

# ChoreoNet: Towards Music to Dance Synthesis with Choreographic Action Unit

Zijie Ye\*<sup>†</sup>

yzjscwy@gmail.com  
Department of Computer Science  
and Technology, Tsinghua University,  
Beijing 100084, China

Haozhe Wu\*<sup>†</sup>

wuhz19@mails.tsinghua.edu.cn  
Department of Computer Science  
and Technology, Tsinghua University,  
Beijing 100084, China

Jia Jia\*<sup>†‡</sup>

jjia@tsinghua.edu.cn  
Department of Computer Science  
and Technology, Tsinghua University,  
Beijing 100084, China

Yaohua Bu\*<sup>†</sup>

18610082539@163.com  
Department of Computer Science  
and Technology, Tsinghua University,  
Beijing 100084, China

Wei Chen

chenweibj8871@sogou-inc.com  
Sogou, Inc.

Fanbo Meng

mengfanbosi0935@sogou-inc.com  
Sogou, Inc.

Yanfeng Wang

wangyanfeng@sogou-inc.com  
Sogou, Inc.

## ABSTRACT

Dance and music are two highly correlated artistic forms. Synthesizing dance motions has attracted much attention recently. Most previous works conduct music-to-dance synthesis via directly music to human skeleton keypoints mapping. Meanwhile, human choreographers design dance motions from music in a two-stage manner: they firstly devise multiple choreographic dance units (CAUs), each with a series of dance motions, and then arrange the CAU sequence according to the rhythm, melody and emotion of the music. Inspired by these, we systematically study such two-stage choreography approach and construct a dataset to incorporate such choreography knowledge. Based on the constructed dataset, we design a two-stage music-to-dance synthesis framework ChoreoNet to imitate human choreography procedure. Our framework firstly devises a CAU prediction model to learn the mapping relationship between music and CAU sequences. Afterwards, we devise a spatial-temporal inpainting model to convert the CAU sequence into continuous dance motions. Experimental results demonstrate that the proposed ChoreoNet outperforms baseline methods (0.622 in terms of CAU BLEU score and 1.59 in terms of user study score).

\*Key Laboratory of Pervasive Computing, Ministry of Education.

<sup>†</sup>Beijing National Research Center for Information Science and Technology.

<sup>‡</sup>Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

MM '20, October 12–16, 2020, Seattle, WA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7988-5/20/10...\$15.00

<https://doi.org/10.1145/3394171.3414005>

## CCS CONCEPTS

• **Applied computing** → **Media arts**;

## KEYWORDS

Dance Synthesis, Choreography, Motion Synthesis

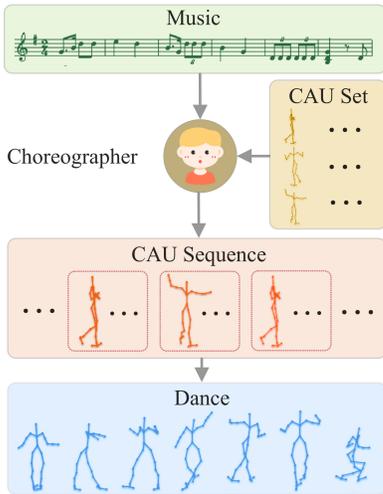
### ACM Reference Format:

Zijie Ye, Haozhe Wu, Jia Jia, Yaohua Bu, Wei Chen, Fanbo Meng, and Yanfeng Wang. 2020. ChoreoNet: Towards Music to Dance Synthesis with Choreographic Action Unit. In *Proceedings of the 28th ACM International Conference on Multimedia (MM '20)*, October 12–16, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3394171.3414005>

## 1 INTRODUCTION

A famous choreographer named Martha Graham once said, 'Dance is the hidden language of the soul'. As an artistic form with a long history, dance is an important medium for people to express their feelings. Conventionally, dance is always accompanied by music. Dancers start to dance when the musical atmosphere is going up at the beginning, perform different dance actions according to the rhythm, melody and emotion of music clips, and take a bow at the end. The complicated mapping relationship between dance and music has prompted researchers to investigate dance-to-music synthesis automatically.

Several previous research efforts have shown the rationality of music-to-dance synthesis [5, 19–21, 30]. Early works conduct music-to-dance synthesis via solving a similarity-based retrieval problem [5, 21], which shows limited capacity, while recent researches [19, 20, 30] leverage the deep learning methods to automatically learn the mapping between music and dance. Usually, prior methods solve music-to-dance synthesis directly by mapping music to human skeleton keypoints. On account of the highly redundant and noisy nature of skeleton keypoints, the frame by frame keypoints



**Figure 1: Human choreography procedure: firstly recollect CAUs and then arrange CAUs according to the rhythm, melody and emotion of the music.**

prediction can hardly capture the coherent structure of music and dance, resulting in limited synthesis quality.

To address this issue, we propose to fuse human choreography knowledge into the music-to-dance synthesis framework. In choreography knowledge, dance actions are composed of multiple indivisible units. We defined such indivisible unit as choreographic action unit (CAU). A CAU refers to a clip of human motions that lasts for several musical beats and acts as an undividable unit in choreography. During the choreography procedure shown in Figure 1, human choreographers usually devise dance actions in a hierarchical manner, they often recollect CAUs they have seen or used before and arrange CAUs according to the rhythm, melody and emotion of the music to formulate a piece of dance.

Pondering over such characteristics of choreography, we propose the ChoreoNet, a music-to-dance framework that imitates the hierarchical choreographic procedure of human beings. Our ChoreoNet has the following two characteristics: (1) The ChoreoNet firstly applies a CAU prediction model to predict the CAU sequence from musical features. Compared with the dance motion prediction in prior works (predict tens of thousands of frames per music), the prediction of CAU sequence (predict 40-80 CAUs per music) on the one hand prevents the neural network model from learning trivial motion details, on the other hand lessens the burden of predicting overlong sequences. (2) The ChoreoNet leverages a spatial-temporal inpainting model to convert the CAU sequence into continuous dance motions. Since that there always exists motion gaps between adjacent CAU pairs in the CAU sequence, the spatial-temporal inpainting model generates natural transitions between these CAU pairs. The overall implementation of ChoreoNet is two-stage, during each stage, the corresponding model is trained separately.

