

## Music-Driven Group Choreography

Nhat Le<sup>1</sup>, Thang Pham<sup>1</sup>, Tuong Do<sup>1</sup>, Erman Tjiputra<sup>1</sup>, Quang D. Tran<sup>1</sup>, Anh Nguyen<sup>2</sup>

<sup>1</sup>AIOZ, Singapore <sup>2</sup>University of Liverpool, UK



Figure 1. We present the AIOZ-GDANCE dataset with in-the-wild videos, music audio, and 3D group dance motion.

### Abstract

*Music-driven choreography is a challenging problem with a wide variety of industrial applications. Recently, many methods have been proposed to synthesize dance motions from music for a single dancer. However, generating dance motion for a group remains an open problem. In this paper, we present AIOZ – GDANCE, a new large-scale dataset for music-driven group dance generation. Unlike existing datasets that only support single dance, our new dataset contains group dance videos, hence supporting the study of group choreography. We propose a semi-autonomous labeling method with humans in the loop to obtain the 3D ground truth for our dataset. The proposed dataset consists of 16.7 hours of paired music and 3D motion from in-the-wild videos, covering 7 dance styles and 16 music genres. We show that naively applying single dance generation technique to creating group dance motion may lead to unsatisfactory results, such as inconsistent movements and collisions between dancers. Based on our new dataset, we propose a new method that takes an input music sequence and a set of 3D positions of dancers to efficiently produce multiple group-coherent choreographies. We propose new evaluation metrics for measuring group dance quality and perform intensive experiments to demonstrate the effectiveness of our method. Our project facilitates future research on group dance generation and is available at <https://aioz-ai.github.io/AIOZ-GDANCE/>.*

### 1. Introduction

Dancing is an important part of human culture and remains one of the most expressive physical art and communication forms [17, 25]. With the rapid development of digital social media platforms, creating dancing videos has gained significant attention from social communities. As the consequence, millions of dancing videos are created and watched daily on online platforms. Recently, studies of how to create natural dancing motion from music have attracted great attention in the research community [8]. The outcome of dancing generation techniques can be applied to various applications such as animation [39], virtual idol [47], meta-verse [35], or in dance education [4, 51].

Although there is some progress towards synthesizing realistic dancing motion from music in recent literature [26, 39, 47, 53] creating natural 3D dancing motions from the input audio remains an open problem [39]. This is mainly due to (i) the complex structure and the non-linear correlation between continuous human motion and the accompanying music audio, (ii) the variety in the repertoire of dancing motions for an expressive choreography performance, (iii) the difficulty in generating long motion sequences, and (iv) the complication of capturing the correspondences between the dancing motion and audio such as dancing styles or music rhythms. Furthermore, recent works focus on generating dancing motion for solo dancer [16, 20, 39, 61] while

producing dancing motion for a group of dancers have not been well-investigated by the community yet.

Compared to the single dance generation task, the group dance generation poses a more challenging problem [13]. In practice, group dance may contain complicated and different choreographies between attended dancers while also retaining the relationship between the motion and the input music rhythmically. Furthermore, group dance has the communication between dancers through physical contact, hence performing a correlation between motion series in a group is essential and challenging. Currently, most of the music-to-dance datasets [37, 39, 47, 55, 58, 66] only contain solo dance videos. Hence, they are not able to capture essential aspects that only occurred in group dance scenarios, including multiple-person motions, synchronization, and interaction between dancers.

Since learning to synthesize dancing motion from music is a highly challenging task, one of the primary requirements is the availability of large-scale datasets. Some works employ professional choreographers to obtain music-dance datasets under a highly complex Motion Capture (MoCap) system [10, 37, 58, 63]. Although the captured motions are accurate, it is challenging to scale up and increase the diversity of the data with several dance styles and music genres. To overcome the limitation of the MoCap system, another line of works leverage existing pose estimation algorithm to generate the pseudo-ground truths for in-the-wild dancing videos [20, 34, 55]. However, these aforementioned datasets are designed originally for the single motion generation task and provide only paired music and single dancing motion [37, 39, 47], thus they cannot be applied to facilitate generating multiple motions within a group of dancers.

Motivated by these shortcomings, this paper introduces AIOZ-GDANCE, a new large-scale dataset to advance the study of group choreography. Unlike existing choreography datasets that only supports single dancer, our dataset consists of group dance videos. As in Figure 1, our dataset has multiple input modalities (i.e., video frames, audio) and multiple 3D human mesh ground truths. To annotate the dataset, we introduce a semi-automatic method with humans in the loop to ensure the data quality. Using the new dataset, we propose the first strong baseline for group dance generation that can jointly generate multiple dancing motions expressively and coherently.

Our contributions are summarised as follows:

- We introduce AIOZ-GDANCE, a new large-scale dataset for group dance generation. To our best knowledge, AIOZ-GDANCE is the *largest* audio-driven group dance dataset.
- Based on our new dataset, we propose a new method, namely GDanceR, to efficiently generate group dancing motion from the input audio.

## 2. Related Work

**Music-Motion Datasets.** Early efforts to create music-motion dataset focus on using motion capture system. Specifically, the authors of [58] use MoCap to record 3D skeletons from dancers to establish a music-to-dance dataset with four dancing types. It is challenging to collect a large dataset using the MoCap method since it is costly to hire performing dancers and equip the required devices. Notice the limitations of MoCap, the authors in [20, 34] propose a music-to-dance dataset by crawling internet videos and use OpenPose [9] to obtain 2D skeleton ground truths. The authors in [47] propose a dataset using MoCap and Virtual Reality data. Instead of generating dance ground truths as human skeletons, AIST++ dataset [39, 56] obtains 3D motion by fitting SMPL model [40] from multi-view videos. Unlike [39] the work in [37] proposes a large-scale music-to-dance dataset from in-the-wild videos. Nevertheless, existing datasets only focus on the single-dance scenario. Therefore, the group motion and the interaction between dancers are not exploited in these datasets. We compare different music-to-dance datasets in Table 1.

