

TM2D: Bimodality Driven 3D Dance Generation via Music-Text Integration

Kehong Gong^{1*}
 Zihang Jiang¹

Dongze Lian^{1*}
 Xinxin Zuo²

Heng Chang²
 Michael Bi Mi²

Chuan Guo²
 Xinchao Wang^{1†}

¹ National University of Singapore

² Huawei Technologies Co., Ltd.

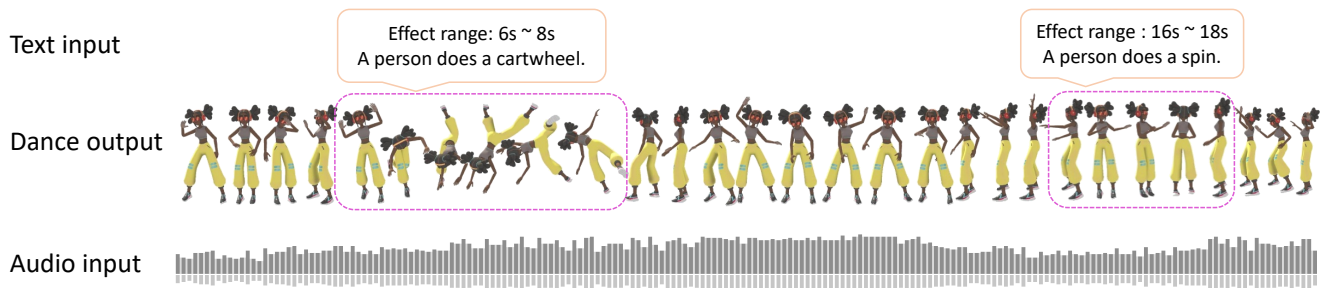


Figure 1: Generated 3D dance examples conditioned on music and text with our method. Given an audio input and text input with a specific starting time and duration, our method is able to generate a sequence of dance motions that fit the music and text instruction. The character is from Mixamo [1].

Abstract

We propose a novel task for generating 3D dance movements that simultaneously incorporate both text and music modalities. Unlike existing works that generate dance movements using a single modality such as music, our goal is to produce richer dance movements guided by the instructive information provided by the text. However, the lack of paired motion data with both music and text modalities limits the ability to generate dance movements that integrate both. To alleviate this challenge, we propose to utilize a 3D human motion VQ-VAE to project the motions of the two datasets into a latent space consisting of quantized vectors, which effectively mix the motion tokens from the two datasets with different distributions for training. Additionally, we propose a cross-modal transformer to integrate text instructions into motion generation architecture for generating 3D dance movements without degrading the performance of music-conditioned dance generation. To better evaluate the quality of the generated motion, we introduce two novel metrics, namely Motion Prediction Distance (MPD) and Freezing Score (FS), to measure the coherence and freezing percentage of the generated motion. Extensive experiments show that our approach can gener-

ate realistic and coherent dance movements conditioned on both text and music while maintaining comparable performance with the two single modalities. Code is available at <https://garfield-kh.github.io/TM2D/>.

1. Introduction

The music-conditioned dance generation has become a topic of great interest in recent years. The ability to generate dance movements that are synchronized with music has numerous applications, such as behavior understanding, simulation, and benefiting the community of dancers and musicians [25, 9, 29, 44]. Although music has been used as a guidance to generate dance movements, another important modality cue, text (or language), which provides richer actions and more flexible motion guidance, and is studied in other tasks such as image classification [38], detection [15], segmentation [50], and text-driven image generation [42], has not been fully explored in dance generation. To this end, we first propose a novel task for generating 3D dance movements that simultaneously incorporate both text and music modalities, enabling the generated human to perform rich dancing movements in accordance with the music and text.

Designing a system pipeline for this bimodality driven 3D dance generation task is non-trivial. There exist two

*Equal contribution: gongkehong@u.nus.edu, dongze@nus.edu.sg

†Corresponding author: xinchao@nus.edu.sg

significant challenges to be considered: i) the existing datasets only cater to either music-driven (music2dance) [46, 5, 57, 29] or text-driven (text2motion) [36, 16] human motion generation, and no paired 3D dance generation dataset exists that takes into account both music and text. While building a new large-scale paired 3D dance dataset based on music and text is possible, it is time-consuming with fully annotated 3D human motion [29]; ii) the integration of text into music-conditioned dance generation requires a suitable architecture. However, existing methods that use music as a driving force to generate dance movements might result in temporal-freezing frames or fail to generalize to in-the-wild scenarios [29, 44]. Therefore, simple integration of text into the existing music-conditioned architecture might pose a risk of degraded dance generation quality in our new task.

To address the first challenge, we take advantage of existing music-dance and text-motion datasets for this new task. However, directly mixing motions from these two datasets would result in inferior performance since the motions from these two datasets are in completely different motion spaces. To overcome this, we propose to utilize a VQ-VAE architecture to project the motions into a consistent and shared latent space. In particular, we build up a shared codebook for all the motions from the training set, and motions from both datasets are now represented as discrete tokens that are implicitly constrained to fall into a shared latent space. For the second challenge, we propose to utilize a cross-modal transformer architecture that formulates both music2dance and text2motion as sequence-to-sequence translation tasks. This architecture directly translates audio and text features into motion tokens and enables bimodality driven ability by introducing a fusion strategy in the latent space with a shared motion decoder for both tasks. With the shared decoder, audio and text information can be efficiently fused during inference. Our entire cross-modal transformer architecture is both effective and efficient, allowing for the integration of text instructions to generate coherent 3D dance motions, as illustrated in Figure 1.

To better evaluate the coherence of generated dance in our task, we propose a new evaluation metric, Motion Prediction Distance (MPD), which measures the distance between the predicted motion and the ground truth at the time of integrating text, thereby providing a more accurate evaluation of the coherence of frames. Additionally, we introduce a Freezing Score (FS) that quantifies the percentage of temporal freezing frames in dance generation, which is a common problem in music-conditioned dance generation. To better evaluate the performance of our method in real-world scenarios, we also collect some in-the-wild music data for evaluation. Our method successfully performs dance generation based on both text and music while maintaining comparable performance on the single modal-

ity tasks (music2dance, text2motion) compared to other state-of-the-art methods.

