

Two-Person Interaction Augmentation with Skeleton Priors

Baiyi Li*

Edmond S. L. Ho[†]

Hubert P. H. Shum[‡]

He Wang[§]

Abstract

Close and continuous interaction with rich contacts is a crucial aspect of human activities (e.g. hugging, dancing) and of interest in many domains like activity recognition, motion prediction, character animation, etc. However, acquiring such skeletal motion is challenging. While direct motion capture is expensive and slow, motion editing/generation is also non-trivial, as complex contact patterns with topological and geometric constraints have to be retained. To this end, we propose a new deep learning method for two-body skeletal interaction motion augmentation, which can generate variations of contact-rich interactions with varying body sizes and proportions while retaining the key geometric/topological relations between two bodies. Our system can learn effectively from a relatively small amount of data and generalize to drastically different skeleton sizes. Through exhaustive evaluation and comparison, we show it can generate high-quality motions, has strong generalizability and outperforms traditional optimization-based methods and alternative deep learning solutions.

1. Introduction

Skeletal motion is a crucial data modality in many applications, such as human activity recognition, motion analysis, security and computer graphics [9, 35, 48, 56, 57, 59]. However, capturing high-quality skeletal motions often requires expensive hardware, professional actors, costly post-processing and laborious trial-and-error processes [40]. Affordable devices such as RGB-D cameras can reduce the cost but usually provide data with jittering and tracking errors [44]. As a result, the majority of available skeletal data is based on single-person [37] or multiple people with short, simple and almost-no-contact interactions [44]. Datasets with close and continuous interactions [15] are rare, limit-

ing the research of motion generation [60], prediction [15], classification [48] within such motions.

One way to tackle the challenge is to carefully capture the motion of actors and retarget it onto different skeletons [16]. With a single skeleton, the problem can be formulated as optimizations with respect to keeping key geometric and dynamic constraints [3, 49]. However, this process quickly becomes intractable with the increase of constraints such as foot contact and hand-environment contact, let alone retargeting two people with close and continuous interactions like wrestling and dancing, where inter-character geometric/topological constraints need to be retained [17, 36]. Consequently, multiple runs of complex optimization with careful hand-tuning of objective function weights are needed [18, 19] for a single motion, which is prohibitively slow and therefore can only be used to generate small amounts of data.

Meanwhile, data-driven approaches for single body retargeting [4], despite being successful, cannot be directly extended to two-character interaction. Methodologically, these methods do not model inter-character geometric constraints, which is key to the semantics of interactions [19]. From the data point of view, these approaches, especially those using deep learning [2, 54], require a large amount of data, which is largely absent for two-character interaction. Existing two-character interaction datasets are for action recognition [8, 43] and low-quality, or only consist of a small amount of data with limited variations in body sizes [15], hardly covering the distribution of possible body variations. Considering the high cost of obtaining interaction data, a method that can learn effectively from limited data and generate interactions with diversified body variations is highly desirable.

We propose a novel lightweight framework for two-character skeletal interaction augmentation, easing the need to capture a large amount of data. Our key insight is the joint relations evolving in time (e.g. relative positions, velocities, etc.) can fully describe an interaction, e.g. hugging always involves wrapping one’s arms around the other’s body. These relations change when the body size changes, but the *distribution* of them should stay similar in the sense that one’s arms should still wrap around the other, such

*University of Leeds, United Kingdom

[†]University of Glasgow, United Kingdom

[‡]Durham University, United Kingdom

[§]corresponding author, he_wang@ucl.ac.uk, University College London, United Kingdom

that the hand-to-body distance is always smaller than e.g. the foot-to-body distance. Meanwhile, this distribution should be very different from other types of interactions e.g. wrestling. Therefore, to generate motions from different skeleton sizes, the key is being able to predict the joint relation distributions based on a given skeleton.

To this end, we propose a conditional motion generation approach, where the generated motions are conditioned on the joint relation distribution which is further conditioned on a skeleton prior, allowing a skeleton change to propagate through the joint relation distribution and finally influence the final motion. We start by modeling the joint probability of two-body motions and proposing a novel factorization to decompose it into three distributions. The three distributions are realized as neural networks, which together form an end-to-end model that conditions two-body motions on one person’s body size. Further, to address the data scarcity challenge, we capture new two-body data and employ an existing optimization-based method for initial data augmentation. After training our model on the data, it can be employed for further motion data augmentation for many downstream tasks.

We evaluate our method in multiple tasks. Since there is no similar method for baselines, we compare our method with adapted baselines and optimization-based approaches, demonstrating that our method is accurate in generating desired motions, can generate diversified interactions while respecting interaction constraints, is much faster for inference and generalizes to large skeletal changes than optimization-based methods. In addition, our model benefits downstream tasks including motion prediction and activity recognition. Formally, our contributions include:

1. A new factorization of two-character interactions that allows for effective modelling of interaction features.
2. a new deep learning method for interaction retargeting/generation to the best of our knowledge, which learns and generalizes effectively from a small number of training samples.
3. A new dataset augmented from single interaction examples, containing interactions with different body sizes and proportions.

2. Related Work

2.1. Deep Learning for Skeletal Motion

Neural networks have been successful in modeling skeletal motions. Convolutional neural networks can learn latent representations for denoising and synthesis [21]. Recurrent neural networks improve the robustness and enable long horizon synthesis [5, 58]. Graph neural networks capture the joint relations [33]. Generative flows combine the style and content in the latent space [63]. Transformers co-embed human motion and body parameters into a la-

tent representation [41]. Diffusion models provide a larger capacity and are less prone to mode collapse in generation [53, 73]. But all the above research is on a single body. While there is some research in modeling human-environment interactions [23, 69], two-body interactions are more complex. Very recent research shows successful synthesis of two interacting characters, but their focus is either on single character control [27, 47], or fix one while generating the other [12, 34]. To our best knowledge, there is no deep-learning method for complex two-character interactions especially under varying body sizes and proportions.

2.2. Motion Retargeting

Motion retargeting adapts a character’s motion to another of a different size while maintaining the motion semantics. Early research employs space-time optimization based on contact [11], purposefully-designed inverse kinematics solver for different morphologies [16], data-driven reconstruction of poses based on end-effectors [4, 46], or physical filters [49] and physical-based solvers [3] considering dynamics constraints. Recently, deep learning has achieved great success, e.g. recurrent neural networks with contact modeling [55], skeleton-aware operators without explicitly pairing the source and target motions [2], and variational autoencoders for motion features preservation during retargeting [14, 28, 54, 72]. Beyond skeletal motions, the skeleton structure is also effective in video based retargeting [68]. Fast deep learning methods are pursued for real-time robotic control [71]. Unlike previous research, we propose a novel deep learning architecture for motion retargeting/generation of two-character interactions, which are intrinsically more complex than single-character retargeting.

2.3. Interaction

Interaction retargeting involving more than one person is more challenging than single-body retargeting, due to their complex motion constraints [25] such as topological constraints [17], but these constraints involve heavy manual designs. As a more general solution, InteractionMesh [19] uses dense mesh structures to represent the spatial relations between two characters and minimizes the mesh change during retargeting [20] and synthesis of character-environment interactions [18]. As it may result in unnatural movements when the skeleton is significantly different from the original one, a prioritization strategy on local relations is proposed [38]. Nevertheless, optimisation-based methods require careful design of constraints, and incur large run-time costs.

Recently, there is a surge of deep learning methods on interactions, including human-object interaction [24, 42, 66], motion generation as reaction [7, 10, 65], from texts [30, 50] and by reinforcement learning [64, 74]. Interaction has

also been investigated in motion forecasting [39, 52, 67]. Among these papers, the closest work is interaction motion generation but existing work either cannot deal with skeletons of different sizes or does not focus on continuous and close interactions. To our best knowledge, there is no deep learning method for interaction modeling as proposed in this research.

Another key bottleneck of two-character interaction re-targeting/generation is the lack of data. Existing datasets focus on action recognition [8, 43, 70] with simple interactions. While some datasets with complex interactions are available [30, 45], they include limited variations of body sizes/proportions and have a limited amount of data. In this research, we present a new dataset and a method that learns efficiently from small amounts of data.

3. Methodology

We denote a motion with T frames as $q = \{q^0, \dots, q^T\}^T \in \mathbb{R}^{T \times N \times 3}$ where q^t is the t^{th} frame, and each frame $q^t = \{p_0^t, \dots, p_N^t\}$ consists of N joints and p_j is the j^{th} joint position. An interaction motion of two characters A and B is represented by $\{q_A, q_B\}$. For a specific interaction, different body sizes and proportions should not change the semantics, e.g. one character always having its arms around the other in hugging. These invariant semantics are often captured by topological/geometric features [17]. Therefore, a skeletal change in B should cause changes in both q_A and q_B to retain the semantics. We represent a B skeleton by its bone length vector $B_s \in \mathbb{R}^n$ where n is the number of bones. The aim is to model the joint probability $p(B_s, q_A, q_B)$. We propose a simple yet effective model, shown in Fig. 1.

3.1. A New Factorization of Interaction Motions

Directly learning $p(B_s, q_A, q_B)$ would need large amounts of data containing different interactions with varying bone lengths. Therefore, we first make it learnable on limited data by introducing a new factorization. First, we represent skeletons with different bone lengths as heterogeneously scaled versions of a *template* skeleton with a bone length scale vector $\hat{B} = \{1, \dots, 1\} \in \mathbb{R}^n$, i.e. we treat the bone lengths of the template skeleton as scale 1. We abuse the notation and denote a skeleton variation by B_s , indicating how each bone is scaled with respect to \hat{B} .

Next, we represent motion data as deviations from some *template* motion $\{\hat{q}_A, \hat{q}_B\}$ with the template skeleton \hat{B} . A skeleton variation B_s corresponds to a distribution of motions $\{q'_A, q'_B\}$, where not only the B motion deviates from \hat{q}_B , the A motion also deviates from \hat{q}_A accordingly to maintain the interaction. So we can split data into template motions and others $\{q_A, q_B\} = \{\hat{q}_A, \hat{q}_A\} \cup \{\hat{q}_B, q'_B\}$, so that $p(B_s, q_A, q_B) = p(q'_A, q'_B, B_s, \hat{q}_A, \hat{q}_B)$. Given $\{\hat{q}_A, \hat{q}_B\}$, $p(q'_A, q'_B, B_s, \hat{q}_A, \hat{q}_B)$ is an easier distribution to

learn than the original $p(B_s, q_A, q_B)$, as $\{\hat{q}_A, \hat{q}_B\}$ serves as an anchor motion with an anchor skeleton, so that all other motion variations can be described by offsets from the template motion, restricting $p(q'_A, q'_B, B_s, \hat{q}_A, \hat{q}_B)$ to only model the distribution of offsets from $\{\hat{q}_A, \hat{q}_B\}$.

There are many ways to factorize $p(q'_A, q'_B, B_s, \hat{q}_A, \hat{q}_B)$ theoretically. Our new factorization follows:

$$\begin{aligned} & p(q'_A, q'_B, B_s, \hat{q}_A, \hat{q}_B) \\ (i) &= p(q'_A|q'_B, B_s, \hat{q}_A, \hat{q}_B)p(q'_B, B_s, \hat{q}_A, \hat{q}_B) \\ (ii) &= p(q'_A|q'_B, \hat{q}_A)p(q'_B|B_s, \hat{q}_B)p(B_s, \hat{q}_A, \hat{q}_B) \\ (iii) &= p(q'_A|q'_B, \hat{q}_A)p(q'_B|B_s, \hat{q}_B)p(B_s) \end{aligned} \quad (1)$$

where (i) gives the conditional probability of $p(q'_A|q'_B, B_s, \hat{q}_A, \hat{q}_B)$, and its prior $p(q'_B, B_s, \hat{q}_A, \hat{q}_B)$. Further, $p(q'_B, B_s, \hat{q}_A, \hat{q}_B)$ can be factorized into $p(q'_B|B_s, \hat{q}_B)p(B_s, \hat{q}_A, \hat{q}_B)$ in (ii), assuming q'_B does not depend on \hat{q}_A . Given the template motion $\{\hat{q}_A, \hat{q}_B\}$ and a changed skeleton B_s , $\{B_s, \hat{q}_A, \hat{q}_B\} \sim p(B_s, \hat{q}_A, \hat{q}_B)$, we can sample a new $q'_B \sim p(q'_B|B_s, \hat{q}_B)$ that satisfies the desired skeleton change, then further sample a new $q'_A \sim p(q'_A|q'_B, \hat{q}_A)$ that maintains the interaction with q'_B . Further, (iii) is obtained when $\{\hat{q}_A, \hat{q}_B\}$ is given.