**Audio-driven Motion Generation.** Generating natural and realistic human motion from audio is a challenging problem [23]. A classical approach is based on the motion graph constructed from a largely captured motion database [30]. To synthesize novel motion, different motion segments are combined by optimizing the transition cost along the path of the graph [50]. Other works apply music-motion similarity matching constraints to further ensure the consistency between motion and music [14, 27, 28, 36, 48, 52]. In recent years, several progresses have been made in the field of music-to-dance motion generation using CNN [2, 55, 63, 66], RNN [5, 20, 55, 58], GCN [16, 49], GAN [34, 55], or Transformer [26, 37, 39, 47, 53]. Typically, these music-to-dance methods are conditioned on multi-modal inputs and then generate the future sequence of human poses. Despite the potential to generate realistic dancing motion from the music, these methods lack the ability to produce coherent movements between multiple dancers altogether.

**Multi-person Motion Prediction.** Learning and forecasting the behavior between multiple people has been a longstanding problem [24, 44, 54]. Alahi *et al.* [3] jointly reasons the trajectories of multiple pedestrians and forecasts their destinations in a scene by using Markov chain model. The authors in [1] combine the visual context of the scene and social interactions between multiple persons to forecast their future motion. Recently, the Multi-Range Transformers proposed by [60] can predict the movements of more than 10 people for social interaction groups. Difference from the motion prediction task that uses motion or visual inputs, we aim to generate human dancing motion from the input music.

Dataset	Hours	Subjects	Captured Environment	Acquisition	Music Genres	Dance Styles	Group Dance	Ground Truth
Dance Melody [58]	1.57	n/a	Lab-control	MoCap	limited	limited	no	3D Joints
DanceNet [66]	0.96	2	Lab-control	MoCap	limited	limited	no	3D Joints
YT-Dance3D [55]	5	n/a	In-the-wild	Fully-automatic	rich	rich	no	3D Joints
Dancing2Music [34]	71	n/a	In-the-wild	Fully-automatic	rich	limited	no	2D Joints
DanceRevolution [20]	12	n/a	In-the-wild	Fully-automatic	rich	limited	no	2D Joints
PMSD [47]	3.84	2	Lab-control	MoCap	varied	limited	no	3D Joints
ShaderMotionVR [47]	10.2	11	In-the-wild	VR Tracking	varied	rich	no	3D Joints
AIST++ [39]	5.2	30	Lab-control	Fully-automatic	limited	varied	no	3D Mesh
PhantomDance [37]	9.6	100+	In-the-wild	Artistic	rich	rich	no	3D Mesh
<b>AIOZ-GDANCE (ours)</b>	<b>16.7</b>	<b>4000+</b>	<b>In-the-wild</b>	<b>Semi-automatic</b>	<b>rich</b>	<b>rich</b>	<b>yes</b>	<b>3D Mesh</b>

Table 1. Comparison between music-to-dance datasets.

### 3. The AIOZ-GDANCE Dataset

Since we want to develop a large-scale dataset with in-the-wild videos, setting up a MoCap system is not feasible. However, manually annotating 3D groundtruth for millions of frames from dancing videos is also an extremely tedious job. Therefore, we propose a semi-automatic labeling method with humans in the loop to produce a large-scale group dance dataset.

#### 3.1. Data Collection and Preprocessing

**Video Collection.** We collect the in-the-wild, public domain group dancing videos along with the music from Youtube, Tiktok, and Facebook. All group dance videos are processed at  $1920 \times 1080$  resolution and 30FPS.

**Human Tracking.** We perform tracking for all humans in the videos using the state-of-the-art multi-object tracker [57] to obtain the tracking bounding boxes. Note that although the tracker can produce reasonable results, there are failure cases in some frames. Therefore, we manually correct the bounding box of the incorrect cases. This tracking correction is crucial since we want the trajectory of each person to be accurately tracked in order to reconstruct their motion in latter stages.

**Pose Estimation.** Given the bounding boxes of each person in the video, we leverage the recent 2D pose estimation method [15] to generate the initial 2D poses for each person. In practice, there exist some inaccurately detected keypoints due to motion blur and partial occlusion. We manually fix the incorrect cases to obtain the 2D keypoints of each human bounding box.

#### 3.2. Group Motion Fitting

To construct 3D group dance motion, we first reconstruct the full body motion for each dancer by fitting the 3D mesh. We then jointly optimize all dancer motions to construct the globally-coherent group motion. Finally, we post-process and remove wrong cases from the optimization results.

**Local Mesh Fitting.** We use SMPL model [40] to represent the 3D human. The SMPL model is a differentiable function that maps the pose parameters  $\theta$ , the shape parameters  $\beta$ , and the root translation  $\tau$  into a set of 3D human body mesh vertices  $\mathbf{V} \in \mathbb{R}^{6890 \times 3}$  and 3D joints  $\mathbf{X} \in \mathbb{R}^{J \times 3}$ , where  $J$  is the number of body joints. Our optimizing motion variables for each individual dancer consist of a sequence of SMPL joint angles  $\{\theta_t\}_{t=1}^T$ , a sequence of the root translation  $\{\tau_t\}_{t=1}^T$ , and a single SMPL shape parameter  $\beta$ . We fit the sequence of SMPL motion variables to the tracked 2D keypoints (obtained from Section 3.1) by extending SMPLify-X [46] across the whole video sequence:

$$E_{\text{local}} = E_J + \lambda_\theta E_\theta + \lambda_\beta E_\beta + \lambda_S E_S + \lambda_F E_F, \quad (1)$$

where  $E_J$ ,  $E_\theta$  and  $E_\beta$  are as in [46] but calculated across every frames of the video sequence. The smoothness term  $E_S = \sum_{t=1}^{T-1} \|\theta_{t+1} - \theta_t\|^2 + \sum_{j=1}^J \sum_{t=1}^{T-1} \|\mathbf{X}_{j,t+1} - \mathbf{X}_{j,t}\|^2$  encourages the temporal smoothness of the motion. The term  $E_F = \sum_{t=1}^{T-1} \sum_{j \in \mathcal{F}} c_{j,t} \|\mathbf{X}_{j,t+1} - \mathbf{X}_{j,t}\|^2$  ensures feet joints to stay stationary when in contact (zero velocity).  $\mathcal{F}$  is the set of feet joint indexes,  $c_{j,t}$  is the feet contact label of joint  $j$  at time  $t$  produced by a contact estimation network [67].

