Simulation and Retargeting of Complex Multi-Character Interactions

Yunbo Zhang* ybzhang3027@gatech.edu Georgia Institute of Technology USA

Jessica Hodgins jkh@cmu.edu Carnegie Mellon University USA Deepak Gopinath* d_gopinath@apple.com Apple USA

Greg Turk turk@cc.gatech.edu Georgia Institute of Technology USA Yuting Ye yuting.ye@meta.com Reality Labs Research, Meta USA

Jungdam Won[†] jungdam@imo.snu.ac.kr Seoul National University South Korea



Figure 1: A variety of two-character interactions animated using on our method.

ABSTRACT

We present a method for reproducing complex multi-character interactions for physically simulated humanoid characters using deep reinforcement learning. Our method learns control policies for characters that imitate not only individual motions, but also the interactions between characters, while maintaining balance and matching the complexity of reference data. Our approach uses a novel reward formulation based on an interaction graph that measures distances between pairs of interaction landmarks. This reward encourages control policies to efficiently imitate the character's motion while preserving the spatial relationships of the interactions in the reference motion. We evaluate our method on a variety of activities, from simple interactions such as a high-five greeting to more complex interactions such as gymnastic exercises, Salsa dancing, and box carrying and throwing. This approach can be used to "clean-up" existing motion capture data to produce physically plausible interactions or to retarget motion to new characters with different sizes, kinematics or morphologies while maintaining the interactions in the original data.

*This work is conducted while authors were working at Meta †corresponding author

SIGGRAPH '23 Conference Proceedings, August 6–10, 2023, Los Angeles, CA, USA © 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0159-7/23/08.

https://doi.org/10.1145/3588432.3591491

CCS CONCEPTS

• Computing methodologies → Physical simulation; Motion capture; Reinforcement learning; Learning from demonstrations.

KEYWORDS

Character Animation, Interactions, Physics Simulation, Physicsbased Characters, Reinforcement Learning

ACM Reference Format:

Yunbo Zhang, Deepak Gopinath, Yuting Ye, Jessica Hodgins, Greg Turk, and Jungdam Won. 2023. Simulation and Retargeting of Complex Multi-Character Interactions. In Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH '23 Conference Proceedings), August 6–10, 2023, Los Angeles, CA, USA. ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3588432.3591491

1 INTRODUCTION

Physical interactions between people are important elements in daily life. For example, we can imagine a simple scenario where people greet each other by shaking hands, and a more complex scenario where two people perform physical activities together such as Yoga or Salsa dancing. A virtual character capable of reproducing such interactions with either autonomous or user-controlled characters in a believable manner would provide an immersive experience in applications such as computer games, movies, or AR/VR platforms.

In this paper, we are interested in transferring complex multicharacter interactions from reference motions to physically simulated characters. Those characters need to be carefully coordinated in both spatial and temporal domains. Interactions among

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

physically simulated characters have been studied far less than single characters, in part because it is very challenging to learn controllers for multiple characters interacting with each other. As with a single character, balance must be maintained, but the interaction constraints also have to be solved simultaneously. Although some breakthrough results were demonstrated in recent studies [Haworth et al. 2020; Liu et al. 2022; Won et al. 2021], the complexity of the demonstrated interactions are still far from what people routinely perform in daily life.

We demonstrate a novel learning-based method that provides a physics-based retargeting of complex interactions for multiple characters. More specifically, given reference motions that capture interactions between people, we learn control policies (a.k.a. controllers) of simulated characters via deep reinforcement learning that imitate not only the motion of the individuals but also the interactions between them. Our learned policies can produce plausible and semantically equivalent interactions when the sizes and kinematics of the characters are varied significantly. If the size of the simulated characters match those in the original motion capture data, the resulting motion is almost indistinguishable from the reference data and any errors from the capture process are eliminated by ensuring that the interactions are now physically plausible. To solve the challenges in learning multi-character interactions, we develop new rewards based on an interaction graph (IG) which measures distances between pairs of specified locations on the characters, and in particular reflects between-character distances. Rewards based on the IG enable control policies to efficiently deploy complex interactions for physically simulated characters while preserving the semantics of the interactions (i.e. spatial relationship) included in the reference data. In our formulation, manual annotation of the interaction in each motion is not necessary except for choosing a set of general interaction landmarks that work for a variety of scenarios.

To show the effectiveness of our method, we record motions that include multi-person interactions at varying levels of difficulty, and test our method with motions that are composed of simple interactions such as a high-five or other greetings, as well as complex interactions such as gymnastic exercises, Salsa dancing, and box moving/throwing. We demonstrate the generality of our system by reproducing interactions not only for simulated characters with different body dimensions than the motion capture subjects, but also for a robot with a different kinematic structure. Finally, we run comparison and ablation studies that justify each choice in the system design.

2 RELATED WORK

We first review papers synthesizing multi-character interaction for kinematic approaches, which inspired our dynamic formulation. We then review recent progress in control of physically simulated characters via deep reinforcement learning.

2.1 Multi-character Interactions for Kinematic Characters

Most approaches for creating or editing multi-character interactions among kinematic characters are data-driven methods, which means that appropriate motion capture data should be obtained in advance. A popular line of work is based on optimization, where the basic idea is to optimize individual character motions with spatio-temporal constraints [Kwon et al. 2008; Liu et al. 2006], game theory [Shum et al. 2007, 2008a; Shum et al. 2012; Wampler et al. 2010] so that the optimized motions have newly synthesized interactions. These methods are suitable for synthesizing motions having sparse interactions, however, the optimization quickly becomes intractable as the complexity of the interactions increases, so it is not suitable for synthesizing dense and complex interactions. Another approach is patch-based methods, where a patch includes a short interaction of multiple characters [Hyun et al. 2013; Lee et al. 2006; Shum et al. 2008b; Won et al. 2014; Yersin et al. 2009]. For this work, motion capture data where multiple actors are recorded simultaneously is required. New motions can be synthesized by connecting boundaries of multiple patches, thus creating multiple interactions that were not performed together in the original data.