In summary, our contributions are as follows: i) We propose an interesting task of utilizing both music and text for 3D dance generation and propose a pipeline named TM2D (Text-Music to Dance) for this task. ii) Rather than collecting a new training set, we effectively combine the existing music2dance and text2motion datasets and employ a VQ-VAE framework to encode motions from all training sets to a shared feature space. iii) We propose a cross-modal transformer as well as a bimodal feature fusion strategy to encode both audio and text features, which is both effective and efficient. iv) We propose two new metrics, MPD and FS, which efficiently reflect the quality of generated motion. v) We successfully generate realistic and coherent dance based on both music and text instructions while maintaining comparable performance on the single modality tasks (music2dance, text2motion).

2. Related Work

2.1. Music to Dance

Music2Dance is typically divided into 2D and 3D dance generation and has been explored for many years. Recent methods model 3D dance generation from the perspective of network architecture. For instance, Lee *et al.* [26] explore the convolutional neural network in casual dilated setting. Lee *et al.* [25] propose a dance unit with Variational Auto-Encoder (VAE). Ren *et al.* [41] employ recurrent neural network (RNN) and the graph convolution network is introduced in [12]. As for the 3D dance generation, [54, 3, 24] implement convolutional neural network, [45, 14] apply adversarial learning in generated dance. [5, 46, 51, 58, 23] implement Long short-term memory (LSTM), and [9] implement motion graph with learned music/dance embedding through matching approach [10]. More recently, transformer has been applied in dance generation [28, 29, 44, 27]. However, the previous methods usually produce temporal freezing frames when generating the long sequences, or are difficult to generalize to in-the-wild music, thereby not satisfiable when directly used for bimodal driven 3D dance generation. In this paper, we design a cross modal transformer which is effective and efficient to integrate text instruction to generate the 3D dance.

2.2. Text to Motion

In addition to music-driven motion generation, text is also utilized as instructions to generate motions. Text2motion can be categorized into action label based motion generation and language description based motion generation. The action label based methods [18, 34] generate motion conditioned on action labels. Action2Motion [18] implement a GRU-VAE to iteratively generate the next

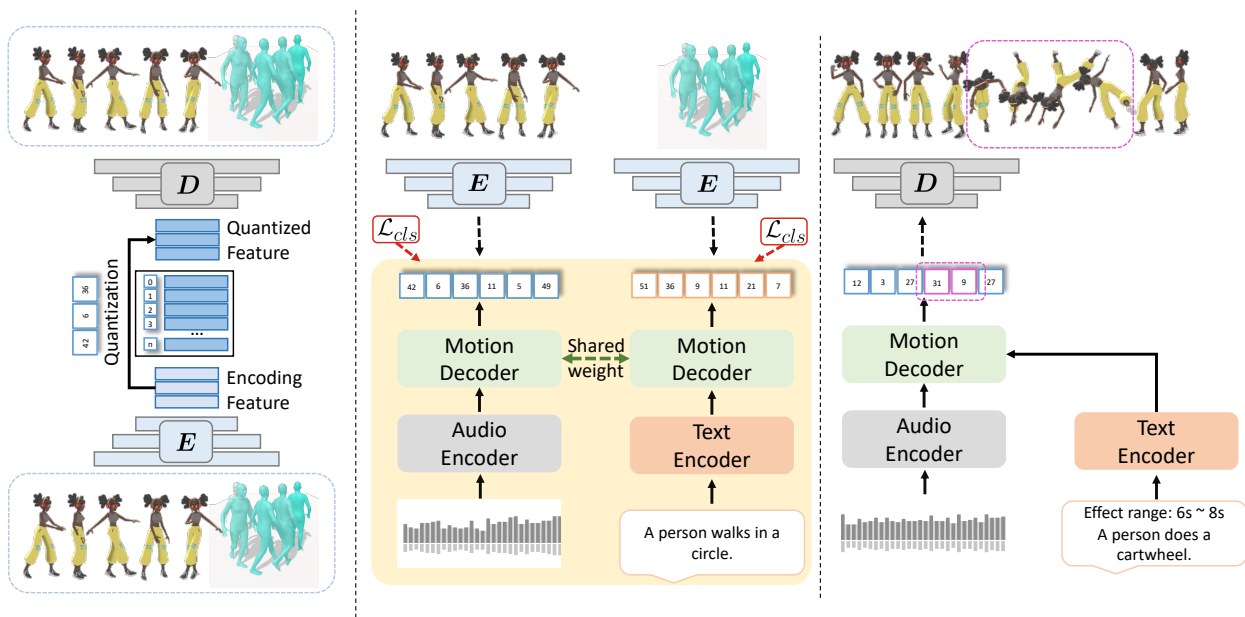


Figure 2: Our proposed pipeline for music-text conditioned 3D dance generation. Three stages from left to right: 3D human motion VQ-VAE, training stage of the cross-modal transformer, and inference stage of our pipeline. In the first stage, a VQ-VAE is trained with both motions from music2dance and text2motion data, which is then used to tokenize all motions. In the second stage, a dual path cross-modal transformer is employed for sequence2sequence translation tasks (*i.e.*, audio to motion tokens, text to motion tokens), with a shared motion decoder. In the third stage, given audio and text inputs, the audio and text encoders first extract the corresponding features, which are then fused (late fusion) in the motion decoder to generate dance condition on both music and text.

frame based on the action label and previous frames. ACTOR [34] implements a transformer-VAE to encode and decode the whole pose sequence in one-shot. Since the action label based methods are restricted in a small set of action labels, language description based methods are proposed for more flexible motion generation. [37, 31, 2, 16, 17] formulate the text-to-motion task as a machine translation problem, others [52, 4, 13, 35] learn a joint embedding space of text and motion. Among them, [37, 31, 2, 52, 4, 13] use RNN encoder-decoder to learn the mapping between text and motion. [17, 16] use transformer encoder with RNN decoder, and TM2T [17] further introduces motion to text as an inverse alignment for text to motion. TEMOS [35] and TEACH [6] use transformer for both encoder and decoder for one-shot generation. Different from these text2motion methods, we focus on how to integrate two modalities (*i.e.*, music and text) together for 3D dance generation.

