GTN-BAILANDO: GENRE CONSISTENT LONG-TERM 3D DANCE GENERATION BASED ON PRE-TRAINED GENRE TOKEN NETWORK

Haolin Zhuang[1,], Shun Lei[1], Long Xiao[1], Weiqin Li[1], Liyang Chen[1],*
Sicheng Yang[1], Zhiyong Wu[1,3,†], Shiyin Kang[2,†], Helen Meng[3]

[1] Shenzhen International Graduate School, Tsinghua University, Shenzhen, China
[2] XVerse Inc., Shenzhen, China
[3] The Chinese University of Hong Kong, Hong Kong SAR, China
{zhuanghl21, leis21}@mails.tsinghua.edu.cn, zywu@sz.tsinghua.edu.cn, kangshiyin@xverse.cn,

ABSTRACT

Music-driven 3D dance generation has become an intensive research topic in recent years with great potential for real-world applications. Most existing methods lack the consideration of genre, which results in genre inconsistency in the generated dance movements. In addition, the correlation between the dance genre and the music has not been investigated. To address these issues, we propose a genre-consistent dance generation framework, GTN-Bailando. First, we propose the Genre Token Network (GTN), which infers the genre from music to enhance the genre consistency of long-term dance generation. Second, to improve the generalization capability of the model, the strategy of pre-training and fine-tuning is adopted. Experimental results on the AIST++ dataset show that the proposed dance generation framework outperforms state-of-the-art methods in terms of motion quality and genre consistency[1].

Index Terms— 3D dance generation, genre token, multi-modal, music-driven

1. INTRODUCTION

Dance genres [1] are the manifestations that arise from dance, each with its specific characteristics. Diversified dance genres increase the variety of dance, making it one of the most well-known forms of artistic expression globally. With the advancement of artificial intelligence (AI), it is possible to generate dance pose sequences for promising applications such as choreography assistance and virtual idol performance. However, producing a satisfactory dance pose sequence using AI choreographies is still challenging due to the difficulty of maintaining genre consistency. In a certain dance genre, not all physically possible dance poses are appropriate, and the choreographed dance poses have stricter positional restrictions as well as a correlation to the given music.

Most existing AI choreographies focus on the alignment between dances and the musical melody. Early methods [2–11] send music and dance directly into a single network to generate dance sequences autoregressively. However, such methods suffer from the accumulation of errors during autoregressive generation and will likely regress non-dancing movements. Meanwhile, some works [12–17] use the retrieval method to divide dances into fixed-length units and choreograph by splicing these units according to

the melody of the music. Although these methods ensure the quality of the generated dance, they are incompatible with different time signatures and beats per minute (BPM). Recently, a great dance generation framework, Bailando [18], uses the vector quantised-variational autoencoder (VQ-VAE) [19] with the Actor-Critic [20] generative pre-trained transformer (GPT) [21] to solve problems in both autoregressive and retrieval dance generation methods.

However, all the studies mentioned above ignore the genre during the dance generation. Without genre, the generated dance will perform multiple genres of dance movements in a single music clip, resulting in performances that are inappropriate with the music (e.g., ballet movements in hip-hop music). Thereby, several latest dance generation frameworks devote efforts to the genre. GCDG [22] uses a one-hot vector to represent the genre and embeds the genre into the transformer decoder to generate genre consistency dance. MNET [23] uses a mapping network to transform a latent code into the style code for multi-genres while generating dance motions. These approaches can generate dance with a particular genre but require auxiliary inputs in inference, such as manually-determined genre labels. In addition, there is a specific correlation between the dance genre and its background music, allowing choreographers to determine the genre of dance based on the music. Therefore, it is preferable for AI choreographies to infer the genre based on the music and generate dances using the genre that has been inferred.

In this paper, we propose a genre consistent long-term dance generation framework based on Genre Token Network (GTN). First, we introduce GTN, to infer the genre by learning the correlation between music and genres. For a given clip of music, the genre can be inferred through GTN and used as a condition in dance generation, ensuring each generated dance pose satisfies the positional restrictions of one genre. Second, we propose a strategy for pre-training GTN to improve its generalization capacity. Due to the insufficiency of music data in existing dance-music-aligned datasets, it is difficult for the GTN to infer genre from music accurately. Thereby, in order to strengthen the correlation between each genre and its corresponding music, we collect a large-scale of dance background music with genre labels to pre-train the GTN. After that, we use a dance-music-aligned dataset to train the dance generation framework and load the pre-trained GTN weights for fine-tuning. Thus, the GTN can infer the genre of the music more effectively, which further enhances genre consistency. Experimental results show that our proposed dance generation framework and pre-training strategy significantly outperforms in both evaluation metrics and visualization judgments.