Methods for adapting existing interactions to new environments and characters have also been studied [Al-Asqhar et al. 2013; Ho et al. 2014, 2010; Jin et al. 2018; Kim et al. 2021, 2014, 2009]. The key idea is to define an interaction descriptor that encodes the spatial and temporal relationship, then to edit the motions while minimizing the semantic difference between the original motion and the edited motions where the difference is measured by the descriptor. This idea has also been used to synthesize hands interacting with objects [Zhang et al. 2021]. Our state representation and reward function for deep reinforcement learning are inspired by one of these descriptor-based approaches [Ho et al. 2010], where they construct an interaction graph by connecting edges among pre-specified markers on the body surface. By utilizing deep reinforcement learning and a novel formulation to measure interaction graph similarities, our method can be applied to dynamic characters having different body shapes instead of generating kinematic interaction motions as was done in [Ho et al. 2010].

2.2 Physically Simulated Characters and Interactions

In many cases, what we refer to as interaction between different characters means physical interaction where physical forces occur between the characters at contacts. By incorporating physics simulation into the motion of the characters, those physical interactions can be synthesized in a plausible manner. Multi-character interactions have been created by solving a quadratic programming problem where the equations of motion for the entire dynamical system are used as either hard or soft constraints [Mordatch et al. 2012; Otani and Bouyarmane 2017; Vaillant et al. 2017]. Although cooperative multi-character interactions could be synthesized by these methods without using reference data, the generated motions are typically slow and less-dynamic due to the quasi-static assumption in their optimization formulation, and they require frame-level specification of all interactions in advance. Combining deep reinforcement learning (DRL) and motion capture data has allowed several breakthroughs in learning imitation controllers [Bergamin et al. 2019; Chentanez et al. 2018; Fussell et al. 2021; Park et al. 2019; Peng et al. 2018, 2021; Won et al. 2020], learning reusable motor skills [Merel et al. 2019; Peng et al. 2019, 2022; Won et al. 2022; Yao

et al. 2022], and motion tracking [Winkler et al. 2022; Ye et al. 2022]. Although there also have been some studies synthesizing dynamic interactions with objects or other characters [Haworth et al. 2020; Liu et al. 2022; Merel et al. 2020; Won et al. 2021], the complexity of the demonstrated interactions are still not comparable to what people routinely perform in daily life. In addition, each of these works developed a task-specific reward function to enforce interactions between multiple entities. In this paper, we aim to synthesize various types of spatially and temporally dense interactions for full-body humanoid characters that are physically simulated. This problem is especially challenging because the motor skills must be sophisticated enough to perform those complex interactions while remaining robust enough to maintain balance.

3 METHOD

Our goal is to build controllers that enable physically simulated characters to perform complex physical interactions with each other. For each behavior, we take a reference motion capture clip representing the desired multi-character interaction and produce controllers that enable the simulated characters to mimic those interactions. Our goal is to generate character interactions that are *semantically* similar to those present in the reference motions. To achieve this, we use multi-agent deep reinforcement learning where the states and rewards are designed based on spatial descriptors inspired by [Ho et al. 2010]. Different from [Ho et al. 2010] where only kinematic motions are generated, our method can be applied to dynamic characters having dramatically different body shapes from the captured actors.

3.1 Environment

Our characters are modeled as articulated rigid body objects by following [Won et al. 2020]. Each character has 22 links and 22 joints, where each joint has three degree-of-freedom and is actuated by stable-PD servos [Tan et al. 2011] given target joint angles. We used an open-source framework [Won et al. 2020] to implement and simulate our characters.

3.2 **Problem Formulation**

We formulate the problem as a multi-agent Markov Decision Process (MDP). Consider *k* controllable agents, we define the tuple $\{S, O_1 \cdots O_k, A_1 \cdots A_k, R_1 \cdots R_k, T, \rho\}$ where *S* is the entire state of our environment, O_i and A_i are the observation and action of *i*-th agent, respectively. The reward function $R_i : O_i \times A_i \to \mathbb{R}$ evaluates the quality of the current state and action of *i*-th agent, the environment is updated by the transition function $T : S \times A_1 \times \cdots \times A_k \to S$ given a set of actions performed by all the agents, and $\rho : S \to [0, 1]$ is the probability distribution of the initial states. We aim to learn a set of optimal control policies $\{\pi_i | i = 1 \cdots k\}$ that maximizes average expected return $\mathbb{E}\left[\sum_{t=0}^T \gamma^t r_{i,t}\right]$ for each agent, where $\gamma \in (0, 1)$ is the discount factor that prevents the sum from being infinity.

3.3 Interaction Graph

To better describe the semantics of the interaction happening between agents (or between an agent and an object) during the motion, we define the notion of an Interaction Graph (IG), a graph-based spatial descriptor where the information on interactions is stored in its vertices and edges. This idea is inspired by [Ho et al. 2010]. To construct an interaction graph, we first place a collection of markers on salient locations on each character (see Figure 2). Fifteen markers are placed in total for each character, where three markers are on each limb in the vicinity of joint locations, one on the pelvis, one on the torso, and one on the head. These markers will be considered as the nodes of the graph, each of which is associated with a 6-dimensional vector $n_i = (p_i, v_i) \in \mathbb{R}^6$, where $p_i \in \mathbb{R}^3$ is the position of the vertex and $v_i \in \mathbb{R}^3$ is the velocity of the vertex. For example, a total of 30 vertices will be used for interactions associated with two characters (see Figure 2). On every time step, we perform a Delauney Tetrahedralization over all the vertices based on the spatial distances between pairs of markers to get a compact collection of edges connecting the vertices. Each edge is assigned a feature vector $e_{ij} = (p_{ij}, v_{ij}) \in \mathbb{R}^6$ that encodes the relative relationship between the two vertices, where $p_{ii} = p_i - p_i \in \mathbb{R}^3$ and $v_{ii} = v_i - v_i \in \mathbb{R}^3$ are the positional and velocity components of the edge features.

The example interaction graph in Figure 2 includes both edges connecting nodes on a single character and edges connecting nodes on different characters. The edges within the character help maintain the motion quality of an individual character, while the edges between the characters act as guides for maintaining the relative position of the body parts of the two characters. Details are discussed later in section 3.4.

There is a major difference between how we compare two spatial descriptors in the interaction graph and how they are compared in the Interaction Mesh (IM) in [Ho et al. 2010]. We perform edge-level (i.e. distance) computation whereas IM computes volumetric deformation on a tetrahedron. We further augment the state of an edge with velocities as they are crucial for a physics simulation. Given the input reference motions clips, we build and store such an IG to capture the spatial relationship across the agents and object at each time-step.