3. Method

3.1. System Pipeline

The overall pipeline is shown in Figure 2, which consists of three stages. The first stage employs a 3D human motion

VQ-VAE to encode both the motions of the music2dance and text2motion datasets to a shared codebook encoded by multiple vectors such that each dance motion can be represented as a discrete motion token. After that, the combinations of music, text, and motion token are fed into a cross-modal transformer for training, which effectively learn how to predict a sequential motion tokens according to the previous those as well as music and text. In the third stage, a random starting motion token is generated and inputted to the cross-modal transformer with the given music and text for the sequential motion token prediction, which will be decoded by the pre-trained 3D human motion VQ-VAE in the first stage to generate 3D dance.

3.2. 3D Human Motion VQ-VAE

Since there are no paired 3D dance motions conditioned on both music and text, we try to tackle our task with the existing music2dance and text2motion datasets. Instead of directly mixing the motions from both datasets for training, we employ a VQ-VAE to project the motions into a consistent and shared latent space, which is represented with a codebook, as illustrated in Figure 2. To be specific, given a 3D human motion $M \in \mathbb{R}^{T \times d_m}$, where T is the time length

and d_m is the dimension of the human motion, an encoder of VQ-VAE that consists of several 1-D convolutional layers projects M to a latent vector $z \in \mathbb{R}^{T' \times d}$, where $T' = \frac{T}{t}$ and t is the time interval for downsampling and d is the dimension of the latent vector. A learnable codebook $e \in \mathbb{R}^{K \times d}$ describes all latent variable features of the whole dataset, where K and d are the length and dimension of the codebook, respectively. A quantized latent vector $z_q \in \mathbb{R}^{T' \times d}$ will record the closest vector from codebook e as follows

$$z_{q,i} = \arg \min_{e_j \in e} \|z_i - e_j\| \in \mathbb{R}^d, \quad (1)$$

and the motion token $t_m \in \mathbb{R}^{T' \times 1}$ stores the index of the closest vector

$$t_{m,i} = \arg \min_j \|z_i - e_j\| \in \mathbb{R}^1. \quad (2)$$

The quantized z_q will be decoded with stacked convolutional layers to reconstruct the human motion \hat{M} , and t_m will be used in the training stage of the cross-modal transformer.

For the training of 3D human motion VQ-VAE, we follow the strategy in [48] and the total loss contains a reconstruction loss for human motion regression, a codebook loss for the dictionary learning, and a commitment loss to stabilize the training process:

$$\mathcal{L}_{vq} = \|\hat{M} - M\|_1 + \|\text{sg}[e] - e_q\|_2^2 + \beta \|e - \text{sg}[e_q]\|_2^2, \quad (3)$$

where $\text{sg}[\cdot]$ is ‘stop gradient’ and β is the factor term to adjust the weight of the commitment loss. A straight-through estimator is also employed to pass the gradient from the decoder to the encoder in back-propagation.

As we will show in Sec. 4, the motion tokens from both datasets encoded by 3D human motion VQ-VAE fall almost in the shared latent space, which shows the feasibility of using the separate music-conditioned and text-conditioned motion datasets for music-text conditioned 3D dance generation task.

3.3. Cross-modal Transformer

The cross-modal transformer contains an audio encoder, a text encoder, and a motion decoder. Since we use two separate datasets for our task, the cross-modal transformer is divided into two branches, where one for music2dance and the other for text2motion. It takes the audio feature, the text feature, and the motion token t_m encoded by 3D human motion VQ-VAE as inputs, and performs the sequence-to-sequence translation task to generate the future motion tokens.

Attention. Attention is introduced in Transformer [49] for natural language processing, and then applied to many domains [30, 8, 20, 39, 11, 53, 32, 55, 40] A L -layer transformer typically consists of L transformer blocks, which

contains a multi-head self-attention (MSA) and a feed forward network (FFN). Given the input $x \in \mathbb{R}^{N \times c}$, where N is the sequence length and c is the dimension of the input, the transformer block first maps it to keys $K \in \mathbb{R}^{N \times c}$, queries $Q \in \mathbb{R}^{N \times c}$, values $V \in \mathbb{R}^{N \times c}$ with the linear projections, and then a self-attention operation is performed by

$$\text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^T}{\sqrt{c}}\right)V. \quad (4)$$

The output is followed by a normalization layer and a FFN. **Audio encoder.** Given a sequence of music, the audio encoder first extracts the raw audio features following Baidando [44], where a public audio processing toolbox, Librosa [22], is employed to obtain the *mel frequency cepstral coefficients (MFCC)*, *MFCC delta*, *constant-Q chromagram*, *tempogram* and *onset strength*. After that, an embedded layer followed by several transformer blocks containing self-attention operation *e.g.*, Eq. (4), is used to generate the processed audio features $f_a \in \mathbb{R}^{T' \times d}$.

Text encoder. Given a text instruction, we first extract the text token with a GloVe [33], and then an embedded layer followed by several transformer blocks is performed to obtain a processed text feature $f_t \in \mathbb{R}^{n \times d}$, where n is the length of the text feature.

Motion decoder. Motion decoder takes the motion tokens encoded by the 3D human motion VQ-VAE as inputs and outputs the future motion tokens, which has the same spirit as the sequence-to-sequence translation task. Since the information after time t is unknown at moment t , a masked MSA is first employed to interact with motions at different times. The mask has the shape of a simple upper triangular matrix and is performed in Eq. (4). For the moment t , only the motions before moment t are able to perform self-attention operations. Motion decoder also maps the extracted audio feature f_a and text feature f_t via audio encoder and text encoder to K and Q , to perform cross-modal attention with motion tokens. To enable the music-text conditional dance generation, (*i.e.*, feature fusion, we apply shared parameters), we use the motion decoder with shared parameters for model efficiency, as shown in Figure 2.