**Global Optimization.** Given the 3D motion sequence of each dancer  $p$ :  $\{\theta_t^p, \tau_t^p\}$ , we further resolve the motion trajectory problems in group dance by solving the following objective:

$$E_{\text{global}} = E_J + \lambda_{\text{pen}} E_{\text{pen}} + \lambda_{\text{reg}} \sum_p E_{\text{reg}}(p) + \lambda_{\text{dep}} \sum_{p,p',t} E_{\text{dep}}(p,p',t) + \lambda_{\text{gc}} \sum_p E_{\text{gc}}(p), \quad (2)$$

where  $E_{\text{pen}}$  is the Signed Distance Function penetration term based on [22] to prevent the overlapping of reconstructed motions between dancers.

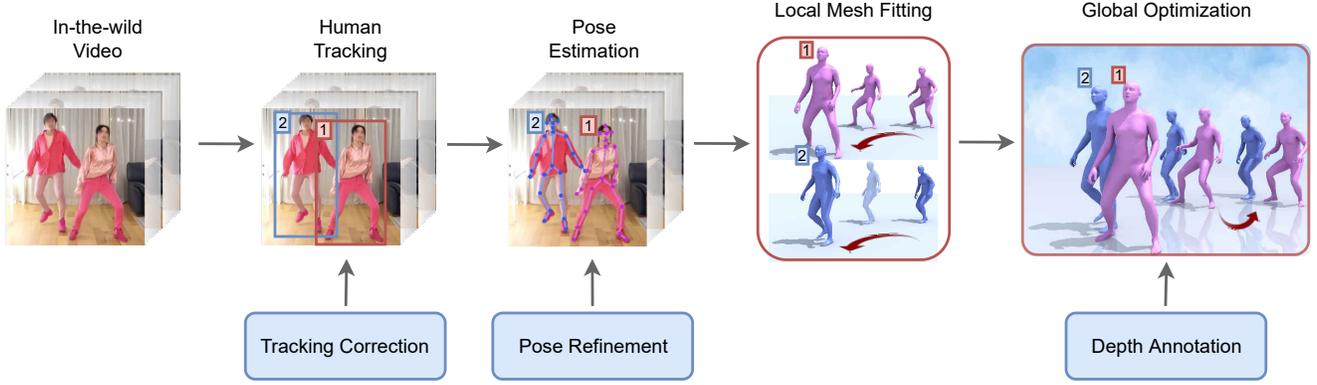


Figure 2. The pipeline of making our AIOZ-GDANCE dataset. Blue boxes denote manual correction/annotation steps.

$E_{\text{reg}}(p) = \sum_{t=1}^T \|\theta_t^p - \hat{\theta}_t^p\|^2$  is the regularization term that prevents the motion from deviating too much from the prior optimized individual motion  $\{\hat{\theta}_t^p\}$  obtained by optimizing Equation 1 for dancer  $p$ .

In practice, we find that the relative depth ordering of dancers in the scene can be inconsistent due to the ambiguity of the 2D projection. To ensure the group motion quality, we watch the videos and manually provide the ordinal depth relation information of all dancers in the scene at each frame  $t$  as follows:

$$r_t(p, p') = \begin{cases} 1, & \text{if dancer } p \text{ is closer than } p' \\ -1, & \text{if dancer } p \text{ is farther than } p' \\ 0, & \text{if their depths are roughly equal} \end{cases} \quad (3)$$

Given the relative depth information provided by human annotators, we derive the depth relation term  $E_{\text{dep}}$  inspired by [11]. This term encourages consistent ordinal depth relation between the motion trajectories of multiple dancers, especially when dancers partially occlude each other:

$$E_{\text{dep}}(p, p', t) = \begin{cases} \log(1 + \exp(z_t^p - z_t^{p'})), & r_t(p, p') = 1 \\ \log(1 + \exp(-z_t^p + z_t^{p'})), & r_t(p, p') = -1 \\ (z_t^p - z_t^{p'})^2, & r_t(p, p') = 0 \end{cases} \quad (4)$$

where  $z_t^p$  is the depth component of the root translation  $\tau_t^p$  of the person  $p$  at frame  $t$ .

Finally, we apply the global ground contact constraint  $E_{\text{gc}}$  to further ensure consistency between the motion of every person and the environment based on the ground contact information. This contact term is also needed to reduce the artifacts such as foot-skating, jittering, and penetration under the ground.

$$E_{\text{gc}}(p) = \sum_{t=1}^{T-1} \sum_{j \in \mathcal{F}} c_{j,t}^p \|\mathbf{X}_{j,t+1}^p - \mathbf{X}_{j,t}^p\|^2 + c_{j,t}^p \|(\mathbf{X}_{j,t}^p - \mathbf{f})^\top \mathbf{n}^*\|^2, \quad (5)$$

where  $\mathcal{F}$  is the set of feet joint indexes,  $\mathbf{n}^*$  is the estimated plane normal and  $\mathbf{f}$  is a 3D fixed point on the ground plane.

The first term in Equation 5 is the zero velocity constraint when the feet are in contact with the ground, while the second term encourages the feet position to stay near the ground when in contact. To obtain the ground plane parameters, we initialize the plane point  $\mathbf{f}$  as the weighted median of all contact feet positions. The plane normal  $\mathbf{n}^*$  is obtained by optimizing a robust Huber objective:

$$\mathbf{n}^* = \arg \min_{\mathbf{n}} \sum_{\mathbf{X}_{\text{feet}}} \mathcal{H} \left( (\mathbf{X}_{\text{feet}} - \mathbf{f})^\top \frac{\mathbf{n}}{\|\mathbf{n}\|} \right) + \|\mathbf{n}^\top \mathbf{n} - 1\|^2, \quad (6)$$

where  $\mathcal{H}$  is the Huber loss function [21],  $\mathbf{X}_{\text{feet}}$  is the 3D feet positions of all dancers across the whole sequence that are labelled as in contact (i.e.,  $c_{j,t}^p = 1$ ).

**Post Processing.** Although our optimization process produces relatively good results, there are some extreme cases that it fails to handle. We recheck all the results and fix the cases with minor problems. Other severely wrong cases are simply discarded. More details can be found in our Supplementary Material.

