


Quantized GAN for Complex Music Generation from Dance Videos

Ye Zhu¹^{*}, Kyle Olszewski², Yu Wu³, Panos Achlioptas², Menglei Chai², Yan Yan¹, and Sergey Tulyakov²

¹ Illinois Institute of Technology, USA

² Snap Inc., USA

³ Princeton University, USA

Abstract. We present Dance2Music-GAN (D2M-GAN), a novel adversarial multi-modal framework that generates complex musical samples conditioned on dance videos. Our proposed framework takes dance video frames and human body motions as input, and learns to generate music samples that plausibly accompany the corresponding input. Unlike most existing conditional music generation works that generate specific types of mono-instrumental sounds using symbolic audio representations (*e.g.*, MIDI), and that usually rely on pre-defined musical synthesizers, in this work we generate dance music in complex styles (*e.g.*, pop, breaking, etc.) by employing a Vector Quantized (VQ) audio representation, and leverage both its generality and high abstraction capacity of its symbolic and continuous counterparts. By performing an extensive set of experiments on multiple datasets, and following a comprehensive evaluation protocol, we assess the generative qualities of our proposal against alternatives. The attained quantitative results, which measure the music consistency, beats correspondence, and music diversity, demonstrate the effectiveness of our proposed method. Last but not least, we curate a challenging dance-music dataset of in-the-wild TikTok videos, which we use to further demonstrate the efficacy of our approach in *real-world* applications – and which we hope to serve as a starting point for relevant future research. Dataset and code at <https://github.com/L-YeZhu/D2M-GAN>.

Keywords: Multimodal Adversarial Learning, Complex Music Generation, Vector Quantized Representation.

1 Introduction

“When the music and dance create with accord, their magic captivates both the heart and the mind.” [†] As a natural form of expressive art, dance and music have enriched our daily lives with a harmonious interplay of melodies, rhythms, and movements, across the millennia. The growing popularity of social media platforms for sharing dance videos such as TikTok has also demonstrated their

^{*}This work was mainly done while the author was an intern at Snap Inc.

[†]Jean-Georges Noverre.




Fig. 1: **Task illustration.** We introduce a Vector-Quantized-based framework for music generation from dance videos, which takes human body motions and visual frames as input, and generates suitable music accordingly. Our proposed model is able to generate complex and rich dance music - in contrast to most existing conditional music generation works that typically output mono-instrumental sounds.

significance as a source of entertainment in our modern society. At the same time, new research works are flourishing following the trend and exploring multi-modal generative tasks between dance motions and music [37,39,38,1].

Although seemingly intuitive, music generation from dance videos has been a challenging task due to two main reasons. First, typical audio music signals are high-dimensional and require sophisticated temporal correlations for overall coherence [28,4]. For example, CD-quality audio has a typical sampling rate of 44.1 kHz, resulting in over 2.5 million data points (“dimensions”) for a one-minute musical piece [9]. In contrast, most dance generation works output the relatively low-dimensional motion data in the form of 2D or 3D skeleton keypoint (*e.g.*, displacement for dozens of joints) conditioned on the music [37,39,55,52], which are then rendered into dance sequences and videos. To tackle the challenge of the high dimensionality of audio data, the research studies on music generation from visual input [16,56,25] often rely on the low-dimensional intermediate symbolic audio representations (*e.g.*, 1D piano-roll or 2D MIDI). The symbolic representations benefit existing learning frameworks with a more explicit audio-visual correlation mapping and more stable training, as well as widely-established standard music synthesizers for decoding the intermediate representations. However, such symbolic-based works suffer from the limitations on the flexibility of the generated music, which brings us to the second challenge of dance video conditioned music generation. Specifically, a separately trained model is usually required for *each* instrument and the generated music is composed with acoustic sounds from a *single predefined* instrument [16,12,46] (*e.g.*, imagine a person dancing hip-hop with piano-based music). These facts make existing conditional music generation works difficult to generalize in complex musical styles and real-world scenarios.

To fill this gap, we propose a novel adversarial multi-modal framework that learns to generate complex musical samples from dance videos via the Vector Quantized audio representations. Inspired by the recent successes of VQ-

VAE [45,51,9] and VQ-GAN [14], we adopt quantized vectors as our intermediate audio representation, and leverage both their increased abstraction ability compared to continuous raw audio signals, as well as their flexibility of better representing complex real-world music compared to classic symbolic representations. Specifically, our framework takes the visual frames and dance motions as input (Fig. 1), which are encoded and fused to generate the corresponding audio VQ representations. After a lookup process of the generated VQ representations in a learned codebook, the retrieved codebook entries are decoded back to the raw audio domains using a fine-tuned decoder from JukeBox [9]. Additionally, we deploy a convolution-based backbone and follow a hierarchical structure with two separate abstraction levels (*i.e.*, different hop-lengths) for the audio signals to test the scalability of our framework. The higher-level model has a larger hop-length and fewer parameters, resulting in faster inference. In contrast, the lower-level model has a lower abstraction level with smaller hop-length, which enables the generation of music with higher fidelity and better quality.

Last but not least, we also contribute a real-world paired dance-music dataset collected from TikTok video compilations. Our dataset contains in total 445 dance videos with 85 songs and an average per-video duration of approximately 12.5 seconds. Unlike existing datasets (e.g., AIST [59,39]), ours is more challenging and better reflects the conditions of real-world scenarios, setting thus a new point for relevant future research.

Tapping on such datasets, we conduct extensive experiments to demonstrate the effectiveness and robustness of the proposed framework. Specifically, we design and follow a rich evaluation protocol to consider its generative quality with respect to the correspondence of dance input in terms of beats, genres and coherence, the general quality of the generated music is also assessed. The attained results (both quantitative and qualitative) show that our model can generate plausible dance music in terms of various musical features, outperforming several competing conditioned music generation methods.