To evaluate the effectiveness of the ChoreoNet framework, we construct a music-to-dance dataset with expert choreographic annotations. Specifically, we collect 164 CAUs from 4 different dance

types (Cha Cha, Waltz, Rumba and Tango) and record 3D motions of each CAU. Then, we collect 62 dance music pieces and invite professional choreographers to annotate the CAU sequence of each music. Totally, we have 94 minutes of music with such annotations. Then we perform quantitative and qualitative experiments on our dataset. Experiments show that compared to baseline methods, our method can generate structured dances with a long duration that match better with the input music and more natural transitions between adjacent dance motions. Specifically, compared to baseline methods, our framework generates CAU sequences with higher BLEU score [24]. Our motion generation model also generates a motion transition closer to groundtruth. The dance animation generated by our framework also scores higher in the user study.

To summarize, our contributions are as two-fold:

- We propose to formalize the human choreography knowledge by defining CAU and introduce it into music-to-dance synthesis.
- We propose a two-stage framework ChoreoNet to implement the music-CAU-skeleton mapping. Experiments demonstrate the effectiveness of our method.

## 2 RELATED WORK

The previous works related to our ChoreoNet are described in two aspects: music to dance synthesis and human motion generation.

### 2.1 Music to Dance Synthesis

Several researches have focused on music to dance synthesis. Early works usually treat this problem as a mechanical template matching problem [5, 21, 28]. Cardle *et al.* [5] modified dance motion according to musical features, while Shiratori *et al.* [28] and Lee *et al.* [21] manually defined musical features and generate dance motions according to musical similarity. These template matching methods have limited capacity on generating natural and creative dance motions. Later, researchers start to address the dance-to-motions synthesis problem with deep learning techniques [7, 20, 27, 30-33]. Crnkovic-Friis *et al.* [7] firstly employ the deep learning methods to generate dance motions, they devise a Chor-RNN framework to predict dance motion from raw motion capture data. Then, Tang *et al.* [30] designed a LSTM-autoencoder to generate 3D dance motion. Previous research [27] also proposed to improve the naturalness of dance motion through perceptual loss [16]. However, the redundant and noisy motion keypoints still limit the quality of synthesized dance. Yalta *et al.* [31] has proposed to solve such issue through weakly supervised learning, whereas, the lack of human choreography experience still makes the generation quality less appealing.

### 2.2 Human Motion Generation

Human motion generation aims to generate natural human motion conditioned on existing motion capture data and plays a key role in music-to-dance synthesis. However, the highly dynamic, non-linear and complex nature of human motion makes this task challenging. Early researchers address this problem with concatenation-based method [2], hidden Markov models [29] and random forests [22]. Then, with the development of deep learning techniques, researchers applied deep learning techniques [1, 9, 10, 13, 15, 23, 26] to solve

motion generation problems. Ghosh *et al.* [10] and Fragkiadaki *et al.* [9] focused on RNN-based models with Euler angle error terms, utilizing the auto-regressive nature of human motion. Later, Pavlo *et al.* [26] addressed the impact of joint angle representation and showed the advantage of quaternion-based representation over Euler angle. Gui *et al.* [12] combined CNN-based model with adversarial training, and achieved better short-term generation quality over previous works. Ruiz *et al.* [13] formulated human motion generation as a spatial-temporal inpainting problem and designed a GAN framework to generate large chunks of missing data. In this work, we devise a similar spatial-temporal inpainting model to generate natural dance motion transitions between adjacent CAUs.

### 3 PROBLEM FORMULATION

Previous researches [7, 20, 27, 30–33] have formulated the music-to-dance synthesis as a music-to-keypoint mapping problem. However, directly predicting human skeleton keypoints raises a series of difficulties. On the one hand, the human skeleton keypoints are usually noisy and redundant, directly mapping music to these noisy keypoints would cause unstable synthesis results. On the other hand, each piece of music is usually accompanied by thousands of motion frames, predicting such an overlong sequence is much too challenging for current sequence model.

To address these issues, in this paper, we propose a two-stage music-to-dance synthesis formulation referring to human choreography knowledge. Conventionally, dance motions are composed of multiple indivisible units. We define such an indivisible units as choreographic action unit (CAU), details of CAU will be illustrated in section 4. During the human choreography procedure, choreographers seldomly improvise dance actions, instead they create dance motion in a two-stage manner: (1) choreographic action unit (CAU) design and (2) CAU sequence arrangement. In the first stage, human choreographers often recollect CAUs they have seen or used before. Then in the second stage, choreographers arrange CAUs to better fit the rhythm, melody and emotion of a given piece of music.

With the observation of such hierarchical choreography procedure, we propose to formulate the music-to-dance synthesis procedure as a two-stage framework. Our framework takes music  $X$  as input to generates human skeleton keypoints sequences  $C$ ,  $C \in \mathbb{R}^{N \times P}$ , where  $N$  is the number of motion frames, and  $P$  is the number of keypoints. The two stages are formally illustrated as follows:

**CAU Choreography Stage.** In this stage, given input music  $X$ , we aim to generate the corresponding CAU sequence  $\{y_1, \dots, y_n\}$ ,  $y_i \in Y$ , where  $Y$  is the overall CAU set.

**Motion Generation Stage.** Having obtained the CAU sequence  $\{y_1, \dots, y_n\}$ , in this stage, we aim to generate the keypoints sequence  $C$ . In this stage, although we have known the keypoints sequence  $C_i$  of particular CAU  $y_i$ , the transition between adjacent CAUs is unknown. Generating smooth and natural transition between adjacent CAUs is the main problem in this stage.

### 4 CHOREOGRAPHIC ACTION UNIT DEFINITION

CAU is short for *choreographic action unit*. One CAU refers to a clip of human motions that lasts for several musical beats and acts

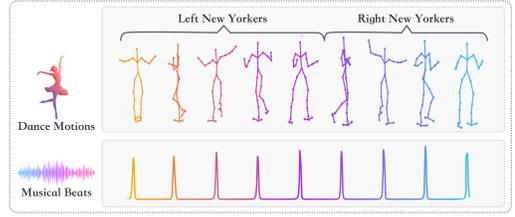


Figure 2: An example of CAUs in an 8 beat motion segment.

as an undividable unit in choreography. Human choreographers recollect CAUs they have seen or used before and arrange them to formulate new dances. We ask professional dancers to design CAUs and record their performance of each CAU with motion capture devices. In this work, we denote a CAU as  $y$ ,  $y$  corresponds to  $len(y)$  musical beats and  $C^y$  is the corresponding motion capture data performed by professional dancers.