### 3.3. How will AIOZ-GDANCE be useful to the community?

We bring up some interesting research directions that can be benefited from our dataset:

- **Group Dance Generation:** While single-person choreography is a hot research topic [10,35,37,39,53], group dance generation has not yet well investigated. We hope that the release of our dataset will foster more this research direction.
- **Human Pose Tracking:** By having SMPL groundtruth motion, our dataset can be used in many human pose/motion tracking tasks such as in [12,57,62].

Apart from these tasks, we believe our dataset can be used in other scenarios such as dance education [16,45],

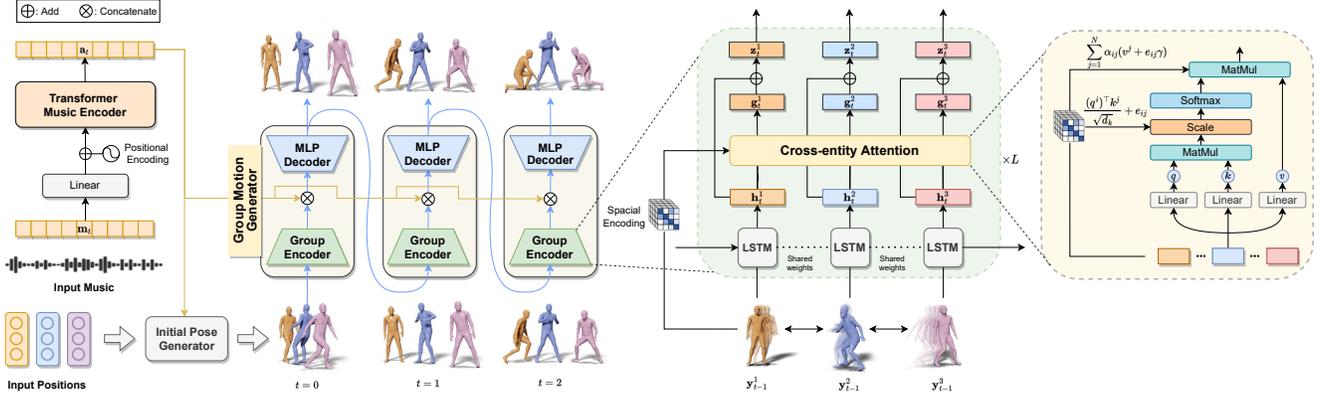


Figure 3. Architecture of our Music-driven 3D **Group Dance** generator (GDanceR). Our model takes in a music sequence and a set of initial positions, and then auto-regressively generates coherent group dance motions that are attuned to the input music.

dance style transfer [48, 64, 65], or human behavior analysis [31–33, 43]. The research community is free to explore other applications of our dataset.

## 4. Audio-driven Group Dance Generation

### 4.1. Problem Formulation

Given an input music audio sequence  $\{m_1, m_2, \dots, m_T\}$  with  $t = \{1, \dots, T\}$  indicates the index of music segments, and the initial 3D positions of  $N$  dancers  $\{\tau_0^1, \tau_0^2, \dots, \tau_0^N\}$ ,  $\tau_0^i \in \mathbb{R}^3$ , our goal is to generate the group motion sequences  $\{y_1^1, \dots, y_T^1; \dots; y_1^n, \dots, y_T^n\}$  where  $y_t^i$  is the generated pose of  $i$ -th dancer at time step  $t$ . Specifically, we represent the human pose as a 72-dimensional vector  $y = [\tau; \theta]$  where  $\tau$ ,  $\theta$  represent the root translation and pose parameters of the SMPL model [40], respectively.

Generally, the generated group dance motion should meet the two conditions: (i) consistency between the generated dancing motion and the input music in terms of style, rhythm, and beat; (ii) the motions and trajectories of dancers should be coherent without cross-body intersection between dancers. Figure 3 shows an overview of our approach.

### 4.2. GDanceR Architecture

**Transformer Music Encoder.** From the raw audio signal of the input music, we first extract music features using the audio processing library [42]. Concretely, we extract the mel frequency cepstral coefficients (MFCC), MFCC delta, constant-Q chromagram, tempogram, onset strength and one-hot beat, which results in a 438-dimensional feature vector. We then encode the music sequence  $\{m_1, m_2, \dots, m_T\}$ ,  $m_t \in \mathbb{R}^{438}$  into a sequence of hidden representation  $\{a_1, a_2, \dots, a_T\}$ ,  $a_t \in \mathbb{R}^{d_a}$ . In practice, we utilize the self-attention mechanism of transformer [59] to effectively encode the multi-scale informa-

tion and the long-term dependency between music frames. The hidden audio at each time step is expected to contain meaningful structural information to ensure that the generated dancing motion is coherent across the whole sequence.

**Initial Pose Generator.** Given the initial positions of all dancers, we generate the initial poses by combing the audio feature with the starting positions. We aggregate the audio representation by taking an average over the audio sequence. The aggregated audio is then concatenated with the input position and fed to a multilayer perceptron (MLP) to predict the initial pose for each dancer:

$$y_0^i = \text{MLP} \left( \left[ \frac{1}{T} \sum_{t=1}^T a_t; \tau_0^i \right] \right), \quad (7)$$

where  $[\cdot]$  is the concatenation operator,  $\tau_0^i$  is the initial position of the  $i$ -th dancer.

**Group Motion Generator.** To generate the group dance motion, we aim to synthesize the coherent motion of each dancer such that it aligns well with the input music. Furthermore, we also need to maintain global consistency between all dancers. In practice, our Group Encoder uses Recurrent Neural Network to capture the temporal motion dynamics of each dancer [38, 41], and attention mechanism to encode the spatial relationship of all dancers [59].

Specifically, at each time step, the pose of each dancer in the previous frame  $y_{t-1}^i$  is sent to an LSTM unit [19] to encode the hidden local motion representation, i.e.,  $h_t^i = \text{LSTM}(y_{t-1}^i, h_{t-1}^i)$ .

To ensure the motions of all dancers have global coherency without strange effects such as cross-body intersection, we introduce the Cross-entity Attention mechanism. In particular, each individual motion representation is first linearly projected into a key vector  $k^i$ , a query vector  $q^i$  and

a value vector  $v^i$  as follows:

$$k^i = h^i W^k, \quad q^i = h^i W^q, \quad v^i = h^i W^v, \quad (8)$$

where  $W^q, W^k \in \mathbb{R}^{d_h \times d_k}$ , and  $W^v \in \mathbb{R}^{d_h \times d_v}$  are parameters that transform the hidden motion  $h$  into a query, a key, and a value, respectively.  $d_k$  is the dimension of the query and key while  $d_v$  is the dimension of the value vector. To encode the relationship between dancers in the scene, our Cross-entity Attention also utilizes the Scaled Dot-Product Attention as in [59].