3.4 Reward Design

We choose to measure the interaction similarity in two ways: an *edge-weighting function* that highlights the importance of interaction regions in the graph and an *edge-similarity function* that measures the similarity between two IGs with the same connectivity.

For the following similarity measurement, we make use of two interaction graphs G^{sim} and G^{ref} with the same connectivity, one from the simulated environment, the other from the reference motion clips. The connectivity of both graphs is the same as computed on the reference motions using the above mentioned method. The interaction graph we defined is a set of spatial descriptors that encode the relative formation among the vertices in the graph.

3.4.1 Edge Weighting Function. We are guided by the intuition that instances where two body parts are close or in contact with each other are particularly important for multi-character interactions. We define a function that dynamically assigns different weights for each edge according to its relative importance to the others. More specifically, for an edge connecting vertices i and j, the weight for the edge w_{ij} is defined as:

SIGGRAPH '23 Conference Proceedings, August 6–10, 2023, Los Angeles, CA, USA Yunbo Zhang, Deepak Gopinath, Yuting Ye, Jessica Hodgins, Greg Turk, and Jungdam Won

$$w_{ij} = 0.5 * \frac{\exp\left(-k_w \|p_{ij}^{sim}\|\right)}{\sum_{ij} \exp\left(-k_w \|p_{ij}^{sim}\|\right)} + 0.5 * \frac{\exp\left(-k_w \|p_{ij}^{ref}\|\right)}{\sum_{ij} \exp\left(-k_w \|p_{ij}^{ref}\|\right)},$$
(1)

where k_w controls how sensitive the weighting function is with respect to the distance of the edges. The first term gives more attention to an edge if the two nodes in the simulation are close to each other, the second term makes sure an edge in the reference motion gets more attention when its nodes stay close. In practice, we found the second term alone is enough for most of our experiments, and we only use the first term for a few examples where it improves the performance.

Normalizing the weights allows our reward function to adapt to various interaction scenarios. For example, when two characters are far away from each other, the edges connecting those two characters do not contribute much to the reward while the edges connecting vertices within individual characters become important. On the other hand, when the two characters are close to each other, some of the connections between their body parts will have large weights. This adjustment based on proximity allows the body parts that are not associated with the close interactions to remain close to the original motion.



Figure 2: Interaction Graph of the reference characters. Higher opacity on an edge indicates a higher weight for the edge when computing the reward.

3.4.2 Edge Similarity Function. Given two interaction graphs *G* and *G'*, we design a distance function measuring their differences. Our distance function measures position and velocity similarity of the two different formations by comparing all corresponding edges in the two graphs.

3.4.3 Positional Graph Similarity. To compare the positional graph similarity between two graphs, we separately consider the similarity of the two graph edges connecting each individual character E_{self} (self-connections) and between characters E_{cross} (cross-connections). The discrepancy of each edge is computed as follows:

$$err_{self,ij} = \|\frac{p_{ij}^{sim} - p_{T,ij}^{sim}}{\|p_{T,ij}^{sim}\|} - \frac{p_{ij}^{ref} - p_{T,ij}^{ref}}{\|p_{T,ij}^{ref}\|}\|$$
(2)

where $p_{T,ij}^{sim}$ and $p_{T,ij}^{ref}$ are edges computed from the first frame of the motion sequence for both simulation and reference motions. In

all the reference motion sequences, the motion capture actors are instructed to start in a T-pose. In other words, we first compute the deviation of an edge from its corresponding T-pose edge, it is then normalized by the length of the T-pose edge. Finally, we compute the difference between the two deviations, one for the simulated characters and the other for the reference motion clips. Note that this similarity measurement is not sensitive to specific body sizes and proportions due to the normalization of the deviations. This formulation is also similar to measuring the Laplacian coordinate difference of all graph nodes between simulation and reference in that they both try to maintain the similarity of the local structure between two graphs, but our formulation gives direct measure to the edge similarity that captures the interaction.

It is challenging to define a reference edge length for the crossconnections because the variance can be extremely high. For example, imagine that the two characters are standing 10m apart versus 0.1m. Instead, we directly penalize the difference in the edge length and direction so that the same cross-connection similarity can be applied to various motions:

$$err_{cross,ij} = 0.5 * \frac{\|p_{ij}^{sim} - p_{ij}^{ref}\|}{\|p_{ij}^{sim}\|} + 0.5 * \frac{\|p_{ij}^{sim} - p_{ij}^{ref}\|}{\|p_{ij}^{ref}\|}$$
(3)

where we normalize the difference by the lengths in the simulation and the reference clips, respectively, then average them so that the similarity becomes symmetric. This symmetry also enables the error to be used for characters having different body shapes. The total error for positional graph similarity is then the sum of the two error terms from all edges:

$$err_{pos_graph} = \sum_{ij \in E_{cross}} w_{ij} err_{cross,ij} + \sum_{ij \in E_{self}} w_{ij} err_{self,ij} \quad (4)$$

3.4.4 Velocity Graph Similarity. To measure the velocity discrepancy between graphs, we simply measure the difference of velocity of all edges in simulation and reference as:

$$err_{vel_graph} = \sum_{ij \in E_{cross} \cup E_{self}} w_{ij} \|v_{ij}^{sim} - v_{ij}^{ref}\|$$
(5)

In contrast to the positional similarities, we observed that the velocities of the graph vertices do not vary much when modifying the body size/proportion of the simulated characters. Thus we do not perform any velocity normalization. We also do not separate the velocity similarities by edge type because we did not find any benefit in doing so.

