

# Bipartite Graph Diffusion Model for Human Interaction Generation

Baptiste Chopin<sup>1</sup>, Hao Tang<sup>2</sup>, Mohamed Daoudi<sup>3,4</sup>  
Inria, Université Côte d’Azur, France <sup>1</sup>,  
Carnegie Mellon University, Pittsburgh, USA<sup>2</sup>

Univ. Lille, CNRS, Centrale Lille, Institut Mines-Télécom, UMR 9189 CRISAL, Lille, F-59000, France<sup>3</sup>,  
IMT Nord Europe, Institut Mines-Télécom, Univ. Lille, Centre for Digital Systems, Lille, F-59000, France<sup>4</sup>

## Abstract

*The generation of natural human motion interactions is a hot topic in computer vision and computer animation. It is a challenging task due to the diversity of possible human motion interactions. Diffusion models, which have already shown remarkable generative capabilities in other domains, are a good candidate for this task. In this paper, we introduce a novel bipartite graph diffusion method (BiGraphDiff) to generate human motion interactions between two persons. Specifically, bipartite node sets are constructed to model the inherent geometric constraints between skeleton nodes during interactions. The interaction graph diffusion model is transformer-based, combining some state-of-the-art motion methods. We show that the proposed achieves new state-of-the-art results on leading benchmarks for the human interaction generation task. Code, pre-trained models and additional results are available at <https://github.com/CRISTAL-3DSAM/BiGraphDiff>.*

## 1. Introduction

Modeling the dynamics of human motion is at the core of many applications in computer vision and computer graphics. However, the production of real human motion requires sophisticated equipment (e.g., expensive motion capture systems) and domain experts. In order to remove the skill requirements for users and potentially reach the general public, it is essential to create a human motion generation model capable of producing diverse motion sequences. Human motion generation is an active research area in robotics, computer graphics, and machine learning. It involves developing algorithms and techniques that enable machines to generate human-like motion for various applications such as animation, virtual reality, and robotics. Ongoing research efforts in this field aim to improve the quality, realism, and efficiency of motion generation methods, paving the way for innovative applications and breakthroughs in related fields.

Several approaches have been proposed to predict or

generate human motion for different tasks. Zhou *et al.* [42] proposed an LSTM-based approach to generate long human motions. However the method suffers from error accumulation, a common problem with LSTMs. The authors also reported that their generative process sometimes converges toward a mean pose and stops generating motion. To avoid these issues some approaches used convolutional neural networks. Kosmas *et al.* [18] use CNNs to create a generative model for dance motions from music. However, using the samples generated by the method suffer from shaking caused by limited temporal coherency. More recently ACTOR [26], a Transformer-based variational autoencoder (VAE) was designed to generate human motion based on the label. This approach limits the learned distribution to normal latent distribution since it mainly uses VAEs [17]. From this point of view, diffusion models are a better candidate for human motion generation, as they are free of assumptions about the target distribution. Motion-Diffuse [41] is a diffusion model-based text-driven motion generation framework. EDGE [36] uses a diffusion-based model that generates dance sequences conditioned on music. Tevet *et al.* [35] proposed a motion diffusion model (MDM) for generating human motion animation given an arbitrary condition or no condition. However, most works on human motion generation ignore human interactions and focus instead on the generation of motion of a single person. However, Human motion interaction modeling is a key component of video understanding and is indispensable because it is frequently observed in real video. It has various applications, such as a human-computer interface, sports, dance, game development, and an understanding of human behavior, and prediction. Generating interaction instead of single-person motion is also more challenging. Indeed, in addition to modeling the spatial and temporal of each skeleton for single-person motion generation, we need to model the interaction between the two skeletons. The motion of each skeleton influences the motion of the other. This is especially challenging since the interaction must be temporally coherent and require the motion of both persons to be

synchronous if we want the interaction to be realistic.

In this paper, we address the challenge of generating high-quality 3D human motion interactions. The complexity lies in the nonlinear nature of human motion interactions and the diverse range of interactions between individuals. To tackle these challenges, we explore several key questions: how to effectively represent interactions between humans and how to model and generate diverse motion interactions. To address these questions, we propose *BiGraphDiff*, a graph Transformer denoising diffusion model that leverages the Transformer architecture. This approach enables us to generate diverse motion interactions while capturing the temporal dependencies of skeleton motion through the use of attention mechanisms. Additionally, the diffusion process incorporated in *BiGraphDiff* facilitates the generation of high-quality and diverse motions. In *BiGraphDiff*, we represent skeleton interactions using a bipartite graph [34]. The bipartite graph serves as a means to capture the interactions between humans, with each human represented by a skeleton. By formulating the motion interaction generation as a reverse diffusion process, we can effectively generate realistic and diverse motion interactions. Overall, our contributions are summarized as follows:

- We propose the first Bipartite graph denoising diffusion model (*BiGraphDiff*) for human interaction generation. Our *BiGraphDiff* is able to generate motion interaction in a stochastic way, naturally leading to high diversity, and is able to generate very long motion sequences (>1000 frames).
- *BiGraphDiff* is a denoising diffusion process that learns not only the denoising of the motion but also it learns a Bipartite graph. The Bipartite graph aims to capture the relations between the two persons.
- *BiGraphDiff* achieves state-of-the-art quantitatively and qualitatively in action interaction and dance tasks. A user study shows that the generated sequences are better qualitatively than the sequences generated by state-of-the-art methods.