---

[1]Generated demo: https://im1eon.github.io/ICASSP23-GTNB-DG/
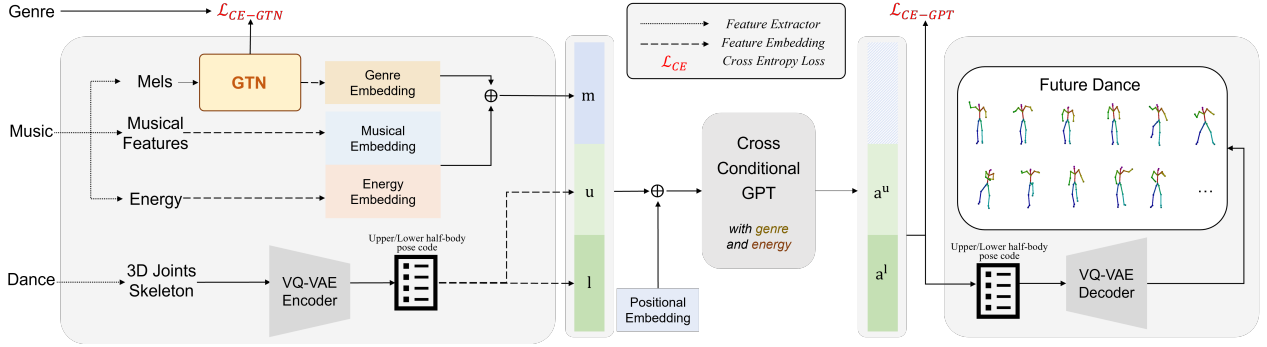
Fig. 1. GTN-Bailando. Mel-spectrogram of music is fed to Genre Token Network (GTN). The generated genre embedding is utilized as a condition in GPT to guarantee that each generated dance is consistent with its genre. The parameters are learned via cross-entropy loss $\mathcal{L}_{CE-GTN}$ with genre label and $\mathcal{L}_{CE-GPT}$ with ground-truth dance pose code.
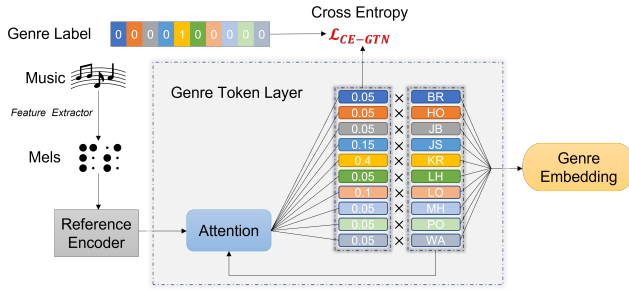
Fig. 2. Architecture of the proposed Genre Token Network (GTN).

2. PROPOSED METHOD

Problem Definition. We aim to infer genre from music and generate dances based on the inferred genre. Formally, given a clip of music $X = \{x_t\}_{t=1}^N$, where N is the music length, a initial dance pose $y_{t=1}$, our goals are, first, infer the genre $\hat{g} \in G$ from X, where G is the set of pre-defined genre categories. Second, use $\hat{g}$ as a condition and X to generate dance $\hat{Y} = \{\hat{y}_t\}_{t=2}^N$.

We intend to generate dances with genre consistency. In this section, we first describe GTN and show how to infer genre from the music. Second, we introduce our genre-consistent dance generation framework. After that, the strategy of pre-training and fine-tuning for GTN is explained. The architecture of our proposed GTN-Bailando is shown in Fig. 1.

2.1. Genre Token Network

Inspired by the work on emotion style tokens based speech synthesis [24,25], we propose supervised training of genre token to achieve the correlation between genre and music. As shown in Fig. 2, the architecture of the proposed GTN is composed of three modules: reference encoder, genre token layer, and genre embedding.

The reference encoder is adopted from [24], which is applied to compress the audio signal into a vector of set length. In this work, the mel-spectrogram of the music clip is fed to the reference encoder, and compressed into the learnable reference embedding.

The genre token layer includes a set of genre token embeddings and an attention module, which uses the reference embedding as the query vector. The attention module learns a measure of similarity

between the reference embedding and each token in a set of randomly initialized embeddings. This set of embeddings, also referred as genre tokens, are shared across all music clips.