In practice, we find that people having closer positions to each other tend to have higher correlation in their movement. Therefore, we use Spatial Encoding to encode the spacial relationship between two dancers. The Spatial Encoding between two entities based on their distance in the 3D space is defined as follows:

$$e_{ij} = \exp\left(-\frac{\|\tau^i - \tau^j\|^2}{\sqrt{d_\tau}}\right), \quad (9)$$

where  $d_\tau$  is the dimension of the position vector  $\tau$ . Considering the query  $q^i$ , which represents the current entity information, and the key  $k^j$ , which represents other entity information, we inject the spatial relation information between these two entities onto their cross attention coefficient:

$$\alpha_{ij} = \text{softmax}\left(\frac{(q^i)^\top k^j}{\sqrt{d_k}} + e_{ij}\right). \quad (10)$$

To preserve the spatial relative information in the attentive representation, we also embed them into the hidden value vector and obtain the global-aware representation  $g^i$  of the  $i$ -th entity as follows:

$$g^i = \sum_{j=1}^N \alpha_{ij}(v^j + e_{ij}\gamma), \quad (11)$$

where  $\gamma \in \mathbb{R}^{d_v}$  is the learnable bias and scaled by the Spatial Encoding. Intuitively, the Spatial Encoding acts as the bias in the attention weight, encouraging the interactivity and awareness to be higher between closer entities. Our attention mechanism can adaptively attend to each dancer and others temporally and spatially, thanks to the encoded motion as well as the spatial information.

We then fuse both the local and global motion representation by adding  $h^i$  and  $g^i$  to obtain the final latent motion  $z^i$ . Our final global-local representation of each entity is expected to carry the comprehensive information of their own past motion as well as the motion of every other entity, enabling the MLP Decoder to generate coherent group dancing sequences. Finally, we generate the next movement  $y_t^i$  based on the final motion representation  $z_t^i$  as well as the hidden audio representation  $a_t$ , and thus can capture the fine-grained correspondence between music feature sequence and dance movement sequence:

$$y_t^i = \text{MLP}([z_t^i; a_t]). \quad (12)$$

Criteria	Train	Validate	Test	Total
Duration (hours)	13.5	1.6	1.6	16.7
Total Frames	1,459K	175K	174K	1,808K

Table 2. Train/val/ test split of our AIOZ-GDANCE dataset.

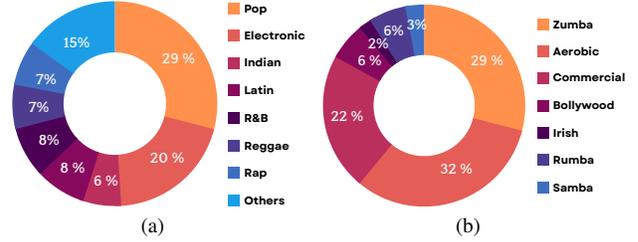


Figure 4. Distribution (%) of music genres (a) and dance styles (b) in our dataset.

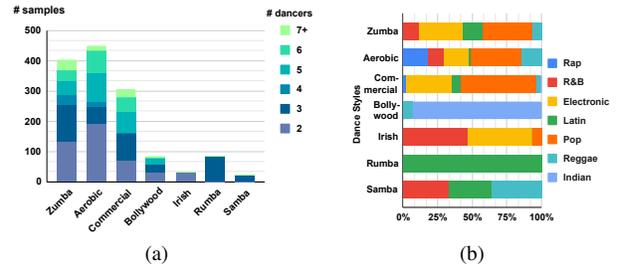


Figure 5. The correlation between dance styles and number of dancers (a); and between dance styles and music genres (b).

## 5. Experiments

### 5.1. AIOZ-GDANCE Statistic

**Dataset Split.** AIOZ-GDANCE comprises 16.7 hours of whole-body motion and music audio of group dancing. The duration of each video in our dataset is ranging from 15 to 60 seconds. We decode all videos at 30 FPS. We randomly sample all videos into train, validation and test sets with 80%; 10%; and 10% of total videos, respectively. Table 2 shows the details about the training, validation, and testing splits of our dataset.

**Dataset Analysis.** In Figure 4, we show the distribution of music genres and dance styles in our dataset. As illustrated in Figure 4a, Pop and Electronic are popular music genres while other music genres nearly share the same distribution. In Figure 4b, Zumba, Aerobic, and Commercial are dominant dance styles.

Figure 5a shows the number of dancers in each dance styles. Naturally, we see that Zumba, Aerobic, and Commercial have more dancers. Figure 5b shows the correlation between music genres and dance styles. In Figure 6, we show an example sequence of a dancing mo-



Figure 6. Example motion sequence from our dataset from front-view and top-view.

tion from our AIOZ-GDANCE dataset. We recommend the readers to check our Supplementary Material and Demonstration Video for more detailed analysis and illustration.

## 5.2. Group Dance Generation Result

### 5.2.1 Implementation Details

The MLP in Equation 7 and 12 has the same architecture with three hidden layers of 512 neurons each. We apply layer normalization [6] and ReLU non-linearity at each hidden layer. The Transformer Music Encoder has 2 transformer layers with 8 attention heads. Both the hidden audio and hidden motion have dimension  $d_a = d_h = 1024$ . We stack  $L = 3$  identical Group Encoder layers in the Group Motion Generator to enhance the learning capacity of the model. For the Cross-entity Attention, we also employ multi-head attention strategy with 8 heads and the dimension of query, key, and value for each head is set to  $d_k = d_v = 64$ . During training, we randomly sample the dance motion with the sequence length  $T = 240$  frames and train the model using  $L2$  loss as in [39]. We also use scheduled sampling [7] to improve the model robustness and enable long-term generation. The whole model is trained end-to-end using Adam optimizer [29] with batch size of 16 and learning rate of  $10^{-4}$ . At test time, the group dance motions are generated in an auto-regressive manner based on the given inputs.