In summary, our main contributions are:

- We propose *D2M-GAN*, a novel adversarial multi-modal framework that generates complex and free-form music from dance videos via *Vector Quantized (VQ) representations*.
- Specifically, the proposed model, using a VQ generator and a multi-scale discriminator, is able to effectively capture the temporal correlations and rhythm for the musical sequence to generate complex music.
- To assess our model we introduce a comprehensive *evaluation protocol* for conditionally generated music and demonstrate how the proposed *D2M-GAN* is able to generate more complex and plausible accompanying music compared to existing approaches.
- Last but not least, we create a novel real-world dataset with dance videos captured *in the wild* – and use it to establish a new more challenging setup for conditioned music generation, which further demonstrates the superiority of our framework.

2 Related Work

Audio, Vision and Motion. Combining data from audio, vision, and motion has been a popular research topic in recent years within the field of multi-modal learning [67,62,68,52,15]. Research focusing on general audio visual learning typically assumes that the two modalities are intrinsically correlated based on the synchronization nature of the audio and visual signals [34,47,48,68,2,3]. Such jointly learned audio-visual representations thus can be applied in multiple downstream tasks like sound source separation [17,18,19,20,66], audio-visual captioning [50,61], audio-visual action recognition [21,31], and audio-visual event localization and parsing [58,64,68,63].

On the other hand, another branch of studies related to our work has been investigating the correlations between motions and sounds [37,70,39,16,38]. A large portion of the research works here, aim to generate human motions based on the audio signals, either in the form of 2D pose skeletons [55,52,37] or direct 3D motions [39,57,29]. For the inverse direction that seeks to generate audio from motions, Zhao *et al.* [66] introduces an end-to-end model to generate sounds from motion trajectories. Gan *et al.* [16] propose a graph-based transformer framework to generate music from performance videos using raw movements as input. Di *et al.* [10] propose to generate video background music conditioned on the motion and special timing/rhythmic features of the input videos. In contrast to these previous works, our work combines three modalities, which takes the vision and motion data as input and generates music accordingly.

Music Generation. Raw music generation is a challenging task due to the high dimensionality of the audio data and sophisticated temporal correlations. Therefore, the existing music generation approaches usually adopt an intermediate audio representation for learning the generative models to reduce the computational demand and simplify the learning task [35,12,25,9,44,69]. Classic audio representations mainly include the symbolic and continuous categories. Musegan [12] introduces a multi-track GAN-based model for instrumental music generation via the 1D piano-roll symbolic representations. Music Transformer [25] aims to improve the long-term coherence of generated musical pieces using the 2D event-based MIDI-like audio representations [46]. Melgan [35] is a generative model for music in form of the audio mel-spectrogram features. Recently, JukeBox [9] introduces a generic music generation model based on the novel Vector Quantized (VQ) representations. Our proposed framework adopts the VQ representations for music generation.

Vector Quantized Generative Models. VQ-VAEs [45,51] are firstly proposed as a variant of the Variational Auto-Encoder (VAE) [32] with discrete codes and learned priors. Following works have demonstrated the potential of VQ-based framework in multiple generative tasks such as image and audio synthesis [14,9,26]. Specifically, the VQ-VAE [45] is initially tested for generating images, videos, and speech. An improved version of VQ-VAE [51] is proposed with a multi-scale hierarchical organization. Esser *et al.* [14] apply the VQ representations in the GAN-based framework for generating high-resolution images. Dhariwal *et al.* [9] introduce the JukeBox as a large-scale generative model for

music synthesis based on VQ-VAE. Compared to the symbolic and continuous representations, VQ representations leverage the benefits of flexibility (*i.e.*, the ability to represent complex music genres with a unified codebook in contrast to symbolic representations) and high compression levels (*i.e.*, the learned codebooks largely reduce the data dimensionality compared to raw waveform or spectrogram). Our proposed framework combines both the GAN [23] and VAE [32], which uses the GAN-based learning to generate VQ representations from the dance videos, and adopts the VAE-based decoder for synthesizing music.

3 Method

An overview of the architecture of the proposed D2M-GAN is shown in Fig. 2. Our approach entails a hierarchical structure with two levels of models that are independently trained with a similar pipeline for flexible scalability. For each level, the model consists of four modules: the motion module, the visual module, the VQ module consisting of a VQ generator and the multi-scale discriminators, and the music synthesizer. Our hierarchical structure amplifies the flexibility to choose between the trade-off of the music quality and computational costs according to practical application scenarios. A detailed description of these modules is given below while further architectural details and model-selection-tuning are included in the supplementary.

3.1 Data Representations

During the inference, the input to our proposed *D2M-GAN* come from two major modalities: the visual frames of the dance videos and human body motions of dance performers. The ground-truth music audio is also used as the supervision for the discriminators during the training stage. For the human body motions, several different forms of data representations such as 3D Skinned Multi-Person Linear model (SMPL) [41] or 2D body keypoints [6,5] can be utilized by our framework. We use SMPL and 2D body keypoints for different datasets in our experiments. To encode the visual frames, we extract I3D features [7] using a model pre-trained on Kinetics [30]. For the musical data, we adopt the VQ as the intermediate audio representation. To leverage the strong representation ability of codebooks trained on the large-scale musical dataset, we use the pre-learned codebooks from JukeBox [9], which are trained on a dataset of 1.2 million songs.