## 2. Related Work

We discuss the relevant literature from two perspectives, namely, previous methods of Human interaction motion synthesis and the literature on diffusion models.

**Human Interaction Motion Generation.** Recently there has been an increase in motion generation based on different modalities, [40] use control signals such as the global trajectory of the person to generate human motion in long-term horizons while [1] and [10] generate motion based on speech audio. Meanwhile, others use only knowledge of the past motion which allows them to work in real-time but on shorter motion [7, 25, 32]. More recently, several works have been dedicated to human pose and motion generation from text or action labels, as well as its reciprocal

task [5, 9, 24, 28]. These papers focus only on one person, while our approach is dedicated to the generation of two-person interactions. [4] propose a multimodal variational recurrent neural network to predict the future motion of both participants in an interaction based on past sequences of motion. In contrast, we propose to generate human interaction between two persons.

**Generative Diffusion Models.** Diffusion models [14, 33] have shown great promise in terms of generative modeling by showing impressive results in synthesis applications ranging from image generation [11], audio-drive motion synthesis [2], molecule generation [16], to text-driven motion generation [29]. More recently, some concurrent work in the field of text-to-motion introduces a diffusion-based method for generating text-conditioned motion. For example, Zhang *et al.* [41] propose *MotionDiffuse*, a diffusion model-based text-driven motion generation framework. Tseng *et al.* [36] propose *EDGE*, a method for generating editable dances that is able to create a realistic dance while remaining faithful to the original music. Dabral *et al.* [8] introduce *MoFusion*, a denoising-diffusion-based framework for high-quality conditional human motion synthesis that can generate long and temporally plausible motions conditioned based on music or text. Despite achieving impressive performance, these methods use a diffusion-based method for generating the motion of only one person. In contrast, our proposed method *BiGraphDiff* proposes to generate the interaction between two persons and propose to learn a bipartite graph during the diffusion process. In addition, *BiGraphDiff* is applied for both text-to-motion and text-to-dance, and it is able to generate a long sequence of dance motions.

## 3. Background

Denoising diffusion models (DDMs) [33] have emerged as powerful generative models, achieving outstanding results not only on image synthesis [13] but also for video synthesis [31]. DDMs consist of two separate processes called forward diffusion and reverse diffusion. During the forward diffusion process, the data gradually is perturbed by Gaussian noise repeatedly until the data becomes Gaussian noise, while a neural model learns the reverse process of gradually denoising the sample.

Formally, the forward process on a real sample from a real data distribution  $x_0 \sim q(x)$  consists in a Markov chain that gradually adds noise following a variance schedule  $\beta_t$  to obtain the posterior  $q(x_{1:T}|x_0)$  with  $x_1$  to  $x_T$  the latent data:

$$\begin{aligned} q(x_{1:T}|x_0) &:= \prod_{t=1}^T q(x_t|x_{t-1}), \\ q(x_t|x_{t-1}) &:= \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}). \end{aligned} \tag{1}$$

The reverse diffusion process,  $p_\theta(x_{0:T})$ , is a Markov chain that eliminates the noise from  $x_T$  recursively until we obtain  $x_1$ . The reverse process can also be conditioned on an arbitrary condition  $c$ . With  $p(x_T) = \mathcal{N}(x_T; \mathbf{0}, \mathbf{I})$ :

$$p(x_{0:T}) := p(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t), \quad (2)$$

$$p_\theta(x_{t-1}|x_t) := \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t, c), \Sigma_\theta(x_t, t, c)).$$

During the denoising process the goal is to estimate  $\mu_\theta(x_t, t, c)$  and  $\Sigma_\theta(x_t, t, c)$ . Based on [15] we can set  $\Sigma_\theta(x_t, t, c) = \sigma_t^2 \mathbf{I}$  with  $\sigma_t$  a constant and replace  $\mu_\theta(x_t, t, c)$  as follow:

$$\mu_\theta(x_t, t, c) = \frac{1}{\sqrt{\alpha_t}} \left( x_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(x_t, t, c) \right), \quad (3)$$

this means that we only need to estimate  $\epsilon_\theta(x_t, t, c)$  to be able to denoise the latent data since we can recover  $x_{t-1}$  using:

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( x_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(x_t, t, c) \right) + \sigma_t \gamma, \quad (4)$$

with  $\gamma \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ . In our model we set  $\sigma_t = \log(\beta_t \frac{1 - \alpha_{t-1}}{1 - \alpha_t})$  following [15] recommendation. To estimate  $\epsilon_\theta(x_t, t, c)$ , we will train a Bipartite Graph Interaction Transformer (defined in section 4.2) to minimize the loss:

$$L := E_{t \in [1, T], x_0 \sim q(x_0), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\| \epsilon - \epsilon_\theta(x_t, t, c) \|^2] \\ := E_{t \in [1, T], x_0 \sim q(x_0), \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\| \epsilon - \epsilon_\theta(\sqrt{\alpha_t} x_0 + \epsilon \sqrt{1 - \alpha_t}, t, c) \|^2]. \quad (5)$$