3.4. Details of Music-text Fusion

The detailed architecture is listed in Figure 3. Given the audio feature and text feature, the motion decoder first processes the past motion tokens with a self-attention layer followed by addition, normalization, a cross-attention with audio and text features separately, addition, normalization, and a feed-forward layer are performed. Such a procedure is repeated L ($L = 6$) times to build a typical L -layer transformer. Then we adopt a late fusion strategy to perform a weighted sum of features of audio and text at a specific time (*i.e.*, effect range). The weight curve is shown in Figure 3,

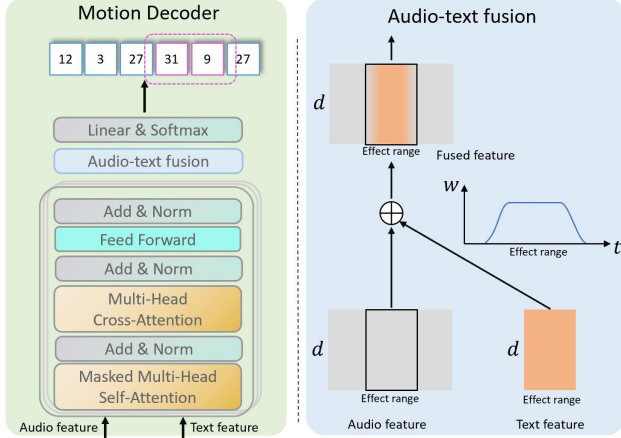


Figure 3: Details of music-text feature fusion: Given the audio and text features from the encoder, the decoder processes them separately in the early layers. Then the audio-text fusion layer is applied by the weighted sum at the effect range with weight cure. Finally, a linear projection layer and a softmax operation are applied to predict the music-text conditioned motion tokens.

where we slightly increase the weight of the text feature by a half cosine curve until a peak value of 0.8 at the beginning (20% time of effect range), and decrease it by a half cosine curve at the end (20% time of effect range). The weight of audio feature W_{audio} is $1 - W_{text}$ to ensure the feature keeps the same scale. With the fused feature, a linear projection layer and a softmax operation are applied to predict the music-text conditioned motion tokens, which are then decoded by the decoder of human motion VQ-VAE to obtain the music-text conditioned 3D dance sequences.

3.5. Training and Inference

Training. To train the 3D human motion VQ-VAE, we crop 64 frames with a sliding window from the original motion sequences as inputs, *i.e.*, $T = 64$, for both music2dance and text2motion datasets. We use a three-layer encoder and decoder so that the time interval of downsampling t is 8. We randomly sample motions from both datasets and employ the Adam optimizer with a batch size of 128 and a learning rate of $1e^{-4}$ to optimizer our 3D human motion VQ-VAE.

To train the cross-modal transformer, we split two branches into two streams to train the music2dance and text2motion tasks on both datasets. For the music2dance task, we use the same frequency as the motion token to sample the music vector so that the music feature has the same time length as the motion token. We perform the sequence to sequence translation with the motion decoder to obtain the prediction of motion tokens conditioned by the music input. As for the text2motion task, since the text and motion token have different lengths, we use a maximum text length

of 84 for text2motion translation with a padding strategy. We employ 6 self-attention layers for the audio encoder, text encoder, and motion decoder with the hidden dimension of 512 and 8 heads. The cross entropy loss is adopted for both music2dance and text2motion tasks

$$\mathcal{L}_{cls} = -\frac{1}{m} \sum_{i=0}^{m-1} \sum_{j=0}^{C-1} y_{ij} \log \hat{y}_{ij}, \quad (5)$$

where \hat{y} is the prediction of motion token and y is the ground-truth. m is the length of motion tokens and C is the number of classes of motion tokens, *i.e.*, the length of codebook of 3D human motion VQ-VAE. Both tasks are simultaneously optimized with a batch size of 64 and a learning rate of $1e^{-4}$.

Inference. Our aim is to generate 3D dance with music-text integration. Therefore, at the inference stage, we first fuse the music feature and text feature extracted by the audio encoder and text encoder with a weighted sum. Specifically, given an audio feature with a time length of T' and a text feature with the length of n , we first feed them into the motion decoder for future motion token prediction. Then we adopt the late fusion strategy to have a weighted sum of the generated motion tokens from both audio and text features at the duration of the integrated text, followed by a linear projection layer and a softmax operation, to obtain the combined motion tokens. Finally, the combined motion tokens are decoded into a 3D dance sequence after going through the decoder of our 3D human motion VQ-VAE model.

4. Experiments

4.1. Experimental Settings

Datasets. Since there are no paired 3D dance motions conditioned on both music and text for our task, we combine the existing music2dance and text2motion datasets. For the music2dance dataset, we employ the AIST++ dataset [29]. For the text2motion dataset, we employ the HumanML3D [16] dataset. Additionally, we also collect a new in-the-wild music dataset to evaluate the generalization ability of our method, which contains 82 wild music clips from Youtube. The total of 53 minutes duration (8x larger size than the AIST++ test set) and various styles and content of music, which faithfully lies out of the distribution of AIST++ (explained in supplementary material). The evaluation on our dataset better reflects the generalization ability of the models in real-world scenarios.

Baselines. Since we are the first to propose this bimodality driven 3D dance generation task, there are no existing methods for comparisons. We implement a traditional motion editing algorithm, slerp [43], which aligns the last generated frame of the dance with the first frame of the motion by text with a transition window of 10 frames, and

applies spherical interpolation of quaternion to fill in the transition in between. To validate the effectiveness of our method on music2dance task, we compare our method to the previous music2dance works [28, 57, 21, 29, 44]. To validate the effectiveness of our method on text2motion task, we compare our method to the previous text2motion works [31, 4, 7, 13, 47, 25, 17].

Evaluation metrics. We adopt the same evaluation settings as suggested by FACT [29] and Bailando [44] to evaluate the dance generation quality, including Fréchet Inception Distances (FID) [19], Diversity, and Beat Align Score. For the text2motion quality, we adopt the same evaluation settings as suggested by TM2T [17], including R-precision, Multimodal-Dist, FID, Diversity, and MultiModality.