The three distributions in Eq. (1) have explicit meanings. $p(B_s)$ is the *skeleton prior* which captures skeletal variations that are likely to be observed; $p(q'_B|B_s, \hat{q}_B)$ is for *motion re-targeting*, i.e. modeling the distribution of possible B motions w.r.t. \hat{q}_B , given a skeletal variation B_s ; $p(q'_A|q'_B, \hat{q}_A)$ is for *motion adaptation*, i.e. modeling the possible A motions w.r.t. \hat{q}_A , given a specific B motion q'_B . Among many possible ways of factorization, our particular choice in Eq. (1) conforms to a plausible workflow where user input can be injected at multiple stages. The input can be a skeletal change B_s to $p(q'_B|B_s, \hat{q}_B)$, or a keyframed new motion q'_B to $p(q'_A|q'_B, \hat{q}_A)$. Alternatively, the B_s can be drawn from $p(B_s)$ for unlimited motion generation.

To keep our model small, inspired by the recent research in human motions [29, 31], we learn a generative model by assuming $p(B_s)$, $p(q'_B|B_s, \hat{q}_B)$ and $p(q'_A|q'_B, \hat{q}_A)$ to have well-behaved latent distribution, e.g. Gaussian, shown in Fig. 1. Compared with other alternative networks such as flows and Transformers, our model is especially suitable since our data is limited. We introduce the general architecture and refer the readers to the Supplementary for details.

3.2. Network Architecture

In Fig. 1, MLP1 and MLP2 are a five-layer (16-32-64-128-256) fully-connected (FC) network, and a five-layer (256-128-64-32-dim(B_s)) FC network, respectively. As B_s is a simple n -dimensional vector with fixed structural information, i.e. each dimension representing the scale of a bone, simple MLPs work well in projecting B_s into a latent space where it conforms to a Normal distribution.

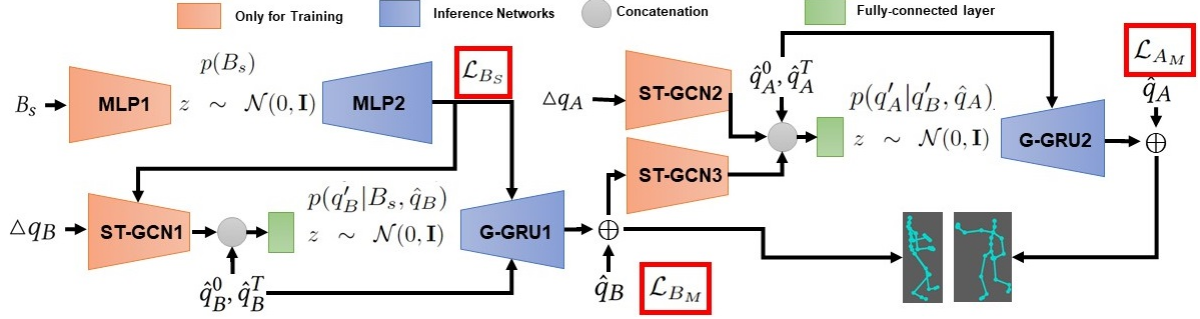


Figure 1. Overview of our model. The key components include Spatial-temporal Graph Convolution Networks (ST-GCN), Multi-layer perceptrons (MLP) and G-GRU networks. Details are in the supplementary material.

Next, we choose two types of networks as key components of our model to learn motion dynamics and interactions. First, spatio-temporal Graph Convolution Networks (ST-GCN) extract features by conducting spatial and temporal convolution on graph data and have been proven effective in analyzing human motions [9, 57]. We use ST-GCNs as encoders to extract reliable features. The other network is a Recurrent Neural Network named Graph Gated Recurrent Unit or G-GRU [29]. G-GRU models time-series data by Gated Recurrent Unit on graph structures and have the ability to stably unroll into the future on predicting human motions [29]. We use it as decoders in our model. This choice is again for reducing the required amount of data for training, which would be much larger if other networks, e.g. ST-GCNs are used as decoders based on our experiments.

Instead of directly learning the distribution of q'_B , learning the distribution of the differences $\Delta q_B = q'_B - \hat{q}_B$ is easier [51, 58]: $p(q'_B|B_s, \hat{q}_B) = p(\Delta q_B|B_s)$, which is easier as it becomes learning the distribution of offsets from the template motion \hat{q}_B and a skeleton variation B_s . We encode Δq_B into a latent space and then decode it back to the data space by:

$$\begin{aligned} z &= \text{FC}(\text{Concat}(\text{ST-GCN1}(\Delta q_B, B_s), \hat{q}_B^0, \hat{q}_B^T)) \\ \Delta q'_B &= \text{G-GRU1}(z, \hat{q}_B^0, \hat{q}_B^T, B_s) \\ &\text{subject to } z \sim \mathcal{N}(0, \mathbf{I}) \end{aligned} \quad (2)$$

where in both the encoding and decoding processes, we also incorporate the first and last frame of the template motion \hat{q}_B^0, \hat{q}_B^T because they help stabilize the dynamics based on our results. After decoding, we add the predicted $\Delta q'_B$ back to the template motion to get the new motion $q'_B = \hat{q}_B + \Delta q'_B$.

Next, given a motion q'_B , character A needs to adjust its motions to keep the interaction, leading to a distribution of possible q'_A . Similarly, we focus on learning $\Delta q_A = q'_A - \hat{q}_A$ by an autoencoder:




Figure 2. The architecture of ST-GCN1 and G-GRU1. More details are in the supplementary material.

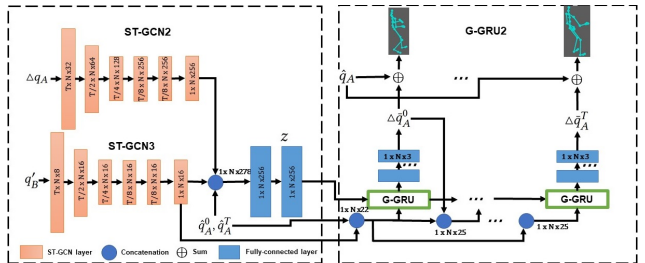


Figure 3. The architecture of ST-GCN2, ST-GCN3 and G-GRU2. More details are in the supplementary material.

$$\begin{aligned} z &= \text{FC}(\text{Concat}(\text{ST-GCN2}(\Delta q_A), \hat{q}_A^0, \hat{q}_A^T, \text{ST-GCN3}(q'_B)) \\ \Delta q'_A &= \text{G-GRU2}(z, \hat{q}_A^0, \hat{q}_A^T) \text{ subject to } z \sim \mathcal{N}(0, \mathbf{I}) \end{aligned} \quad (3)$$

where after decoding we compute the new motion $q'_A = \hat{q}_A + \Delta q'_A$.

We give more detailed architectures of ST-GCN1 and G-GRU1 in Figure 2, and the detailed architectures of ST-GCN2, ST-GCN3 and G-GRU2 in Figure 3.

3.3. Loss functions

Training our model involves three loss terms for the three autoencoders:

$$\mathcal{L} = \mathcal{L}_{B_s} + \mathcal{L}_{B_M} + \mathcal{L}_{A_M}. \quad (4)$$

Minimizing \mathcal{L}_{B_s} learns MLP1 and MLP2 to learn the distribution of possible skeleton variations B_s :

$$\mathcal{L}_{B_s} = \frac{1}{M} \sum \|B'_s - B_s\|_2^2 + D_{KL}[z|\mathcal{N}(0, \mathbf{I})], \quad (5)$$

where z is the output of MLP1, B'_s is the output of MLP2, B_s is the ground-truth skeleton variation and D_{KL} is the KL-divergence.

Next, \mathcal{L}_{B_M} is for training ST-GCN1 and G-GRU1:

$$\mathcal{L}_{B_M} = \frac{1}{M} \sum \{\omega_1 \|\tilde{q}'_B - q'_B\|_1 + \omega_2 \|\dot{\tilde{q}}'_B - \dot{q}'_B\|_1 + \omega_3 BL(\tilde{q}'_B, q'_B)\} + \omega_4 D_{KL}[z|\mathcal{N}(0, \mathbf{I})], \quad (6)$$

where z is the latent variable, $\omega_4 = 1 - \omega_1 - \omega_2 - \omega_3$, M is the total number of motions. \tilde{q}'_B and q'_B are the predicted and the ground-truth B motion. $\omega_1 = 0.75$, $\omega_2 = 0.1$ and $\omega_3 = 0.05$. $\|\cdot\|_1$ is the l_1 norm and $p(z|c) \sim \mathcal{N}(0, \mathbf{I})$. $BL(\tilde{q}'_B, q'_B)$ is the bone-length loss between \tilde{q}'_B and q'_B :

$$BL(\tilde{q}, q) = \sum_t \|bone_len(\tilde{q}^t) - bone_len(q^t)\|_2^2, \quad (7)$$

where $bone_len$ computes the bone lengths of frame t of \tilde{q} and q . Note we minimize the difference between the ground-truth and prediction on the *zero-order* and *first-order* derivative in Eq. 6.

Summarily for \mathcal{L}_{A_M} :

$$\mathcal{L}_{A_M} = \frac{1}{M} \sum [\omega_1 \|\tilde{q}'_A - q'_A\|_1 + \omega_2 \|\dot{\tilde{q}}'_A - \dot{q}'_A\|_1 + \omega_3 BL(\tilde{q}'_A, q'_A)] + \omega_4 D_{KL}[z|\mathcal{N}(0, \mathbf{I})], \quad (8)$$

where z is the latent variable. $\omega_4 = 1 - \omega_1 - \omega_2 - \omega_3$, M is the total number of motions. \tilde{q}'_A and q'_A are the predicted and the ground-truth A motion. $\omega_4 = 1 - \omega_1 - \omega_2 - \omega_3$, and $\omega_1 = 0.75$, $\omega_2 = 0.1$ and $\omega_3 = 0.05$. $BL(\tilde{q}'_A, q'_A)$ is the same bone length loss as in Eq. 7.

4. A New Interaction Dataset

To our best knowledge, there are few public datasets focusing on close and continuous interactions except [15]. To construct our dataset, we first obtain base motions and augment them. The base motion details are shown in the supplementary. We obtain ‘‘Judo’’. From CMU [1], we choose ‘‘Face-to-back’’, ‘‘Turn-around’’ and ‘‘Hold-body’’. From ExPI [15], we choose ‘‘Around-the-back’’, ‘‘Back-flip’’, ‘‘Big-ben’’, ‘‘Noser’’ and ‘‘Chandelle’’. These interactions are sufficiently complex to fully evaluate the robustness and generalizability of our model. They show the need for automated motion retargeting/generation as it requires hiring professional actors. Also, these motions contain rich and sustained contacts and close and continuous interactions, where single-body motion retargeting methods can

easily lead to breach of contact and severe body penetrations.

After obtaining the base motions, a number of variations of each motion are collected to form a dataset. Our method is independent of how the variations are obtained. One may consider motion capture with actors of different body sizes, or manual keyframing with different characters. We employ a semi-automated approach. We manually change the skeleton to generate variations, after which we adapt an iterative and interactive optimization approach called InteractionMesh [19] to generate new motions based on the changed skeletons. This allows us to precisely control the bone sizes for rigorous and consistent evaluation.

For each base motion, we vary the bones by scales within [0.75, 1.25] with a 0.05 spacing, where the original skeleton is used as the scale-1 template skeleton. This spans the +25% range of the original skeleton, covering most of the population. The process is semi-automatic, involving the use of an optimisation engine to carefully retarget an interaction to different body sizes, with manual adjustment of constraint weights and inspection of results. Synthesizing a few seconds of interaction generally requires around 2 minutes of computation. This is done multiple times for one variation of a base motion, due to the need for manual weighting tuning.

5. Experiments

5.1. Tasks, Metrics and Generalization Settings

Tasks. Since our model can generate motions with or without user input to specify a skeleton variation, we test different model variants for motion augmentation. Specifically, we evaluate our model on motion augmentation via retargeting and generation. If B_s is given, we refer to the task as *retargeting* where we only use G-GRU1 and G-GRU2 for inference; if B_s is not given, we use the full model (MLP2+G-GRU1+G-GRU2) and refer to it as *generation*.