## 4. Bipartite Graph Diffusion Model

### 4.1. Framework Overview

As shown in Figure 1, our goal is to generate a human motion interaction  $x^{1:N}$  given an arbitrary condition  $c$ . Let us consider  $x^{1:N} = \{x^1, \dots, x^N\}$  an arbitrary sequence of joints that compose the two skeletons,  $x^i \in \mathbb{R}^{k \times 3 \times 2}$ , where  $k$  is the number of joints (3 the number of dimensions, 2 the number of skeletons). The motion generation is formulated as a reverse diffusion process that requires sampling a random noise  $x_T^{1:N}$  from noise distribution to generate a motion sequence as explained in section 3. We propose Transformers to learn the denoising function and a bipartite graph to represent the relationship between the joints of the skeleton. The Bipartite Graph Interaction Transformer used by BiGraphDiff is based on the original Transformer [39]. It is composed of a text encoder, embedding and positional encoding layers, self-attention modules, cross-attention modules, a Bipartite graph module, feed-forward modules, and

a final linear layer. The architecture is composed of a text encoder to encode the class label and a motion decoder to learn the interaction motion. This decoder deals with each skeleton separately with self-attention and cross-attention then our bipartite graph network learns a graph that represents the interaction and encodes this information in each skeleton. The proposed Transformer learns not only the denoising function but also the bipartite graph.

### 4.2. Bipartite Graph Interaction Transformer

The text encoder is used to encode the class label  $c$ . We use a simple four layers Transformer encoder as described in [39] that uses multi-head self-attention. To avoid training the encoder from scratch, we initialize the weight with those of CLIP [27]. By using a text encoder instead of a simpler label encoder working with one hot vector, we allow our architecture to be directly used on more complex interaction datasets where motion is composed of several sub-motions when this kind of dataset will be available. The motion decoder uses  $x_t^{1:N}$  and the output of the text encoder to obtain  $\epsilon_\theta(x_t, t, c)$ . First we split  $x_t^{1:N}$  into  $x_{1,t}^{1:N}$  and  $x_{2,t}^{1:N}$  which represent the first and second skeleton, respectively. Each skeleton passes through an embedding layer followed by a positional encoding layer introduced by [39] that encodes the temporal information from each frame of the sequence. Then the data goes through self-attention and cross-attention layers. Attention is used to find correlations within the data. To reduce the complexity of attention layers, we use efficient attention [30]. Each attention layer (self and cross) is followed by a stylization block. This module, introduced by [41], allows the generative process to keep track of the current diffusion timestep  $t$  improving the generation. The output of this module is added to the input of the attention through a residual connection. We then use a Bipartite graph module to learn the interaction between the two skeletons. After the bipartite graph module, the data of each skeleton goes through a feed-forward network. It is composed of linear projections, dropout, and GELU activation functions. It is followed by a stylization block to ensure that the information about the current timestep is not lost. The output is added to the input of the feed-forward network thanks to a residual connection. The Motion decoder contains eight identical layers and the input of layer  $m$  is the output of layer  $m - 1$ . Following those eight layers, the data of the two skeletons are concatenated and goes through a final linear projection to obtain  $\epsilon_\theta(x_t, t, c)$  that we can use in our loss and to retrieve  $x_{t-1}^{1:N}$ .

#### 4.2.1 Bipartite Graph

Here we describe in more detail the Bipartite graph module used in the motion decoder. Following the self-attention and cross-attention module, we obtain embeddings for both