In addition, we also propose two new evaluation metrics, Motion Prediction Distance (MPD) and Freezing Score. The former reflects the coherence of frames at the time of integrating text, and the latter reacts to the percentage of temporal freezing frames in the generated dance. For the Freezing Score, we introduce the Percentage of Freezing Frame (PFF) and AUC_f , where the PFF is defined by measuring the percentage of frozen frames with two criteria: 1) the maximum joint velocity blew a threshold (0.015m/s). 2) its duration exceeds a certain period (3s). The AUC_f is defined by area under curve of PFF in the threshold range from 0 to 0.03. As for MPD, it is defined as

$$MPD = \min_i \|f_i(M_{t_0 \rightarrow t_1}) - M_{t_1 \rightarrow t_2}\|_2, \quad (6)$$

where M is the dance motions, f is a motion prediction model, $f_i(M_{t_0 \rightarrow t_1})$ is the i -th predicted possible future motion, and t_0 , t_1 , and t_2 are the timestamps. It means that the model predicts various potential possible motion from t_1 to t_2 with the motion from t_0 to t_1 . If this distance is small (*i.e.*, the generated dance $M_{t_1 \rightarrow t_2}$ lays in the possible future), then the generated dance motion $M_{t_1 \rightarrow t_2}$ is more coherent at the time of integrating text. We adopt the motion prediction model DLow [56] in this evaluation metric as it is designed for diverse potential future motion prediction.

Perceptual Evaluation. Besides the above metric measuring, we also conduct extensive user studies on Amazon Mechanical Turk (AMT) to perceptually evaluate the visual effects of our generated 3D dance results. Particularly, given each pair of dance movements sampled from our method and others with the same music clip, we request 3 distinct users on AMT to present their preference regarding the music-dance alignment, motion realism, and mobility. We further set the bar of involved users that only the ones with Master recognition who also have finished more than 1,000 tasks with over 98% approve rate are considered. Overall, there are 55 users employed in our user studies that come from various regions, ages, races and gender. The results of the user study are more representative to show the effect of

the generated dance in practice.

4.2. Evaluation on Music-text Conditioned Dance Generation

We validate the results of music-text conditioned dance generation on the test set of AIST++ and our dataset, as illustrated in Table 2. Here we mainly evaluate the coherence of generated dance with MPD metric at the time where music meets text, *i.e.*, the time point where text starts to take effect. We measure the coherence of pure music2dance generation as a baseline named a2d to reflect the influence of with/without text instruction. We use the past 25 motion frames to predict the future 30 frames, and calculate the MPD from future frame ($ft = 10$ to $ft = 30$, respectively). Our method consistently outperforms slerp [43] and a2d baseline in both datasets, and gains a similar result compared to the ground-truth shown in Table 2, which verifies the naturalness of our generated motion. More importantly, user study experiments show that our method generates more realistic 3D dance compared to dance generation only conditioned music, and 70.0% Win rate even compared to the ground-truth shown in Table 2. The results of our user study indicate that the music-text conditioned dance generation received high ratings from participants, highlighting the importance of considering audience preferences in evaluating dance quality.

4.3. Evaluation on Music Conditioned Dance Generation

To validate the effectiveness of our architecture, we quantitatively compare our method with state-of-the-art those for music conditioned 3D dance generation. The results on the test set of AIST++ are shown in Table 1. We can find that our method outperforms the previous ones in terms of motion diversity (Div_k and Div_g), while the performance of the motion quality (FID_g) and beat align score is inferior under the condition that we do not use seed motion from ground-truth compared with FACT [29] and Bailando [44].

The existing evaluation metrics are not sufficient to reflect the quality of generated dance in practice, which motivates us to propose a PFF and AUC_f to evaluate the percentage of freezing frames. It is worth noting that our architecture outperforms FACT [29] and Bailando [44] in terms of PFF and AUC_f , which shows that our method rarely generates frozen frames. Meanwhile, Li *et al.* [28] and DanceNet [57] also gain nearly zero freezing in PFF and AUC_f . The reason is that the generated dances of Li *et al.* [28] are highly jittery making its velocity variation extremely high, which is also reported in [29, 44], and leads to the non-freezing issue. And for DanceNet [57], it generates dance in a repeat motion pattern, nearly zero freezing but with low diversity. Furthermore, the results of user study show that the generated 3D dance is more visually re-

Data	Method	Motion Quality		Motion Diversity		Freezing		Beat Align Score \uparrow	User Study
		FID $_k$ \downarrow	FID $_g^\dagger$ \downarrow	Div $_k$ \uparrow	Div $_g^\dagger$ \uparrow	PFF \downarrow	AUC $_f$ \downarrow		Our Method Wins
AIST++	Ground-truth	17.10	10.60	8.19	7.45	0.00	0.00	0.2374	41.9%
	Li <i>et al.</i> [28]	86.43	43.46	6.85*	3.32	0.00	0.00	0.1607	98.3%
	DanceNet [57]	69.18	25.49	2.86	2.85	0.00	<u>0.98</u>	0.1430	87.2%
	DanceRevolution [21]	73.42	25.92	3.52	4.87	<u>11.01</u>	12.22	0.1950	70.5%
	FACT [29]	35.35	22.11	5.94	6.18	25.29	21.59	<u>0.2209</u>	83.3%
	Bailando [44]	28.16	<u>9.62</u>	<u>7.83</u>	<u>6.34</u>	14.91	13.25	0.2332	65.0%
	ours (only dance data)	<u>23.94</u>	9.53	7.69	4.53	0.00	0.00	0.2127	-
	ours	19.01	20.09	9.45	6.36	0.00	0.00	0.2049	-
Wild Audio	FACT [29]	70.36	20.23	7.33	6.34	32.64	27.21	<u>0.2211</u>	89.6%
	Bailando [44]	50.56	22.55	3.80	<u>6.04</u>	17.55	14.14	0.2166	70.8%
	ours (only dance data)	<u>43.85</u>	13.08	8.52	4.76	<u>0.78</u>	<u>0.76</u>	0.1998	-
	ours	27.65	<u>20.34</u>	<u>7.88</u>	5.27	0.12	0.11	0.2290	-

Table 1: Music Conditioned Dance Generation: quantitative results on AIST++ and Wild Audio test set. The best and runner-up values are bold and underlined, respectively. Among compared methods, “Li *et al.*”, DanceNet and FACT are multiplexing the same results of AIST++ benchmark [28], while DanceRevolution [21] is followed from Bailando [44]. \dagger FID $_k$ and DIV $_k$ are fetched from [29] while FID $_g$ and DIV $_g$ are fetched from [44].