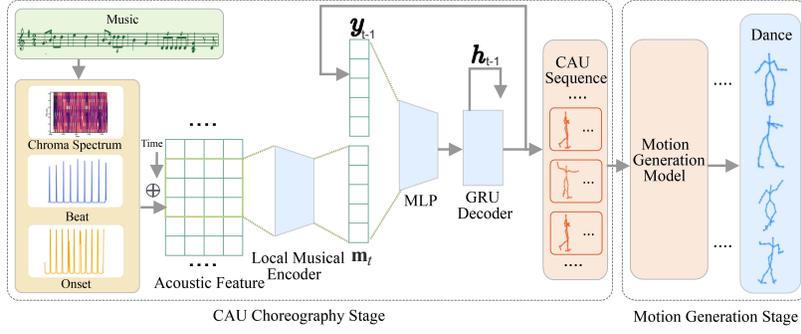
Figure 2 shows the procedure of performing CAUs from a motion segment that lasts for 8 musical beats. At beat 1, the dancer pivots on the left foot and conducts a quarter turn. At beat 2, the dancer checks forward and spread arms. Afterwards, the dancer moves backward and turn to the original position at beat 3 and 4. Because the dance motions in the first 4 beats make up a whole action unit and can not be divided into smaller units in choreography, it is recognized as a CAU named the *Left New Yorkers*. After the *Left New Yorkers*, the dancer performs the *Right New Yorkers* in the last 4 beats of this motion segment. The dance motions in the last 4 beats are also recognized as a CAU. So totally we have two CAUs from this motion segment: the *Left New Yorkers* CAU and the *Right New Yorkers* CAU.

## 5 METHODOLOGY

Following the formulation proposed in section 3, we propose a two-stage framework ChoreoNet, as shown in Figure 3. In the first stage, we generate appropriate choreographic action unit (CAU) sequences from musical features with an encoder-decoder model. Based on the predicted CAU sequence, in the second stage, we leverage a spatial-temporal inpainting model to convert CAU sequences to dance motions (human skeleton keypoint coordinates). In the next two subsections, we will introduce the CAU prediction model and the motion generation model respectively.

### 5.1 CAU Prediction Model

In the first stage of ChoreoNet, we design our CAU prediction model to predict CAU sequence from the input music. As explained above, CAUs should match with music clips and adjacent CAUs, so our CAU prediction model needs to consider both musical context and CAU context when deciding the next appropriate CAU. In order to model musical context and CAU context simultaneously, we adopt an encoder-decoder model as the CAU prediction model. In this model, we first encode musical context with a temporal convolution network, then we feed the encoded musical features to the decoder model. Our decoder model takes musical features and the CAU history as input and outputs the next CAU. To model the time dependency, we apply gated recurrent unit (GRU) [6] as our



**Figure 3: The pipeline of the ChoreoNet.** In CAU choreography stage, firstly we extract the deep chroma spectrum, beat and onset of input music. We concatenate the extracted features to formulate acoustic features. Then the local musical encoder computes the encoded musical features  $m_t$  in a slide window scheme. The GRU decoder combines  $m_t$  and CAU history to predict the next CAU  $y_t$ . In motion generation stage, the motion generation model recovers motion data from CAU sequences. Details of the motion generation model would be illustrated in Figure 4.  $\oplus$  denotes the concatenation operator.

decoder model. The procedure of CAU prediction is shown in Figure 3. Next, we’ll explain each part of the CAU prediction model in detail.

For the encoder of the CAU prediction model, we leverage a temporal CNN to encode music frames. We propose to design the encoder as a local music encoder (encode a local clip of music) rather than a global encoder (encode the whole input music). This is because the CAU sequence is much sparser than music frames. One piece of music is usually coupled with only dozens of CAUs but has thousands of music frames. The imbalance between CAUs and music frames causes difficulty in training a global musical encoder between the encoder needs much more capacity than the decoder and is hard to converge. Afterwards, we adopt the following three musical features as the input of the encoder: (1) deep chroma spectrum [18] (2) beat [4] and (3) onset [8]. We respectively denote chroma spectrum as  $X^c$ , beat as  $X^b$  and onset as  $X^s$ . At time  $t$ , the encoder applies temporal convolution on raw musical features within a fixed-length time window  $[t - w, t + w]$ :

$$m_t = \text{encode}(X_{[t-w, t+w]}^c, X_{[t-w, t+w]}^b, X_{[t-w, t+w]}^s), \quad (1)$$

where  $m_t$  represents the encoded musical feature at time  $t$ .

Having obtained the encoded musical feature  $m_t$ , we now model the CAU generation process of the decoder as a probability distribution conditioned on musical context and CAU context. Given the CAU history  $y_1, \dots, y_{t-1}$  and the encoded musical feature  $m_t$ , the distribution of next CAU  $y_t$  at time  $t$  is described as:

$$p(y_t) = p(y_t | y_1, \dots, y_{t-1}, m_t). \quad (2)$$

We model the conditional distribution with a GRU, as shown in Figure 3.  $y_{t-1}$  and  $m_t$  are fused with an MLP before feeding to the GRU.  $m_t$  represents the musical context, while  $y_{t-1}$  and the GRU’s hidden state contains the information of CAU history. In this way, our model has access to both musical context and CAU context.

Combining the two parts of the CAU prediction model, we adopt a sliding-window scheme to generate a CAU sequence from the input music: (1) raw musical features within the time window  $[t - w, t + w]$  is fed to the encoder and we get the encoded musical feature  $m_t$ , (2) last predicted CAU  $y_p$  and  $m_t$  are fed to the decoder

to generate  $\hat{y}_t$ , (3)  $t$  is updated by  $t \leftarrow t + \text{len}(\hat{y}_t)$ .  $t$  is initially set to 0 and the whole generation process is repeated until an end annotation is generated by the decoder or we meets the end of the input music. Algorithm 1 shows the whole procedure of the CAU generation.