Metrics. We employ four metrics for evaluation: joint position reconstruction error (E_r), bone-length error (E_b), Fréchet Inception Distance (FID), and joint-pair distance error (JPD). E_r , E_b and JPD are based on l_2 distance. FID is used to compare the distributional difference between the generated motions and the data. JPD measures the key joint-pair distance error. The key joint pairs are the body parts in continuous contact. It is to investigate the key spatial relations between joint pairs in different motions (Judo: A’s right hand to B’s spine; Face-to-back: A’s left hand to B’s right hand; Turn-around: A’s left hand to B’s right hand; Hold-Body: A’s right hand to B’s spine; Around-the-back: A’s left hand to B’s right hand; Back-flip: A’s left hand to B’s right hand; Big-ben: A’s right hand to B’s right hip; Noser: A’s right hand to B’s right hip; Chandelle: A’s right hand to B’s right hip). All results reported are per joint re-

Base Motion	M1	M2	M3	M4	M5	M6	M7	M8	M9	Total
Original frames	91	536	561	488	294	248	238	518	345	3,319
Augmented motion	160	119	119	119	90	90	90	90	90	967
Augmented frames	14,560	63,784	66,759	58,072	26,460	22,320	21,420	46,620	31,050	351,045

Table 1. M1: Judo, M2 Face-to-back, M3 Turn-around, M4: Hold-body, M5 Around-the-back, M6 Back-flip, M7 Big-ben, M8 Noser, M9 Chandelle. More details are in the Supplementary.

sults averaged over A and B.

Generalization Settings. Our dataset has two different skeletal topologies shown in the Supplementary. Therefore, we divide them into two datasets: D1 (M1-4) and D2 (M5-M9) and conduct experiments on them separately. We employ four different settings to evaluate our model: *random*, *cross-scale*, *cross-interaction* and *cross-scale-interaction*:

1. *Random* means a random split on the data for training and testing where we keep 20% data for testing.
2. *Cross-scale* means we train on moderate bone scales but predict on larger skeleton variations. Our training data is within the scale [0.95, 1.05] and our testing data is both much smaller [0.75, 0.85] and larger [1.15, 1.25]. Note the testing varies up to +/- 25% of the bone lengths covering a wide range of bodies.
3. *Cross-interaction* is splitting the data by interaction types, e.g. training on Judo and tested dancing. When we choose one or several interactions for testing, the other interactions are used for training the model. Specifically, in D1, we split the data into two sets: M1-M2 and M3-M4; in D2, we split them into two sets: M5-M7 and M8-M9. In both, when one group is used for training, the other is used for testing.
4. *Cross-scale-interaction* is both cross-scale and cross-interaction, which is the hardest setting. This means that the scale [0.95, 1.05] of some interactions are used for training, and the scale [0.75, 0.85] and [1.15, 1.25] in the other interactions are for testing. For instance, in D1, when the scale [0.95, 1.05] of M1-M2 is used for training, the scale [0.75, 0.85] and [1.15, 1.25] in M3-M4 are for testing.

5.2. Evaluation

5.2.1 Retargeting and Generation

We present the main results here and refer the readers to the Supplementary for more results and details.

We first show quantitative evaluation in Tab. 2. Across the two tasks, generation is harder than retargeting, as the bone scales are not given in generation. Naturally, the bone length error E_b is almost always slightly worse than Retargeting and so is JPD. But even the worst case is 330% in E_b and 206.89% worse in JPD which suggests the model generalizability on unseen scales and interactions in general is strong. We show visual results in Fig. 4 and the video. To-

		E_r	E_b	JPD	FID	E_b	JPD
M1	Random	1.069	0.171	3.008	2.934	0.18	3.421
	Cross-scale	2.017	0.304	4.248	3.973	0.354	4.304
	Cross-interaction	2.843	0.476	4.443	4.071	0.492	4.903
	Cross-scale-interaction	3.021	0.679	4.754	4.369	0.753	5.067
M2	Random	0.067	0.004	0.104	1.719	0.005	0.101
	Cross-scale	0.344	0.018	0.241	2.364	0.023	0.645
	Cross-interaction	0.671	0.087	0.625	3.077	0.097	1.004
	Cross-scale-interaction	1.051	0.131	0.845	3.256	0.143	1.317
M3	Random	1.076	0.02	2.274	5.573	0.03	2.134
	Cross-scale	1.563	0.066	2.948	6.556	0.094	2.872
	Cross-interaction	1.644	0.089	3.147	6.712	0.127	3.095
	Cross-scale-interaction	1.928	0.13	3.493	6.863	0.153	3.317
M4	Random	0.191	0.017	0.264	1.579	0.03	0.297
	Cross-scale	0.471	0.079	0.418	2.148	0.087	1.071
	Cross-interaction	0.617	0.104	0.589	2.648	0.111	1.347
	Cross-scale-interaction	0.897	0.112	0.624	3.094	0.129	1.915
M5	Random	1.975	0.003	0.398	0.69	0.01	0.604
	Cross-scale	2.674	0.016	0.837	1.283	0.031	1.157
	Cross-interaction	3.067	0.034	1.672	1.431	0.05	1.894
	Cross-scale-interaction	3.864	0.067	2.268	1.897	0.094	3.068
M6	Random	1.878	0.008	0.448	0.688	0.013	0.624
	Cross-scale	3.615	0.022	0.997	1.22	0.028	1.273
	Cross-interaction	4.013	0.031	1.923	1.523	0.039	2.024
	Cross-scale-interaction	4.876	0.076	2.641	1.667	0.083	3.264
M7	Random	2.746	0.006	0.495	0.645	0.015	0.702
	Cross-scale	5.204	0.017	1.163	1.153	0.03	2.14
	Cross-interaction	5.648	0.029	2.32	1.492	0.042	2.32
	Cross-scale-interaction	5.757	0.066	2.759	1.475	0.069	3.762
M8	Random	2.272	0.006	0.402	0.676	0.012	0.634
	Cross-scale	3.124	0.021	0.964	1.349	0.038	1.374
	Cross-interaction	3.389	0.04	1.534	1.671	0.057	1.862
	Cross-scale-interaction	3.971	0.103	2.341	2.965	0.103	2.675
M9	Random	2.234	0.005	0.403	0.634	0.009	0.561
	Cross-scale	2.935	0.01	0.934	1.412	0.043	1.259
	Cross-interaction	3.256	0.023	1.674	1.842	0.051	1.903
	Cross-scale-interaction	3.623	0.064	2.842	2.854	0.114	2.971

Table 2. Retargeting (left) and Generation (right). Here is the result of D1 (M1-4) and D2 (M5-9).

gether with the scaled skeleton, the poses are automatically adapted on both characters to keep the geometric relations of the interaction.

In terms of generation settings, the overall difficulty should be Cross-scale-interaction > Cross-interaction > Cross-scale > Random, as more and more information is included in the training data from Cross-scale-interaction to Random. The metrics in Tab. 2 are consistent with this expectation. Cross-scale-interaction is the most challenging task which is testing the model on both unseen bone sizes and interactions simultaneously. Its metrics are worse than the other three in general as expected. Despite the worse results, the visual results of cross-scale-interaction are of good quality. We show one example (with the worst met-




Figure 4. In the original Judo motion (top), the red character is augmented for a bigger body (middle) and a smaller body (bottom), while retaining the key features of the interaction semantics. The black boxes in column **a** highlight how the “Judo holding” semantics, i.e., the red character holding the blue one, are adapted. The black boxes in column **b** show a similar example.




Figure 5. Comparison between ground-truth (top) and cross-scale interaction (bottom). The skeleton of the red character is changed. Both of them are Back-flip on scale 0.85.

rics) in Fig. 5 in comparison with ground-truth.

5.2.2 Extrapolating to Large Unseen Scales

We predict larger scales. The scales are beyond our dataset (including the testing data). We show one example of Turn-around on 0.65 and 1.3 in the Supplementary, which shows that our model can extrapolate to larger skeletal variations when trained only using data on scales [0.95, 1.05]. More examples can be found in the video. Although larger scale variations e.g. 0.5 and 1.5 might lead to unnatural motions, the Supplementary already demonstrate the generalizability of our model.

	Hold-Body			Judo		
	Our method	[41]	[13]	Our method	[41]	[13]
FID	0.412	2.257	40.351	0.267	1.998	28.459
Eb	0.002	0.541	0.389	0.118	0.334	0.311
JPD	0.168	1.463	4.903	3.401	4.532	5.648

Table 3. Results at Scale 1.25, averaged over 10 randomly generated motions.




Figure 6. Scale 1.25 comparison. Left: ground-truth, mid: ours, right: [41]. [41] generates unnatural poses and break contact (enlarged parts). Zoom-in for better visualization.

5.3. Comparison

To our best knowledge, it is new for deep learning to be employed for interaction augmentation with varying body sizes. So there is no similar research. Therefore, we adapt two single-body methods ([13, 41]) which provide conditioned generation and are the only methods we know that could potentially be adapted for handling varying bone lengths, i.e. we train the model by labelling different scales as different conditions and train the model on scale [0.75, 1.25]. More specifically, both models require action type (i.e. a class label) as input, so we label data at different scales as different classes. Note [41] and [13] cannot generate motions for unseen action types, which means they cannot predict on unseen scales like our method.

We show the metrics in Tab. 3. After trying our best to train [13], it still generates jittering motions. It can preserve the bone-length better than [41] but its FID and JPD are much worse. [41] generate better results but it is still much worse than our method. We show one example of Hold-Body in Fig. 6 in comparison with [41]. Overall, single-body methods even when adapted cannot easily generate interactions.

We also compare with InteractionMesh [19]. Since our ground-truth is from InteractionMesh, comparisons on the aforementioned evaluation metrics would be meaningless. Instead, we compare the speed and motion quality on unseen extreme scales. The inference time of our model is 0.323 seconds, while InteractionMesh needs ~ 120 seconds on average per optimization, plus the time needed for manual tuning of the weighting. Admittedly, our model needs overheads for training. However, once trained, it is very fast and can be used for interactive applications. Further, InteractionMesh needs to do optimization for every given B_s , while our model is trained once then does inference for any B_s . Last but not least, InteractionMesh sometimes fails to

Predict(sec)	0.2	0.4	0.6	0.8	1.0	
M5	JME	0.234/0.449	0.427/0.771	0.593/1.073	0.722/1.365	0.848/1.594
	AME	0.417/0.605	0.750/1.100	1.036/1.499	1.250/1.877	1.474/2.176
M6	JME	0.520/0.552	0.848/0.874	1.098/1.187	1.485/1.533	1.670/1.799
	AME	0.671/0.682	1.170/1.168	1.530/1.579	1.958/1.968	2.253/2.326
M7	JME	0.538/0.565	0.971/0.959	1.298/1.302	1.720/1.708	1.926/1.848
	AME	0.708/0.727	1.334/1.367	1.809/1.823	2.319/2.290	2.608/2.573
M8	JME	0.507/0.562	0.927/0.985	1.137/1.284	1.648/1.692	1.886/2.013
	AME	0.673/0.695	1.315/1.330	1.796/1.830	1.908/1.968	2.353/2.483
M9	JME	0.505/0.590	0.834/0.920	1.263/1.312	1.567/1.725	1.904/2.201
	AME	0.721/0.723	1.469/1.634	1.848/1.923	2.031/2.224	2.415/2.657
M5	JME	0.278/0.507	0.444/0.767	0.652/1.122	0.763/1.299	0.867/1.641
	AME	0.467/0.668	0.748/1.094	1.085/1.603	1.345/1.894	1.551/2.230
M6	JME	0.538/0.548	0.856/0.880	1.096/1.180	1.488/1.586	1.622/1.793
	AME	0.683/0.690	1.194/1.196	1.528/1.566	1.960/1.973	2.256/2.335
M7	JME	0.584/0.579	1.023/1.049	1.322/1.315	1.645/1.648	1.937/1.940
	AME	0.723/0.746	1.466/1.489	1.896/1.900	2.391/2.379	2.608/2.612
M8	JME	0.597/0.605	1.036/1.068	1.204/1.315	1.701/1.767	1.892/2.148
	AME	0.710/0.748	1.348/1.347	1.808/1.810	2.064/2.101	2.332/2.425
M9	JME	0.524/0.528	0.862/0.892	1.378/1.392	1.674/1.702	1.923/2.046
	AME	0.718/0.713	1.486/1.497	1.867/1.901	2.067/2.209	2.523/2.672

Table 4. Motion prediction of [15] (top) and [62] (bottom) in JME (joint mean error) and AME (aligned mean error) from D2 (M5-9). In each test, xx/xx is with/without data augmentation.

converge due to its optimization set up, resulting in either numerical explosion or very unnatural motions (see video). This requires careful manual tuning. Comparatively, our model does not need manual intervention.