3.2 Generator

The generator $G = \{G_m, G_v, G_{vq}\}$ includes the motion module G_m , the visual module G_v , and the principal VQ generator G_{vq} in the VQ module, which takes the fused motion-visual data as input and outputs the desired VQ audio representations.

$$f_{vq} = G_{vq}(G_m(x_m), G_v(x_v)) = G(x_m, x_v), \quad (1)$$

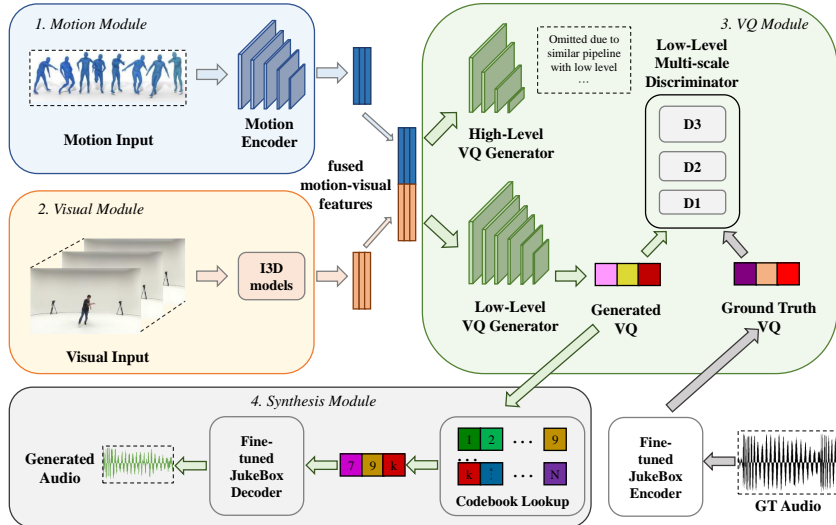


Fig. 2: **Overview of the proposed architecture of the $D2M\text{-GAN}$.** Our model takes the motion and visual data from the dance videos as input and process them with the motion and visual modules, respectively. It then forwards the concatenated representation containing information from both modalities to ground the generation of audio VQ-based representations with the VQ module. The resulting features are calibrated by a multi-scale GAN-based discriminator and are used to perform a *lookup* in the pre-learned codebook. Last, the retrieved codebook entries are decoded to raw musical samples via by a pre-trained and fine-tuned decoder, responsible for synthesizing music.

where x_m and x_v represent the motion and visual input data, respectively. f_{vq} is the output VQ representations. All these modules are implemented as convolution-based feed-forward networks. For the principal VQ generator, we use leaky rectified activation functions [65] for its hidden layers and a tanh activation for its last layer before output to promote the stability of GAN training [49].

It is also worth noting that we find that using batch normalization and the aforementioned activation function designs [42,49,54] is crucial for a stable GAN training in our framework. However, the application of the tanh activation will also restrict the output VQ representations within the data range between -1 and $+1$. We choose to scale activation after the last tanh activation by multiplying by a factor σ . The hyper-parameter σ enlarges the data range of VQ output and makes it possible to perform the lookup of pre-learned large-scale codebooks $\text{LookUp}(f'_{vq})$ with $f'_{vq} = \sigma f_{vq}$. Another significant observation regarding the generator’s design is using a wide receptive field. Music has long temporal dependencies and correlations compared to images, therefore, the principal VQ generator with a larger receptive field is beneficial for generating music sam-

ples with better quality, which is consistent with the findings from previous works [35,11]. To this end, we design our generator with relatively large kernel sizes in the convolutional layers, and we also add residual blocks with dilations after the convolutional layers. All previously described sub-modules within our generator G are jointly optimized.

3.3 Multi-Scale Discriminator

Similar to the generator, the discriminator in the D2M-GAN is also expected to capture the long-term dependencies of musical signals encoded in the generated sequence of VQ features. However, different from the generator design that focuses on increasing the receptive fields of the neural networks, we address this problem in the discriminator design by using a multi-scale architecture. The multi-scale discriminator design has been studied in previous works within the field of audio synthesis and generation [60,35,33]. The discriminator $D = \{D_1, D_2, D_3\}$ in the VQ module of our D2M-GAN is composed of 3 discriminators that operate on the sequence of generated VQ representations and its downsampled features by a factor of 2 and 4, respectively. Specifically, different from the multi-scale discriminators proposed in previous works that directly take the raw audio as input, we reshape the VQ representations f'_{vq} along the temporal dimension before feeding them into the discriminators, which is also important for *D2M-GAN* to reach a stable adversarial training since music is a temporal audio sequence. Finally, we use the window-based objectives [35] (Markovian window-based discriminator analog to image patches in [27]). Instead of learning to distinguish the distributions between two entire sequences, window-based objective learns to classify between distributions of small chunks of VQ sequences to further enhance the overall coherence as illustrated in Fig. 3.




Fig. 3: **Illustration of the reshape operation and the window-based discriminator for our *D2M-GAN*.**

3.4 Lookup and Synthesis

After generating the VQ representations, we perform a codebook lookup operation similar to other VQ-based generative models [45,51,14,9] to retrieve the corresponding entries with closest distance. Finally, we fine-tune the decoder from the JukeBox [9] without modifying the codebook entries as the music synthesizer for our learned VQ representations. Specifically, we also adopt the GAN-based technique for fine-tuning the music synthesizer, where the generator is replaced by the decoder of JukeBox and the discriminator follows the similar architecture as described in the previous subsection.