3.4.5 *Final Reward Design.* We define our reward function based on the errors computed from the interaction graphs. In addition, we add two more error terms measuring the tracking of the root joint and center-of-mass, which are frequently used in learning imitation controllers for physically simulated characters. As a result, our reward function is composed of four terms

$$r = r_{pos_graph} \cdot r_{vel_graph} \cdot r_{root} \cdot r_{com} \tag{6}$$

Simulation and Retargeting of Complex Multi-Character Interactions

SIGGRAPH '23 Conference Proceedings, August 6-10, 2023, Los Angeles, CA, USA

$$r_{pos_graph} = \exp(-k_1 * err_{pos_graph})$$

$$r_{vel_graph} = \exp(-k_2 * err_{vel_graph})$$

$$r_{root} = \exp(-k_3 * err_{root})$$

$$r_{com} = \exp(-k_4 * err_{com})$$
(7)

where r_{pos_graph} , r_{vel_graph} measure the difference between the two interaction graphs and r_{root} , r_{com} encourage the tracking of the root joint and the center-of-mass projected on the ground, k_1, \dots, k_4 are the sensitivities of the terms, respectively. The errors for the tracking are defined as follows:

$$err_{root} = w_p \|\bar{p}_{sim} - \bar{p}_{ref}\|^2 + w_q \|log(q_{sim}^{-1} \cdot q_{ref})\|^2 + w_v \|\bar{v}_{sim} - \bar{v}_{ref}\|^2 + w_\omega \|\omega_{sim} - \omega_{ref}\|^2$$
(8)

$$err_{com} = w_{com,x} ||x_{sim} - x_{ref}|| + w_{com,\dot{x}} ||\bar{x}_{sim} - \dot{x}_{ref}||$$
 (9)

where \bar{p}, \bar{v} are the position and velocity of the root joint excluding the height components, and q,ω are the orientation and angular velocity of the root joint, respectively, x and \dot{x} are the center-of-mass position and velocity of the simulated character excluding their height components. $w_p, w_q, w_v, w_w, w_{com,x}$, and $w_{com,\dot{x}}$ are the relative weights of the terms. Note that we ignore the height components of the linear positions and velocities so that the relevant errors are not directly affected by the absolute size of the character. In contrast to [Ho et al. 2010], where tetrahedron volumes are used to measure similarities of meshes, our edges-based reward is more sensitive to point-to-point physical interactions. In addition, it is not trivial to design an adaptive weight function in a volume-based setting, which makes sure that individual characters motion quality is preserved, even when characters are staying far apart, making our reward a good substitute for motion imitation.

3.5 Observation and Action Spaces

The observation space of our environment is inspired by the design from prior work [Won et al. 2020, 2021] where the observation of an agent $o_i = (o_{sim}, o_{ref})$ consists of the states of the simulated characters and objects, which are computed from the simulation and the reference motion clips. For simulated observation space $o_{sim} = (o_{sim,self}, o_{sim,other}, o_{sim,object})$, we include the position, orientation, linear and angular velocity for each link of the characters and the objects. To make sure the state is invariant to the global position and orientation of the agent, all values are transformed to the facing frame of the controlled character. The facing frame of the character is computed by projecting the global transformation of the character root to the ground.

The reference observation $o_{ref} = (o_{ref}^0, o_{ref}^{0.05}, o_{ref}^{0.15})$ contains the reference information 0, 0.05, and 0.15 seconds in the future to the current time. For each future reference observation frame $o_{ref}^* = (o_{ref,self}^*, o_{ref,other}^*, o_{ref,object}^*)$, we include the position, orientation, linear and angular velocity for each link of the characters and the objects in the facing frame of the reference character.

Our action *a* is the change of pose Δq from the pose q_{ref} given the reference frame at each time-step. A new reference pose $q_{ref} + \Delta q$ (i.e. a set of joint angles) is given to the stable PD servos attached to our simulated character and then joint torques are computed accordingly.

4 RESULTS

In this section, we show that our formulation is robust enough to be applied to a variety of motions with multiple characters and objects. The dynamically adjusting weights of our method focus the adjustments to the motion on the physical interactions. This approach results in higher quality motion than existing work in scenarios with complex interactions. Further, our formulation is able to preserve interaction when the body sizes, kinematics, and skeletons of the characters differ from the reference motion sequences.

4.1 Experiment Setup

The structure of our policy follows a encoder-decoder style as presented in [Won et al. 2021], where the encoder is a fully connected neural network with two hidden layers with 256 and 128 units respectively. The encoder takes the full observation and projects it onto a 32 dimensional latent vector z. The decoder is another fully connected network with two hidden layers with 256 units, and it takes as input the concatenated vector $z_{decoder} = (o_{sim,self}, z)$ and outputs the action of the policy. To speed up the learning for all the experiments below, we pre-train an imitation policy of a single character on sequences that can be done alone (e.g. high five, greetings, and push ups). When training an interaction-graph based policy, we reuse the pre-trained decoder and allow its weights to be updated during the training. Note that the encoder trained simultaneously with the pre-trained decoder is not reusable due to differences in input dimensions. The decoder is reusable because the latent dimensions are unchanged. This design makes it easier for the policy to maintain balance at the initial phase of learning, and hence results in faster training. The training time of a policy varies based on the difficulty of the sequence. For easier sequences, it takes about 300 million to 500 million samples to train one policy. For harder sequences, it could take more than 2 billion samples to train a policy. All experiments are run using 640 CPUs, and it could take from 3 days to 9 days to train a policy based on the sequence difficulty. We demonstrate the success of our method over several different scenarios.

4.2 Human-Human Interaction

In the human-human interaction scenarios, we aim to show that our method is capable of reproducing imitation policies with similar motion quality as existing works such as [Fussell et al. 2021; Won et al. 2020; ?] while the interaction is better preserved. We show a variety of scenarios ranging from sparsely interacting motions to continuously interacting motions between the two human characters.

4.2.1 Light Interaction. Figure 3a and 3b shows light physical interactions. In *Rapper-Style Greetings*, the two characters touch their hands, elbows, shoulders, and legs in sequence to greet each other, an action which has been shown in many hip-hop music videos [sinestesia3000 2012]. In *Jumpover*, one character jumps over the other character. In these scenarios, physical interactions are of short duration with little physical forces, and the interactions are well-preserved semantically when the interacting body parts are close enough at the right timing.