### 5.2.2 Evaluation Protocol

We use the following metrics to evaluate the quality of single dancing motion: Frechet Inception Distance (FID) [18, 39], Motion-Music Consistency (MMC) [39], Generation Diversity (GenDiv) [20, 34, 39]. To evaluate the group dancing quality, we propose three new metrics: Group Motion Realism (GMR), Group Motion Correlation (GMC), and Trajectory Intersection Frequency (TIF). Please see our Supplementary Material for more discussion.

Method	Single-dance Metric			Group-dance Metric			
	FID↓	MMC↑	GenDiv↑	GMR↓	GMC↑	TIF↓	
FACT [39]	56.20	0.222	8.64	101.52	62.68	0.321	
<b>GDanceR</b> (ours)	w/o CA	63.83	0.218	8.99	109.80	68.47	0.379
	w CA	<b>43.90</b>	<b>0.250</b>	<b>9.23</b>	<b>51.27</b>	<b>79.01</b>	<b>0.217</b>

Table 3. The generation results on our dataset. *w/o* CA denotes without using Cross-entity Attention.

### 5.2.3 Experimental Results

We compare our method with FACT [39]. FACT is designed for single dance generation, thus giving our method an advantage. However, it is still the closest competing method as we propose a new group dance dataset that is not available for benchmarking before. We also analyse our method with and without using the Cross-entity Attention. We train all methods with mini-batch containing all dancers within the group instead of sampling each dancer independently as in FACT’s original implementation.

**Cross-entity Attention Analysis.** Table 3 shows the method comparison between the baseline FACT [39] and our proposed GDanceR with and without Cross-entity Attention. The results show that GDanceR, especially with the Cross-entity Attention, out-performs the baseline by a large margin in all metrics. In Figure 7, we also visualize the example outputs of FACT and GDanceR. It is clear that FACT does not handle well the intersection problem. This is understandable as FACT is not designed for group dance generation, while our method with the Cross-entity Attention can deal with this problem better.

**Number of Dancers Analysis.** Table 4 demonstrates the generation results of our method when we want to generate different numbers of dancers. In general, the FID, GMR, and GMC metric do not show much correlation with the numbers of generated dancers since the results are varied. On the other hand, MMC shows its stability among all se-

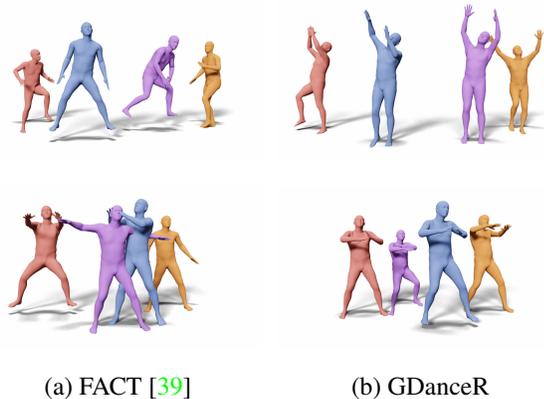


Figure 7. Comparison between FACT [39] and our GDanceR. Our method handles better the consistency and cross-body intersection problem between dancers.

$N$ Generated Dancers	Single-dance Metric			Group-dance Metric		
	FID↓	MMC↑	GenDiv↑	GMR↓	GMC↑	TIF↓
2	48.82	0.248	9.66	53.83	75.44	0.086
3	44.47	0.245	9.46	52.85	74.07	0.104
4	47.32	0.248	9.24	58.79	77.71	0.162
5	44.19	0.249	9.38	55.05	78.72	0.218
6	50.95	0.250	9.25	59.05	75.24	0.319
7	48.86	0.250	9.19	56.23	76.01	0.367

Table 4. Performance of our proposed method when we increase the number of generated dancers.

tups ( $\sim 0.248$ ), which indicates that our network is robust in generating motion from given music regardless of the changing of initial positions. The generation diversity (GenDiv) decreases while the intersection frequency (TIF) increases when more dancers are generated. These results show that dealing with the collision during the group generation process is worth further investigation.

**Dance Style Analysis.** Different dance styles exhibit different challenges in group dance generation. As shown in Table 5, Aerobic and Zumba are quite similar for generating choreography as they usually focus on workout and sporty movements. Besides, while Commercial and Irish are easier for the model to generate, Bollywood and Samba contain highly skilled movements that are challenging to capture and represent accurately. In Figure 8, we show the generated results of GDanceR with different dance styles. Our Supplementary Material and Demonstration Video also provide more examples.

**Latent Motion Fusion Analysis.** We investigate different fusion strategies between the local motion  $h^i$  and global-aware motion  $g^i$  to obtain the final motion representation  $z^i$ . Specifically, we experiment with three settings: (i) No Fusion: the final motion is the global-aware motion obtained from our Cross-entity Attention ( $z^i = g^i$ ); (ii) Concatenate:



Figure 8. Examples of generated group motions from our method.

Dance Styles	Single-dance Metric			Group-dance Metric		
	FID↓	MMC↑	GenDiv↑	GMR↓	GMC↑	TIF↓
Zumba	45.86	0.268	9.77	50.97	72.70	0.133
Aerobic	38.68	0.252	6.57	63.62	75.12	0.249
Commercial	46.22	0.232	8.58	51.18	81.02	0.056
Bollywood	81.89	0.211	2.14	101.49	74.00	0.377
Irish	42.02	0.219	8.56	42.73	82.00	0.083
Rumba	69.62	0.273	3.91	68.00	71.85	0.228
Samba	71.00	0.228	7.77	98.83	67.76	0.441

Table 5. The results of different dance styles. Note that these results are obtained by training the model on each dance style.

Fusion Strategy	Single-dance Metric			Group-dance Metric		
	FID↓	MMC↑	GenDiv↑	GMR↓	GMC↑	TIF↓
No Fusion	47.19	0.242	9.14	57.84	69.67	0.221
Concatenate	52.25	0.223	9.23	54.23	72.46	0.242
<b>Add</b>	<b>43.90</b>	<b>0.250</b>	<b>9.23</b>	<b>51.27</b>	<b>79.01</b>	<b>0.217</b>

Table 6. Ablation study on different fusion strategies for the motion representation.

the final motion is the concatenation of the local and global-aware motion ( $z^i = [h^i; g^i]$ ); (iii) Add: the final motion is the addition between local and global ( $z^i = h^i + g^i$ ). Table 6 summarizes the results. We find that fusing the motion by adding both the local and global motion features achieves the best results.