---

#### Algorithm 1 CAU Generation

---

```

1:  $t \leftarrow 0$ 
2:  $Y_{gen} = []$ 
3:  $y_p = \text{StartOfDance}$ 
4: while  $y_p \neq \text{EndOfDance}$  AND the music is not ended do
5:    $m_t = \text{encode}(X_{[t-w, t+w]}^c, X_{[t-w, t+w]}^b, X_{[t-w, t+w]}^s)$ 
6:    $\hat{y} = \text{decode}(m_t, y_p)$ 
7:   Add  $\hat{y}$  to  $Y_{gen}$ 
8:    $y_p \leftarrow \hat{y}$ 
9:    $t \leftarrow t + \text{len}(\hat{y})$ 
return  $Y_{gen}$ 

```

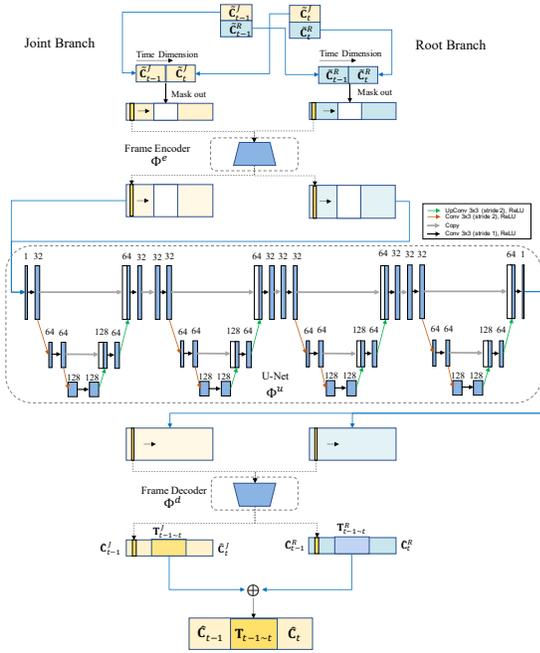
---

Having obtained the predicted distribution  $p(y_t)$  at each step, we train the CAU prediction model by maximizing the log-likelihood of the expert annotations in the dataset. Given expert annotation  $y_t^p$  at time  $t$ , we minimize the negative log-likelihood loss:

$$\mathcal{L}_{nll} = - \sum_{t=1}^{t=N} \log(p(y_t = y_t^p)). \quad (3)$$

## 5.2 Motion Generation Model

Having obtained the CAU sequence generated from the input music, we convert the CAU sequence to dance motions. Here, dance motions refer to human skeleton coordinates that can be directly used to drive humanoid animations. For each CAU  $y_i$ , we collect a clip of motion capture data  $C_i = \{C_i^1, \dots, C_i^{N_i}\}$  performed by professional dancers,  $C_i$  consists of  $N_i$  frames and each pose frame is denoted as  $C_i^j$ . Each frame  $C_i^j$  is made up of two parts: (1) the rotation of each joint related to its parent joint, (2) the translation and rotation of the root point (the hip) related to the world coordinate. We adopt quaternion-based joint rotation representation



**Figure 4: The motion generation model. We generate human joint rotation and root point respectively. The frame encoder firstly projects each frame into the latent vector space. Then the U-Net inpaints the transition gap. Finally, the frame decoder projects each frame back into the original parameter space. Note that the models of the two branches do not share parameters.**

according to Pavlo *et al.* [26]. Thus  $C_i^J$  consists of  $P = P_r + P_j$  parameters, where  $P_r$  represents the number of root point parameters while  $P_j$  represents the number of joint rotation parameters. The motion generation problem is formulated as follows: given  $\{C_1, \dots, C_n\}$ ,  $C_i \in \mathbb{R}^{N_i \times P}$ , we aim to generate  $C \in \mathbb{R}^{N \times P}$ ,  $N = \sum_{i=1}^n N_i$ .

For each clip of motion capture data, a large part of  $C_i$  can be directly used to drive humanoid animation. However, the motion transition between adjacent CAUs needs to be natural and smooth. To address this issue, we devise a motion generation model to inpaint dance motions. Before feeding motion capture data to our model, we preprocess  $C_i$  in two steps: (1) align the root points of adjacent CAUs to avoid sudden change of root position and root rotation, (2) align kinematic beats (the beats of dance motions) with musical beats. After these steps,  $\{C_1, \dots, C_n\}$  is converted to  $\tilde{C} = \{\tilde{C}_1, \dots, \tilde{C}_n\}$ . Afterwards, we propose to inpaint  $\tilde{C}$  to generate natural and smooth motion transition between adjacent CAUs. One simple solution could be linear blending. However, linear blending generates unnatural transition when there's huge gap between adjacent CAUs. The reason behind is that human motion is highly complicated and non-linear in nature. There exists sudden acceleration, deceleration and body turns in dance motions. To address this issue, we design a spatial-temporal inpainting model shown in Figure 4 inspired by Ruiz *et al.* [13].

The general idea of spatial-temporal inpainting is similar to image inpainting. Given two clips of motion capture data  $\tilde{C}_{t-1}, \tilde{C}_t$ , the model generates motion transition and its contexts:  $\{\hat{C}_{t-1}, T_{t-1-t}, \hat{C}_t\}$ .  $T_{t-1-t}$  is the generated motion transition, while  $\hat{C}_{t-1}$  and  $\hat{C}_t$  are predicted contexts. Usually,  $\hat{C}_{t-1}$  and  $\hat{C}_t$  are a slightly modified version of  $\tilde{C}_{t-1}$  and  $\tilde{C}_t$ . After generating transition for each pair of adjacent CAUs, we get  $C$  by concatenating  $\hat{C}_i$  and  $R_i$ .

Our spatial-temporal inpainting model is made up of two sub-models: (1) human joint rotation inpainting model and (2) root point inpainting model. The two sub-models share the same structure but with different parameters. The workflow of our model is shown in Figure 4. The human joint rotation inpainting model takes only the joint rotation parameters as input, denoted as  $\tilde{C}_{t-1}^J$  and  $\tilde{C}_t^J$  in Figure 4. We first concatenate these two matrices and mask out the transition window between the two motion segments by setting parameters within this window to zeros. Then we project each frame of the motion segment to a latent vector space  $\mathbb{R}^M$  through a frame encoder  $\Phi^e$ . After that, a 2D U-Net  $\Phi^u$  is applied to inpaint the masked matrix. At last, a frame decoder  $\Phi^d$  projects each frame of the inpainted matrix from  $\mathbb{R}^M$  to  $\mathbb{R}^J$ , where  $J$  is the number of joint rotation parameters. The output of  $\Phi^d$  consists of  $\hat{C}_{t-1}^J$ ,  $T_{t-1-t}^J$  and  $\hat{C}_t^J$ . Similarly, the root point inpainting model takes the rotation and the velocity of the root point as input, denoted as  $\tilde{C}_{t-1}^R$  and  $\tilde{C}_t^R$ . The workflow of the root point inpainting model is the same as that of the joint rotation inpainting model.  $\hat{C}_{t-1}$ ,  $T_{t-1-t}$  and  $\hat{C}_t$  are the concatenation of the outputs from the human joint inpainting model and root inpainting model.