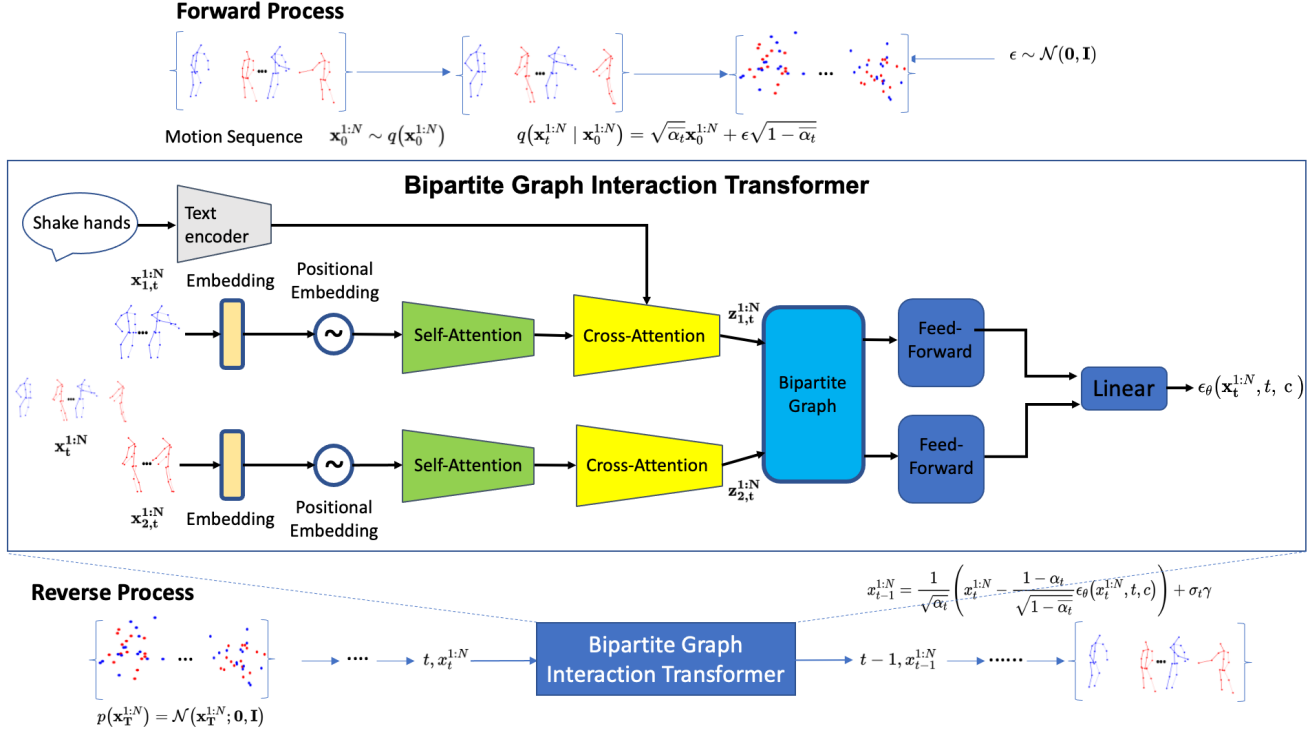


Figure 1. **BiGraphDiff overview.** Top: the forward diffusion process to add noise to the motion sequence. Middle: the proposed Bipartite Graph Interaction Transformer to learn the denoising function. Bottom reverse diffusion process to generate motion sequence from noise.




Figure 2. **Illustration of the proposed bipartite graph method.**

skeletons, i.e.,  $z_{1,t}^{1:N}$  and  $z_{2,t}^{1:N}$  (size  $N \times d_l$ ,  $d_l$  size of the embedding of a single skeleton). These embeddings go through the bipartite graph module. The proposed bipartite graph aims to capture the long-range cross relations between the two embeddings  $S_a = z_{1,t}^{1:N}$  and  $S_b = z_{2,t}^{1:N}$  in a bipartite graph via GCNs. Each node in  $S_a$  is connected to all the nodes in  $S_b$ , as shown in Figure 2.

Firstly,  $S_a$  and  $S_b$  are separately fed into two encoders to obtain the feature  $F_a$  and  $F_b$ , respectively. We then reduce the dimension of  $F_a$  with the function  $\varphi_a(F_a) \in \mathbb{R}^{C \times D_a}$ , where  $C$  is the number of feature map channels, and  $D_a$  is the number of nodes of  $F_a$ . Meanwhile, we reduce the dimension of  $F_b$  with the function  $\theta_b(F_b) = H_b^T \in \mathbb{R}^{D_b \times C}$ , where  $D_b$  is the number of nodes of  $F_b$ . Next, we project  $F_a$

to a new feature  $V_a$  in a bipartite graph using the projection function  $H_b^T$ . Thus we have:

$$V_a = H_b^T \varphi_a(F_a) = \theta_b(F_b) \varphi_a(F_a), \quad (6)$$

where both functions  $\theta_b(\cdot)$  and  $\varphi_a(\cdot)$  are implemented using a  $1 \times 1$  convolutional layer. This results in a new feature  $V_a \in \mathbb{R}^{D_b \times D_a}$  in the bipartite graph, which represents the cross relations between the nodes of the skeleton  $F_b$  and the skeleton  $F_a$ .

After projection, we employ a fully connected bipartite graph with adjacency matrix  $A_a \in \mathbb{R}^{D_b \times D_b}$ . We then use a graph convolution to learn the long-range cross relations between the nodes from both skeletons, which can be represented as:

$$M_a = (I - A_a) V_a W_a, \quad (7)$$

where  $W_a \in \mathbb{R}^{D_a \times D_a}$  denotes the trainable edge weights. We use Laplacian smoothing [6, 20] to propagate the node features over the bipartite graph. The identity matrix  $I$  can be viewed as a residual sum connection to alleviate optimization difficulties. We randomly initialize both the adjacency matrix  $A_a$  and the weights  $W_a$  and then train them by gradient descent.

After the cross-reasoning process, the new updated feature  $M_a$  is mapped back to the original coordinate space for further processing. Next, we add the result to the original



Table 1. Classification score on NTU-26.

Method	GT	ACTOR [26]	MotionDiffuse [41]	BiGraphDiff
Classification Accuracy ↑				
Punching	76.0%	1.0%	43.0%	<b>49.0%</b>
Kicking	86.0%	14.0%	61.0%	<b>86.0%</b>
Pushing	97.0%	77.0%	<b>86.0%</b>	74.0%
Pat on back	88.0%	4.0%	72.0%	<b>80.0%</b>
Point Finger	83.0%	0.0%	52.0%	<b>76.0%</b>
Hugging	97.0%	59.0%	90.0%	<b>97.0%</b>
Giving object	91.0%	34.0%	68.0%	<b>86.0%</b>
Touch pocket	93.0%	35.0%	81.0%	<b>84.0%</b>
Shaking hands	89.0%	16.0%	80.0%	<b>90.0%</b>
Walking toward	93.0%	72.0%	98.0%	<b>99.0%</b>
Walking apart	95.0%	<b>90.0%</b>	<b>90.0%</b>	<b>90.0%</b>
Hit with object	44.0%	8.0%	23.0%	<b>28.0%</b>
Wield knife	50.0%	7.0%	31.0%	<b>41.0%</b>
Knock over	85.0%	4.0%	<b>61.0%</b>	<b>61.0%</b>
Grab stuff	74.0%	0.0%	57.0%	<b>62.0%</b>
Shoot with gun	57.0%	1.0%	<b>46.0%</b>	44.0%
Step on foot	89.0%	5.0%	85.0%	<b>90.0%</b>
High five	90.0%	4.0%	75.0%	<b>78.0%</b>
Cheers and drink	90.0%	16.0%	69.0%	<b>92.0%</b>
Carry object	96.0%	<b>98.0%</b>	92.0%	95.0%
Take a photo	87.0%	19.0%	63.0%	<b>80.0%</b>
Follow	94.0%	68.0%	<b>90.0%</b>	81.0%
Whisper	83.0%	0.0%	72.0%	<b>79.0%</b>
Exchange things	88.0%	6.0%	65.0%	<b>78.0%</b>
Support somebody	94.0%	<b>100.0%</b>	94.0%	92.0%
Rock paper scissor	91.0%	6.0%	75.0%	<b>91.0%</b>
Average	84.6%	30.7%	70.0%	<b>77.0%</b>