5.4. Downstream Tasks

Motion augmentation can benefit various downstream tasks. Here we show two downstream tasks: motion prediction and activity recognition. In motion prediction we train two models [62] and [15] on the ExPI dataset [15] with/without our data augmentation, following their settings. The testing protocols and evaluation metrics follow [15]. The results are shown in Tab. 4, where 90 of 100 metrics are improved by our augmentation, with a maximum 47.88% improvement on JME (M5-AB-0.2sec) and a maximum 47.74% improvement on AME (M5-AB-0.6sec).

In activity recognition, we train three latest activity classifiers HD-GCN [26], STGAT [22] and TCA-GCN [61] on ExPI with/without data augmentation, following two data splits: 80/10/10 and 50/20/30 split on training/validation/testing data. The results are shown in Tab. 5. The data augmentation improves the accuracy across all models and all split settings. As the training data is reduced from 80% to 50%, the results with data augmentation have a small deterioration (less than 1.49%). Without data augmentation, it quickly drops by as much as 3.42%.

We further show the quality of the augmented motions via a trained classifier. If a trained classifier can correctly recognize the generated motions, then it suggests the generated features have similar features to the original data. We train the aforementioned classifiers on the original ExPI data and use the generated motions as testing data. Tab. 6 shows the action recognition result. Our method outperforms the other two methods in all three action recognition

Settings/Classifiers	HD-GCN [26]	STGAT [22]	TCA-GCN [61]
80/10/10	94.80/94.36	94.27/94.10	94.68/94.62
50/20/30	93.92/92.65	93.66/92.40	93.27/91.38

Table 5. Activity recognition accuracy on 3 different classifiers from ExPI [15]. In each test, xx/xx is with/without data augmentation.

Methods/Classifiers	HD-GCN [26]	STGAT [22]	TCA-GCN [61]
ACTOR[41]	97.68	98.03	97.22
Action2motion[13]	97.43	96.90	96.45
Our method	98.64	98.53	97.93

Table 6. Activity recognition accuracy on 3 different methods from D2 (M5-9). Training on the ground-truth and testing on generated 200 motions.

classifiers, which shows that our generated data has more similar features to the ground-truth. Given close interaction data is new [15] and its limited variety and amounts, our method provide an efficient way of augmenting such data for activity recognition.

5.5. Alternative Architectures

Our model combines existing network components in a novel way for interaction augmentation, so a natural question is if there are other better alternative architectures. We test several alternative network architectures inspired by existing research. The selection criteria is they need to be data efficient for learning, so we exclude some data-demanding architectures such as Transformers or Diffusion models. The details and results are shown in the Supplementary, but overall our model outperforms the alternative architectures.

6. Conclusion, Limitations & Discussion

To our best knowledge, our research is the very first deep learning model for interaction augmentation. It has high accuracy in generating desired skeletal changes, great flexibility in generating diversified motions, strong generalizability to unseen and large skeletal scales, and benefits to multiple downstream tasks. One limitation is that we need some data samples to start and require the same skeletal topology to do cross-motion motion augmentation. Also, the model does not explicitly enforce contacts between two characters now. However, considering the difficulties of interaction motion capture, our method provides a new and fast way of iteratively augmenting a single captured motion then learning to generate infinite number of variations. Next, although we use InteractionMesh to generate training data, our method can easily incorporate other data sources such as captured motions from different subjects as well as manually created motions by animators. Given the small number of motions needed by our method, this is still a fast pipeline to acquire a large number of interactions with varying body sizes.

Two-Person Interaction Augmentation with Skeleton Priors

Supplementary Material

1. More Details on Dataset

One instance of the nine motions (Judo, Face-to-back, Turn-around, Hold-body, Around-the-back, Back-flip, Big-ben, Noser and Chandelle) was captured from different subjects by different systems. Therefore, we have two skeletons, with 25 joints and 24 bones (Judo, Face-to-back, Turn-around and Hold-body), and 17 joints and 16 bones (Around-the-back, Back-flip, Big-ben, Noser and Chandelle), shown in Fig. 7. The motions in D1 and D2 are shown in Fig. 8 and Fig. 9.

For each captured motion, we vary bones with scales within $[0.75, 1.25]$ with a 0.05 spacing, where the original skeleton is used as the template skeleton and labeled as scale 1. An exhaustive permutation of all possible scaling is impractical. Therefore, we only use full-body uniform scaling and single-bone scaling on the upper-body bones which are heavily involved in interactions. We manually specify the skeleton variations and use InteractionMesh [19] to generate motions.

InteractionMesh is an optimization framework where the required input is the original motion and the scaled target skeleton. InteractionMesh make a mesh structure by connecting every pair of points between two characters, called interaction mesh. When adapting the motion for a desired scaled skeleton, it minimizes the Laplacian energy, i.e. a deformation energy term of the interaction mesh, to keep the spatial relations as much as possible for every pair of joints. Using InteractionMesh, instead of hiring more actors, allows us to: (1) have exact control over the bone lengths; (2) explore atypical skeleton/body sizes, e.g. left arm longer than right arm. However, the optimization process is sensitive to initialization and weight tuning of the object function. For each skeleton variation, we manually conduct several rounds of optimizations and visually inspect the quality of the generated motion, until it become satisfactory.

Admittedly, compared with the only dataset for interactions [15], the number of interactions in our dataset is smaller (9 vs 16), but our emphasis is the diversity of body sizes. Overall, we have 9 base motions, a total of 967 body variations with 351045 frames, which is larger than [15] in terms of the number of sequences and frames.

2. Additional Results and Details

2.1. Detailed experiments

The full comparison results of different methods for both retargeting and generation are shown in Tab. 7-Tab. 15.




Figure 7. Two skeletons in our dataset. Left: 25 joints, Right: 17 joints




Figure 8. The base motion in D1(M1-M4). From top to bottom: Judo, Face-to-back, Turn-around and Hold-body.

2.2. Skeletal Visualization vs Body Visualization.

Skeletal visualization is widely adopted in existing research (e.g. character animation, motion prediction, activity recognition, etc.), but we do notice a recent trend of showing body shapes with skeletal motions. Theoretically, it is possible to generate body meshes e.g. via SMPL [32]. However, for our problem, this is not the case because generating/adapting body meshes for varying bone lengths is non-trivial and is itself an entirely different topic. Not only is there no body geometry in the data we used, but the motion contains rich contacts between characters. Therefore, generated body meshes could easily lead to penetration so

Metric	F-CNNs	F-GCNs	M-CNNs	M-GCNs	Ours
E_r	0.673/2.569	0.834/5.642	0.517/2.771	0.681/3.447	0.298/1.840
	0.821/4.225	1.684/8.018	1.206/4.060	1.372/3.530	0.463/3.571
	1.251/6.031	3.604/6.858	1.363/4.823	1.622/5.166	0.942/4.744
	1.521/6.455	4.002/6.840	1.684/5.690	1.812/6.110	1.130/4.912
E_b	0.093/0.355	0.120/0.484	0.109/0.835	0.136/0.650	0.072/0.270
	0.127/1.097	0.135/1.067	0.164/1.014	0.158/1.270	0.102/0.506
	0.234/1.214	0.262/1.42	0.273/1.449	0.278/1.516	0.189/0.763
	0.448/1.368	0.403/1.653	0.428/1.506	0.418/1.834	0.305/1.053
JPD	4.078	4.358	6.654	6.877	3.008
	4.938	4.877	6.821	7.239	4.248
	6.674	7.034	8.345	8.003	4.443
	7.894	8.234	8.861	8.642	4.754
FID	3.574/4.928	6.784/14.304	5.421/11.483	3.136/8.09	2.134/3.734
	4.841/8.103	7.541/24.021	7.412/11.838	4.158/9.046	3.824/4.122
	6.854/7.112	7.984/25.080	8.025/15.867	4.278/10.846	4.033/4.109
	7.931/8.761	12.841/26.721	9.541/16.207	6.418/10.654	4.214/4.524
E_b	0.254/0.350	0.365/0.569	0.315/0.549	0.228/0.518	0.176/0.184
	0.621/0.925	0.421/0.825	0.512/1.480	0.862/0.926	0.285/0.423
	0.687/1.763	1.654/2.276	1.862/2.820	1.923/1.947	0.532/0.452
	1.325/3.081	2.284/3.022	2.684/3.424	2.047/5.267	0.737/0.769
JPD	7.542	8.844	6.543	7.832	3.421
	8.043	9.641	7.965	8.239	4.304
	8.821	10.632	8.517	9.632	4.903
	9.852	12.245	9.786	10.985	5.067

Table 7. Comparison on Judo retargeting (top) and generation (bottom). XX/XX are Character A/B results. All results are per joint results. The four rows in each cell are results of Random, Cross-scale, Cross-interaction and Cross-scale-interaction respectively.

Metric	F-CNNs	F-GCNs	M-CNNs	M-GCNs	Ours
E_r	0.227/0.328	0.445/1.474	0.102/0.232	0.424/2.760	0.058/0.076
	0.234/0.335	0.544/1.554	0.124/0.372	1.732/3.714	0.263/0.425
	0.297/0.451	0.548/1.573	0.156/0.434	1.988/3.876	0.352/0.990
	0.725/1.812	0.641/1.785	0.921/1.932	4.412/5.689	0.630/1.472
E_b	0.009/0.018	0.040/0.107	0.035/0.048	0.120/0.727	0.002/0.006
	0.022/0.053	0.082/0.241	0.056/0.078	0.312/0.739	0.012/0.024
	0.054/0.081	0.103/0.357	0.841/0.959	0.327/0.884	0.089/0.085
	0.245/0.432	0.584/0.633	1.294/2.230	0.972/1.064	0.045/0.217
JPD	0.599	0.517	0.330	0.465	0.104
	0.658	0.505	0.414	0.302	0.241
	0.892	0.703	0.678	0.526	0.625
	1.284	1.724	1.595	0.951	0.845
FID	2.637/7.218	21.238/37.530	4.825/14.917	1.379/2.782	1.134/2.304
	3.118/8.745	20.457/35.483	5.215/14.551	1.751/3.451	1.824/2.904
	3.331/8.286	25.844/38.517	5.466/19.651	2.154/3.756	2.533/3.621
	5.542/9.844	26.723/40.425	7.983/24.842	2.831/4.237	2.814/3.698
E_b	0.032/0.046	0.186/0.200	1.157/1.965	0.125/0.753	0.001/0.009
	0.043/0.062	0.267/0.352	1.305/2.021	0.163/0.847	0.018/0.028
	0.107/0.108	0.349/0.514	2.687/2.984	0.195/0.954	0.053/0.141
	0.342/0.504	0.652/0.721	3.864/4.030	0.642/1.231	0.073/0.213
JPD	1.017	3.916	2.469	0.369	0.101
	1.157	4.148	2.672	0.454	0.645
	1.872	6.216	2.896	0.648	1.004
	1.904	7.385	4.542	2.034	1.317

Table 8. Comparison on Face-to-back retargeting (top) and generation (bottom). XX/XX are on Character A/B. All results are per joint results. The four rows in each cell are results of Random, Cross-scale, Cross-interaction and Cross-scale-interaction respectively.