The output of the genre token layer is the probability that the input music belongs to each genre category. In order to improve the robustness of the GTN, a soft embedding method [26] is used to represent the genre. In whis way, the tokens are then weighted summed by the probability to form a genre embedding.

In order to enhance the correlation between music and genre, the number of tokens is set to be consistent with the number of genre categories. Meanwhile, the genre label is converted to one-hot embedding and introduced into the genre token layer to serve as targets for token weights. Thereby, GTN is optimized via supervised training with cross-entropy loss between the genre labels and the genre token weights as follow,

$$\mathcal{L}_{CE-GTN} = \frac{1}{T} \sum_{t=0}^{T-1} \text{CE}\left(g_t, \hat{g}_t\right), \tag{1}$$

where g_t and $\hat{g}_t$ denote the genre label vector and the genre token weight vector of the music for t-th time's clip. $\text{CE}\left(\cdot\right)$ refers to the cross entropy loss function, T represents the number of music clips.

2.2. Dance Generation Framework

Inspired by Bailando [18], we generate dances based on VQ-VAE with cross-conditional GPT. Since there is a correlation between the velocity of the dance movement and the musical energy, we further consider the energy feature to improve the motion quality of the generated dance.

Precisely, given a piece of music as the input, we first extract the energy and musical features, then embed them to learnable vector $\boldsymbol{Z}_e$ and $\boldsymbol{Z}_m$, respectively. In the meantime, we extract the mel-spectrogram of the input music clip, then send it to the GTN to produce the genre embedding $\boldsymbol{Z}_g$. Then, we concatenate $\boldsymbol{Z}_m$ and $\boldsymbol{Z}_e$ on the temporal dimension, and add $\boldsymbol{Z}_g$, forming $\mathbf{m}$. As for the dance, we first feed the skeleton joints position of the dance clip into the VQ-VAE encoder to generate the upper and lower half-body pose codes, then embed them to learnable vector $\mathbf{u}$ and $\mathbf{l}$, respectively.

After getting $\mathbf{m}, \mathbf{u}, \mathbf{l}$, we concatenate them on the temporal dimension and add a learned positional embedding, then feed it to the cross-conditional GPT. At last, we get the outputs of the GPT, which are the probabilities of the upper and lower half-body pose codes. We get the predicted upper and lower body pose code pairs from

the probabilities and feed them into the VQ-VAE decoder to get the future dance.

Note that, we use the teacher-forcing scheme for GTN during training to improve the overall genre consistency of the dance generation framework. The cross-conditional GPT is optimized via supervised training with cross-entropy loss between the predicted action probability $\mathbf{a}$ and ground-truth pose codes p as follow:

$$\mathcal{L}_{CE-GPT} = \frac{1}{T'} \sum_{t=0}^{T'-1} \sum_{h=u,l} \text{CE}\left(\mathbf{a}_t^h, p_{t+1}^h\right). \qquad (2)$$

Finally, the loss of dance generation framework $\mathcal{L}_{DG}$ can be computed as:

$$\mathcal{L}_{DG} = \alpha \cdot \mathcal{L}_{CE-GTN} + \beta \cdot \mathcal{L}_{CE-GPT}. \qquad (3)$$

2.3. Pre-training and Fine-tuning

Although GTN is capable of inferring the genre, it still suffers from poor generalization capability, making it difficult to establish the correlation between the genre and the music nor to ensure genre consistency in generated dances. It is due to the fact that the music training data in the dance-music-aligned dataset is not large enough (less than 1 hour in length for all the ten genres).

Therefore, we propose a pre-training method to enhance the generalization capability for GTN and improve genre consistency for generated dances. First, we collect a large amount of dance background music data from the Internet, tag corresponding dance genre labels, and then use this data to train the GTN. After that, we load the parameters of the pre-trained GTN for fine-tuning while training the cross-conditional GPT of the dance generation framework. Note that, in order to prevent GTN from overfitting during fine-tuning, the GTN will freeze when it reaches the threshold epoch.

3. EXPERIMENTS

3.1. Datasets

3.1.1. Dataset for pre-training GTN

In pre-training GTN, we collect the background music of the dance corresponding to genres from the Internet. We collect 1-hour music data for each genre, constituting a music dataset of ten genres of 10 hours in length. The genres are the same as the ten genres in the dance-music-aligned dataset AIST++ [9], which are Break (BR), Ballet Jazz (JB), House (HO), Street Jazz (JS), Krump (KR), LA style Hip-hop (LH), Lock (LO), Middle Hip-hop (MH), Pop (PO), and Waack (WA).