feature  $F_a$  to form a residual connection, as follows:

$$\tilde{F}_a = \phi_a(H_b M_a) + F_a, \quad (8)$$

where we reuse the projection matrix  $H_b$  and apply a linear projection  $\phi_a(\cdot)$  to project  $M_a$  back to the original coordinate space. Therefore, we obtain the feature  $\tilde{F}_a$ , which has the same dimension as the original one  $F_a$ .

Similarly, we can obtain the new feature  $\tilde{F}_b$ . Overall, the proposed method reasons the cross relations between feature maps of different skeletons using a bipartite graph.

## 5. Experiments

### 5.1. Datasets

There are few 3D motion two-person interaction datasets. Therefore we focus on two complementary datasets. The NTU RGB+D 120 dataset [21], among its 120 classes, contains 26 classes labeled as “Mutual Actions / Two Person Interactions” which show two persons performing simple interaction motions. We take this 26-class subset that we call NTU-26 and split each class randomly to obtain our training and testing set. The testing set contains 2,600 samples (100 per class), and the training set is 19,787 samples. The framerate for NTU-26 is 30fps. The second dataset is DuetDance [19], which contains five classes of two-person dance motions for a total of 406 sequences. The motions are more complex than the one from NTU-26 and harder to classify, even for a human observer. The original dataset contains motions with great variations in lengths

from 100 frames to more than 4,000. The average length is 483 frames with a median of 360 frames. While our model can generate very long motions it causes a problem when obtaining quantitative results and lower the quality of the generation due to the limited presence of some sequence of certain lengths. We decided to split the sequences into sub-sequences of 300 frames or less. This increases the number of samples to train our network with. This increased number of samples will also help the diffusion model since it needs a lot of data. This leaves us with 698 training samples and 125 test samples (25 per class randomly selected). The skeletons from DuetDance are extracted from YouTube videos. The frame rate is either 25 fps or 30 fps, depending on the video. We use the same normalization for both datasets. First, we compute the means of the x,y, and z values of both skeletons. Then we subtract these means from the x,y, and z values. We divide the results by the Frobenius norm. Finally, we center the two skeletons around a point that is the middle of a segment that links the torso joints of both skeletons.

### 5.2. Implementation Details

We train our model on an NVIDIA A100 80Go GPU with PyTorch with a batch size of 128 for NTU and 64 for DuetDance. We train on NTU for 1,500 epochs and for 30,000 epochs on DuetDance.

### 5.3. Baselines

As there are currently limited interaction generation methods available, we compared our architecture to single-person motion generation methods. To ensure a fair comparison, we focused on generation methods conditioned on text or class labels and excluded methods conditioned on music, such as EDGE [36]. We compared our BiGraphDiff method to two state-of-the-art methods: MotionDiffuse [41] and ACTOR [26]. MotionDiffuse is a recent architecture based on diffusion and transformers that generate single-person motion from the text. We used the code provided by the authors and ran the model with the recommended parameters. To ensure consistency, we used the same batch size and number of epochs as for our BiGraphDiff method. We chose MotionDiffuse over other diffusion methods because it has available code and its results are among the best, comparable to, or better than MoFusion [8]. ACTOR is a Transformer VAE method that also generates single-person motion. We used the code provided by the authors and retrained it on our datasets with recommended parameters. We had to deactivate the SMPL [22] loss function because it is not available in the NTU RGB+D and DuetDance datasets. These two methods are mainly intended to generate the motion of a single individual. However, we can use them to generate the motion of two individuals by modifying their inputs. These methods assume that the motion is

Table 2. FVD and Multimodality on NTU-26.

Method	FVD↓	Multimodality↓
ACTOR [26]	25298.73	34.91
MotionDiffuse [41]	1292.32	14.94
BiGraphDiff	<b>1048.13</b>	<b>11.28</b>

Table 3. Classification score on DuetDance.

Method	GT	ACTOR [26]	MotionDiffuse [41]	BiGraphDiff
Classification Accuracy ↑				
Cha-cha	28.0%	<b>36.0%</b>	32.0%	32.0%
Jive	52.0%	16.0%	<b>20.0%</b>	16.0%
Rumba	56.0%	16.0%	48.0%	<b>68.0%</b>
Salsa	88.0%	0.0%	64.0%	<b>76.0%</b>
Samba	52.0%	<b>80.0%</b>	32.0%	52.0%
Average	55.2%	29.6%	39.2%	<b>48.8%</b>

represented by a matrix with dimensions of  $3k \times N$  ( $k$  is the number of joints, and  $N$  is the number of frames). Instead, we introduce a matrix twice as large, with dimensions of  $6k \times N$ , which consists of the concatenation of two skeletons. By doing so, neural networks observe an increase in the number of joints and can generate interacting motion.