Note that there are two characteristics in the design of our motion generation model. First, we adopt two sub-models for human joint rotation inpainting and root point inpainting respectively rather than a single model. This is because joint rotation and root velocity are of two different parameter spaces, it's hard for one inpainting model to operate on two parameter spaces. Second, we project each frame of motion to a larger space of dimension  $M$  before feeding the frame into the 2D-U-Net. That's because 2D-U-Net would conduct 2D convolution on the input matrix, while the motion frame matrix  $\tilde{C}_i$ , different from the traditional image matrix, is continuous on the time dimension but not on the motion parameter dimension. Thus we cannot directly perform 2D convolution on the frame matrix. To address this issue, we leverage a frame encoder  $\Phi^e$  to project each frame to a motion embedding space, so as to guarantee the continuation property in the embedding space. Extensive experiments verify the effectiveness of the two designs.

During training, we adopt the motion capture data as groundtruth. Each time, we randomly clip two consecutive clips of motion capture data denoted as  $\{C_1, C_2\}$ ,  $C_i \in \mathbb{R}^{N_i \times P}$ . Then we feed  $\{C_1, C_2\}$  to the motion generation model, the human joint rotation inpainting model outputs  $\hat{C}^J = \text{concat}(\hat{C}_1^J, T_{1-2}^J, \hat{C}_2^J) \in \mathbb{R}^{(N_1+N_2) \times P_j}$ , the root point inpainting model outputs  $\hat{C}^R = \text{concat}(\hat{C}_1^R, T_{1-2}^R, \hat{C}_2^R) \in \mathbb{R}^{(N_1+N_2) \times P_r}$ . Afterwards, we minimize the distance between the outputs and the groundtruth  $C_{gnd} = \text{concat}(C_1, C_2) \in \mathbb{R}^{(N_1+N_2) \times P}$ . The two sub-models define two different distance function (i.e. loss function) respectively: (1) joint rotation loss  $\mathcal{L}_{joint}$ , (2) root point loss  $\mathcal{L}_{root}$ . For joint rotation loss, we adopt geodesic loss proposed

by Gui *et al.* [12]. Given two rotation matrices  $R$  and  $\hat{R}$ , the geodesic distance is defined using the logarithm map in  $SO(3)$ :

$$d_{geo}(R, \hat{R}) = \|\log(R\hat{R}^T)\|_2. \quad (4)$$

Summing up geodesic distances between the inpainted frames and the groundtruth frames, we obtain joint rotation loss:

$$\mathcal{L}_{joint} = \sum_{k=1}^{k=N_1+N_2} \sum_{j=1}^{j=J} d_{geo}(R_k^j, \hat{R}_k^j), \quad (5)$$

where  $R_k^j$  represents the rotation matrix of the  $j$ -th joint in the  $k$ -th frame of  $C_{gnd}$ , while  $\hat{R}_k^j$  represents the rotation matrix of the  $j$ -th joint in the  $k$ -th frame of  $\hat{C}^J$ . For root point loss, we adopt L-1 distance as loss function:

$$\mathcal{L}_{root} = \sum_{i=1}^{i=N_1+N_2} \|C_i^R - \hat{C}_i^R\|, \quad (6)$$

where  $C_i^R$  is the  $i$ -th frame of  $C_{gnd}^R$ , while  $\hat{C}_i^R$  represents the  $i$ -th frame of  $\hat{C}^R$ . The two sub-models are trained separately.

To summarize, we devise a two-stage framework to apply choreography experience into music-dance synthesis. In the first stage, a CAU prediction model arranges a CAU sequence according to the input music. The CAU prediction model is designed as a CAU prediction model so as to consider both the musical context and the CAU context when arranging CAUs. In the second stage, we convert the CAU sequence to human motions that can directly drive humanoid animations. A spatial-temporal inpainting model is used to generate natural and smooth transition between adjacent dance motions.

## 6 EXPERIMENT

In this section, we conduct extensive experiments to demonstrate the effectiveness of our framework. We evaluate our framework on the CAU annotation dataset. Our framework has shown better performance compared with baseline methods, both qualitatively and quantitatively. Afterwards, we verify the effectiveness of our motion generation model with quantitative experiments.

### 6.1 Experiment Setup

**Dataset.** The expert choreographic annotations we collect come as pairs of music pieces and CAU sequences. Specifically, we collect 62 pieces of dancing music and 164 types of CAUs. For each piece of music, we invite professional choreographers to choreograph for each piece of music and record their choreography as CAU sequences. We annotate each CAU in the CAU sequence with its start time and end time in the music. These annotations include four types of dance (Waltz, Tango, Cha Cha, and Rumba) with a total of 94 minutes of music.

The CAU annotations used in this model consist of all the types of CAU we collected from expert choreographers and three other special annotations: [SOD], [EOD] and [NIL], summing up to a total of 167 annotations. [SOD] and [EOD] represent *start of dance* and *end of dance* respectively, performing similar functions to those of [SOS] and [EOS] annotations used in NLP problems. [NIL] annotation is introduced to imitate the decision of start time in choreography. One [NIL] annotation represents waiting for one musical

beat. One possible CAU sequence generated by our model could be: [NIL, NIL, C-1-3, C-18-1]. The first two [NIL] annotations represents that the dance should start after the first two beats, while [C-1-3] and [C-18-1] are normal CAUs collected from expert choreographers.

Meanwhile, we ask professional dancers to perform dances consists of the aforementioned CAUs and record their performances through motion capture devices. Then we crop segments corresponding to each CAU from the motion capture data to form a motion dataset. The FPS (frame per second) of our motion capture data is 80, we collect a total of 12688 frames of motion capture data. During training, each time, we clip a segment of 192 frames randomly from the motion capture data and set the motion blending window size to 64.