3.1.2. Dataset for dance generation

During training and evaluating the dance generation framework with fine-tuning GTN, we use the AIST++ [9] dataset, which is the largest publicly available 3D dance-music-aligned dataset with genre labels to our knowledge. The AIST++ dataset includes 992 high-performance 60-FPS 3D pose sequences in SMPL format, split into 952 and 40 as training and validation sets, respectively.

3.2. Experiment Setup

For GTN, we extract an 80-dimensional mel-spectrogram from the music by Librosa. For the dance generation framework, the musical features are extracted by Librosa, including mel frequency cepstral coefficients (MFCC), MFCC delta, constant-Q chromagram, tempogram, and onset strength, resulting in a 438-dimensional

musical feature. Meanwhile, we extracted energy from the music by referring to Fastspeech2 [27], resulting in a 1-dimensional energy feature.

The settings of dance generation framework and GTN follow previous studies, and we set $\alpha = 1$, $\beta = 0.001$ for equation (3). The num of tokens is set to 10 in the genre token layer in GTN. The VQ-VAE in the framework is adopted from the choreographic memories in Bailando [18]. All dance and music data are cropped to the length of 4 seconds. We first pre-train GTN for 250 epochs and VQ-VAE for 500 epochs, then load the pre-trained VQ-VAE to generate the pose codes corresponding to the dance motions. After that, we load the pre-trained GTN and train the dance generation framework for 400 epochs, and after 90 epochs, we freeze the GTN.

We compare our proposed method with FACT [9] and Bailando [18], which are among the state-of-the-art dance generation methods. For each method, we generate dances in the condition of ten different AIST++ test music and two different starting pose codes or sequences, which results in 20 dance clips. We cut the generated dances into the length of 20 seconds for further experiments.

3.3. Experimental Results and Analysis

3.3.1. Genre Embedding Visualization

In Fig. 3, we illustrate genre embedding of different genres visualized by t-SNE method. We use the AIST++ test set to verify the GTN in the dance generation framework, which shows that the different genre embeddings are well separated from each other, proving that the GTN can infer genre from music properly.


**Fig. 3**. T-SNE visualization of genre embedding. Different colors mark different dance genres.

3.3.2. Subjective Evaluation

Due to the lack of objective quantitative metrics for genre consistency, we conduct subjective evaluation referring to the method in emotional speech synthesis [25, 26], to further evaluate the visual performance of dances generated by the proposed dance generation framework. The test is conducted with 24 subjects, who have participated in the dance training and have a certain understanding of the included dance genres. The subjects are asked to evaluate the dance quality and genre consistency, and rate the dances on a scale of 1-5 (the poorest score is 1, the best score is 5) with a 1-point interval.

The last two columns of Table 1 report the mean opinion scores (MOS) on dance quality and genre consistency. The proposed model outperforms all baseline models, indicating that GTN can establish the correlation between music and genres. Also, with the inferred genre to serve as a condition, the proposed dance generation framework can generate higher-quality and genre-consistent dances.

Table 1. Evaluation results of different dance generation frameworks. Dance quality and genre consistency are results of MOS with 95% confidence intervals. '*' denotes the proposed model. 'w/o' is short for 'without' in ablation study.

	Motion Quality		Motion Diversity			User Study	
	$FID_k \downarrow$	$FID_g \downarrow$	$DIV_k \uparrow$	$DIV_g \uparrow$	$BAS \uparrow$	Quality$\uparrow$	Consistency$\uparrow$
Ground-Truth	17.10	10.60	8.19	7.45	0.2484	-	-
FACT	37.31	34.87	5.75	5.47	0.2175	2.15±0.10	2.19±0.11
Bailando	32.11	9.95	6.09	5.70	0.2299	3.19±0.11	3.11±0.11
Proposed*	**29.51**	**8.57**	**6.15**	**6.65**	**0.2352**	**3.65±0.10**	**3.60±0.10**
w/o $\mathcal{L}_{CE-GTN}$	34.35	10.92	5.51	5.95	0.2292	3.35±0.13	3.25±0.13
w/o teacher-forcing	41.77	14.13	5.96	3.54	0.2298	3.49±0.13	3.29±0.14
w/o pre-trained GTN	30.55	7.03	7.03	6.60	0.2466	3.10±0.14	3.06±0.15
w/o energy	32.66	6.23	6.23	6.94	0.2366	3.40±0.15	3.33±0.15

Fig. 4. Genre consistent visualization. Different colors mark dances of different genres.