#### 5.4. Quantitative Results

We perform the quantitative evaluation by using classification accuracy, Fréchet Video Distance (FVD) score [37], and Multimodality.

**Fréchet Video Distance (FVD)** adapts the Fréchet Inception distance (FID) [12] for video sequences [38]. FVD computes the distance between the generated data distribution and the ground truth using deep features.

**MultiModality.** Multimodality measures the diversity of generated samples in each class. It is defined as the average deep features distance of the samples generated by a method compared to the average deep features distance of the ground truth on a specific class. The deep features are extracted from the classifier used for classification accuracy. To compute the average deep features distance we split the set of features of each class into two equal sets and compare the Euclidean norm between the pairs formed by a member of each set and compute the average over the size of the subsets. The Multimodality score is similar the multimodality from [41] but in our experimental results, we directly compare it with the multimodality of the ground truth.

**The classification accuracy** is obtained using a simple Transformer encoder followed by an MLP. The classifier is trained and tested on the same set as the generative methods.

**NTU-26.** Table 1 shows that our method outperforms the two the-state-of-art methods in terms of average accuracy. BiGraphDiff outperforms MotionDiffuse by 7.0% and ACTOR by 46.3%. We are also very close to the accuracy of the classifier on the ground truth. This shows that the se-

Table 4. FVD and Multimodality on DuetDance.

Method	FVD↓	Multimodality↓
ACTOR [26]	2641.08	67.79
MotionDiffuse [41]	1133.51	12.24
BiGraphDiff	<b>997.92</b>	<b>4.33</b>

quences generated by our method are realistic and correspond to the input class. In more detail, we can see that we outperform or equate the other methods on 22 classes out of 26. MotionDiffuse and ACTOR are both better in 2 classes. However, we can see that ACTOR results being actually better is debatable as some classes have very low accuracy, down to 0%. We can also see that the classes in which we perform the worse (i.e., “Hit with object” 28%, “Wield knife” 41%, and “Shoot with gun” 44%) are the ones where the results are also low for the ground truth. Those are classes where the main difference is the object used which is something we can not see using 3D skeleton data. Table 2 shows the FVD and multimodality results. In terms of FVD and Multimodality, our method also outperforms the two other methods indicating that our method produces sequences closer to the real data. One issue when using the NTU dataset is that it is very noisy (see the ground truth in the qualitative results). This means that it is harder to generate noiseless sequences but also that a method that generates samples without noise might be disadvantaged in the quantitative results since they are compared with the ground truth and the classifier is trained on the noisy data.

**DuetDance.** Table 3 shows the classification results on the DuetDance dataset. We can see that, as on NTU-26, we have the best performance on average. The accuracy of our method is 9.6% higher than on MotionDiffuse and only 6.4% lower than on the ground truth. We can note that the accuracy for the ground truth is much lower than for NTU-26. This is due to the nature of the motion in DuetDance. The dance motions are much harder to recognize even for a human and also longer so it is not surprising that the results are worse. ACTOR, on the other only achieves results slightly higher than chance (20%) we will see in the qualitative results that on DuetDance, ACTOR does not produce any motion, we will discuss this in detail in the qualitative results. In the “jive” class, all methods only achieve chance level or lower accuracy. But the results are not so low for the ground truth with means that all methods have trouble generating motion of the “jive class”. In Table 4, we show that we outperform the other methods on both metrics meaning that the results of our method are more realistic.

#### 5.5. Qualitative Results

**NTU-26.** Figure 3 shows visuals of sequences generated for the “Cheers and drinks” class. This class of motion is more complex than others because it is composed of two




Figure 3. Examples of diverse motion generation for a given text prompt “Cheers and Drink” action from NTU.

separate motions “cheers” and “drink”. All methods generate a proper motion but ACTOR shows a low intensity for “cheers” and does not really generate the “drink” motion. MotionDiffuse generates a good motion with both “cheers” and “drink” but there is some noise and the arm length grows over time. Our method generates the proper motion with the two steps and does not produce the noise that is present in the ground truth. In our case, one character drinks while grabbing the glass with one hand while the other uses both hands showing the diversity in the generated motions. Overall we see that our motion is more realistic, temporally, and spatially coherent and manages well to keep the interaction coherent.

**DuetDance.** Figure 4 shows examples of motion generation of the “salsa” class. The dance motions are more complex and the sequences are longer than NTU-26 sequences. ACTOR does not produce motion. We believe this to be due to the great variability of motions from the same class. ACTOR converges to a mean and finds that an unmoving pair of skeletons is the best generation for its losses. We see that MotionDiffuse produces a dance motion without noise. This is because there is less noise in DuetDance than in NTU-26. Our method also generates a dance motion but is better than MotionDiffuse, we reproduce the motion of characters changing sides that is present in the ground truth and the interaction is better as the arm of both characters does not overlap.

## 5.6. User Study

The user study compared BiGraphDiff with two leading methods (i.e., ACTOR [26], MotionDiffuse [41]) and the ground truth sequence. For both datasets, we randomly select 20 samples for each class from the test data. For each comparison, 30 participants are asked to answer two questions, i.e., ‘Q1: Which skeleton sequence is more realistic?’, and ‘Q2: Which skeleton sequence matches the input text better?’. The numbers indicate the preference percentage of users who favor the results of the corresponding methods or

Table 5. User study results (%).