3.5 Training Objectives

GAN Loss. We use the hinge loss version of GAN objective [40,43] adopted for our music generation task to train the proposed *D2M-GAN*.

$$\begin{aligned} L_{adv.}(D; G) &= \sum_k L_{adv.}(D_k; G) \\ &= \sum_k (\mathbb{E}_{\phi(x_a)}[\min(0, 1 - D_k(\phi(x_a)))] \\ &\quad + \mathbb{E}_{(x_m, x_v)}[\min(0, 1 + D_k(G(x_m, x_v)))]), \end{aligned} \quad (2)$$

$$L_{adv.}(G; D) = \mathbb{E}_{x_m, x_v} [\sum_k -D_k(G(x_m, x_v))], \quad (3)$$

where x_a is the original music in a waveform, ϕ represents the fine-tuned encoder from JukeBox [9]. k indicates the number of multi-scale discriminators, which is empirically chosen to be 3 in our case.

Feature Matching Loss. To encourage the construction of subtle details in audio signals, we also include a feature matching loss [36] in the overall training objective. Similar to the audio generation works [35,33], the feature matching loss is defined as the L_1 distance between the discriminator feature maps of the real and generated VQ features.

$$L_{FM}(G; D) = \mathbb{E}_{(x_m, x_v)} \left[\sum_{i=1}^T \frac{1}{N_i} \|D^i(\phi(x_a)) - D^i(G(x_m, x_v))\|_1 \right]. \quad (4)$$

Codebook Commitment Loss. The codebook commitment loss [45,51] is defined as the L_1 distance between the generated VQ features and the corresponding codebook entries of the ground truth VQ features after the codebook lookup process.

$$L_{code}(G) = \mathbb{E}_{(x_m, x_v)} [\|Lookup(\phi(x_a)) - G(x_m, x_v)\|_1]. \quad (5)$$

Audio Perceptual Losses. To further improve the perceptual auditory quality, we consider the perception losses of the raw audio signals from both time and frequency domains. Specifically, the perceptual losses are calculated as the L_1 distance between the original audio and the generated audio samples:

$$L_{wav}(G) = \mathbb{E}_{(x_m, x_v)} [\|x_a - G(x_m, x_v)\|_1]. \quad (6)$$

$$L_{Mel}(G) = \mathbb{E}_{(x_m, x_v)} [\|\theta(x_a) - \theta(G(x_m, x_v))\|_1]. \quad (7)$$

where θ is the function to compute the mel-spectrogram features for the audio signals in waveform.

Final Loss. The final training objective for the entire generator module is defined as follows:

$$L_G = L_{adv.}(G; D) + \lambda_{fm} L_{FM}(G; D) + \lambda_c L_{code} + \lambda_w L_{wav} + \lambda_m L_{mel}, \quad (8)$$

where the λ_{fm} , λ_c , λ_a , and λ_{mel} are set to be 3, 15, 40 and 15, respectively during our experiments for both levels.



Fig. 4: **Examples of dance videos from our TikTok dance-music dataset.** Different from the AIST dataset [59] where dancing is performed by professional dancers in a studio environment, our dataset consists of real-world videos collected “in the wild”.




Fig. 5: **Qualitative example of rhythm evaluations and beat correspondence.** The lower-abstraction level model (D2M-Low) appears to align better than its high-counterpart (D2M-High) with the ground-truth (GT), which is consistent with the quantitative scores from the Table 1.

4 Experiments

4.1 Experimental Setup

Datasets. We validate the effectiveness of our method by conducting experiments on two datasets with paired dance video and music: the AIST++ [39] and our proposed TikTok dance-music dataset. The AIST++ dataset [39] is a subset of AIST dataset [59] with 3D motion annotations. We adopt the official cross-modality data splits for training, validation, and testing, where the videos are divided without overlapping musical pieces between the training and the validation/testing sets. The number of videos in each split is 980, 20, and 20, respectively. The videos from this dataset are filmed in professional studios with clean backgrounds. There are in total 10 different dance genres and corresponding music styles, which include breaking, pop, lock and etc. The number of total songs is 60, with 6 songs for each type of music. We use this dataset for the main experiments and evaluations. We also collect and annotate a **TikTok dance-music dataset** which contains 445 dance videos, with an average length of 12.5 seconds. This dataset utilizes 85 different songs, with the majority of videos having a single dance performer, and a maximum of five performers. The training-testing splits contain 392 and 53 videos, respectively, without overlapping songs. Fig. 4 shows example frames of the dance videos and makes apparent the key differences compared to the professional studio filmed dance video from AIST [59]. Our videos have wildly different backgrounds, and oftentimes contain incomplete human body skeleton data, which increases significantly the difficulty of the learning problem. For the TikTok music dataset, we use 2D human skeleton data as the underlying motion representation.

Implementation Details. For the presented experiments, we adopt a sampling rate of 22.5 kHz for all audio signals. We use the video and audio segments in the length of 2 seconds for training and standard testing in the main experiments. The generation of longer sequences is also investigated in Section 4.3. The hop lengths for the high and low level are 128 and 32, respectively. During the GAN training, we adopt the Adam optimizer with a learning rate of 1e-4 with $\beta_1 = 0.5$ and $\beta_2 = 0.9$ for the generators and discriminators. We define the scaling factor $\sigma = 100$ for the VQ generators. The number of discriminators k is 3 for the multi-scale structure. The batch size is set to be 16 for all experiments. During the fine-tuning of the JukeBox synthesizer, we use the Adam optimizer with a learning rate of 1e-5 with $\beta_1 = 0.5$ and $\beta_2 = 0.9$ for the synthesizer and multi-scale discriminators. We perform a denoising process [53] on the generated raw music data for better audio quality.