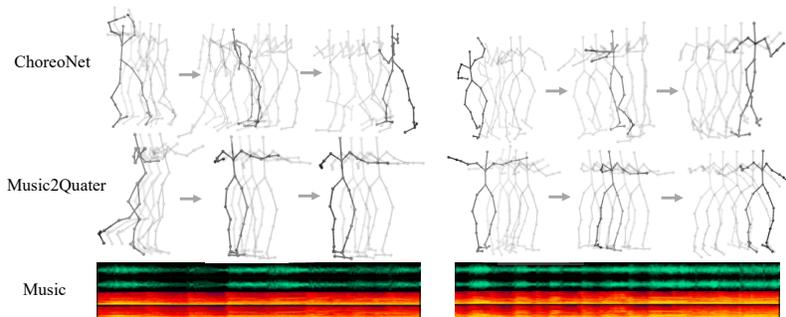
**Implementation Details.** We first extract acoustic features from raw input music using *Madmom* [3] toolkit. Specifically, we extract deep chroma spectrum [18]  $X^c$ , beat feature [4]  $X^b$  and onset feature [8]  $X^s$ . The frame-per-second (FPS) of  $X^c$  is 10 while the FPS of  $X^b$  and  $X^s$  is 100. The dimension of each frame of  $X^c$  is 12, while the other two features are both 1D vectors.

Then for the network architecture, the encoder of the CAU prediction model contains 5 1D convolutional layers, and we choose ReLU as activation function. The sliding window size of this local musical feature encoder is set to 10 seconds. The acoustic features are firstly convolved through the 5 convolutional layers and then fed to an MLP layer to output a 64-dimension local musical feature. The embedding dimension of CAU is set to 128, and the decoder of the CAU prediction model consists of an MLP layer and a GRU with 64 units.

Before feeding the recorded motion capture segments of the generated CAU sequence to our motion generation model, we align the kinematic beats with musical beats. A kinematic beat refer to the sudden motion deceleration. We detect the sudden deceleration of human limbs and mark the times of such deceleration as kinematic beats. Afterwards, we align them with musical beats detected by *Madmom* [3].

The frame encoder  $\Phi^e$  and the frame decoder  $\Phi^d$  of the motion generation model are both 1D convolutional networks with 6 convolutional layers and we adopt ReLU as their activation function. The dimension of motion embedding space is set to 84. The 2D U-Net  $\Phi^u$  consists of 4 blocks, each block is made up of 4 downsample layers and 4 upsample layers. The number of hidden states of each block is set to 32. In practice, we mask out the middle part of a clip of motion capture data and then feed the masked clip to the motion generation model to reconstruct the original motion. The size of  $G_i$  are set to 64 frames, i.e. the model takes 192 frames of motion data and outputs 192 frames of inpainted motion data.

We adopt RMSprop algorithm [11] to train our CAU prediction model for 1000 epochs with an initial learning rate at  $10^{-3}$ . We apply *ReduceOnPlateau* learning rate decay strategy provided by PyTorch [25] with the patience set to 8 and the decay factor set to 0.9. For the motion generation model, we adopt Adam algorithm [17] as optimizer and train the model for 400 epochs with a mini-batch size of 48 samples and a initial learning rate at  $10^{-3}$ . *ReduceOnPlateau* strategy is also applied, with the patience set to 5 and the decay factor set to 0.7.



**Figure 5: Comparison of Generated Dances.** The dances generated by Music2Quater (directly music-to-dance mapping baseline) collapse to a static pose after the first few seconds of music, while our ChoreoNet can generate different motions according to the input music.

**Table 1: Comparison of Different Methods**

Methods	BLEU-4	Music Matchness	Motion Naturalness	Motion Multimodality
Music2Quater	-	2.47	2.31	2.09
One-to-one CAU Prediction	0.0823	-	-	-
ChoreoNet Non-recurrent	0.539	3.78	<b>4.05</b>	3.80
ChoreoNet	<b>0.704</b>	<b>3.88</b>	3.94	<b>3.82</b>

## 6.2 Metrics

To evaluate the effectiveness of our framework, we adopt different metrics for the CAU prediction model and the motion generation model respectively. For the CAU prediction model, BLEU score [24] is used for quantitative comparison. BLEU score is originally designed to evaluate the quality of machine-translated text. In our experiment, we adopt it to evaluate the quality of the generated CAU sequences. For the motion generation model, we evaluate the geodesic distance and FID [14]. For geodesic distance, we evaluate the human joint rotation distance between the generated transition and the groundtruth. For FID, we evaluate the distance between the feature of the generated human motions and the feature of the real human motions. We train a motion auto-encoder on our motion capture dataset as the feature extractor for our FID evaluation.

We also conduct a user study to evaluate the quality of the dances generated by our framework. Participants are required to rate the following factors from 0 to 5: (1) the matchness between dance and music, (2) the naturalness of the dance motions, (3) the multimodality of the dances.

## 6.3 Comparison with Baselines

In this section, we compare the following baseline methods: (1) Music2Quater. This method directly maps musical features to human skeleton keypoints. It’s similar to the LSTM-autoencoder proposed by Tang *et al.* [30], but replaces LSTM with GRU and adopts quaternion as motion representation. (2) One-to-one CAU prediction. This method takes musical features as input and outputs CAU sequences without referring to the previous CAU context. It’s implemented by replacing the decoder of our CAU prediction model with a simple MLP layer. (3) ChoreoNet non-recurrent. This method

replaces the decoder of our CAU prediction model with a CNN and takes the last predicted CAU as input. Because the new decoder does not have hidden state, it’s only capable to capture the local CAU context rather than the whole CAU context. (4) The proposed ChoreoNet framework.

**Quantitative Evaluations.** We perform a quantitative evaluation of the quality of the CAU sequences generated from the input music. Specifically, we select 5 pieces of dancing music that are not in our training dataset. Afterwards, we evaluate the BLEU scores of CAU sequences generated from these music. The BLEU score measures the similarity between the generated CAU sequences and the CAU sequences we collected from professional choreographers. Table 1 shows the evaluation results. The BLEU score of Music2Quater is not given because it does not produce CAU sequences in the dance motion synthesis process. Overall, we can see the CAU context is crucial to CAU prediction. The one-to-one CAU prediction method ignores CAU context produces poor results. Replacing the GRU decoder with a CNN would also cause performance degradation. Our framework outperforms the ChoreoNet non-recurrent baseline by 0.165 with 0.704 BLEU score on our CAU annotation dataset.