Method	NTU-26		DuetDance	
	Q1	Q2	Q1	Q2
ACTOR [26]	6.1	7.8	5.6	5.9
MotionDiffuse [41]	22.4	24.3	21.8	23.7
BiGraphDiff	<b>31.6</b>	<b>32.7</b>	<b>28.5</b>	<b>30.1</b>
GT	39.9	35.2	44.1	40.3

the GT skeleton sequence. The results highlight the quality of the sequence generated by our method.

## 5.7. Ablation Study

We report ablation results in Table 6 on the NTU-26 dataset. We compare a simple two-stream Transformer (S1), a two-stream Transformer in a diffusion process (S2), a two-stream Transformer in a diffusion process with a simple GCN (S3), and finally our method with bipartite graphs (S4).

The results of S1 are extremely bad. It is explained by the fact the Transformer is a deterministic method and has a low generation diversity which explains the very high FVD. Furthermore, the noisy data from the NTU dataset makes it even harder to provide well-generated sequences. S2 provides much better results both in classification accuracy and FVD, the results are similar to the results obtained by MotionDiffuse. With S3 the simple GCN helps enhance the generation leading to better accuracy and FVD. This highlights the ability of the GCN to model more accurately the spatio-temporal dependencies from each skeleton. Adding a bipartite graph network in S4 provides a stronger increase in performance. It shows that modeling the interactions between the two skeletons is more important than trying to refine the interactions inside each skeleton as S3 did. It validates the use of the bipartite graph network in BiGraphDiff architecture.




Figure 4. Examples of diverse motion generation for a given text prompt “Salsa” action from DuetDance.




Figure 5. Example of generation of very long sequences on the “rumba” class from DuetDance. Under each skeleton the frame number.

Table 6. Ablation study on NTU.

Method	Classification $\uparrow$	FVD $\downarrow$
S1: Two Stream Transformer	3.9%	21215.21
S2: Two Stream Transformer + Diffusion	69.3%	1406.09
S3: S2 + Simple GCN	73.2%	1123.88
S4: S2 + Bipartite Graph	<b>77.0%</b>	<b>1048.16</b>

## 6. Very Long Generation

Long-term motion generation plays an important role in real-world applications. Our method is able to generate longer sequences, as shown in Figure 5. We train the network on the original DuetDance dataset with a maximum sequence length of 4050 frames. We use 376 samples for training and 40 (8 per class) for testing. Figure 5 shows an example of 1580 frames from the “rumba” class. We can see that we generate dance-like motion for the entire duration of the sequence. However, it is very noticeable that we generate better motion for the first few hundred frames, we see that the motion quality around 300 frames is good, but then around 600 frames, we see deterioration that gradually becomes worse. This is due to the length of the sequences in the DuetDance dataset distribution which are usually not very long (average: 483 frames, median: 360 frames). Because of this, we do not use many very long positional encodings during training, preventing a good generation of very long motion.

## 7. Conclusion, Limitations and Future Work

We introduce the first approach for 3D human motion interactions based on denoising diffusion models. Both quan-

titative and qualitative evaluations show that BiGraphDiff outperforms state-of-the-art methods. The proposed BiGraphDiff framework generates coherent human motion sequences that are longer and more diverse than the results of previous approaches. The proposed BiGraphDiff suffers however from the common limitations of diffusion models: the need for large datasets and the long training and testing duration. While the use of Bipartite graph improve performance by modeling the interaction, it also increase the complexity of the model. Our method has 221,476,320 trainable parameters while MotionDiffuse only has 86,990,688 and ACTOR 14,795,866. Recently modifications to the diffusion process have been investigated to improve the complexity of diffusion models [3, 23]. The method is also still slightly sensitive to noise in the training data and can sometimes generate deformed skeletons. This is due in part to the quality of the data used but also because we do not set any constraint related to the input data, e.g., bone length or relative position of joint for 3D skeletons. This means that BiGraphDiff can be used for tasks other than human interaction generation. As long as the input data can be split into two sets and has a temporal or positional component BiGraphDiff can be used for generation.

## Acknowledgments

This project has received financial support from the CNRS through the 80—Prime program, from the French State, managed by the National Agency for Research (ANR) under the Investments for the future program with reference ANR- 16-IDEX-0004 ULNE.