Metric	F-CNNs	F-GCNs	M-CNNs	M-GCNs	Ours
E_r	0.454/0.874	0.622/1.121	0.334/1.244	1.735/2.714	0.398/1.754
	0.534/0.925	0.751/1.334	0.453/1.348	1.956/2.819	0.263/2.863
	0.796/2.071	0.728/1.127	0.879/2.941	2.001/2.771	0.352/2.936
	1.296/2.842	1.121/2.254	1.641/3.263	2.942/3.234	0.530/3.326
E_b	0.020/0.038	0.075/0.082	0.363/0.473	0.320/0.773	0.003/0.037
	0.050/0.079	0.098/0.112	0.383/0.536	0.334/0.801	0.028/0.104
	0.112/0.135	0.102/0.133	0.349/0.551	0.503/0.978	0.059/0.119
	0.234/0.524	0.221/0.508	0.641/0.897	0.842/1.235	0.105/0.155
JPD	3.359	2.291	3.765	2.155	2.274
	3.507	3.814	4.202	2.261	2.948
	3.741	4.001	4.268	3.054	3.147
	4.542	6.123	4.964	4.637	3.493
FID	6.806/7.702	9.037/10.487	9.830/11.495	4.407/8.824	3.214/7.932
	7.023/8.112	10.148/12.046	11.049/16.839	4.466/8.847	3.854/9.258
	7.214/8.849	12.645/20.984	12.057/18.213	5.121/9.157	3.708/9.716
	8.678/10.845	13.412/23.582	14.325/21.842	6.051/9.821	3.923/9.803
E_b	0.413/0.454	0.315/0.445	0.940/1.986	0.332/0.776	0.006/0.054
	0.464/0.457	0.486/0.781	1.001/2.068	0.348/0.816	0.025/0.163
	0.516/0.604	0.715/1.033	1.104/2.211	0.401/0.849	0.052/0.202
	0.605/0.840	1.254/1.930	1.529/2.842	0.645/1.731	0.137/0.169
JPD	3.678	4.120	4.399	3.206	2.134
	3.845	4.368	4.501	4.025	2.872
	4.008	4.808	4.815	4.419	3.095
	4.845	5.614	5.325	5.004	3.317

Table 9. Comparison on Turn-around retargeting (top) and generation (bottom). XX/XX are on Character A/B. All results are per joint results. The four rows in each cell are results of Random, Cross-scale, Cross-interaction and Cross-scale-interaction respectively.

Metric	F-CNNs	F-GCNs	M-CNNs	M-GCNs	Ours
E_r	0.230/0.258	0.504/1.178	0.077/0.289	0.339/0.860	0.098/0.284
	0.257/0.291	0.604/1.258	0.125/0.291	0.458/0.909	0.163/0.779
	0.304/0.345	0.771/1.541	0.201/0.294	0.517/0.931	0.252/0.982
	0.651/0.837	1.204/1.976	0.604/0.849	0.915/1.677	0.430/1.364
E_b	0.007/0.014	0.049/0.055	0.045/0.059	0.127/0.569	0.003/0.031
	0.015/0.041	0.051/0.059	0.057/0.169	0.199/0.605	0.007/0.151
	0.049/0.064	0.074/0.098	0.099/0.203	0.232/0.771	0.012/0.196
	0.184/0.251	0.142/0.194	0.204/0.531	0.671/0.949	0.025/0.199
JPD	0.617	2.076	0.807	1.685	0.264
	0.824	2.148	0.814	1.694	0.418
	0.835	2.548	1.215	1.805	0.589
	1.542	4.287	2.674	3.004	0.624
FID	3.585/8.344	20.815/24.261	0.721/3.867	0.322/3.513	0.214/2.944
	3.748/9.424	21.784/28.454	0.915/2.245	1.751/2.158	0.854/3.442
	3.982/9.458	22.511/30.368	1.052/2.244	1.981/2.752	0.712/4.584
	4.874/12.828	23.074/29.241	2.452/3.657	3.642/3.777	0.923/5.265
E_b	0.056/0.101	0.504/0.591	0.097/0.519	0.137/0.587	0.006/0.054
	0.077/0.125	0.607/0.614	0.128/0.684	0.252/0.640	0.025/0.149
	0.098/0.130	0.701/0.848	0.157/0.745	0.425/0.672	0.052/0.170
	0.249/0.341	1.204/1.899	0.531/1.112	0.822/1.054	0.067/0.191
JPD	1.105	2.217	1.304	1.707	0.297
	1.235	2.148	1.365	1.735	1.071
	1.442	3.331	1.317	1.844	1.347
	2.140	4.640	2.384	2.896	1.915

Table 10. Comparison on Hold-body retargeting (top) and generation (bottom). XX/XX are on Character A/B. All results are per joint results. The four rows in each cell are results of Random, Cross-scale, Cross-interaction and Cross-scale-interaction respectively.

Metric	F-CNNs	F-GCNs	M-CNNs	M-GCNs	Ours
E_r	1.586/1.614	2.277/4.159	3.120/4.384	1.759/4.045	1.153/2.797
	1.662/4.770	2.282/4.342	3.252/4.844	2.201/4.349	1.851/3.497
	1.976/4.824	2.044/4.157	3.924/4.121	2.471/4.174	2.252/3.882
	2.782/5.452	3.451/5.735	4.812/6.328	3.421/5.418	3.453/4.275
E_b	0.141/0.225	0.088/0.226	0.030/0.063	0.003/0.031	0.001/0.005
	0.156/0.267	0.135/0.287	0.105/0.161	0.010/0.061	0.003/0.029
	0.225/0.305	0.197/0.334	0.210/0.370	0.017/0.120	0.018/0.050
	0.647/0.812	0.729/0.964	0.748/0.792	0.079/0.234	0.055/0.079
JPD	1.844	3.841	3.525	0.307	0.398
	2.217	3.428	3.191	0.941	0.837
	2.618	3.627	3.224	1.715	1.672
	3.542	4.521	4.751	2.642	2.114
FID	0.349/0.746	5.142/8.672	1.824/1.971	0.337/0.584	0.214/1.166
	0.662/0.997	5.771/9.071	2.054/2.642	0.417/0.742	0.854/1.712
	1.087/1.671	6.041/9.817	2.511/2.912	0.661/0.942	1.212/1.650
	1.574/1.942	8.452/10.122	2.981/3.514	0.967/1.345	1.723/2.071
E_b	0.170/0.317	0.565/0.953	0.105/0.251	0.006/0.040	0.006/0.014
	0.204/0.391	0.642/1.074	0.287/0.354	0.038/0.084	0.025/0.037
	0.396/0.504	0.699/1.611	0.487/0.515	0.051/0.191	0.042/0.058
	0.925/1.213	1.077/1.921	1.073/1.258	0.254/0.293	0.047/0.141
JPD	2.485	4.253	3.671	0.345	0.604
	2.671	4.506	3.851	1.414	1.157
	3.211	5.011	4.514	1.892	1.894
	4.359	5.824	5.942	2.487	2.268

Table 11. Comparison on Around-the-back motion retargeting (top) and generation (bottom). XX/XX are on Character A/B. All results are per joint results. The four rows in each cell are results of Random, Cross-scale, Cross-interaction and Cross-scale-interaction respectively.



Figure 9. The base motion in D2 (M5-M9). From top to bottom: Around-the-back, Back-flip, Big-ben, Noser and Chandelle.

manually created meshes are needed. Furthermore, since we sample different bone lengths, manual creation of body geometry for every scaled skeleton would be required, as naive non-uniform scaling on the body mesh designed for a template skeleton would easily cause mesh deformation artefacts or contact breach. Methods such as SMPL might help but with no guarantee, because arbitrary bone scaling easily leads to out-of-distribution skeletons deviating from their training data. We tested SMPL and show one such example in Fig. 10. But this does not mean our motion quality is low. The motion quality can be visually inspected in the video.

2.3. Generation Diversity

Our model contains 3 learned Gaussian distributions and therefore is intrinsically stochastic. We show a Judo motion sampled multiple times (in different colors) using the same skeleton in Fig. 11 (zoom-in for better visualization). While there are motion diversity, we do realize that the motions do not visually show big variations. Note that this is due to the fact that the skeleton is exactly the same for all motions, and more importantly the key interaction features such as contacts need to be maintained in different samples. These contacts implicitly act as constraints for augmentation. However, as shown before, when the bone sizes

Metric	F-CNNs	F-GCNs	M-CNNs	M-GCNs	Ours
E_r	1.402/4.493	2.015/3.483	2.641/3.783	1.501/3.142	1.541/2.215
	1.427/4.025	2.421/3.214	2.453/2.924	1.701/3.542	2.511/4.719
	1.481/4.406	2.812/4.082	2.125/3.421	1.412/3.199	3.052/4.974
	3.214/6.643	3.542/6.547	3.895/6.852	4.624/8.954	3.453/6.299
E_b	0.051/0.063	0.122/0.272	0.031/0.092	0.001/0.031	0.002/0.014
	0.071/0.094	0.228/0.309	0.077/0.108	0.008/0.071	0.005/0.039
	0.081/0.104	0.320/0.481	0.334/0.471	0.014/0.191	0.012/0.050
	0.171/0.307	0.422/0.554	0.445/0.575	0.089/0.201	0.041/0.111
JPD	0.637	4.123	4.241	0.480	0.495
	1.734	4.187	4.651	1.712	1.163
	2.794	4.914	4.987	2.923	2.320
	4.045	5.514	5.612	3.818	3.762
FID	0.283/0.806	5.849/6.246	2.421/3.841	0.305/0.512	0.424/0.952
	0.310/0.884	5.244/6.121	2.451/3.874	0.540/0.917	0.878/1.562
	0.711/0.976	6.018/6.924	3.084/4.312	0.749/1.034	0.912/2.134
	1.874/1.854	6.684/7.896	4.548/4.845	1.342/1.837	1.027/2.307
E_b	0.112/0.389	0.398/1.662	0.248/0.745	0.003/0.064	0.006/0.020
	0.162/0.401	0.407/1.823	0.425/0.945	0.008/0.118	0.015/0.041
	0.227/0.454	0.487/2.132	0.504/1.003	0.031/0.216	0.022/0.056
	0.421/0.645	0.722/2.972	0.924/1.781	0.135/0.421	0.037/0.129
JPD	1.510	5.204	4.312	1.613	0.624
	1.601	5.405	4.894	1.819	1.273
	2.718	5.827	5.181	3.003	2.024
	4.248	5.922	6.247	4.252	3.941

Table 12. Comparison on Back-flip motion retargeting (top) and generation (bottom). XX/XX are on Character A/B. All results are per joint results. The four rows in each cell are results of Random, Cross-scale, Cross-interaction and Cross-scale-interaction respectively.

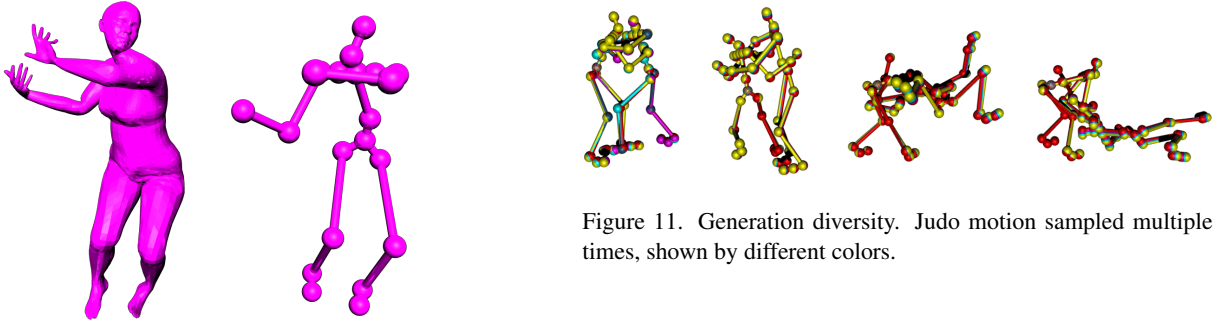


Figure 10. SMPL results on our skeleton. Left: the SMPL generated mesh. Right: the skeleton we captured in Judo motion for Character A. Due to the skeleton differences, e.g. different number of joints and different lengths of bones, severe distortion (both hands and left foot) exists in the body shape.

change, bigger diversities can be seen.

2.4. Generalizability on Reduced Training Samples

Since high-quality interaction motion is hard to capture and data augmentation is not easy, it is highly desirable if augmentation can work on as few training samples as possible. To test this, we choose Face-to-back (M2) and Big-ben (M7) under Cross-scale-interaction, and reduce the training

Figure 11. Generation diversity. Judo motion sampled multiple times, shown by different colors.

samples to 24, to 12 and 6. More specifically, when using the scale [0.75, 0.85] and [1.15, 1.25] of M2 as the testing data, we randomly select 24, 12 and 6 training samples from the scale [0.95, 1.05] of M3-M4 for training. Similarly, when choosing the scale [0.75, 0.85] and [1.15, 1.25] of M7 as the testing data, we randomly select 24, 12 and 6 training samples from the scale [0.95, 1.05] of M8-M9 for training. Note this is a very challenging setting.