Comparisons. We compare our proposed method with several baselines. *Foley Music [16]*: Foley Music model generates MIDI musical representations based on keypoints motion data and then converts the MIDI back to raw waveform using a pre-defined MIDI synthesizer. Specifically, the MIDI audio representation is unique for each musical instrument, and therefore the Foley music model can only generate musical samples with mono-instrumental sound. *Dance2Music [1]*: Similar to [16], the generated music with this method is also monotonic in terms of the musical instrument. *Controllable Music Transformer (CMT) [10]*: CMT is a Transformer-based model proposed for video background music generation using MIDI representation. In addition to the above cross-modality models that are closely related to our work, we also consider *Ground Truth*: GT samples are the original music from dance videos. *JukeBox [9]*: music samples generated or reconstructed via the JukeBox model.

4.2 Music Evaluations

We design a comprehensive evaluation protocol that incorporates objective (*i.e.*, metrics that can be automatically calculated) and subjective (*i.e.*, scores given by human testers) metrics to evaluate the generated music from various perspectives. Specifically, the evaluations are divided into two categories: the first category, which is also the focus of our work, measures correlations between the generated music and the input dance videos, for which we compare our proposed model with other cross-modality music generation works [16,1,10] and a random baseline from JukeBox [9]. The second category focuses on the quality of the music in general, for which we use the reconstructed samples using JukeBox [9] given the original audio as input and GT samples for comparisons.

Rhythm. Musical rhythm accounts for an important characteristic of the generated music samples, especially given the dance video as input. To evaluate the correspondence between the dance beats and generated musical rhythm, we adopt two objective scores as evaluation metrics, which are the Beats Coverage Scores and the Beats Hit Scores similar to [37,8]. Previous works [37,8] have demonstrated the kinematic dance and musical beats (*i.e.*, rhythm) are generally aligned, we can therefore reasonably evaluate the musical rhythm by

Table 1: Evaluation protocol and the corresponding results for the experiments on the AIST++ dataset [39]. *Obj.* stands for *Objective*, which means the scores are automatically calculated. *Subj.* stands for *Subjective*, which means the scores are given by human evaluators.

Category	Features	Type	Metric	Methods	Scores
Dance-Music	Rhythm	Obj.	Beats Coverage & Beats Hit	Dance2Music [1]	83.5 & 82.4
				Foley Music [16]	74.1 & 69.4
				CMT [10]	85.5 & 83.5
				Ours High-level	88.2 & 84.7
				Ours Low-level	92.3 & 91.7
Dance-Music Genre&Diversity		Obj.	Genre Accuracy (Retrieval-based)	Dance2Music [1]	7.0
				Foley Music [16]	8.1
				CMT [10]	11.6
				Ours High-level	24.4
				Ours Low-level	26.7
Dance-Music	Coherence	Subj.	Mean Opinion Scores	Random JukeBox [9]	2.0
				Dance2Music [1]	2.8
				Foley Music [16]	2.8
				CMT [10]	3.0
				Ours High-level	3.5
				Ours Low-level	3.3
				GT	4.6
Music	Overall quality	Subj.	Mean Opinion Scores	JukeBox [9]	3.5
				Ours High-level	3.5
				Ours Low-level	3.7
				GT	4.8

comparing the beats from the generated music and those from the GT music samples as shown in Fig. 5. We detect the musical beats by the second-level onset strength [13], which can be considered as the start of an acoustic event. We define the number of detected beats from the generated music samples as B_g , the total beats from the original music as B_t , and the number of aligned beats from the generative samples as B_a . The Beats Coverage Scores B_g/B_t measure the ratio of overall generated beats to the total musical beats. The Beats Hit Scores B_a/B_t measure the ratio of aligned beats to the total musical beats. The quantitative results are presented in Table 1. We observe that both levels of our proposed *D2M-GAN* achieve better scores compared to competing methods.

Genre and Diversity. Dance and music are both diverse in terms of genres. The generated music samples are expected to be diverse and harmonious with the given dance style (*e.g.*, breaking dance with strong beats to pair with music in fast rhythm). Therefore, we calculate the genre accuracy for evaluating whether the generated music samples have a consistent genre with the dance style. The calculation of this objective metric requires the annotations of dance and music genres, we thus use the retrieved musical samples from the AIST++ [39] for this evaluation setting. Specifically, we retrieve the musical samples with the highest similarity scores from the segment-level database formed by original audio samples with the same sequence length. The similarities scores are defined

Table 2: Evaluations for the experiments on the TikTok dataset.

Models	Beats Coverage	Beats Hit
High w/o M	85.5	72.4
High w/o V	86.3	81.7
High (full)	88.4	82.3
Low w/o M	83.8	74.6
Low w/o V	85.2	81.7
Low (full)	87.1	83.9

Table 3: Results for ablation studies in terms of sequence length.

Length	Beats Coverage	Beats Hit	Genre Acc.
High - 2s	88.2	84.7	24.4
High - 3s	88.2	85.3	25.6
High - 4s	87.1	83.0	23.3
Low - 2s	92.3	91.7	26.7
Low - 3s	90.1	88.2	25.6
Low - 4s	88.2	84.7	23.3

as the euclidean distance between the audio features extracted via a VGG-like network [24] pre-trained on AudioSet [22]. In case that the retrieved musical sample has the same genre as the given dance style, we consider the segment to be genre accurate. The genre accuracy is then calculated by S_c/S_t , where S_c counts the number of genre accurate segments and S_t is the total number of segments from the testing split.

We observe in Table 1 that the genre accuracy scores of our *D2M-GAN* are considerably higher compared to the competing methods. This is due to the reason that the competing methods rely on MIDI events as audio representations, which require a specific synthesizer for each instrument, and thus can only generate music samples with mono-instrumental sound. In contrast, our generated VQ audio representations can represent complex dance music similar to the input music types, which helps to increase the diversity of the generated music samples. It also makes the generated samples to be more harmonious with the dance videos compared to acoustic instrumental sounds from [16,1], as shown in the next evaluation protocol for the coherence test.