3.3.3. Objective Evaluation

For objective evaluation, following [18], we evaluate the quality and the diversity of the generated dance motions, and the alignment of the generated motions to the music beats. Specifically, as for the motion quality, we calculate the distribution distance between all the motion sequences of AIST++ and the generated dance motions through Fréchet Inception Distance (FID) [28] on the kinetic feature (noted as 'FID_k') and the geometric feature (noted as 'FID_g'). The lower the FID is, the closer the generated dances are to the ground-truth. As for the motion diversity, we calculate the average Euclidean distance on the kinetic feature (noted as 'DIV_k') and the geometric feature (noted as 'DIV_g'). The higher the DIV is, the larger the distance, meaning that the generated dances have more various movements. As for the alignment, we calculate the Beat Align Score (BAS) between the music beats and the motion beats. The higher the BAS, the more the dance is on the beat. The results of different methods as shown in Table 1.

It can be seen that our proposed framework outperforms baseline frameworks on all aspects. Specifically, our proposed dance generation framework achieves the lowest FID_k and FID_g, as well as the highest DIV_k, DIV_g and BAS. The proposed method generates higher-quality, more diverse dances by taking genre and energy into account, as well as improves the alignment between dance movements and the musical melody.

3.3.4. Ablation Study

Moreover, we conduct ablation study to explore the ability of our model to produce dances that are realistic. The results are shown in Table 1 and visual comparisons of the ablation study can also be found in the demo webpage.

As shown in Table 1, when the framework is optimized without the $\mathcal{L}_{CE-GTN}$ and the teacher-forcing scheme is not used during the training, the generated dances will be much dissimilar to the ground truth dances and have lower diversity. When the correlation between the velocity of the dance movement and the musical energy is not

considered, the expressiveness of the dance will reduce. When the pre-training and fine-tuning strategy is not used, although the framework can generate dances that are similar to ground truth, it will have limitations in genre consistency due to the poor generalization capability of GTN.

3.3.5. Genre Consistent Visualization

To further evaluate the genre consistency of the generated dances, we visualize the qualitative results from the joint skeletons of dances generated by our proposed framework and Bailando [18]. We randomly select a 20-second clip from the "LO" (Locking) genre dances generated by each framework, and sample the result at a frequency of 1 FPS.

As shown in Fig. 4, given a music clip, the dance generated by Bailando performs multiple dance genres, which is inappropriate with the music. While, our framework can infer the genre and generate a dance that matches the melody of the music and is consistent with the "LO" genre. More genre-consistent visualization results can be found in the demo-page.

4. CONCLUSION

In this paper, to address the challenge of genre consistency, we propose a dance generation framework, GTN-Bailando, which is able to establish the correlation between music and genres, and generate high-quality, genre-consistent dances on the basis of energy, music, and the genre inferred through GTN. Meanwhile, the proposed strategy of pre-training GTN on the large-scale dataset effectively improves its generalization ability. Experimental results on AIST++ show that our method can generate high-quality and genre-consistent dances based on the music.

5. REFERENCES

[1] Shuhei Tsuchida, Satoru Fukayama, Masahiro Hamasaki, and Masataka Goto, "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*, Delft, Netherlands, Nov. 2019, pp. 501–510.

[2] Omid Alemi, Jules Françoise, and Philippe Pasquier, "Groovenet: Real-time music-driven dance movement generation using artificial neural networks," *networks*, vol. 8, no. 17, pp. 26, 2017.

[3] Taoran Tang, Jia Jia, and Hanyang Mao, "Dance with melody: An lstm-autoencoder approach to music-oriented dance synthesis," *acm multimedia*, 2018.

[4] Nelson Yalta, Shinji Watanabe, Kazuhiro Nakadai, and Tetsuya Ogata, "Weakly-supervised deep recurrent neural networks for basic dance step generation," *international joint conference on neural network*, 2019.

[5] Hsuan-Kai Kao and Li Su, "Temporally guided music-to-body-movement generation," *acm multimedia*, 2020.