Data	Method	Motion Prediction Distance			User Study
		ft=10	ft=20	ft=30	Our Method Wins
AIST++	GT	0.048	0.072	0.088	70.0%
	a2d	<u>0.052</u>	<u>0.084</u>	<u>0.108</u>	66.6%
	slerp [43]	0.088	0.122	0.135	73.3%
	ours	0.049	0.080	0.102	-
Wild Audio	a2d	<u>0.052</u>	<u>0.088</u>	<u>0.113</u>	60.0%
	slerp [43]	0.104	0.148	0.171	70.0%
	ours	0.048	0.079	0.107	-

Table 2: Music-text conditioned Dance Generation: quantitative results on AIST++ and wild audio test set. The best and runner-up values are bold and underlined, respectively.

alistic compared to other methods. Even in comparison to the ground truth, 41.9% of our generated dance is voted as the better in average. We also report the result trained with dance-only data, which shows comparable performance.

In addition to the evaluation on the test set of AIST++, we also show the experimental results on our in-the-wild dataset in Table 1. Our method outperforms FACT [29] and Bailando [44] in almost all metrics except for FID $_g$ Div $_g$. We can observe that both FACT [29] and Bailando [44] shows a large performance drop in terms of FID $_k$, while ours maintain a small change. This is because FACT [29] and Bailando [44] requires seed motion, however, there is no ground-truth for in-the-wild scenario. With random sampled seed motion or token, their methods are not adapted

well for the in-the-wild scenarios. One can also notice that there exists freezing in our method on the in-the-wild dataset but the percentage of frozen frames is zero on the test set of AIST++, which results from the different distributions of AIST++ and our dataset. It also shows that our dataset is more challenging. We also report the result trained with dance-only data, which shows the advantage of our mix-training strategy in the wild scenario.

4.4. Evaluation on Text Conditioned Motion Generation.

We evaluate our text2motion approach in two different scenarios: inference with text only (t2m), and inference with text and music feature fusion. Table 3 shows the results. In the text-only setting, our approach achieves comparable performance with TM2T baseline(T) [17], which demonstrates that the mixed data/tasks training does not affect the quality of text2motion generation. We then apply our late fusion method by randomly sampling a music clip (from the whole AIST++ test set) and a text (from the whole t2m test set), and evaluate the generated dance clips following the same protocol as TM2T [17]. As Table 3 indicates, the text-dance consistency remains acceptable with a late fusion ratio of 0.8. For more details on the relation between the late fusion rate and text2motion result, please refer to the supplementary material.

4.5. Qualitative Results

We also show the qualitative results of our method for music-text conditioned 3D dance generation in Figure 4, Figure 5 and Figure 6. Specifically, Figure 4 shows the gen-

Methods	R Precision \uparrow	FID \downarrow	MM Dist \downarrow	Diversity \rightarrow	MModality \uparrow
Real motions	0.511	0.002	2.974	9.503	-
Seq2Seq [31]	0.180	11.75	5.529	6.223	-
Language2Pose [4]	0.246	11.02	5.296	7.676	-
Text2Gesture [7]	0.165	5.012	6.030	6.409	-
Hier [13]	0.301	6.532	5.012	8.332	-
MoCoGAN [47]	0.037	94.41	9.643	0.462	0.019
Dance2Music [25]	0.033	66.98	8.116	0.725	0.043
TM2T baseline(T) [17]	0.351	1.669	4.046	9.632	4.352
TM2T [17]	0.424	1.501	3.467	8.589	2.424
TM2D (t2m)	0.319	1.021	4.098	9.513	4.139
TM2D (LFR 0.8)	0.300	<u>1.105</u>	4.307	8.887	4.443

Table 3: The evaluation of text2motion on the HumanML3D dataset, where the results are averaged from 20 runs.

erated 3D dance with the same music sequence but with different text instructions. Figure 5 shows the generated 3D dance with the same text instruction but with a different start time. Figure 6 shows the generated 3D dance with the same text instruction but with different durations. As one can see, our approach maintains plausible visual results according to the text instructions for all three cases, which confirms that our approach is more flexible.

4.6. The Impact of Mixed Data

We employ two datasets (AIST++ and HumanML3D) to train our 3D human motion VQ-VAE so that the motion tokens drop in the same latent space. To empirically show this point, we conduct experiments to count the number of shared motion tokens. For a trained 3D human motion VQ-VAE with AIST++ and HumanML3D, there are 855 vectors and 912 vectors from codebook used to construct the motion token in AIST++ and HumanML3D. The total number of vectors contained in the codebook is 1024. among them, 846 vectors (98.9% in AIST++ and 92.8% in HumanML3D) are shared to generate the motion tokens, which shows the feasibility of using the separate music-conditioned and text-conditioned motion datasets for music-text conditioned 3D dance generation task. We further verify the latent space of the motions by VQ-VAE, and we perform t-SNE visualizations of the raw motion distribution (before VQ-VAE) and feature distribution encoded by VQ-VAE (after VQ-VAE) from a2d motion and t2m motion in Figure 7. Two motion datasets are mixed successfully, which provides the potential integration of bimodality dance generation. Refer to the supplementary materials for more details.

4.7. Model Efficiency

We also show the efficiency of our architecture in Table 4. In the inference stage, the complete model parameters and the inference time for 1s music are compared. Our architecture requires only about half model parameters and




Figure 4: **Same music, different actions.** These two generated dance share the same music and effect range (from 6s to 8s), with different text instructions: “A person does a cartwheel” (top), “A person is spinning with arms spread out” (bottom).




Figure 5: **Same action, different start time points.** These two generated dance share the same music and text instruction (“A person jumps up and down”) but different effect ranges: from 6 to 8s (top), from 7 to 9s (bottom).




Figure 6: **Same action, different durations.** These two generated dance share the same music and text instruction (“A person is keeping jumping”) and action start time point (*i.e.*, 6s) but different effect duration: 2s (top), 3s (bottom, *i.e.*, jumped twice).

inference time compared to Bailando [44], which attributes




Figure 7: The t-SNE visualizations of motions before and after VQ-VAE (orange: motions in AIST++, blue: motions in HumanML3D).