Coherence. Since we generate music samples conditioned on the dance videos, the dance video input and the output are expected to be harmonious and coherent when combined together. Specifically, a given dance sequence could be accompanied by multiple appropriate songs. However, the evaluation of the dance-music coherence is very subjective, therefore we conduct the Mean Opinion Scores (MOS) human test for assessing the coherence feature. During the evaluation process, the human testers are asked to give a score between 1 and 5 to evaluate the coherence between the dance moves and the music given a video with audio sounds. The higher scores indicate the fact the tester feels the given dance and music are more coherent. We prepare the videos with original visual frames and fused generated music samples for testing. In addition to the previously cross-modality generation methods [1,16,10], we also include the GT samples and the randomly generated music from JukeBox [9] for comparison. Our *D2M-GAN* achieves better scores compared to other baselines, which validates the fact that our proposed framework is able to catch the correlations with the given dance video and generates rather complex music that well matches the input. Details about the human evaluations are included in the supplementary.

Overall Quality. Although our main research focus is to learn the dance-music correlations in this work, we also look at the general sound quality of the gen-

Table 4: Results for ablation studies in terms of input modalities on the AIST++ dataset. M means the motion data, and V means the visual data.

Models	Beats Coverage	Beats Hit	Genre Acc.
High w/o M	83.5	82.9	15.1
High w/o V	87.1	88.2	16.3
High (full)	88.2	84.7	24.4
Low w/o M	89.4	87.6	15.1
Low w/o V	90.6	90.0	17.4
Low (full)	92.3	91.7	26.7

Table 5: Results for ablation studies in terms of losses on the AIST++ dataset. The mel loss is especially helpful for beats scores since the beats are characteristic by high frequencies.

Losses	Beats Coverage	Beats Hit	Genre Acc.
High w/o L_{FM}	85.3	84.7	23.3
High w/o L_{wav}	85.9	84.7	23.3
High w/o L_{mel}	77.6	76.5	18.6
High (full)	88.2	84.7	24.4
Low w/o L_{FM}	91.7	90.1	24.4
Low w/o L_{wav}	89.4	88.8	23.3
Low w/o L_{mel}	78.8	77.1	17.4
Low (full)	92.3	91.7	26.7

erated samples. We conduct the subjective MOS tests similar to the coherence evaluation, where the human testers are asked to give a score between 1 to 5 for the general quality of the music samples. During this test, only audio signals are played to the testers. The JukeBox samples are obtained by directly feeding the GT samples as input. The MOS tests show that our $D2M-GAN$ is able to generate music sample with plausible sound quality comparable to the JukeBox. JukeBox has multiple variants with different hop lengths, we compare with samples obtained from the model with same audio hop length for fairness (*i.e.*, the hop lengths for our high and low levels are 128 and 32, respectively). It is worth noting that synthesizing high quality audio itself is a vary challenging and computational demanding research topic, for example, it takes $3\ hrs$ to sample a 20-seconds high-quality music sample with a hop length of 8 [9].

Results on the TikTok Dataset. Compared to the AIST++ [39], our TikTok dance-music dataset is a more challenging dataset with “in the wild” video settings that contains various occlusions and noisy backgrounds. Table 2 shows the quantitative evaluation results for the experiments on the TikTok dataset, which demonstrates the overall robustness of the proposed $D2M-GAN$.

4.3 Ablation Studies

Sequence Length. In the main experiments, we use the 2-second samples for experiments with reference to other similar cross-modality generation tasks [39]. However, our model can also be effectively trained and tested with a longer sequence length as shown in Table 3 via a larger network with more parameters.

Data Modality. We perform ablation studies in terms of the input data modalities, by removing either the dance motion or the visual frame from the input data. Table 4 lists the corresponding experimental results. We observe that both motion and visual data contribute to our conditioned music generation task. Specifically, the motion data impose a larger impact on the musical rhythm, which is consistent with our expectations since the musical rhythm is closely correlated with the dance motions.

Table 6: Results for ablation studies in terms of model architectures on the AIST++ dataset. *D.* means discriminators.

Models	Beats Coverage	Beats Hit	Genre Acc.
High 1-layer D.	75.3	72.9	9.3
High 2-layer D.	85.3	82.9	21.0
High w/o scaling	72.9	71.8	14.0
High w/o reshape	73.5	70.1	11.6
High w/o fine-tune	87.0	84.7	24.4
High (full)	88.2	84.7	24.4
Low 1-layer D.	73.5	71.8	8.1
Low 2-layer D.	87.0	85.9	22.1
Low w/o scaling	72.4	70.1	12.8
Low w/o reshape	73.5	71.8	12.8
Low w/o fine-tune	92.3	91.2	26.7
Low (full)	92.3	91.7	26.7

Loss function. We analyze the impact of different losses included in the overall training objective. The results from Table 5 show the contributions of each loss term. Specifically, we observe the audio perceptual loss from the frequency domain L_{mel} helps with the generation of musical rhythm, it is reasonable due to the fact that mel-spectrogram features help to capture the high frequencies from the audio signals, which is closely related to the dance beats.

Model Architecture. We also test various variants of our *D2M-GAN* in terms of the model architecture and proposed model design techniques as in Tabel 6. The experimental results show that the multi-scale layer for the discriminators, the scaling operation in the generator, as well as the reshape techniques for discriminators are crucial.