Tab. 16 shows a quantitative comparison. Note metrics have different scales and cross-metric comparison is not meaningful. Unsurprisingly, all metrics become worse when the number of training samples decreases. However, the increase of errors is slow compared with the corresponding experiments in retargeting and generation part, showing our method has high data efficiency. We show more results

Metric	F-CNNs	F-GCNs	M-CNNs	M-GCNs	Ours
E_r	1.691/3.948	2.371/4.434	3.013/3.212	1.926/4.337	1.621/3.871
	3.802/8.071	3.427/8.724	3.412/8.894	4.016/9.711	2.054/8.354
	3.914/9.221	3.624/9.217	4.642/10.945	5.174/10.915	2.752/8.544
	5.544/10.221	4.052/9.826	4.952/10.954	5.204/11.065	2.453/9.061
E_b	0.049/0.134	0.159/0.316	0.031/0.140	0.001/0.033	0.003/0.009
	0.108/0.227	0.207/0.375	0.099/0.204	0.004/0.091	0.012/0.022
	0.172/0.271	0.271/0.405	0.123/0.306	0.017/0.194	0.022/0.036
	0.211/0.294	0.301/0.534	0.325/0.452	0.036/0.205	0.031/0.101
JPD	0.956	3.569	3.356	0.737	0.495
	0.917	3.453	3.541	1.571	1.163
	2.127	4.485	3.941	2.584	2.320
	2.354	4.755	4.842	2.948	3.762
FID	0.301/1.816	6.833/8.771	0.504/2.051	0.472/0.520	0.407/0.883
	0.651/1.971	7.661/8.875	0.571/2.364	0.841/1.117	0.841/1.465
	0.806/2.011	8.054/9.404	0.604/2.781	0.894/1.199	0.934/2.050
	1.068/2.325	8.725/9.891	1.262/2.934	1.288/1.824	0.939/2.011
E_b	0.072/0.216	1.084/1.497	0.105/0.310	0.0017/0.0606	0.004/0.026
	0.109/0.312	1.400/1.425	0.184/0.412	0.018/0.107	0.009/0.051
	0.206/0.337	1.701/1.832	0.208/0.577	0.037/0.401	0.017/0.067
	0.332/0.521	2.054/2.641	0.355/0.851	0.109/0.484	0.022/0.116
JPD	2.072	5.972	3.451	0.895	0.702
	2.424	6.422	3.411	1.718	2.140
	2.672	7.051	4.823	1.896	2.320
	2.791	7.455	6.271	2.942	2.759

Table 13. Comparison on Big-ben motion retargeting (top) and generation (bottom). XX/XX are on Character A/B. All results are per joint results. The four rows in each cell are results of Random, Cross-scale, Cross-interaction and Cross-scale-interaction respectively.

Metric	F-CNNs	F-GCNs	M-CNNs	M-GCNs	Ours
E_r	1.256/4.045	1.853/3.997	2.645/4.241	2.736/4.360	0.953/3.591
	1.864/4.689	2.907/5.084	2.756/5.601	2.862/5.336	1.638/4.610
	2.362/4.898	3.015/5.915	3.808/6.032	2.991/5.617	1.937/4.841
	2.623/5.185	3.530/6.240	4.810/6.937	3.244/6.054	2.223/5.719
E_b	0.088/0.105	0.141/0.353	0.286/0.422	0.144/0.216	0.002/0.010
	0.582/0.698	0.164/0.391	0.231/0.488	0.189/0.271	0.009/0.033
	0.612/0.706	0.200/0.446	0.409/0.521	0.217/0.595	0.017/0.063
	0.620/0.705	0.220/0.450	0.426/0.554	0.237/0.607	0.031/0.175
JPD	3.557	3.792	5.451	5.670	0.402
	3.804	3.984	5.669	6.265	0.964
	4.602	4.205	6.324	6.618	1.534
	5.552	5.434	6.729	6.887	2.341
FID	0.624/4.578	10.764/12.042	3.011/11.084	0.831/2.471	0.297/1.055
	2.642/3.637	11.684/14.587	3.512/12.986	1.986/3.076	0.685/2.013
	3.186/5.804	15.545/22.688	4.336/14.745	3.957/4.225	0.907/2.435
	4.169/7.804	17.550/23.821	5.306/16.075	4.580/4.904	1.274/4.656
E_b	0.176/0.278	0.121/0.350	0.334/1.147	0.186/0.286	0.004/0.020
	0.532/0.758	0.225/0.421	0.418/1.379	0.167/0.399	0.009/0.067
	0.685/0.721	0.345/0.584	0.514/1.536	0.231/0.763	0.017/0.097
	0.688/0.842	0.385/0.604	0.595/1.623	0.243/0.789	0.052/0.154
JPD	1.116	2.207	2.914	2.843	0.634
	2.513	4.741	5.045	5.157	1.374
	4.895	5.068	6.861	6.854	1.862
	5.542	6.068	7.861	7.560	2.675

Table 14. Comparison on Noser retargeting (top) and generation (bottom). XX/XX are Character A/B results. All results are per joint results. The four rows in each cell are results of Random, Cross-scale, Cross-interaction and Cross-scale-interaction respectively.

Metric	F-CNNs	F-GCNs	M-CNNs	M-GCNs	Ours
E_r	0.754/4.548	1.696/4.302	0.857/4.241	1.733/4.346	0.735/3.733
	0.930/5.288	1.957/4.553	1.263/4.914	1.869/4.866	1.328/4.542
	1.513/5.585	2.013/4.9014	1.881/4.937	1.909/5.670	1.863/4.649
	2.062/6.018	2.415/6.001	2.384/6.237	2.841/6.287	2.197/5.049
E_b	0.018/0.102	0.086/0.203	0.082/0.222	0.043/0.196	0.003/0.007
	0.064/0.203	0.184/0.334	0.258/0.547	0.134/0.220	0.008/0.012
	0.127/0.259	0.222/0.446	0.299/0.687	0.205/0.335	0.015/0.031
	0.156/0.372	0.252/0.474	0.394/0.697	0.229/0.385	0.034/0.094
JPD	2.645	3.762	4.552	4.850	0.403
	3.512	3.874	5.589	5.125	0.934
	4.214	4.978	6.872	6.051	1.674
	4.985	5.541	7.085	6.452	2.971
FID	0.587/2.584	6.542/7.255	3.214/6.211	0.610/2.714	0.384/0.884
	1.524/3.450	7.225/8.254	3.5124/7.986	1.226/3.274	0.571/2.253
	2.269/4.804	10.542/12.457	3.303/9.524	2.957/4.545	0.694/2.990
	3.542/5.274	13.275/17.681	4.656/12.865	4.033/5.125	1.250/4.458
E_b	0.076/0.278	0.071/0.357	0.124/0.254	0.128/0.208	0.006/0.012
	0.142/0.305	0.122/0.402	0.361/0.537	0.146/0.409	0.015/0.071
	0.285/0.421	0.205/0.408	0.484/0.596	0.223/0.658	0.017/0.085
	0.435/0.527	0.321/0.568	0.590/0.605	0.338/0.763	0.022/0.206
JPD	4.436	4.258	5.454	4.954	0.561
	5.452	5.751	6.592	6.334	1.259
	5.494	6.006	6.881	8.454	1.903
	6.899	8.158	7.461	9.046	2.842

Table 15. Comparison on Chandelle retargeting (top) and generation (bottom). XX/XX are Character A/B results. All results are per joint results. The four rows in each cell are results of Random, Cross-scale, Cross-interaction and Cross-scale-interaction respectively.

Training samples		E_r	E_b	JPD	FID	E_b	JPD
M2	24	1.158	0.142	0.892	3.485	0.167	1.428
	12	1.347	0.186	0.963	3.676	0.192	1.667
	6	1.657	0.224	1.305	3.983	0.258	2.017
M7	24	5.861	0.082	3.035	1.897	0.094	3.923
	12	6.025	0.104	3.879	1.957	0.123	4.343
	6	6.254	0.173	4.241	2.124	0.205	4.587

Table 16. Result with limited training samples. Here is the result of Face-to-back (M2) and Big-ben (M7).

in the video.

2.5. Extrapolating to Large Unseen Scales

There is one example of Turn-around on 0.65 and 1.3 in the Fig. 12, which shows that our model can extrapolate to larger skeletal variations when trained only using data on scales [0.95, 1.05]. More examples can be found in the video.

3. Methodology Details

3.1. ST-GCN Layers

Spatio-temporal Graph Convolutions (ST-GCNs) are widely used in analyzing human motions. Our construction of it is inspired by [29]. Given $q = \{q^0, \dots, q^T\} \in \mathbb{R}^{T \times N \times 3}$, where T is frame

number of a motion, N is the number of joints and each joint location is represented by its 3D coordinates, we first construct a graph adjacency matrix $A_n \in \mathbb{R}^{n \times n}$ of the skeleton, indicating the connectivity between joints. The spatial graph convolution of a layer can be represented as:

$$X_{i+1}^t = \text{ReLU}(A_n X_i^t W_i + X_i^t U_i) \in \mathbb{R}^{n \times h_i} \quad (9)$$

where the subscript of X is the layer index, t is a frame and h_i is the latent dimension of the layer. W_i and U_i are trainable network weights. Further the temporal convolution can be achieved by using standard 2D convolution on X . In addition, we also add one Batch Normalization layer and a ReLU layer before the 2D convolution and one more Batch Normalization layer and one Dropout layer after the 2D convolution. After combining the spatial and temporal




Figure 12. Large-scale extrapolation results. The skeleton of the red character is changed. The motion is Hold-body on scale 0.65 (top), original scale (mid) and scale 1.3 (bottom).

convolution, we have one ST-GCN layer.

3.2. G-GRU Layers

Graph Gated Recurrent Unit Network, or G-GRU is based on standard GRU network [6], which is a Recurrent Neural Network which can model time-series data. Traditional GRU networks do not consider structured data such as graphs. A combination of GRU and Graph Neural Network can overcome this shortcoming [29]:

$$\begin{aligned}
 r^t &= \sigma(r_{input}(X^t)) + r_{hidden}(A_s H^t W), \\
 u^t &= \sigma(u_{input}(X^t)) + u_{hidden}(A_s H^t W), \\
 c^t &= \tanh(c_{input}(X^t)) + r^t \odot c_{hidden}(A_s H^t W), \\
 H^{t+1} &= u^t H^t + (1 - u^t) \odot c^t \quad (10)
 \end{aligned}$$

where r_{input} , u_{input} , c_{input} , r_{hidden} , u_{hidden} and c_{hidden} are trainable functions. X^t is the input, H^t is the hidden state at t and W is trainable weights. A_s is the adjacency matrix.

3.3. Network Implementation and Training Details

The network implementation details of ST-GCN1 and G-GRU1, including network layer configurations and architecture, are shown in Tab. 17 and Tab. 18. The network details of ST-GCN2, ST-GCN3 and G-GRU2 are shown in Tab. 19 - Tab. 22.

For training, we use a batch size 32 and Adam as the optimizer (learning rate = 0.001) for all our experiments. We train our model on a Nvidia Geforce RTX2080 Ti Graphics Card. The average training time for different models is 243

minutes with training epoch = 50, and the inference time = 0.323s per motion.

4. Alternative Architectures

We use a frame-based Convolution Neural Networks (CNNs) and a frame-based Graph Convolution Networks (GCNs) as the encoders (MLP1, ST-GCN1-3) and decoders (MLP2, G-GRU1-2) in all three VAEs denoted as F-CNNs and F-GCNs. In addition, we also use motion-based CNNs (M-CNNs) and GCNs (M-GCNs). The M-CNNs follow the architecture in [21]. For M-GCNs, we mirror the GCN encoders in ST-GCN1, ST-GCN2 and ST-GCN3, and use them as the decoders. Due to the limited data, we did not choose architectures that require large amounts of data such as Transformers, Flows or Diffusion models.

Totally, there are four baseline networks: Frame-based CNNs (F-CNNs), Frame-based GCNs (F-GCNs), Motion-based CNNs (M-CNNs) and Motion-based GCNs (M-GCNs). The detailed architectures of them are given in Tab. 23 - Tab. 26, respectively. Numerically, our current setting significantly outperforms all the other alternatives by as much as 66.99% in E_r , 49.42% in E_b (retargeting), 56.25% in JPD (retargeting), 72.17% in FID, 74.82% in E_b (generation) and 61.32% in JPD (generation).

Layer index	Output channels	Dimension	Layer	Stride
Input	/	[32,T,n,4]	/	/
1	32	[32,T,n,32]	ST-GCN	1
2	64	[32,T/2,n,64]	ST-GCN	2
3	128	[32,T/4,n,128]	ST-GCN	2
4	256	[32,T/8,n,256]	ST-GCN	2
5	256	[32,T/8,n,256]	ST-GCN	1
6	256	[32,1,n,256]	Temporal Averaging	/
7	262	[32,1,n,262]	Concatenation with \hat{q}_B^0 and \hat{q}_B^T	/
8	256	[32,1,n,256]	Dense	/

Table 17. Detailed architecture of ST-GCN1. T is the motion length. n is the number of joints.