## References

- [1] Chaitanya Ahuja, Dong Won Lee, Ryo Ishii, and Louis-Philippe Morency. No gestures left behind: Learning relationships between spoken language and freeform gestures. In *EMNLP*, 2020. 2
- [2] Simon Alexanderson, Rajmund Nagy, Jonas Beskow, and Gustav Eje Henter. Listen, denoise, action! audio-driven motion synthesis with diffusion models. *arXiv preprint arXiv:2211.09707*, 2022. 2
- [3] Fan Bao, Chongxuan Li, Jiacheng Sun, Jun Zhu, and Bo Zhang. Estimating the optimal covariance with imperfect mean in diffusion probabilistic models. *arXiv preprint arXiv:2206.07309*, 2022. 8
- [4] M. Baruah and B. Banerjee. A multimodal predictive agent model for human interaction generation. In *CVPR Workshops*, 2020. 2
- [5] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18000–18010, 2023. 2
- [6] Yunpeng Chen, Marcus Rohrbach, Zhicheng Yan, Yan Shuicheng, Jiashi Feng, and Yannis Kalantidis. Graph-based global reasoning networks. In *CVPR*, 2019. 4
- [7] Qiongjie Cui, Huaijiang Sun, and Fei Yang. Learning dynamic relationships for 3D human motion prediction. In *CVPR*, 2020. 2
- [8] Rishabh Dabral, Muhammad Hamza Mughal, Vladislav Golyanik, and Christian Theobalt. Mofusion: A framework for denoising-diffusion-based motion synthesis, 2022. 2, 5
- [9] Chuan Guo, Xinxin Zuo, Sen Wang, and Li Cheng. TM2T: stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXV*, pages 580–597, 2022. 2
- [10] Ikhsanul Habibie, Weipeng Xu, Dushyant Mehta, Lingjie Liu, Hans-Peter Seidel, Gerard Pons-Moll, Mohamed Elgharib, and Christian Theobalt. Learning speech-driven 3d conversational gestures from video. In *ACM International Conference on Intelligent Virtual Agents (IVA)*, 2021. 2
- [11] William Harvey, Saeid Naderiparizi, Vaden Masrani, Christian Weilbach, and Frank Wood. Flexible diffusion modeling of long videos. *arXiv preprint arXiv:2205.11495*, 2022. 2
- [12] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017. 6
- [13] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 2
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arxiv:2006.11239*, 2020. 2
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020. 3
- [16] Emiel Hooeboom, Víctor Garcia Satorras, Clément Vignac, and Max Welling. Equivariant diffusion for molecule generation in 3D. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 8867–8887. PMLR, 17–23 Jul 2022. 2
- [17] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. 1
- [18] Kosmas Krissis, Aggelos Gkiokas, Aggelos Pikrakis, and Vassilis Katsouros. Danceconv: Dance motion generation with convolutional networks. *IEEE Access*, 10:44982–45000, 2022. 1
- [19] Jogendra Kundu, Himanshu Buckchash, Priyanka Mandikal, Rahul M V, Anirudh Jamkhandi, and R. Babu. Cross-conditioned recurrent networks for long-term synthesis of inter-person human motion interactions. In *WACV*, 2020. 5
- [20] Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *AAAI*, 2018. 4
- [21] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex Kot. Ntu rgb+d 120: A large-scale benchmark for 3d human activity understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–18, 2019. 5
- [22] 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, Oct. 2015. 5
- [23] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. DPM-solver: A fast ODE solver for diffusion probabilistic model sampling in around 10 steps. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. 8
- [24] Thomas Lucas, Fabien Baradel, Philippe Weinzaepfel, and Grégory Rogez. Posegpt: Quantization-based 3d human motion generation and forecasting. In *European Conference on Computer Vision (ECCV)*, 2022. 2
- [25] Julieta Martinez, Michael J. Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *CVPR*, 2017. 2
- [26] Mathis Petrovich, Michael J. Black, and Gül Varol. Action-conditioned 3D human motion synthesis with transformer VAE. In *International Conference on Computer Vision (ICCV)*, 2021. 1, 5, 6, 7
- [27] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learn-

- ing transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 3
- [28] Zhiyuan Ren, Zhihong Pan, Xin Zhou, and Le Kang. Diffusion motion: Generate text-guided 3d human motion by diffusion model. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, 2023. 2
- [29] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. 2
- [30] Zhuoran Shen, Mingyuan Zhang, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Efficient attention: Attention with linear complexities. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 3531–3539, 2021. 3
- [31] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022. 2
- [32] Theodoros Sofianos, Alessio Sampieri, Luca Franco, and Fabio Galasso. Space-time-separable graph convolutional network for pose forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11209–11218, October 2021. 2
- [33] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. 2
- [34] Hao Tang, Song Bai, Philip H. S. Torr, and Nicu Sebe. Bipartite graph reasoning gans for person image generation. In *31st British Machine Vision Conference 2020, BMVC 2020, Virtual Event, UK, September 7-10, 2020*. BMVA Press, 2020. 2
- [35] Guy Tevet, Sigal Raab, Brian Gordon, Yonatan Shafir, Daniel Cohen-Or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. 1
- [36] Jonathan Tseng, Rodrigo Castellon, and C. Karen Liu. EDGE: editable dance generation from music. *CoRR*, abs/2211.10658, 2022. 1, 2, 5
- [37] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. FVD: A new metric for video generation. In *Deep Generative Models for Highly Structured Data, ICLR 2019 Workshop, New Orleans, Louisiana, United States, May 6, 2019*. OpenReview.net, 2019. 6
- [38] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. FVD: A new metric for video generation. In *Deep Generative Models for Highly Structured Data, ICLR 2019 Workshop, New Orleans, Louisiana, United States, May 6, 2019*. OpenReview.net, 2019. 6
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 3
- [40] Wenjie Yin, Hang Yin, Danica Kragic, and Mårten Björkman. Graph-based normalizing flow for human motion generation and reconstruction, 2021. 2
- [41] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiandiffuse: Text-driven human motion generation with diffusion model. *CoRR*, abs/2208.15001, 2022. 1, 2, 3, 5, 6, 7
- [42] Yi Zhou, Zimo Li, Shuangjiu Xiao, Chong He, Zeng Huang, and Hao Li. Auto-conditioned recurrent networks for extended complex human motion synthesis. In *International Conference on Learning Representations*, 2018. 1