## 6. Conclusion

We have introduced AIOZ-GDANCE, the largest dataset for audio-driven group dance generation. Our dataset contains in-the-wild videos and covers different dance styles and music genres. We then propose a strong baseline along with evaluation protocols for group dance generation task. We hope that the release of our dataset will foster more research on audio-driven group choreography.

## References

- [1] Vida Adeli, Ehsan Adeli, Ian Reid, Juan Carlos Niebles, and Hamid Rezatofighi. Socially and contextually aware human motion and pose forecasting. *RA-L*, 2020. 2
- [2] Hyemin Ahn, Jaehun Kim, Kihyun Kim, and Songhwai Oh. Generative autoregressive networks for 3d dancing move synthesis from music. *RA-L*, 2020. 2
- [3] Alexandre Alahi, Vignesh Ramanathan, and Li Fei-Fei. Socially-aware large-scale crowd forecasting. In *CVPR*, 2014. 2
- [4] Sarah Fdili Alaoui, Cyrille Henry, and Christian Jacquemin. Physical modelling for interactive installations and the performing arts. *International Journal of Performance Arts and Digital Media*, 10(2):159–178, 2014. 1
- [5] Omid Alemi, Jules Françoise, and Philippe Pasquier. Groovenet: Real-time music-driven dance movement generation using artificial neural networks. *networks*, 2017. 2
- [6] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 7
- [7] Samy Bengio, Oriol Vinyals, Navdeep Jaitly, and Noam Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. *Advances in neural information processing systems*, 28, 2015. 7
- [8] Daniel Bisig. Generative dance—a taxonomy and survey. In *Proceedings of the 8th International Conference on Movement and Computing*, 2022. 1
- [9] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 2
- [10] Kang Chen, Zhipeng Tan, Jin Lei, Song-Hai Zhang, Yuan-Chen Guo, Weidong Zhang, and Shi-Min Hu. Choreomaster: choreography-oriented music-driven dance synthesis. *ACM TOG*, 2021. 2, 4
- [11] Weifeng Chen, Zhao Fu, Dawei Yang, and Jia Deng. Single-image depth perception in the wild. *Advances in neural information processing systems*, 29, 2016. 4
- [12] Andreas Doering, Di Chen, Shanshan Zhang, Bernt Schiele, and Juergen Gall. Posetrack21: A dataset for person search, multi-object tracking and multi-person pose tracking. In *CVPR*, 2022. 4
- [13] Jonathan Eisenmann, Benjamin Schroeder, Matthew Lewis, and Rick Parent. Creating choreography with interactive evolutionary algorithms. In *European Conference on the Applications of Evolutionary Computation*, 2011. 2
- [14] Rukun Fan, Songhua Xu, and Weidong Geng. Example-based automatic music-driven conventional dance motion synthesis. *IEEE TVCG*, 2011. 2
- [15] Hao-Shu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. RMPE: Regional multi-person pose estimation. In *ICCV*, 2017. 3
- [16] Joao P Ferreira, Thiago M Coutinho, Thiago L Gomes, José F Neto, Rafael Azevedo, Renato Martins, and Erickson R Nascimento. Learning to dance: A graph convolutional adversarial network to generate realistic dance motions from audio. *Computers & Graphics*, 2021. 1, 2, 4
- [17] Bernhard Fink, Bettina Bläsing, Andrea Ravignani, and Todd K. Shackelford. Evolution and functions of human dance. *Evolution and Human Behavior*, 2021. 1
- [18] 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. *NeurIPS*, 2017. 7
- [19] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 1997. 5
- [20] Ruozi Huang, Huang Hu, Wei Wu, Kei Sawada, and Mi Zhang. Dance revolution: Long sequence dance generation with music via curriculum learning. *CoRR*, abs/2006.06119, 2020. 1, 2, 3, 7
- [21] Peter J Huber. Robust estimation of a location parameter. 1992. 4
- [22] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *CVPR*, 2020. 3
- [23] Manish Joshi and Sangeeta Chakrabarty. An extensive review of computational dance automation techniques and applications. *Proceedings of the Royal Society A*, 2021. 2
- [24] Pushpajit Khaire and Praveen Kumar. Deep learning and rgb-d based human action, human–human and human–object interaction recognition: A survey. *Journal of Visual Communication and Image Representation*, 86:103531, 2022. 2
- [25] Iris Kico, Nikos Grammalidis, Yiannis Christidis, and Fotis Liarokapis. Digitization and visualization of folk dances in cultural heritage: A review. *Inventions*, 3(4), 2018. 1
- [26] Jinwoo Kim, Heeseok Oh, Seongjean Kim, Hoseok Tong, and Sanghoon Lee. A brand new dance partner: Music-conditioned pluralistic dancing controlled by multiple dance genres. In *CVPR*, 2022. 1, 2
- [27] Jae Woo Kim, Hesham Fouad, and James K Hahn. Making them dance. In *AAAI Fall Symposium: Aurally Informed Performance*, 2006. 2
- [28] Tae-hoon Kim, Sang Il Park, and Sung Yong Shin. Rhythmic-motion synthesis based on motion-beat analysis. *ACM TOG*, 2003. 2
- [29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 7
- [30] Lucas Kovar, Michael Gleicher, and Frédéric Pighin. Motion graphs. In *SIGGRAPH*, 2002. 2
- [31] Nhat Le, Khanh Nguyen, Anh Nguyen, and Bac Le. Global-local attention for emotion recognition. *Neural Computing and Applications*, 2022. 5
- [32] Nhat Le, Khanh Nguyen, Quang Tran, Erman Tjiputra, Bac Le, and Anh Nguyen. Uncertainty-aware label distribution learning for facial expression recognition. In *WACV*, 2023. 5
- [33] Dong-Gyu Lee and Seong-Whan Lee. Human interaction recognition framework based on interacting body part attention. *Pattern Recognition*, 128:108645, 2022. 5
- [34] Hsin-Ying Lee, Xiaodong Yang, Ming-Yu Liu, Ting-Chun Wang, Yu-Ding Lu, Ming-Hsuan Yang, and Jan Kautz. Dancing to music. *NeurIPS*, 2019. 2, 3, 7
- [35] Lik-Hang Lee, Zijun Lin, Rui Hu, Zhengya Gong, Abhishek Kumar, Tangyao Li, Sijia Li, and Pan Hui. When creators