**Qualitative Evaluations.** We compare the quality of dances synthesized by different methods. Dances generated by Music2Quater tend to collapse to a static pose after first few seconds. The noisy and highly redundant nature causes difficulty for the model to directly map musical feature to human skeleton keypoints. For ChoreoNet w/o CAU context, the generated dances are made up of some random actions, while ChoreoNet with local CAU context is able to generate valid dances but tend to produce some repeated actions. This proves that CAU context is crucial to our framework. To further investigate the quality of synthesized dances, we also

**Table 2: Comparison with Linear Blending.** We evaluate the geodesic distances and FID of the linear blending baseline and our spatial-temporal inpainting method with different blending window sizes.

Methods	Geodesic Distance	FID
Linear Blending (16 window)	$1.72 \times 10^{-2}$	85.1
Ours (16 window)	$1.24 \times 10^{-3}$	76.9
Linear Blending (32 window)	$1.80 \times 10^{-2}$	75.8
Ours (32 window)	$1.31 \times 10^{-3}$	67.8
Linear Blending (64 window)	$1.95 \times 10^{-2}$	69.7
Ours (64 window)	$1.33 \times 10^{-3}$	63.9
Linear Blending (128 window)	$2.08 \times 10^{-2}$	105.4
Ours (128 window)	$1.97 \times 10^{-3}$	78.4

conduct a user study. 17 participants are asked to rate 'matchness with music', 'motion naturalness' and 'motion multimodality' of each synthesized dances. The results of the user study is shown in Table 1. From the results, we can see that compared to the Music2Quater baseline, our ChoreoNet framework scores 1.41 points higher in *Music Matchness*, 1.63 points higher in *Motion Naturalness* and 1.73 points higher in *Motion Multimodality*. The results confirm our observation that the ChoreoNet framework generates dances of higher quality than directly music-to-skeleton mapping methods.

#### 6.4 Analysis on Motion Generation Model

In order to evaluate the performance of our spatial-temporal inpainting model, in this section, we study it from three aspects. First, we compare it with linear blending baseline. Next, we conduct ablation study to verify the effectiveness of the frame encoder/decoder and the two sub-model design. Finally, we adjust the blending window size to analyze the impact of different blending window sizes.

**Comparison with Baseline.** To verify the effectiveness of our spatial-temporal inpainting model, we compare it with the linear blending baseline. This is a solution used by previous researchers to concatenate two adjacent segments of motion data. Specifically, we conduct linear blending on human joint rotation within the blending window. The results are shown in Table 2. Our method achieves produce motion transition with much smaller geodesic distances than the baseline.

**Ablation Study.** To verify the effectiveness of our spatial-temporal inpainting model, we compare it with the following methods. (1) Ours w/o sub-model. This method ablates the two sub-model design, there is only one model to inpaint both human joint rotation and the root point. (2) Ours w/o frame encoder. This method ablates the frame encoder and decoder, all the motion inpainting is conducted in the original motion parameter space. We calculate the geodesic distance per frame between the generated motion transition and the groundtruth to evaluate the effectiveness of these methods.

Table 3 shows the test geodesic distances of all the methods. The results show that the designs of frame encoder/decoder and sub-models increase the performance by a large margin.

**Table 3: Comparison with Ours w/o frame encoder and Ours w/o sub-model.** The blending window size is set to 64.

Methods	Geodesic Distance
Ours w/o frame encoder	$2.02 \times 10^{-3}$
Ours w/o sub-model	$5.14 \times 10^{-3}$
Ours	$1.33 \times 10^{-3}$

**Motion Blending Window Size Analysis.** The size of motion blending window is an important hyper-parameter. Small blending window would cause quick and unnatural motion transition while big blending window causes difficulty to our model as it needs more capacity to inpaint a larger segment of missing motion transition. The same holds true for the linear blending baseline method. As it has no other parameters, both big or small blending window would cause undesirable results. To evaluate the quality of generated motion transition, we adopt FID [14] to measure the distance between the generated motion transition and the real motion capture data. As there exists no standard feature extractor for motion data, we train a motion auto-encoder on our motion capture dataset as motion feature extractor.

From the results shown in Table 2, we observe that the minimum FID of our model occurs at blending window size of 64 frames. Both small transition window size and large transition window size increase the FID. The results confirm the observation that proper transition window size leads to better transition quality.

## 7 CONCLUSION

In this paper, we formulate music-to-dance synthesis as a two-stage procedure to introduce human choreography experience. Firstly, we define choreographic action unit (CAU) and build a dataset containing 62 pieces of dancing music and 164 types of CAUs. Each piece of music is coupled with human expert CAU annotations. Based on the dataset, we propose a music-to-dance synthesis framework to implement the two-stage music-to-dance synthesis procedure. In the first stage, a CAU prediction model is used to generate CAU sequences from musical features. Then we apply a spatial-temporal inpainting model to generate dance motions from CAU sequences. We conduct extensive experiments to verify the effectiveness of our framework. The results show that compared to baseline methods, our CAU prediction model generate CAU sequences of higher quality and our spatial-temporal inpainting model produce more natural and smoother motion transition.

Overall, our framework improves the music-to-dance synthesis performance by a large margin. Furtherly, the proposed CAU-based formulation paves a new way for future research.

## ACKNOWLEDGEMENTS

This work is supported by National Key R&D Program of China (2019AAA0105200), the state key program of the National Natural Science Foundation of China (NSFC) (No.61831022), and the National Natural, and Science Foundation of China (61521002). We also want to thank Tiangong Institute for Intelligent Computing, Tsinghua University and Beijing Academy of Artificial Intelligence (BAAI) for their support.