Method	Parameters	Inference time
FACT [29]	120M	8.300s/(1s music)
Bailando [44]	152M	0.236s/(1s music)
ours	72M	0.143s/(1s music)

Table 4: Comparisons of parameters and inference time.

that we do not need to encode separate features for the upper and lower half body. Compared to FACT [29], our architectures significantly reduce the inference time. The main reason is that FACT employs two transformers for the music, dance motion, and a cross-modal transformer, but we formulate it as the standard sequence-to-sequence translation task with one transformer only, which establishes the advantage of our architecture.

5. Conclusion

In this paper, a novel task that simultaneously integrates both music and text instruction for 3D dance generation is proposed. Due to the lack of the paired 3D dance motions conditioned on both music and text, we resort to two existing datasets, *i.e.*, music2dance and text2motion, to perform this task and employ a 3D human motion VQ-VAE to project the motions of the two datasets into a shared latent space so that the two datasets with different distributions can be effectively mixed for training. Moreover, we also propose a cross-modal transformer architecture to generate 3D dance without degrading the performance of music-conditioned dance generation. Two new evaluation metrics, MPD and FS, are proposed to reflect the quality of generated motion. Extensive experiments show that our method can generate dance motion that matches both music and text in a realistic and coherent way while maintaining comparable performance on two single modalities.

Acknowledgement. This project is supported by the Ministry of Education Singapore under its Academic Research Fund Tier 2 (Award Number: MOE-T2EP20122-0006),

and Singapore National Research Foundation (“CogniVision – Energy-autonomous always-on cognitive and attentive cameras for distributed real-time vision with milliwatt power consumption” grant NRF-CRP20-2017-0003) – www.green-ic.org/CogniVision.

References

- [1] mixamo. <https://www.mixamo.com/>. Accessed: 2023-03-07. 1
- [2] Hyemin Ahn, Timothy Ha, Yunho Choi, Hwiyeon Yoo, and Songhwai Oh. Text2action: Generative adversarial synthesis from language to action. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5915–5920. IEEE, 2018. 3
- [3] Hyemin Ahn, Jaehun Kim, Kihyun Kim, and Songhwai Oh. Generative autoregressive networks for 3d dancing move synthesis from music. *IEEE Robotics and Automation Letters*, 5(2):3501–3508, 2020. 2
- [4] Chaitanya Ahuja and Louis-Philippe Morency. Language2pose: Natural language grounded pose forecasting. In *2019 International Conference on 3D Vision (3DV)*, pages 719–728. IEEE, 2019. 3, 6, 8
- [5] Omid Alemi, Jules Françoise, and Philippe Pasquier. Groovenet: Real-time music-driven dance movement generation using artificial neural networks. In *Workshop on Machine Learning for Creativity, 23rd ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, volume 8, page 26, 2017. 2
- [6] Nikos Athanasiou, Mathis Petrovich, Michael J. Black, and Gül Varol. Teach: Temporal action compositions for 3d humans. In *International Conference on 3D Vision (3DV)*, 2022. 3
- [7] Uttaran Bhattacharya, Nicholas Rewkowski, Abhishek Banerjee, Pooja Guhan, Aniket Bera, and Dinesh Manocha. Text2gestures: A transformer-based network for generating emotive body gestures for virtual agents. In *IEEE Virtual Reality and 3D User Interfaces (VR)*, pages 1–10. IEEE, 2021. 6, 8
- [8] Hanting Chen, Yunhe Wang, Jianyuan Guo, and Dacheng Tao. VanillaNet: the Power of Minimalism in Deep Learning. *arXiv preprint arXiv:2305.12972*, 2023. 4
- [9] 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 Transactions on Graphics (TOG)*, 40(4):1–13, 2021. 1, 2
- [10] Rukun Fan, Songhua Xu, and Weidong Geng. Example-based automatic music-driven conventional dance motion synthesis. *IEEE transactions on visualization and computer graphics*, 18(3):501–515, 2011. 2
- [11] Gongfan Fang, Xinyin Ma, and Xinchao Wang. Structural pruning for diffusion models. *arXiv preprint arXiv:2305.10924*, 2023. 4
- [12] Joao P Ferreira, Thiago M Coutinho, Thiago L Gomes, José F Neto, and et al. Learning to dance: A graph convolutional adversarial network to generate realistic dance motions from audio. *Computers & Graphics*, 94:11–21, 2021. 2