4.2.2 Heavy Interaction. Figure 9a, 9b, and 9c shows physical interactions where significant forces occur between the two characters. The *Lift-Pushup* example includes interactions where one character needs to lift the other character's legs while that character is performing a push-up exercise. In the first salsa dancing motion (Salsa Grasping), two character's hands are grasped together to form a loop for one character to go under. In another salsa dancing motion (Salsa Support), one character needs to support the other while they lean backward. This type of interaction is more challenging than the light interactions because the two simulated characters need to perform highly coordinated motions with force exchange. For example, the character performing a push-up would not be able to imitate the reference motions successfully unless his legs are grasped by the other character. Furthermore, these heavy interactions make maintaining balance difficult because significant forces are transferred between the two characters. Our method was able to imitate these challenging motions successfully as shown in Figure 9a. Because our characters do not have fingers, we mimic the grasp by adding weld constraints in the physics simulation. More specifically, we examine the reference motion and label a sequence of grasping windows to identify when grasping should be presented on specified body pairs at certain moment. During simulation, when a character's hand is close to a body part it should grasp at that moment, weld constraints between the two body parts are made temporarily. The constraints is removed when the grasping window is over, representing the character releasing their hands. Other than hand grasping, we did not use any additional weld constraints, all the complex physical interactions emerged during the learning process.

These results show that our formulation is capable for the control policies to be aware of the relative formation among various body parts that are close enough in the reference motions regardless of the interacting duration, and those relative formations will be preserved in the simulation.

4.3 Human-Object Interaction

We further demonstrate that our formulation can also handle humanobject interactions where objects are passively simulated. Figure 6a and 6b show the two motions: One includes interactions where two persons are throwing and catching a small box repeatedly, the other includes interactions where two persons are lifting and moving a large box. For both motion sequences, we place an extra marker on every vertex of the box (i.e. 8 markers in total) when constructing the interaction graph. For the edges connecting between characters and the box, we use the same reward formulation following Equation 3 to measure the discrepancy between simulation and reference. In addition, we choose to remove all edges connecting the markers on the box because their relative distances will stay constant throughout the motion. The resulting graph is shown in Figure 12. The control policies learned with those additional markers successfully reproduce both hand-object interaction scenarios, which shows the generality of our method.

4.4 Retageting to different body sizes

Our graph-based formulation is robust to the change in body conditions because we compute features in a normalized manner. As a result, the interactions in the same reference motions can also be applied to simulated characters that have completely different body dimensions from the actors in the reference motions.

Figure 3c, 3d demonstrates motion that include light interaction. In both sequences, we scale all limbs of the yellow and blue characters by a factor of 1.3 and 0.5, respectively, so the yellow character is almost 2 times taller than the blue character. The scaled characters are trained using our framework to track the reference motion. In the *Rapper-style Greeting* motion, for example, we see the taller character deliberately bends down to reach their hand, elbow, and shoulder to the shorter character when the interaction happens, and they straighten their back after they finish the interaction. Similarly, the taller character lowers their waist when the shorter character jumps over their back in the *Jumpover* motion.

Learning how to transfer forces via physical interactions is crucial to imitate motions including heavy interaction as in Figure 9d,9e,and 9f. For the *Lift-Pushup* motion (Figure 9d), we give a 0.5 scaling factor to all limbs of the blue character, for the *Salsa Grasping* motion (Figure 9e), we scale the yellow character's limbs by 0.5, and for *Salsa Support* motion (9f), we scale the yellow character's limbs by 0.8. For this type of motion, our method enables the scaled characters to adjust their motions to preserve the interactions rather than simply minicking the original reference motions, which means that the semantics of the interaction are transferred successfully to the scaled characters. For example, the taller characters in the *Lift-Pushup* motions learned to bend down and reach the target grasping region to form the grasp.

Finally, we also scale the characters and objects for human-object interaction scenarios. Figure 6 shows the control policies learned successfully for the small box throwing-catching and the large box lifting-moving. For both human-object interaction motions, we scale the yellow character's limbs by 0.7.

4.5 Non-human Characters

Our method can also transfer interactions in the reference motions to characters with different kinematic configurations. For example, if we use robots with fewer DoFs than the reference character, our method can still make the robots mimic the interactions existing in the reference motions. As shown in Figure 4, we replace one of the characters by a Baxter robot composed of two industrial manipulators.

Because the robot has a fixed base, we place the robot at the location where the greeting motion is conducted and place a total of eight markers on the upper body of the robot on its head, torso, upper arms, lower arms, and end-effectors to match those of the human character. For the human character, we keep the same 15 markers as described earlier on the human body. We then use the total of 23 markers to construct the interaction graph for the training. During the training, we use two separate reward functions for the character and robot independently. The character will receive the same reward terms as described above, the robot will only receive a reward from $r_{pos\ graph}$ and $r_{vel\ graph}$ because it is not mobile.

In addition, we found that including the first term in Equation 1 become helpful when the robot is immobile. It highlights the edge error when the robot body parts are staying close but the reference characters' body are far away. Because the kinematic structure of the robot is completely different that of the actor in the reference character/motion, we ask the policy to directly output the absolute target joint angles q instead of learning the deviation (i.e. Δq) from the reference q_{ref} for both the human character and the robot. Our framework can successfully generate animations of the Baxter robot performing greetings with a human character (Figure 5a), and perform a highfive with another Baxter robot (Figure 5b). These examples demonstrates the potential of our method as an option to reproduce human-robot interactions.

4.6 Comparison

We conduct comparison and ablation studies to show the effectiveness of our graph-based formulation to reproduce complex interaction for physically simulated characters.

4.6.1 Joint-based Reward. To highlight the necessity of formulating the interaction-graph-based reward, we compare our method with the commonly used joint-based reward formulation for motion imitation. For the sequences that use a joint-based reward, we apply a similar formulation as described in [Peng et al. 2018; Won et al. 2020]. That formulation asks the policy to minimize the positional and angular differences of the joints and links between the simulation and the reference motion. In this formulation, no reward term exists to evaluate the quality of interactions between multiple characters or characters to objects. As a result, when the simulated character has a different body configuration from the reference motion, the characters will only learn to mimic the poses in the reference motion instead of learning to adapt to the other characters (or object) to correctly perform the interaction. Figure 7 shows a comparison for the greeting motions. The control policies trained using a joint-based reward fail to cause the taller character to bend down to meet the shorter character. Similar behaviors are observed in the other motions for the control policies trained using the joint-based reward only.