[6] Hyemin Ahn, Jaehun Kim, Kihyun Kim, and Songhwai Oh, "Generative autoregressive networks for 3d dancing move synthesis from music," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3501–3508, 2020.

[7] 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.

[8] Ruozi Huang, Huang Hu, Wei Wu, Kei Sawada, Mi Zhang, and Daxin Jiang, "Dance revolution: Long-term dance generation with music via curriculum learning," *Learning*, 2020.

[9] Ruilong Li, Shan Yang, David A. Ross, and Angjoo Kanazawa, "Ai choreographer: Music conditioned 3d dance generation with aist++," in *The IEEE International Conference on Computer Vision (ICCV)*, 2021.

[10] Guillermo Valle-Pérez, Gustav Eje Henter, Jonas Beskow, Andre Holzapfel, Pierre-Yves Oudeyer, and Simon Alexanderson, "Transflower: probabilistic autoregressive dance generation with multimodal attention," *ACM Transactions on Graphics (TOG)*, vol. 40, no. 6, pp. 1–14, 2021.

[11] Mingao Zhang, Changhong Liu, Yong Chen, Zhenchun Lei, and Mingwen Wang, "Music-to-dance generation with multiple conformer," in *Proceedings of the 2022 International Conference on Multimedia Retrieval*, 2022, pp. 34–38.

[12] Rukun Fan, Songhua Xu, and Weidong Geng, "Example-based automatic music-driven conventional dance motion synthesis," *IEEE transactions on visualization and computer graphics*, vol. 18, no. 3, pp. 501–515, 2011.

[13] Minho Lee, Kyogu Lee, and Jaeheung Park, "Music similarity-based approach to generating dance motion sequence," *Multimedia tools and applications*, vol. 62, no. 3, pp. 895–912, 2013.

[14] 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*, vol. 32, 2019.

[15] Kang Chen, Tan Zhipeng, Lei Jin, Song-Hai Zhang, Yuan-Chen Guo, Weidong Zhang, and Shi-Min Hu, "Choreomaster: choreography-oriented music-driven dance synthesis," *ACM Transactions on Graphics*, 2021.

[16] 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," *acm multimedia*, 2020.

[17] Ferda Ofli, Engin Erzin, Yücel Yemez, and A Murat Tekalp, "Learn2dance: Learning statistical music-to-dance mappings for choreography synthesis," *IEEE Transactions on Multimedia*, vol. 14, no. 3, pp. 747–759, 2011.

[18] Li Siyao, Weijiang Yu, Tianpei Gu, Chunze Lin, Quan Wang, Chen Qian, Chen Change Loy, and Ziwei Liu, "Bailando: 3d dance generation via actor-critic gpt with choreographic memory," in *CVPR*, 2022.

[19] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu, "Neural discrete representation learning.," *neural information processing systems*, 2017.

[20] Vijay R. Konda and John N. Tsitsiklis, "Actor-critic algorithms," *neural information processing systems*, 2002.

[21] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al., "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, pp. 9, 2019.

[22] Yuhang Huang, Junjie Zhang, Shuyan Liu, Qian Bao, Dan Zeng, Zhineng Chen, and Wu Liu, "Genre-conditioned long-term 3d dance generation driven by music," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 4858–4862.

[23] Jinwoo Kim, Heeseok Oh, Seongjean Kim, Hoseok Tong, and Sanghoon Lee, "A brand new dance partner: Music-conditioned pluralistic dancing controlled by multiple dance genres," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 3490–3500.

[24] Yuxuan Wang, Daisy Stanton, Yu Zhang, RJ Skerry-Ryan, Eric Battenberg, Joel Shor, Ying Xiao, Fei Ren, Ye Jia, and Rif A. Saurous, "Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," *international conference on machine learning*, 2018.

[25] Wu Pengfei, Ling Zhen-Hua, Liu Lijuan, Jiang Yuan, Wu Hongchuan, and Dai Lirong, "End-to-end emotional speech synthesis using style tokens and semi-supervised training," *IEEE Conference Proceedings*, 2019.

[26] Yi Lei, Shan Yang, Xinsheng Wang, and Lei Xie, "Msemotts: Multi-scale emotion transfer, prediction, and control for emotional speech synthesis," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 853–864, 2022.

[27] Yi Ren, Chenxu Hu, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu, "Fastspeech 2: Fast and high-quality end-to-end text to speech," in *International Conference on Learning Representations*, 2020.

[28] 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," *Asian Journal of Applied Science and Engineering*, 2017.