5 Conclusion and Limitations

To conclude, we propose the *D2M-GAN* framework for complex music generation from dance videos via the VQ audio representations. As an early work in the exploitation of VQ based music generation, there are still limitations in the current work from two major aspects: the audio quality and inference speed. As we employ a learning-based encoder-decoder model for raw music (JukeBox [9]), its performance is the major bottleneck for the quality of our generated music. Though JukeBox can synthesize relatively high-quality audio signals, there is a tradeoff between computational cost and quality. Achieving fast inference requires increasing the hop length for the generated waveform, which limits the audio quality and introduces noise. On the other hand, another direction to balance the above two goals would be investigating a proper approach to automatically compose multiple instruments into a single performance based on video input via MIDI musical representations.

Acknowledgements This work is partially supported by NSF ECCS-2123521 research grant and Snap unrestricted gift research grant. This article solely reflects the opinions and conclusions of its authors and not the funding agents.

References

1. Aggarwal, G., Parikh, D.: Dance2music: Automatic dance-driven music generation. arXiv preprint arXiv:2107.06252 (2021)
2. Arandjelovic, R., Zisserman, A.: Look, listen and learn. In: ICCV (2017)
3. Aytar, Y., Vondrick, C., Torralba, A.: Soundnet: Learning sound representations from unlabeled video. NeurIPS (2016)
4. Briot, J.P., Hadjeres, G., Pachet, F.D.: Deep learning techniques for music generation, vol. 1. Springer (2020)
5. Cao, Z., Hidalgo Martinez, G., Simon, T., Wei, S., Sheikh, Y.A.: Openpose: Realtime multi-person 2d pose estimation using part affinity fields. IEEE TPAMI (2019)
6. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: CVPR (2017)
7. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: CVPR (2017)
8. Davis, A., Agrawala, M.: Visual rhythm and beat. In: ACM Transactions on Graphics (TOG) (2018)
9. Dhariwal, P., Jun, H., Payne, C., Kim, J.W., Radford, A., Sutskever, I.: Jukebox: A generative model for music. arXiv preprint arXiv:2005.00341 (2020)
10. Di, S., Jiang, Z., Liu, S., Wang, Z., Zhu, L., He, Z., Liu, H., Yan, S.: Video background music generation with controllable music transformer. In: ACM MM (2021)
11. Donahue, C., McAuley, J., Puckette, M.: Adversarial audio synthesis. In: ICLR (2019)
12. Dong, H.W., Hsiao, W.Y., Yang, L.C., Yang, Y.H.: Musegan: Multi-track sequential generative adversarial networks for symbolic music generation and accompaniment. In: AAAI (2018)
13. Ellis, D.P.: Beat tracking by dynamic programming. *Journal of New Music Research* (2007)
14. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: CVPR (2021)
15. Ferreira, J.P., Coutinho, T.M., Gomes, T.L., Neto, J.F., Azevedo, R., Martins, R., Nascimento, E.R.: Learning to dance: A graph convolutional adversarial network to generate realistic dance motions from audio. *Computers & Graphics* (2021)
16. Gan, C., Huang, D., Chen, P., Tenenbaum, J.B., Torralba, A.: Foley music: Learning to generate music from videos. In: ECCV (2020)
17. Gan, C., Huang, D., Zhao, H., Tenenbaum, J.B., Torralba, A.: Music gesture for visual sound separation. In: CVPR (2020)
18. Gao, R., Feris, R., Grauman, K.: Learning to separate object sounds by watching unlabeled video. In: ECCV (2018)
19. Gao, R., Grauman, K.: 2.5 d visual sound. In: CVPR (2019)
20. Gao, R., Grauman, K.: Co-separating sounds of visual objects. In: ICCV (2019)
21. Gao, R., Oh, T.H., Grauman, K., Torresani, L.: Listen to look: Action recognition by previewing audio. In: CVPR (2020)
22. Gemmeke, J.F., Ellis, D.P., Freedman, D., Jansen, A., Lawrence, W., Moore, R.C., Plakal, M., Ritter, M.: Audio set: An ontology and human-labeled dataset for audio events. In: ICASSP. IEEE (2017)
23. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NeurIPS (2014)