We further contrast the performance for the dense interaction example between the interaction graph and joint-based rewards. Figure 8 shows such a comparison on *Lift-Pushup* sequence with a scaled character. When using the interaction graph reward, the taller character actively bends forward to reach its hands to the shorter character's lower leg to form the grasping constraints and lift the shorter character. When using a joint-based reward, on the other hand, there is no reward based on the relative poses between the two characters and the taller character cannot grasp the shorter character's leg and the interaction semantics is not preserved.

Furthermore, we show that a joint-based reward also produces lower quality motions when re-targeting motions for human-object interactions. Figure 10 shows a comparison for a small box throw and catch motion trained with the interaction graph reward and Joint-based reward. The two characters are able to perform the throw and catch motion sequence in the joint-based reward because of the presence of the additional object observation and reward as described above. However, it fails to preserve the interaction semantics because the shorter character should catch the box by holding on two opposite faces of the box instead of supporting the box on its bottom. 4.6.2 Edge Weighting Function. We do an ablation on the edge weighting function (Equation 1) to understand how this helps the training selectively pay attention to more important edges and ignore irrelevant edges. Our experiments demonstrate that this design can help generating more natural-looking motions. In Figure 11, we compare the resulting policy trained with (left) and without (right) the weighting function for the greeting motion. When the edge weighting function is present, the taller character learns to bend its waist to reduce its height when greeting the shorter character. However, when all the edges have the same weight during training, the taller character instead learns to walk and complete all the greetings with the legs bent at an unnatural angle. This unnatural behavior is created because the policy tries to get a low error on every edge of the graph regardless of the distances of the nodes.

5 DISCUSSION

We demonstrated a method of simulating and retargeting of complex multi-Character interactions by using deep reinforcement learning where novel state and rewards that are character-agnostic are developed based on an *Interaction Graph*. Our formulation is applicable to a variety of interactions among people ranging from sparse interactions (e.g. greeting, jumpover) to complex ones (e.g. exercise motion, Salsa dancing) regardless of whether body sizes, kinematics or skeletons of the simulated characters are the same as that of the actors who recorded the reference motions. Those interactions even also able to be transferred to a robot to generate human-robot interactions.

While we demonstrate many successful examples, there are some limitations to our method. First, there are some limitations of our reward function design. Because the action space of our policy is not directly associated with the reward function, our training usually requires more samples to converge compared to the joint-based reward function. In addition, due to the lack of supervision on the joint angles, the motion generated from our policy could contain artifacts on joints that have little impact to the interaction. For example, sometimes the character may tilt the head or the waist at an unnatural angle because this deviation from the reference will not affect the positions of the interaction graph's node, and therefore it does not decrease the reward. Adding more markers would be an immediate remedy but this would also increase computational cost. Another limitation is that our controllers are still imitation controllers which cannot perform interactions that do not exist in the reference motions. Further the controllers only work for the specific body configuration that it was trained on, so one policy cannot easily be generalized to work on a character with a different body configuration. We also observe that the variability of our result is limited by the dissimilarity of the character and the difficulty of the task. The characters scaled extremely or with drastically different body plans could fail to imitate the interactions due to their physical limits. For example, in challenging interaction scenarios such as box throwing compared to the greeting motions, our method could only work with slight scaling on the characters but will fail when replacing one human character with a robot to perform the task with another human character.

We envision several future directions to improve the limitations. For better variability, we can build a stronger motion prior that SIGGRAPH '23 Conference Proceedings, August 6-10, 2023, Los Angeles, CA, USA Yunbo Zhang, Deepak Gopinath, Yuting Ye, Jessica Hodgins, Greg Turk, and Jungdam Won

contains a larger variation of motion types. Further training on top of the motion prior could be more sample efficient, and allow the policy to explore better in the motion space to find a valid solution when trained character shapes undergo extreme scalings. To improve the generalization of our method, better observation space would be helpful. Currently we are still using the commonly used joint-based observation space for the policy to account for the environmental context. A better formulation of the observation space that directly interprets the interaction graph (instead of joint angles) would be a first step towards a universal policy. A good candidate would be utilizing a Graph Neural Network to directly observe the interaction graph in the simulation to better understand the spatial relationship of different objects in the scene. After the creation of a comprehensive model that can better understand the interaction observations, a subsequent future direction would be building a interaction prior model that can sense the surrounding environment and create motions to complete general interactions without requiring specific reference motions as demonstrated in [Peng et al. 2022; Won et al. 2022] for a single character motion. We believe that our method has made it possible to generate complex multi-character interactions of simulated characters for many behaviors and will serve as a stepping stone for future research.

ACKNOWLEDGMENTS

Jungdam Won was partially supported by the New Faculty Startup Fund from Seoul National University, ICT(Institute of Computer Technology) at Seoul National University, and the MSIT (Ministry of Science and ICT), Korea, under the ITRC (Information Technology Research Center) support program (IITP-2023-2020-0-01460) supervised by the IITP(Institute for Information Communications Technology Planning Evaluation).

REFERENCES

- Rami Ali Al-Asqhar, Taku Komura, and Myung Geol Choi. 2013. Relationship Descriptors for Interactive Motion Adaptation. In Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '13). 45–53. https://doi.org/10.1145/2485895.2485905
- Kevin Bergamin, Simon Clavet, Daniel Holden, and James Richard Forbes. 2019. DReCon: Data-driven Responsive Control of Physics-based Characters. ACM Trans. Graph. 38, 6, Article 206 (2019). http://doi.acm.org/10.1145/3355089.3356536
- Nuttapong Chentanez, Matthias Müller, Miles Macklin, Viktor Makoviychuk, and Stefan Jeschke. 2018. Physics-based motion capture imitation with deep reinforcement learning. In Motion, Interaction and Games, MIG 2018. ACM, 1:1–1:10. https://doi.org/10.1145/3274247.3274506
- Levi Fussell, Kevin Bergamin, and Daniel Holden. 2021. SuperTrack: Motion Tracking for Physically Simulated Characters using Supervised Learning. ACM Trans. Graph. 40, 6, Article 197 (2021). https://dl.acm.org/doi/10.1145/3478513.3480527
- Brandon Haworth, Glen Berseth, Seonghyeon Moon, Petros Faloutsos, and Mubbasir Kapadia. 2020. Deep Integration of Physical Humanoid Control and Crowd Navigation. In *Motion, Interaction and Games (MIG '20)*. Article 15. https://doi.org/10. 1145/3424636.3426894
- Edmond SL Ho, He Wang, and Taku Komura. 2014. A multi-resolution approach for adapting close character interaction. In *Proceedings of the 20th ACM symposium on virtual reality software and technology*. 97–106.
- Edmond S. L. Ho, Taku Komura, and Chiew-Lan Tai. 2010. Spatial Relationship Preserving Character Motion Adaptation. ACM Trans. Graph. 29, 4, Article 33 (2010). https://doi.org/10.1145/1778765.1778770
- K. Hyun, M. Kim, Y. Hwang, and J. Lee. 2013. Tiling Motion Patches. IEEE Transactions on Visualization and Computer Graphics 19, 11 (2013), 1923–1934. https://doi.org/ 10.1109/TVCG.2013.80
- Taeil Jin, Meekyoung Kim, and Sung-Hee Lee. 2018. Aura Mesh: Motion Retargeting to Preserve the Spatial Relationships between Skinned Characters. *Computer Graphics Forum* 37, 2 (2018), 311–320. https://doi.org/10.1111/cgf.13363
- Jongmin Kim, Yeongho Seol, and Taesoo Kwon. 2021. Interactive multi-character motion retargeting. Computer Animation and Virtual Worlds 32, 3-4 (2021), e2015.