- meet the metaverse: A survey on computational arts. *CoRR*, abs/2111.13486, 2021. 1, 4
- [36] Minh Lee, Kyogu Lee, and Jaeheung Park. Music similarity-based approach to generating dance motion sequence. *Multimedia tools and applications*, 2013. 2
- [37] Buyu Li, Yongchi Zhao, and Lu Sheng. Dancenet3d: Music based dance generation with parametric motion transformer. In *AAAI*, 2022. 2, 3, 4
- [38] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *CVPR*, 2018. 5
- [39] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *CVPR*, 2021. 1, 2, 3, 4, 7, 8
- [40] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics*, 2015. 2, 3, 5
- [41] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion prediction. In *ICCV*, 2019. 5
- [42] Brian McFee, Colin Raffel, Dawen Liang, Daniel P Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. librosa: Audio and music signal analysis in python. In *Proceedings of the 14th python in science conference*, 2015. 5
- [43] Qianhui Men, Hubert PH Shum, Edmond SL Ho, and Howard Leung. Gan-based reactive motion synthesis with class-aware discriminators for human-human interaction. *Computers and Graphics*, 2022. 5
- [44] Binh X Nguyen, Binh D Nguyen, Tuong Do, Erman Tjiputra, Quang D Tran, and Anh Nguyen. Graph-based person signature for person re-identifications. In *CVPRW*, 2021. 2
- [45] Mathilde Papillon, Mariel Pettee, and Nina Miolane. Pirounet: Creating intentional dance with semi-supervised conditional recurrent variational autoencoders. *arXiv*, 2022. 4
- [46] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 3
- [47] Guillermo Valle Perez, Jonas Beskow, Gustav Henter, Andre Holzapfel, Pierre-Yves Oudeyer, and Simon Alexanderson. Transflower: probabilistic autoregressive dance generation with multimodal attention. 2021. 1, 2, 3
- [48] Trong-Thang Pham, Nhat Le, Tuong Do, Hung Nguyen, Erman Tjiputra, Quang D Tran, and Anh Nguyen. Style transfer for 2d talking head animation. *arXiv preprint*, 2023. 2, 5
- [49] Xuanchi Ren, Haoran Li, Zijian Huang, and Qifeng Chen. Self-supervised dance video synthesis conditioned on music. In *ACM MM*, 2020. 2
- [50] Alla Safonova and Jessica K Hodgins. Construction and optimal search of interpolated motion graphs. In *SIGGRAPH*, 2007. 2
- [51] Jianxing Shi. Application of 3d computer aided system in dance creation and learning. In *International Conference on Machine Learning and Big Data Analytics for IoT Security and Privacy*, 2021. 1
- [52] Takaaki Shiratori, Atsushi Nakazawa, and Katsushi Ikeuchi. Dancing-to-music character animation. In *Comput. Graph. Forum*, 2006. 2
- [53] Li Siyao, Weijiang Yu, Tianpei Gu, Chunze Lin, Quan Wang, Chen Qian, Chen Change Loy, and Ziwei Liu. Bailando: 3d dance generation by actor-critic gpt with choreographic memory. *CVPR*, 2022. 1, 2, 4
- [54] Alexandros Stergiou and Ronald Poppe. Analyzing human-human interactions: A survey. *Computer Vision and Image Understanding*, 188:102799, 2019. 2
- [55] Guofei Sun, Yongkang Wong, Zhiyong Cheng, Mohan S. Kankanhalli, Weidong Geng, and Xiangdong Li. Deepdance: Music-to-dance motion choreography with adversarial learning. *IEEE TMM*, 2020. 2, 3
- [56] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *CVPR*, 2019. 2
- [57] Peize Sun, Jinkun Cao, Yi Jiang, Zehuan Yuan, Song Bai, Kris Kitani, and Ping Luo. Dancetrack: Multi-object tracking in uniform appearance and diverse motion. In *CVPR*, 2022. 3, 4
- [58] Taoran Tang, Jia Jia, and Hanyang Mao. Dance with melody: An lstm-autoencoder approach to music-oriented dance synthesis. In *ACM MM*, pages 1598–1606, 2018. 2, 3
- [59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 5, 6
- [60] Jiashun Wang, Huazhe Xu, Medhini Narasimhan, and Xiaolong Wang. Multi-person 3d motion prediction with multi-range transformers. *NeurIPS*, 2021. 2
- [61] Nelson Yalta, Shinji Watanabe, Kazuhiro Nakadai, and Tetsuya Ogata. Weakly-supervised deep recurrent neural networks for basic dance step generation. In *IJCNN*, 2019. 1
- [62] Yiding Yang, Zhou Ren, Haoxiang Li, Chunluan Zhou, Xinchao Wang, and Gang Hua. Learning dynamics via graph neural networks for human pose estimation and tracking. In *CVPR*, 2021. 4
- [63] Zijie Ye, Haozhe Wu, Jia Jia, Yaohua Bu, Wei Chen, Fanbo Meng, and Yanfeng Wang. Choreonet: Towards music to dance synthesis with choreographic action unit. In *ACM MM*, 2020. 2
- [64] Wenjie Yin, Hang Yin, Kim Baraka, Danica Kragic, and Márten Björkman. Dance style transfer with cross-modal transformer. *arXiv*, 2022. 5
- [65] Xinjian Zhang, Yi Xu, Su Yang, Longwen Gao, and Huyang Sun. Dance generation with style embedding: Learning and transferring latent representations of dance styles. *arXiv*, 2021. 5
- [66] Wenlin Zhuang, Congyi Wang, Jinxiang Chai, Yangang Wang, Ming Shao, and Siyu Xia. Music2dance: Dancenet for music-driven dance generation. *ACM MM*, 2022. 2, 3
- [67] Yuliang Zou, Jimei Yang, Duygu Ceylan, Jianming Zhang, Federico Perazzi, and Jia-Bin Huang. Reducing footskate in human motion reconstruction with ground contact constraints. In *WACV*, 2020. 3