## REFERENCES

- [1] Emre Aksan, Manuel Kaufmann, and Otmar Hilliges. 2019. Structured prediction helps 3d human motion modelling. In *Proceedings of the IEEE International Conference on Computer Vision*. 7144–7153.
- [2] Okan Arıkan, David A Forsyth, and James F O’Brien. 2003. Motion synthesis from annotations. In *ACM SIGGRAPH 2003 Papers*. 402–408.
- [3] Sebastian Böck, Filip Korzeniowski, Jan Schlüter, Florian Krebs, and Gerhard Widmer. 2016. Madmom: A new python audio and music signal processing library. In *Proceedings of the 24th ACM international conference on Multimedia*. 1174–1178.
- [4] Sebastian Böck and Markus Schedl. 2011. Enhanced beat tracking with context-aware neural networks. In *Proc. Int. Conf. Digital Audio Effects*. 135–139.
- [5] Marc Cardle, Loic Barthe, Stephen Brooks, and Peter Robinson. 2002. Music-driven motion editing: Local motion transformations guided by music analysis. In *Proceedings 20th Eurographics UK Conference*. IEEE, 38–44.
- [6] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 1724–1734.
- [7] Luka Crnkovic-Friis and Louise Crnkovic-Friis. 2016. Generative Choreography using Deep Learning. In *Proceedings of the Seventh International Conference on Computational Creativity, UPMC, Paris, France, June 27 - July 1, 2016*, François Pachet, Amílcar Cardoso, Vincent Coruble, and Fiammetta Ghedini (Eds.). Sony CSL Paris, France, 272–277. <http://www.computationalcreativity.net/iccc2016/wp-content/uploads/2016/01/Generative-Choreography-using-Deep-Learning.pdf>
- [8] Florian Eyben, Sebastian Böck, Björn Schuller, and Alex Graves. 2010. Universal onset detection with bidirectional long-short term memory neural networks. In *Proc. 11th Intern. Soc. for Music Information Retrieval Conference, ISMIR, Utrecht, The Netherlands*. 589–594.
- [9] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. 2015. Recurrent network models for human dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*. 4346–4354.
- [10] Partha Ghosh, Jie Song, Emre Aksan, and Otmar Hilliges. 2017. Learning human motion models for long-term predictions. In *2017 International Conference on 3D Vision (3DV)*. IEEE, 458–466.
- [11] Alex Graves. 2013. Generating Sequences With Recurrent Neural Networks. *arXiv preprint arXiv:1308.0850* (2013).
- [12] Liang-Yan Gui, Yu-Xiong Wang, Xiaodan Liang, and José MF Moura. 2018. Adversarial geometry-aware human motion prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 786–803.
- [13] Alejandro Hernandez, Jurgen Gall, and Francese Moreno-Noguer. 2019. Human motion prediction via spatio-temporal inpainting. In *Proceedings of the IEEE International Conference on Computer Vision*. 7134–7143.
- [14] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett (Eds.). 6626–6637. <http://papers.nips.cc/paper/7240-gans-trained-by-a-two-time-scale-update-rule-converge-to-a-local-nash-equilibrium>
- [15] Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. 2016. Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5308–5317.
- [16] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*. Springer, 694–711.
- [17] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1412.6980>
- [18] Filip Korzeniowski and Gerhard Widmer. 2016. Feature Learning for Chord Recognition: The Deep Chroma Extractor. In *Proceedings of the 17th International Society for Music Information Retrieval Conference, ISMIR 2016, New York City, United States, August 7-11, 2016*, Michael I. Mandel, Johanna Devaney, Douglas Turnbull, and George Tzanetakis (Eds.). 37–43.
- [19] Hsin-Ying Lee, Xiaodong Yang, Ming-Yu Liu, Ting-Chun Wang, Yu-Ding Lu, Ming-Hsuan Yang, and Jan Kautz. 2019. Dancing to Music. In *Advances in Neural Information Processing Systems*. 3581–3591.
- [20] Juheon Lee, Seohyun Kim, and Kyogu Lee. 2018. Listen to dance: Music-driven choreography generation using autoregressive encoder-decoder network. *arXiv preprint arXiv:1811.00818* (2018).
- [21] Minho Lee, Kyogu Lee, and Jaehung Park. 2013. Music similarity-based approach to generating dance motion sequence. *Multimedia tools and applications* 62, 3 (2013), 895–912.
- [22] Andreas M Lehrmann, Peter V Gehler, and Sebastian Nowozin. 2014. Efficient nonlinear markov models for human motion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1314–1321.
- [23] Julieta Martinez, Michael J Black, and Javier Romero. 2017. On human motion prediction using recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2891–2900.
- [24] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*. ACL, 311–318. <https://www.aclweb.org/anthology/P02-1040/>
- [25] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.). 8024–8035.
- [26] Dario Pavllo, David Grangier, and Michael Auli. 2018. QuaterNet: A Quaternion-based Recurrent Model for Human Motion. In *British Machine Vision Conference 2018, BMVC 2018, Northumbria University, Newcastle, UK, September 3-6, 2018*. BMVA Press, 299.
- [27] Xuanchi Ren, Haoran Li, Zijian Huang, and Qifeng Chen. 2019. Music-oriented Dance Video Synthesis with Pose Perceptual Loss. *arXiv preprint arXiv:1912.06606* (2019).
- [28] Takaaki Shiratori, Atsushi Nakazawa, and Katsushi Ikeuchi. 2006. Dancing-to-music character animation. In *Computer Graphics Forum*, Vol. 25. Wiley Online Library, 449–458.
- [29] Luis Molina Tanco and Adrian Hilton. 2000. Realistic synthesis of novel human movements from a database of motion capture examples. In *Proceedings Workshop on Human Motion*. IEEE, 137–142.
- [30] Taoran Tang, Jia Jia, and Hanyang Mao. 2018. Dance with melody: An lstm-autoencoder approach to music-oriented dance synthesis. In *Proceedings of the 26th ACM international conference on Multimedia*. 1598–1606.
- [31] Nelson Yalta, Shinji Watanabe, Kazuhiro Nakadai, and Tetsuya Ogata. 2019. Weakly-Supervised Deep Recurrent Neural Networks for Basic Dance Step Generation. In *2019 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [32] Yuan Yao and Ruochen Wen. [n.d.]. Follow the Music: Dance Motion Synthesis Corresponding to Arbitrary Music. ([n. d.]).
- [33] Wenlin Zhuang, Congyi Wang, Siyu Xia, Jinxiang Chai, and Yangang Wang. 2020. Music2Dance: Music-driven Dance Generation using WaveNet. *arXiv preprint arXiv:2002.03761* (2020).