- Jongmin Kim, Yeongho Seol, Taesoo Kwon, and Jehee Lee. 2014. Interactive Manipulation of Large-Scale Crowd Animation. ACM Trans. Graph. 33, 4, Article 83 (2014). https://doi.org/10.1145/2601097.2601170
- Manmyung Kim, Kyunglyul Hyun, Jongmin Kim, and Jehee Lee. 2009. Synchronized Multi-Character Motion Editing. ACM Trans. Graph. 28, 3, Article 79 (2009). https: //doi.org/10.1145/1531326.1531385
- T. Kwon, Y. Cho, S. I. Park, and S. Y. Shin. 2008. Two-Character Motion Analysis and Synthesis. *IEEE Transactions on Visualization and Computer Graphics* 14, 3 (2008), 707–720. https://doi.org/10.1109/TVCG.2008.22
- Kang Hoon Lee, Myung Geol Choi, and Jehee Lee. 2006. Motion Patches: Building Blocks for Virtual Environments Annotated with Motion Data. ACM Trans. Graph. 25, 3 (2006), 898–906. https://doi.org/10.1145/1141911.1141972
- C. Karen Liu, Aaron Hertzmann, and Zoran Popović. 2006. Composition of Complex Optimal Multi-Character Motions. In Proceedings of the 2006 ACM SIG-GRAPH/Eurographics Symposium on Computer Animation (SCA '06). 215–222.
- Siqi Liu, Guy Lever, Zhe Wang, Josh Merel, S. M. Ali Eslami, Daniel Hennes, Wojciech M. Czarnecki, Yuval Tassa, Shayegan Omidshafiei, Abbas Abdolmaleki, Noah Y. Siegel, Leonard Hasenclever, Luke Marris, Saran Tunyasuvunakool, H. Francis Song, Markus Wulfmeier, Paul Muller, Tuomas Haarnoja, Brendan Tracey, Karl Tuyls, Thore Graepel, and Nicolas Heess. 2022. From motor control to team play in simulated humanoid football. *Science Robotics* 7, 69 (2022), eabo0235. https://doi.org/10.1126/scirobotics.abo0235
- Josh Merel, Leonard Hasenclever, Alexandre Galashov, Arun Ahuja, Vu Pham, Greg Wayne, Yee Whye Teh, and Nicolas Heess. 2019. Neural Probabilistic Motor Primitives for Humanoid Control. In 7th International Conference on Learning Representations, ICLR 2019. https://openreview.net/forum?id=BJIGTjRcY7
- Josh Merel, Saran Tunyasuvunakool, Arun Ahuja, Yuval Tassa, Leonard Hasenclever, Vu Pham, Tom Erez, Greg Wayne, and Nicolas Heess. 2020. Catch Carry: Reusable Neural Controllers for Vision-Guided Whole-Body Tasks. ACM Trans. Graph. 39, 4, Article 39 (2020). https://doi.org/10.1145/3386569.3392474
- Igor Mordatch, Emanuel Todorov, and Zoran Popović. 2012. Discovery of Complex Behaviors through Contact-Invariant Optimization. ACM Trans. Graph. 31, 4, Article 43 (jul 2012), 8 pages. https://doi.org/10.1145/2185520.2185539
- Kazuya Otani and Karim Bouyarmane. 2017. Adaptive whole-body manipulation in human-to-humanoid multi-contact motion retargeting. In 2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids). 446–453. https://doi. org/10.1109/HUMANOIDS.2017.8246911
- Soohwan Park, Hoseok Ryu, Seyoung Lee, Sunmin Lee, and Jehee Lee. 2019. Learning Predict-and-simulate Policies from Unorganized Human Motion Data. ACM Trans. Graph. 38, 6, Article 205 (2019). http://doi.acm.org/10.1145/3355089.3356501
- Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. 2018. Deep-Mimic: Example-guided Deep Reinforcement Learning of Physics-based Character Skills. ACM Trans. Graph. 37, 4, Article 143 (2018). http://doi.acm.org/10.1145/ 3197517.3201311
- Xue Bin Peng, Michael Chang, Grace Zhang, Pieter Abbeel, and Sergey Levine. 2019. MCP: Learning Composable Hierarchical Control with Multiplicative Compositional Policies. In Advances in Neural Information Processing Systems 32. 3681–3692.
- Xue Bin Peng, Yunrong Guo, Lina Halper, Sergey Levine, and Sanja Fidler. 2022. ASE: Large-scale Reusable Adversarial Skill Embeddings for Physically Simulated Characters. ACM Trans. Graph. 41, 4, Article 94 (July 2022).
- Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. 2021. AMP: Adversarial Motion Priors for Stylized Physics-Based Character Control. *ACM Trans. Graph.* 40, 4, Article 1 (July 2021), 15 pages. https://doi.org/10.1145/3450626.3459670
- Hubert P. H. Shum, Taku Komura, Masashi Shiraishi, and Shuntaro Yamazaki. 2008b. Interaction Patches for Multi-Character Animation. ACM Trans. Graph. 27, 5 (2008). https://doi.org/10.1145/1409060.1409067
- Hubert P. H. Shum, Taku Komura, and Shuntaro Yamazaki. 2007. Simulating Competitive Interactions Using Singly Captured Motions. In Proceedings of the 2007 ACM Symposium on Virtual Reality Software and Technology (VRST '07). 65–72. https://doi.org/10.1145/1315184.1315194
- Hubert P. H. Shum, Taku Komura, and Shuntaro Yamazaki. 2008a. Simulating Interactions of Avatars in High Dimensional State Space. In Proceedings of the 2008 Symposium on Interactive 3D Graphics and Games (I3D '08). 131–138. https: //doi.org/10.1145/1342250.1342271
- H. P. H. Shum, T. Komura, and S. Yamazaki. 2012. Simulating Multiple Character Interactions with Collaborative and Adversarial Goals. *IEEE Transactions on Vi*sualization and Computer Graphics 18, 5 (2012), 741–752. https://doi.org/10.1109/ TVCG.2010.257
- sinestesia3000. 2012. Hip-hop Greetings. Youtube. https://www.youtube.com/watch? v=VnyvqSgiZz0
- Jie Tan, Karen Liu, and Greg Turk. 2011. Stable Proportional-Derivative Controllers. IEEE Computer Graphics and Applications 31, 4 (2011), 34–44. https://doi.org/10. 1109/MCG.2011.30
- Joris Vaillant, Karim Bouyarmane, and Abderrahmane Kheddar. 2017. Multi-Character Physical and Behavioral Interactions Controller. *IEEE Transactions on Visualization* and Computer Graphics 23, 6 (2017), 1650–1662. https://doi.org/10.1109/TVCG.2016. 2542067