24. Hershey, S., Chaudhuri, S., Ellis, D.P., Gemmeke, J.F., Jansen, A., Moore, R.C., Plakal, M., Platt, D., Saurous, R.A., Seybold, B., et al.: Cnn architectures for large-scale audio classification. In: ICASSP. IEEE (2017)
25. Huang, C.Z.A., Vaswani, A., Uszkoreit, J., Shazeer, N., Simon, I., Hawthorne, C., Dai, A.M., Hoffman, M.D., Dinculescu, M., Eck, D.: Music transformer: Generating music with long-term structure. In: ICLR (2019)
26. Iashin, V., Rahtu, E.: Taming visually guided sound generation. In: British Machine Vision Conference (BMVC) (2021)
27. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017)
28. Ji, S., Luo, J., Yang, X.: A comprehensive survey on deep music generation: Multi-level representations, algorithms, evaluations, and future directions. arXiv preprint arXiv:2011.06801 (2020)
29. Kao, H.K., Su, L.: Temporally guided music-to-body-movement generation. In: ACMMM (2020)
30. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017)
31. Kazakos, E., Nagrani, A., Zisserman, A., Damen, D.: Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In: ICCV (2019)
32. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: ICLR (2014)
33. Kong, J., Kim, J., Bae, J.: Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. In: NeurIPS (2020)
34. Korbar, B., Tran, D., Torresani, L.: Cooperative learning of audio and video models from self-supervised synchronization. In: NeurIPS (2018)
35. Kumar, K., Kumar, R., de Boissiere, T., Gestin, L., Teoh, W.Z., Sotelo, J., de Brébisson, A., Bengio, Y., Courville, A.C.: Melgan: Generative adversarial networks for conditional waveform synthesis. In: NeurIPS (2019)
36. Larsen, A.B.L., Sønderby, S.K., Larochelle, H., Winther, O.: Autoencoding beyond pixels using a learned similarity metric. In: International conference on machine learning. PMLR (2016)
37. Lee, H.Y., Yang, X., Liu, M.Y., Wang, T.C., Lu, Y.D., Yang, M.H., Kautz, J.: Dancing to music. In: NeurIPS (2019)
38. Li, B., Zhao, Y., Sheng, L.: Dancenet3d: Music based dance generation with parametric motion transformer. arXiv preprint arXiv:2103.10206 (2021)
39. Li, R., Yang, S., Ross, D.A., Kanazawa, A.: Ai choreographer: Music conditioned 3d dance generation with aist++. In: ICCV (2021)
40. Lim, J.H., Ye, J.C.: Geometric gan. arXiv preprint arXiv:1705.02894 (2017)
41. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM Trans. Graphics (Proc. SIGGRAPH Asia) **34** (2015)
42. Lucic, M., Kurach, K., Michalski, M., Gelly, S., Bousquet, O.: Are gans created equal? a large-scale study. In: NeurIPS (2018)
43. Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. arXiv preprint:1802.05957 (2018)
44. Oord, A.v.d., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., Kavukcuoglu, K.: Wavenet: A generative model for raw audio. In: ICLR (2016)
45. Oord, A.v.d., Vinyals, O., Kavukcuoglu, K.: Neural discrete representation learning. In: NeurIPS (2017)

46. Oore, S., Simon, I., Dieleman, S., Eck, D., Simonyan, K.: This time with feeling: Learning expressive musical performance. *Neural Computing and Applications* pp. 955–967 (2020)
47. Owens, A., Efros, A.A.: Audio-visual scene analysis with self-supervised multisensory features. In: *ECCV* (2018)
48. Owens, A., Wu, J., McDermott, J.H., Freeman, W.T., Torralba, A.: Ambient sound provides supervision for visual learning. In: *ECCV* (2016)
49. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434* (2015)
50. Rahman, T., Xu, B., Sigal, L.: Watch, listen and tell: Multi-modal weakly supervised dense event captioning. In: *ICCV* (2019)
51. Razavi, A., van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with vq-vae-2. In: *NeurIPS* (2019)
52. Ren, X., Li, H., Huang, Z., Chen, Q.: Self-supervised dance video synthesis conditioned on music. In: *ACM MM* (2020)
53. Sainburg, T., Thielk, M., Gentner, T.Q.: Finding, visualizing, and quantifying latent structure across diverse animal vocal repertoires. *PLoS computational biology* **16**(10), e1008228 (2020)
54. Salimans, T., Kingma, D.P.: Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In: *NeurIPS* (2016)
55. Shlizerman, E., Dery, L., Schoen, H., Kemelmacher-Shlizerman, I.: Audio to body dynamics. In: *CVPR* (2018)
56. Su, K., Liu, X., Shlizerman, E.: Audeo: Audio generation for a silent performance video. In: *NeurIPS* (2020)
57. Tang, T., Jia, J., Mao, H.: Dance with melody: An lstm-autoencoder approach to music-oriented dance synthesis. In: *ACMMM* (2018)
58. Tian, Y., Shi, J., Li, B., Duan, Z., Xu, C.: Audio-visual event localization in unconstrained videos. In: *ECCV* (2018)
59. Tsuchida, S., Fukayama, S., Hamasaki, M., Goto, M.: Aist dance video database: Multi-genre, multi-dancer, and multi-camera database for dance information processing. In: *Proceedings of the 20th International Society for Music Information Retrieval Conference, (ISMIR)* (2019)
60. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: *CVPR* (2018)
61. Wang, X., Wang, Y.F., Wang, W.Y.: Watch, listen, and describe: Globally and locally aligned cross-modal attentions for video captioning. In: *NAACL* (2018)
62. Wu, Y., Jiang, L., Yang, Y.: Switchable novel object captioner. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 1–1 (2022). <https://doi.org/10.1109/TPAMI.2022.3144984>
63. Wu, Y., Yang, Y.: Exploring heterogeneous clues for weakly-supervised audio-visual video parsing. In: *CVPR* (2021)
64. Wu, Y., Zhu, L., Yan, Y., Yang, Y.: Dual attention matching for audio-visual event localization. In: *ICCV* (2019)
65. Xu, B., Wang, N., Chen, T., Li, M.: Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853* (2015)
66. Zhao, H., Gan, C., Ma, W.C., Torralba, A.: The sound of motions. In: *ICCV* (2019)
67. Zhu, X., Zhu, Y., Wang, H., Wen, H., Yan, Y., Liu, P.: Skeleton sequence and rgb frame based multi-modality feature fusion network for action recognition.

- ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM) **18**(3), 1–24 (2022)
68. Zhu, Y., Wu, Y., Latapie, H., Yang, Y., Yan, Y.: Learning audio-visual correlations from variational cross-modal generation. In: ICCASP (2021)
 69. Zhu, Y., Wu, Y., Olszewski, K., Ren, J., Tulyakov, S., Yan, Y.: Discrete contrastive diffusion for cross-modal and conditional generation. arXiv preprint arXiv:2206.07771 (2022)
 70. Zhuang, W., Wang, C., Xia, S., Chai, J., Wang, Y.: Music2dance: Dancenet for music-driven dance generation. arXiv preprint arXiv:2002.03761 (2020)