Layer Index	Input	Dimension	Layer
1	Hidden state at time t	[32,1,n,256]	/
2	$B_s, \hat{q}_B^0, \hat{q}_B^T, \text{ and } \Delta \bar{q}_B^t$	[32,1,n,10]	Concatenation
3	output of 1, 2	[32,1,n,256]	G-GRU
4	output of 3	[32, 1, n, 256]	Dense
5	output of 4	[32, 1, n, 256]	Dense
6	output of 5	[32,1,n,3]	Dense

Table 18. Detailed architecture of G-GRU1. It takes as input z, \hat{q}_B^0 and \hat{q}_B^T and outputs $\Delta \bar{q}_B$. n is the number of joints.

Layer Index	Output channels	Dimension	Layer	Stride
Input	/	[32,T,n,3]	/	/
1	32	[32,T,n,32]	ST-GCN	1
2	64	[32,T/2,n,64]	ST-GCN	2
3	128	[32,T/4,n,128]	ST-GCN	2
4	256	[32,T/8,n,256]	ST-GCN	2
5	256	[32,T/8,n,256]	ST-GCN	1
6	256	[32,1,n,256]	Temporal Averaging	/

Table 19. Detailed architecture of ST-GCN2. T is the motion length and n is the number of joints.

Layer Index	Output channels	Dimension	Layer	Stride
Input	/	[32,T,n,8]	/	/
1	16	[32,T,n,16]	ST-GCN	1
2	16	[32,T/2,n,16]	ST-GCN	2
3	16	[32,T/4,n,16]	ST-GCN	2
4	16	[32,T/8,n,16]	ST-GCN	2
5	16	[32,T/8,n,16]	ST-GCN	1
6	16	[32,1,n,16]	Temporal Averaging	/

Table 20. Detailed architecture of ST-GCN3. T is the motion length and n is the number of joints.

Layer Index	Output channels	Dimension	Layer	Stride
Input	278	[32,1,n,278]	Concatenation of outputs from the Δq_A and q'_B branches, \hat{q}_A^0 and \hat{q}_A^T	/
1	256	[32,1,n,256]	Dense	/
2	256	[32,1,n,256]	Dense	/

Table 21. Detailed architecture after the ST-GCN2 and ST-GCN3. n is the number of joints. The network finally outputs z .

Layer Index	Input	Dimension	Layer
1	Hidden state at time t	$[32,1,n,256]$	/
2	encoded $q'_B, \hat{q}_A^0, \hat{q}_A^T$, and $\Delta \bar{q}_A^t$	$[32,1,n,10]$	Concatenation
3	output of 1, 2	$[32,1,n,256]$	G-GRU
4	output of 3	$[32, 1, n, 256]$	Dense
5	output of 4	$[32, 1, n, 256]$	Dense
6	output of 5	$[32,1,n,3]$	Dense

Table 22. Detailed architecture of G-GRU2. It takes as the first input z , encoded q'_B, \hat{q}_A^0 and \hat{q}_A^T and outputs $\Delta \bar{q}_A$. n is the number of joints.

Index	Output channels	Feature Shape	Operation	Stride
Input	/	$[32,n,3]$	/	/
1	32	$[32,n,32]$	Conv	1
2	64	$[32,n/2,64]$	Conv and Maxpooling	1
3	128	$[32,n/4,128]$	Conv and Maxpooling	1
4	256	$[32,n/8,256]$	Conv and Maxpooling	1
5	260	$[32,n/8,260]$	Concatenate B_s and \hat{q}_B	/
6	256	$[32,n/8,256]$	Dense	/
Index	Output channels	Feature Shape	Operation	Stride
7	/	$[32,n/8,260]$	Concatenate B_s and \hat{q}_B	/
8	256	$[32,n/8,256]$	Dense	/
9	256	$[32,n/4,256]$	ConvTranspose	2
10	128	$[32,n/2,128]$	ConvTranspose	2
11	32	$[32,n,32]$	ConvTranspose	2
Output	3	$[32,n,3]$	Dense	/
Index	Output channels	Feature Shape	Operation	Stride
Input	/	$[32,n,3]$	/	/
1	32	$[32,n,32]$	Conv	1
2	64	$[32,n/2,64]$	Conv and Maxpooling	1
3	128	$[32,n/4,128]$	Conv and Maxpooling	1
4	256	$[32,n/8,256]$	Conv and Maxpooling	1
5	264	$[32,n/8,264]$	Concatenate encoding \hat{q}_A and q'_B	1
6	256	$[32,n/8,256]$	Dense	/
Index	Output channels	Feature Shape	Operation	Stride
7	/	$[32,n/8,264]$	Concatenate encoding \hat{q}_A and q'_B	/
8	256	$[32,n/8,256]$	Dense	/
9	128	$[32,n/4,128]$	ConvTranspose	2
10	64	$[32,n/2,64]$	ConvTranspose	2
11	32	$[32,n,32]$	ConvTranspose	2
Output	3	$[32,n,3]$	Dense	/

Table 23. F-CNNs detailed architecture in Character B (top) and Character A (bottom)

Index	Output channels	Feature Shape	Operation	Stride
Input	/	[32,n,3]	/	/
1	32	[32,n,32]	GCN	1
2	64	[32,n,64]	GCN	1
3	84	[32,n,84]	Concatenate encoding B_s and \hat{q}_B	1
4	128	[32,n,128]	GCN	1
5	256	[32,n,256]	GCN	1
Index	Output channels	Feature Shape	Operation	Stride
6	/	[32,n,276]	Concatenate encoding B_s and \hat{q}_B	/
7	256	[32,n,256]	GCN	1
8	128	[32,n,128]	GCN	1
9	64	[32,n,64]	GCN	1
10	32	[32,n,32]	GCN	1
11	3	[32,n,3]	GCN	1
Output	3	[32,n,3]	Dense	1
Index	Output channels	Feature Shape	Operation	Stride
Input	/	[32,n,3]	/	/
1	32	[32,n,32]	GCN	1
2	64	[32,n,64]	GCN	1
3	80	[32,n,80]	Concatenate encoding q'_B and \hat{q}_A	1
4	128	[32,n,128]	GCN	1
5	256	[32,n,256]	GCN	1
Index	Output channels	Feature Shape	Operation	Stride
6	/	[32,n,272]	Concatenate encoding q'_B and \hat{q}_A	/
7	256	[32,n,256]	GCN	1
8	128	[32,n,128]	GCN	1
9	64	[32,n,64]	GCN	1
10	32	[32,n,32]	GCN	1
11	3	[32,n,3]	GCN	1
Output	3	[32,n,3]	Dense	/

Table 24. F-GCNs detailed architecture in Character B (top) and Character A (bottom)

Index	Output channels	Feature Shape	Operation	Stride
Input	/	[32,T,n,4]	/	/
1	16	[32,T,n,16]	Conv and Maxpooling	1
2	32	[32,T,n,32]	Conv and Maxpooling	1
3	64	[32,T,n,64]	Conv and Maxpooling	1
Index	Output channels	Feature Shape	Operation	Stride
4	/	[32,T,n,65]	Concatenate B_s	/
5	64	[32,T,n,64]	Dense	/
6	32	[32,T,n,32]	ConvTranspose	1
7	16	[32,T,n,16]	ConvTranspose	1
Output	3	[32,T,n,3]	ConvTranspose	1
Output	3	[32,T,n,3]	Dense	/
Index	Output channels	Feature Shape	Operation	Stride
Input	/	[32,T,n,16]	Concatenate encoding \hat{q}_A and q'_B	/
1	32	[32,T,n,32]	Conv and Maxpooling	1
2	64	[32,T,n,64]	Conv and Maxpooling	1
Index	Output channels	Feature Shape	Operation	Stride
3	/	[32,T,n,72]	Concatenate encoding \hat{q}_A and q'_B	/
4	64	[32,T,n,64]	Dense	/
5	32	[32,T,n,32]	ConvTranspose	1
6	16	[32,T,n,16]	ConvTranspose	1
Output	3	[32,T,n,3]	Dense	/

Table 25. M-CNNs detailed architecture in Character B (top) and Character A (bottom)

Index	Output channels	Feature Shape	Operation	Stride
Input	/	[32,T,n,4]	/	/
1	32	[32,T,n,32]	ST-GCN	1
2	64	[32,T,n,64]	ST-GCN	1
3	128	[32,T,n,128]	ST-GCN	1
4	128	[32,T,n,128]	Dense	1
Index	Output channels	Feature Shape	Operation	Stride
5	/	[32,T,n,129]	Concatenate B_s	/
6	128	[32,T,n,128]	ST-GCN	1
7	64	[32,T,n,64]	ST-GCN	1
8	32	[32,T,n,32]	ST-GCN	1
9	16	[32,T,n,16]	ST-GCN	1
Output	3	[32,T,n,3]	ST-GCN	1
Index	Output channels	Feature Shape	Operation	Stride
Input	/	[32,T,n,3]	/	/
1	32	[32,T,n,32]	ST-GCN	1
2	64	[32,T,n,64]	ST-GCN	1
3	128	[32,T,n,128]	ST-GCN	1
4	144	[32,T,n,144]	Concatenate encoding q'_B	1
5	128	[32,T,n,128]	Dense	1
Index	Output channels	Feature Shape	Operation	Stride
6	/	[32,T,n,144]	Concatenate encoding q'_B	/
7	128	[32,T,n,128]	ST-GCN	1
8	64	[32,T,n,64]	ST-GCN	1
9	32	[32,T,n,32]	ST-GCN	1
Output	3	[32,T,n,3]	Dense	/

Table 26. M-GCNs detailed architecture in Character B (top) and Character A (bottom)

References

- [1] CMU graphics lab motion capture database. <http://mocap.cs.cmu.edu>, accessed 2022. 5
- [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. 1, 2
- [3] Mazen Al Borno, Ludovic Righetti, Michael J. Black, Scott L. Delp, Eugene Fiume, and Javier Romero. Robust physics-based motion retargeting with realistic body shapes. *Computer Graphics Forum*, 37(8):81–92, 2018. 1, 2
- [4] Jinxiang Chai and Jessica K. Hodgins. Performance animation from low-dimensional control signals. *ACM Trans. Graph.*, 24(3):686–696, 2005. 1, 2
- [5] Wenheng Chen, He Wang, Yi Yuan, Tianjia Shao, and Kun Zhou. *Dynamic Future Net: Diversified Human Motion Generation*, page 2131–2139. Association for Computing Machinery, New York, NY, USA, 2020. 2
- [6] Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014. 8
- [7] Baptiste Chopin, Hao Tang, Naima Otberdout, Mohamed Daoudi, and Nicu Sebe. Interaction transformer for human reaction generation. *IEEE Transactions on Multimedia*, pages 1–13, 2023. 2
- [8] C. Coppola, S. Cosar, D. Faria, and N. Bellotto. Automatic detection of human interactions from rgb-d data for social activity classification. In *IEEE Int. Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 871–876, 2017. 1, 3
- [9] Yunfeng Diao, Tianjia Shao, Yongliang Yang, Kun Zhou, and He Wang. Basar:black-box attack on skeletal action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 4
- [10] Anindita Ghosh, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Philipp Slusallek. Remos: 3d motion-conditioned reaction synthesis for two-person interactions, 2024. 2
- [11] Michael Gleicher. Retargeting motion to new characters. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, page 33–42, New York, NY, USA, 1998. Association for Computing Machinery. 2
- [12] Aman Goel, Qianhui Men, and Edmond S. L. Ho. Interaction Mix and Match: Synthesizing Close Interaction using Conditional Hierarchical GAN with Multi-Hot Class Embedding. *Computer Graphics Forum*, 2022. 2
- [13] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2021–2029, 2020. 7, 8
- [14] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5152–5161, 2022. 2
- [15] Wen Guo, Xiaoyu Bie, Xavier Alameda-Pineda, and Francesc Moreno-Noguer. Multi-person extreme motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13053–13064, 2022. 1, 5, 8
- [16] Chris Hecker, Bernd Raabe, Ryan W. Enslow, John DeWeese, Jordan Maynard, and Kees van Prooijen. Real-time motion retargeting to highly varied user-created morphologies. *ACM Trans. Graph.*, 27(3):1–11, 2008. 1, 2
- [17] Edmond S.L. Ho and Taku Komura. Indexing and retrieving motions of characters in close contact. *IEEE Transactions on Visualization and Computer Graphics*, 15(3):481–492, 2009. 1, 2, 3
- [18] Edmond S. L. Ho and Hubert P. H. Shum. Motion adaptation for humanoid robots in constrained environments. In *2013 IEEE International Conference on Robotics and Automation*, pages 3813–3818, 2013. 1, 2
- [19] Edmond S. L. Ho, Taku Komura, and Chiew-Lan Tai. Spatial relationship preserving character motion adaptation. *ACM Trans. Graph.*, 29(4), 2010. 1, 2, 5, 7
- [20] Edmond S. L. Ho, He Wang, and Taku Komura. A multi-resolution approach for adapting close character interaction. In *Proceedings of the 20th ACM Symposium on Virtual Reality Software and Technology*, page 97–106, New York, NY, USA, 2014. Association for Computing Machinery. 2
- [21] Daniel Holden, Jun Saito, and Taku Komura. A deep learning framework for character motion synthesis and editing. *ACM Trans. Graph.*, 35(4), 2016. 2, 8
- [22] Lianyu Hu, Shenglan Liu, and Wei Feng. Spatial temporal graph attention network for skeleton-based action recognition. *arXiv preprint arXiv:2208.08599*, 2022. 8
- [23] 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 *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 13274–13285, 2022. 2
- [24] Nan Jiang, Tengyu Liu, Zhexuan Cao, Jieming Cui, Yixin Chen, He Wang, Yixin Zhu, and Siyuan Huang. Full-body articulated human-object interaction. In *ICCV*, 2023. 2
- [25] Jongmin Kim, Yeongho Seol, and Taesoo Kwon. Interactive multi-character motion retargeting. *Computer Animation and Virtual Worlds*, 32(3-4):e2015, 2021. 2
- [26] Jungho Lee, Minhyeok Lee, Dogyoon Lee, and Sangyoon Lee. Hierarchically decomposed graph convolutional networks for skeleton-based action recognition. *arXiv preprint arXiv:2208.10741*, 2022. 8
- [27] Kyungho Lee, Seyoung Lee, and Jehee Lee. Interactive character animation by learning multi-objective control. *ACM Trans. Graph.*, 37(6), 2018. 2
- [28] Sunmin Lee, Taeho Kang, Jungnam Park, Jehee Lee, and Jungdam Won. Same: Skeleton-agnostic motion embedding for character animation. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–11, 2023. 2