Simulation and Retargeting of Complex Multi-Character Interactions

- Kevin Wampler, Erik Andersen, Evan Herbst, Yongjoon Lee, and Zoran Popović. 2010. Character Animation in Two-Player Adversarial Games. ACM Trans. Graph. 29, 3, Article 26 (2010). https://doi.org/10.1145/1805964.1805970
- Alexander Winkler, Jungdam Won, and Yuting Ye. 2022. QuestSim: Human Motion Tracking from Sparse Sensors with Simulated Avatars. In SIGGRAPH Asia 2022 Conference Papers. 1–8.
- Jungdam Won, Deepak Gopinath, and Jessica Hodgins. 2020. A Scalable Approach to Control Diverse Behaviors for Physically Simulated Characters. ACM Trans. Graph. 39, 4, Article 33 (2020). https://doi.org/10.1145/3386569.3392381
- Jungdam Won, Deepak Gopinath, and Jessica Hodgins. 2021. Control Strategies for Physically Simulated Characters Performing Two-Player Competitive Sports. ACM Trans. Graph. 40, 4, Article 146 (2021). https://doi.org/10.1145/3450626.3459761
- Jungdam Won, Deepak Gopinath, and Jessica Hodgins. 2022. Physics-Based Character Controllers Using Conditional VAEs. ACM Trans. Graph. 41, 4 (2022). https: //doi.org/10.1145/3528223.3530067
- Jungdam Won, Kyungho Lee, Carol O'Sullivan, Jessica K. Hodgins, and Jehee Lee. 2014. Generating and Ranking Diverse Multi-Character Interactions. ACM Trans. Graph. 33, 6, Article 219 (2014). https://doi.org/10.1145/2661229.2661271
- Heyuan Yao, Zhenhua Song, Baoquan Chen, and Libin Liu. 2022. ControlVAE: Model-Based Learning of Generative Controllers for Physics-Based Characters. ACM Transactions on Graphics (TOG) 41, 6 (2022), 1–16.
- Yongjing Ye, Libin Liu, Lei Hu, and Shihong Xia. 2022. Neural3Points: Learning to Generate Physically Realistic Full-body Motion for Virtual Reality Users. arXiv preprint arXiv:2209.05753 (2022).
- Barbara Yersin, Jonathan Maïm, Julien Pettré, and Daniel Thalmann. 2009. Crowd Patches: Populating Large-Scale Virtual Environments for Real-Time Applications. In Proceedings of the 2009 Symposium on Interactive 3D Graphics and Games (I3D '09). 207-214. https://doi.org/10.1145/1507149.1507184
- He Zhang, Yuting Ye, Takaaki Shiratori, and Taku Komura. 2021. ManipNet: Neural Manipulation Synthesis with a Hand-Object Spatial Representation. ACM Trans. Graph. 40, 4, Article 121 (2021), 14 pages. https://doi.org/10.1145/3450626.3459830

SIGGRAPH '23 Conference Proceedings, August 6-10, 2023, Los Angeles, CA, USA Vunbo Zhang, Deepak Gopinath, Yuting Ye, Jessica Hodgins, Greg Turk, and Jungdam Won



(a) Rapper-Style Greeting



(c) Scaled Rapper-Style Greeting

Figure 3: Light Interaction

(b) Jumpover

(d) Scaled Jumpover



(a) Box Throw and Catch





(c) Scaled Box Throw and Catch

(d) Scaled Box Carry

Figure 6: Human-Object Interaction



Figure 4: Baxter robot with markers





(a) Joint-based Reward

(b) Interaction Graph Reward

Figure 7: Reward comparison for Greetings.





(a) Human-Baxter Greeting

(b) Baxter-Baxter Highfive

Figure 5: Robot Interaction





(a) Joint-based Reward

(b) Interaction Graph Reward

Figure 8: Reward comparison for Lift Pushup.

Simulation and Retargeting of Complex Multi-Character Interactions

SIGGRAPH '23 Conference Proceedings, August 6-10, 2023, Los Angeles, CA, USA



(a) Lift Pushup



(d) Scaled Lift Pushup



(b) Salsa Grasping



(e) Scaled Salsa Grasping

Figure 9: Heavy Interaction



(c) Salsa Support



(f) Scaled Salsa Support



(a) Joint-based Reward



(b) Interaction Graph Reward

Figure 10: Box Throw comparison



Figure 12: An Interaction Graph that includes an object (the top box) as well as the characters.





(a) Without Edge Weights

(b) With Edge Weights

Figure 11: Comparison showing the effect of edge weights