- [13] Anindita Ghosh, Noshaba Cheema, Cennet Oguz, Christian Theobalt, and Philipp Slusallek. Synthesis of compositional animations from textual descriptions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1396–1406, 2021. 3, 6, 8
- [14] Shiry Ginosar, Amir Bar, Gefen Kohavi, Caroline Chan, Andrew Owens, and Jitendra Malik. Learning individual styles of conversational gesture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3497–3506, 2019. 2
- [15] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *International Conference on Learning Representations*, 2022. 1
- [16] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5152–5161, 2022. 2, 3, 5
- [17] Chuan Guo, Xinxin Zuo, Sen Wang, and Li Cheng. Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *European Conference on Computer Vision*, pages 580–597, 2022. 3, 6, 7, 8
- [18] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2021–2029, 2020. 2
- [19] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, volume 30, 2017. 6
- [20] Duc N.M Hoang, Shiwei Liu, Radu Marculescu, and Zhangyang Wang. Revisiting pruning at initialization through the lens of ramanujan graph. In *International Conference on Learning Representations*, 2023. 4
- [21] Ruozi Huang, Huang Hu, Wei Wu, Kei Sawada, Mi Zhang, and Daxin Jiang. Dance revolution: Long-term dance generation with music via curriculum learning. In *International Conference on Learning Representations*, 2021. 6, 7
- [22] Yanghua Jin, Jiakai Zhang, Minjun Li, Yingtao Tian, Huachun Zhu, and Zhihao Fang. Towards the automatic anime characters creation with generative adversarial networks. *arXiv preprint arXiv:1708.05509*, 2017. 4
- [23] Hsuan-Kai Kao and Li Su. Temporally guided music-to-body-movement generation. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 147–155, 2020. 2
- [24] Kosmas Krtsis, Aggelos Gkiokas, Aggelos Pikrakis, and Vassilis Katsourous. Danceconv: Dance motion generation with convolutional networks. *IEEE Access*, 10:44982–45000, 2022. 2
- [25] Hsin-Ying Lee, Xiaodong Yang, Ming-Yu Liu, Ting-Chun Wang, Yu-Ding Lu, Ming-Hsuan Yang, and Jan Kautz. Dancing to music. *Advances in neural information processing systems*, 32, 2019. 1, 2, 6, 8
- [26] Juheon Lee, Seohyun Kim, and Kyogu Lee. Listen to dance: Music-driven choreography generation using autoregressive encoder-decoder network. *arXiv preprint arXiv:1811.00818*, 2018. 2
- [27] Buyu Li, Yongchi Zhao, Shi Zhelun, and Lu Sheng. Danceformer: Music conditioned 3d dance generation with parametric motion transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 1272–1279, 2022. 2
- [28] Jiaman Li, Yihang Yin, Hang Chu, Yi Zhou, Tingwu Wang, Sanja Fidler, and Hao Li. Learning to generate diverse dance motions with transformer. *arXiv preprint arXiv:2008.08171*, 2020. 2, 6, 7
- [29] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13401–13412, 2021. 1, 2, 5, 6, 7, 9
- [30] Dongze Lian, Daquan Zhou, Jiashi Feng, and Xinchao Wang. Scaling & shifting your features: A new baseline for efficient model tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. 4
- [31] Angela S. Lin, Lemeng Wu, Rodolfo Corona, Kevin W. H. Tai, Qixing Huang, and Raymond J. Mooney. Generating animated videos of human activities from natural language descriptions. 2018. 3, 6, 8
- [32] Xinyin Ma, Gongfan Fang, and Xinchao Wang. LLM-Pruner: On the Structural Pruning of Large Language Models. *arXiv preprint arXiv:2305.11627*, 2023. 4
- [33] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014. 4
- [34] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10985–10995, 2021. 2, 3
- [35] Mathis Petrovich, Michael J Black, and Gül Varol. Temos: Generating diverse human motions from textual descriptions. In *European Conference on Computer Vision*, page 480–497, 2022. 3
- [36] Matthias Plappert, Christian Mandery, and Tamim Asfour. The kit motion-language dataset. *Big data*, 4(4):236–252, 2016. 2
- [37] Matthias Plappert, Christian Mandery, and Tamim Asfour. Learning a bidirectional mapping between human whole-body motion and natural language using deep recurrent neural networks. *Robotics and Autonomous Systems*, 109:13–26, 2018. 3
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 1

- [39] Sucheng Ren, Xingyi Yang, Songhua Liu, and Xinchao Wang. SG-Former: Self-guided Transformer with Evolving Token Reallocation. In *IEEE International Conference on Computer Vision*, 2023. 4
- [40] Sucheng Ren, Daquan Zhou, Shengfeng He, Jiashi Feng, and Xinchao Wang. Shunted Self-Attention via Multi-Scale Token Aggregation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 4
- [41] Xuanchi Ren, Haoran Li, Zijian Huang, and Qifeng Chen. Self-supervised dance video synthesis conditioned on music. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 46–54, 2020. 2
- [42] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 1
- [43] Ken Shoemake. Animating rotation with quaternion curves. In *Proceedings of the 12th annual conference on Computer graphics and interactive techniques*, pages 245–254, 1985. 5, 6, 7
- [44] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11050–11059, 2022. 1, 2, 4, 6, 7, 8, 9
- [45] Guofei Sun, Yongkang Wong, Zhiyong Cheng, Mohan S Kankanhalli, Weidong Geng, and Xiangdong Li. Deepdance: music-to-dance motion choreography with adversarial learning. *IEEE Transactions on Multimedia*, 23:497–509, 2020. 2
- [46] Taoran Tang, Jia Jia, and Hanyang Mao. Dance with melody: An lstm-autoencoder approach to music-oriented dance synthesis. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 1598–1606, 2018. 2
- [47] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018. 6, 8
- [48] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017. 4
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4
- [50] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. Groupvit: Semantic segmentation emerges from text supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18134–18144, 2022. 1
- [51] Nelson Yalta, Shinji Watanabe, Kazuhiro Nakadai, and Tetsuya Ogata. Weakly-supervised deep recurrent neural networks for basic dance step generation. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2019. 2
- [52] Tatsuro Yamada, Hiroyuki Matsunaga, and Tetsuya Ogata. Paired recurrent autoencoders for bidirectional translation between robot actions and linguistic descriptions. *IEEE Robotics and Automation Letters*, 3(4):3441–3448, 2018. 3
- [53] Xingyi Yang, Daquan Zhou, Songhua Liu, Jingwen Ye, and Xinchao Wang. Deep model reassembly. In *Advances in neural information processing systems*, 2022. 4
- [54] 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 *Proceedings of the 28th ACM International Conference on Multimedia*, pages 744–752, 2020. 2
- [55] Weihao Yu, Chenyang Si, Pan Zhou, Mi Luo, Yichen Zhou, Jiashi Feng, Shuicheng Yan, and Xinchao Wang. Metaformer Baselines for Vision. *arXiv preprint arXiv:2210.13452*, 2022. 4
- [56] Ye Yuan and Kris Kitani. Dlow: Diversifying latent flows for diverse human motion prediction. In *European Conference on Computer Vision*, pages 346–364. Springer, 2020. 6
- [57] Wenlin Zhuang, Congyi Wang, Jinxiang Chai, Yangang Wang, Ming Shao, and Siyu Xia. Music2dance: Dancenet for music-driven dance generation. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 18(2):1–21, 2022. 2, 6, 7
- [58] Wenlin Zhuang, Yangang Wang, Joseph Robinson, Congyi Wang, Ming Shao, Yun Fu, and Siyu Xia. Towards 3d dance motion synthesis and control. *arXiv preprint arXiv:2006.05743*, 2020. 2