- [29] Maosen Li, Siheng Chen, Yangheng Zhao, Ya Zhang, Yanfeng Wang, and Qi Tian. Dynamic multiscale graph neural networks for 3d skeleton based human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 214–223, 2020. 3, 4, 7, 8
- [30] Han Liang, Wenqian Zhang, Wenxuan Li, Jingyi Yu, and Lan Xu. Intergen: Diffusion-based multi-human motion generation under complex interactions. *International Journal of Computer Vision*, 2024. 2, 3
- [31] Hung Yu Ling, Fabio Zinno, George Cheng, and Michiel van de Panne. Character controllers using motion vaes. *ACM Trans. Graph.*, 39(4), 2020. 3
- [32] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015. 1
- [33] Qianhui Men, Edmond S. L. Ho, Hubert P. H. Shum, and Howard Leung. A quadruple diffusion convolutional recurrent network for human motion prediction. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(9): 3417–3432, 2021. 2
- [34] Qianhui Men, Hubert P. H. Shum, Edmond S. L. Ho, and Howard Leung. Gan-based reactive motion synthesis with class-aware discriminators for human-human interaction. *Computers and Graphics*, 2021. 2
- [35] Lucas Mourot, Ludovic Hoyet, François Le Clerc, François Schnitzler, and Pierre Hellier. A survey on deep learning for skeleton-based human animation. *Computer Graphics Forum*, 41(1):122–157, 2022. 1
- [36] Meinard Müller and Tido Röder. Motion templates for automatic classification and retrieval of motion capture data. In *Proceedings of the 2006 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, page 137–146, Goslar, DEU, 2006. Eurographics Association. 1
- [37] M. Müller, T. Röder, M. Clausen, B. Eberhardt, B. Krüger, and A. Weber. Documentation Mocap Database HDM05. Technical Report CG-2007-2, Universität Bonn, 2007. ISSN: 1610-8892. 1
- [38] Maryam Naghizadeh and Darren Cosker. Multi-character motion retargeting for large-scale transformations. In *Advances in Computer Graphics*, pages 94–106, Cham, 2019. Springer International Publishing. 2
- [39] Xiaogang Peng, Siyuan Mao, and Zizhao Wu. Trajectory-aware body interaction transformer for multi-person pose forecasting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17121–17130, 2023. 3
- [40] Maksym Peregichka, Daniel Holden, Sudhir P. Mudur, and Tiberiu Popa. Robust marker trajectory repair for mocap using kinematic reference. In *Motion, Interaction and Games*, New York, NY, USA, 2019. Association for Computing Machinery. 1
- [41] Mathis Petrovich, Michael J. Black, and Gül Varol. Action-conditioned 3D human motion synthesis with transformer VAE. In *Proc. International Conference on Computer Vision (ICCV)*, pages 10985–10995, 2021. 2, 7, 8
- [42] Huaijin Pi, Sida Peng, Minghui Yang, Xiaowei Zhou, and Hujun Bao. Hierarchical generation of human-object interactions with diffusion probabilistic models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15061–15073, 2023. 2
- [43] M. S. Ryoo and J. K. Aggarwal. UT-Interaction Dataset, ICPR contest on Semantic Description of Human Activities (SDHA). http://cvrc.ece.utexas.edu/SDHA2010/Human_Interaction.html, 2010. 1, 3
- [44] Amir Shahrudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. NTU RGB+D: A large scale dataset for 3D human activity analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1010–1019, 2016. 1
- [45] Yijun Shen, Longzhi Yang, Edmond S. L. Ho, and Hubert P. H. Shum. Interaction-based human activity comparison. *IEEE Transactions on Visualization and Computer Graphics*, 26(8):115673–115684, 2020. 3
- [46] Hubert P. H. Shum, Edmond S. L. Ho, Yang Jiang, and Shu Takagi. Real-time posture reconstruction for microsoft kinect. *IEEE Transactions on Cybernetics*, 43(5):1357–1369, 2013. 2
- [47] Sebastian Starke, Yiwei Zhao, Fabio Zinno, and Taku Komura. Neural animation layering for synthesizing martial arts movements. *ACM Trans. Graph.*, 40(4), 2021. 2
- [48] Zehua Sun, Qihong Ke, Hossein Rahmani, Mohammed Bennamoun, Gang Wang, and Jun Liu. Human action recognition from various data modalities: A review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–20, 2022. 1
- [49] Seyoon Tak and Hyeong-Seok Ko. A physically-based motion retargeting filter. *ACM Trans. Graph.*, 24(1):98–117, 2005. 1, 2
- [50] Mikihiro Tanaka and Kent Fujiwara. Role-aware interaction generation from textual description. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15999–16009, 2023. 2
- [51] Xiangjun Tang, Linjun Wu, He Wang, Bo Hu, Xu Gong, Yuchen Liao, Songnan Li, Qilong Kou, and Xiaogang Jin. Rsmt: Real-time stylized motion transition for characters. In *ACM SIGGRAPH*, pages 1–10, 2023. 4
- [52] Julian Tanke, Linguang Zhang, Amy Zhao, Chengcheng Tang, Yujun Cai, Lezi Wang, Po-Chen Wu, Juergen Gall, and Cem Keskin. Social diffusion: Long-term multiple human motion anticipation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9601–9611, 2023. 3
- [53] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Amit H Bermano, and Daniel Cohen-Or. Human motion diffusion model. *arXiv preprint arXiv:2209.14916*, 2022. 2
- [54] Seong Uk Kim, Hanyoung Jang, and Jongmin Kim. A variational u-net for motion retargeting. *Computer Animation and Virtual Worlds*, 31(4-5):e1947, 2020. 1, 2
- [55] Ruben Villegas, Duygu Ceylan, Aaron Hertzmann, Jimei Yang, and Jun Saito. Contact-aware retargeting of skinned motion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9720–9729, 2021. 2

- [56] He Wang, Edmond SL Ho, and Taku Komura. An energy-driven motion planning method for two distant postures. *IEEE transactions on visualization and computer graphics*, 21(1):18–30, 2015. [1](#)
- [57] He Wang, Feixiang He, Zhexi Peng, Tianjia Shao, Yongliang Yang, Kun Zhou, and David Hogg. Understanding the robustness of skeleton-based action recognition under adversarial attack. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [1](#), [4](#)
- [58] He Wang, Edmond S. L. Ho, Hubert P. H. Shum, and Zhanxing Zhu. Spatio-temporal manifold learning for human motions via long-horizon modeling. *IEEE Transactions on Visualization and Computer Graphics*, 27(1):216–227, 2021. [2](#), [4](#)
- [59] He Wang, Yunfeng Diao, Zichang Tan, and Guodong Guo. Defending black-box skeleton-based human activity classifiers, 2022. [1](#)
- [60] Jingbo Wang, Yu Rong, Jingyuan Liu, Sijie Yan, Dahua Lin, and Bo Dai. Towards diverse and natural scene-aware 3d human motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20460–20469, 2022. [1](#)
- [61] Shengqin Wang, Yongji Zhang, Fenglin Wei, Kai Wang, Minghao Zhao, and Yu Jiang. Skeleton-based action recognition via temporal-channel aggregation. *arXiv preprint arXiv:2205.15936*, 2022. [8](#)
- [62] Mao Wei, Liu Miaomiao, and Salzemann Mathieu. History repeats itself: Human motion prediction via motion attention. In *ECCV*, 2020. [8](#)
- [63] Yu-Hui Wen, Zhipeng Yang, Hongbo Fu, Lin Gao, Yanan Sun, and Yong-Jin Liu. Autoregressive stylized motion synthesis with generative flow. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13607–13607, 2021. [2](#)
- [64] Jungdam Won, Deepak Gopinath, and Jessica Hodgins. Control strategies for physically simulated characters performing two-player competitive sports. *ACM Transactions on Graphics (TOG)*, 40(4):1–11, 2021. [2](#)
- [65] Liang Xu, Yizhou Zhou, Yichao Yan, Xin Jin, Wenhan Zhu, Fengyun Rao, Xiaokang Yang, and Wenjun Zeng. Regennet: Towards human action-reaction synthesis, 2024. [2](#)
- [66] Sirui Xu, Zhengyuan Li, Yu-Xiong Wang, and Liang-Yan Gui. Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In *ICCV*, 2023. [2](#)
- [67] Sirui Xu, Yu-Xiong Wang, and Liang-Yan Gui. Stochastic multi-person 3d motion forecasting. In *ICLR*, 2023. [3](#)
- [68] Zhuoqian Yang, Wentao Zhu, Wayne Wu, Chen Qian, Qiang Zhou, Bolei Zhou, and Chen Change Loy. Transmomo: Invariance-driven unsupervised video motion retargeting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#)
- [69] Hongwei Yi, Chun-Hao P. Huang, Dimitrios Tzionas, Muhammed Kocabas, Mohamed Hassan, Siyu Tang, Justus Thies, and Michael J. Black. Human-aware object placement for visual environment reconstruction. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3959–3970, 2022. [2](#)
- [70] Kiwon Yun, Jean Honorio, Debaleena Chattopadhyay, Tamara L. Berg, and Dimitris Samaras. Two-person interaction detection using body-pose features and multiple instance learning. In *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 28–35, 2012. [3](#)
- [71] Haodong Zhang, Weijie Li, Jiangpin Liu, Zexi Chen, Yuxiang Cui, Yue Wang, and Rong Xiong. Kinematic motion retargeting via neural latent optimization for learning sign language. *IEEE Robotics and Automation Letters*, 7(2):4582–4589, 2022. [2](#)
- [72] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Shaoli Huang, Yong Zhang, Hongwei Zhao, Hongtao Lu, and Xi Shen. T2m-gpt: Generating human motion from textual descriptions with discrete representations, 2023. [2](#)
- [73] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiandiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. [2](#)
- [74] Yunbo Zhang, Deepak Gopinath, Yuting Ye, Jessica Hodgins, Greg Turk, and Jungdam Won. Simulation and retargeting of complex multi-character interactions. In *ACM SIGGRAPH 2023 Conference Proceedings*, New York, NY, USA, 2023. Association for Computing Machinery. [2](#)