

Scalable Group Choreography via Variational Phase Manifold Learning

Nhat Le¹, Khoa Do^{2,3}, Xuan Bui^{2,3}, Tuong Do^{1,4}, Erman Tjiputra¹,
Quang D. Tran¹, and Anh Nguyen⁴

¹ AIOZ, Singapore

² University of Science, Ho Chi Minh City, Vietnam

³ Vietnam National University, Ho Chi Minh City, Vietnam

⁴ University of Liverpool, Liverpool L69 3BX, United Kingdom

Abstract. Generating group dance motion from the music is a challenging task with several industrial applications. Although several methods have been proposed to tackle this problem, most of them prioritize optimizing the fidelity in dancing movement, constrained by predetermined dancer counts in datasets. This limitation impedes adaptability to real-world applications. Our study addresses the scalability problem in group choreography while preserving naturalness and synchronization. In particular, we propose a phase-based variational generative model for group dance generation on learning a generative manifold. Our method achieves high-fidelity group dance motion and enables the generation with an unlimited number of dancers while consuming only a minimal and constant amount of memory. The intensive experiments on two public datasets show that our proposed method outperforms recent state-of-the-art approaches by a large margin and is scalable to a great number of dancers beyond the training data.



Fig. 1: We present a new group dance generation method that can generate a large number of dancers within a fixed resource consumption. The illustration shows a generated group dance sample with 100 dancers.

1 Introduction

The widespread availability of digital social media platforms has led to a significant increase in the popularity of creating dance videos. This heightened interest has resulted in the daily production and viewing of millions of dance videos on various online platforms [13, 32]. Recently, researchers from the computer vision and computer graphics community have focused on developing methods for generating authentic dance movements in response to music [6, 94]. These advancements have broad implications and find applications in diverse areas,

including animation [49, 50], virtual idols [59], virtual metaverses [42], and dance education [2, 67, 72]. Such techniques empower artists, animators, and educators by providing powerful tools to enhance their creative pursuits and improve the overall dance experience for performers and audiences alike.

Recently, significant strides have been made in creating dance motions for solo dance performers [12, 26, 34, 49, 59, 71, 83]. However, generating synchronized group dance movements that are both lifelike and in harmony with music remains a complex challenge [9, 27, 87]. The introduction of AIOZ-GDance [40] stands as the extensive dataset to aid in generating group choreography. Different aspects of group dance motion such as the study of consistency and diversity among the movements of dancers are explored in [39]. However, despite recent advances in the field, existing methods can either only generate dances for a limited number of dancers given the input audio [40, 86, 90], or would consume an expensively large amount of memory due to their design nature [39, 40]. Typically, these methods are usually based on specialized collaborative mechanisms such as cross-entity attention [40, 51], group global attention [39], or residual communication block [66] achieve coherent multi-person synthesis. In other words, their innate architectural design requires the network to process and synthesize all motions simultaneously, leading to excessively growing computational/memory costs with the increasing number of inputs, and hence extremely difficult to scale up.

Despite the comprehensive exploration of group dance motion generation, the predominant emphasis remains on optimizing the execution of the generated movements [14, 19, 31, 37, 52, 93]. These existing methodologies function within a confined scope, often tethered to a predetermined maximum number of dancers depicted in group choreography videos outlined by the predefined dataset [39]. This inherent limitation poses a potential constraint on the adaptability of trained models when transitioning to real-world applications [47]. In our pursuit, we center our efforts on addressing the critical aspect of *scalability in group dance*. Our primary objective revolves around not only expanding the number of dancers but also ensuring the preservation of the innate naturalness in the dance motions and the seamless synchronization between dancers throughout the entirety of their performance.

In the field of character motion control and synthesis, existing approaches can be broadly divided into two categories: deterministic [16, 26, 40, 49, 55, 86] and probabilistic [19, 39, 59, 61, 83]. The former tries to learn direct and deterministic mapping from the input conditioning signal (such as text, audio, or user controls) to the desired output motion. However, these methods usually regress towards the mean pose as they are typically trained by minimizing a regression objective under different outputs for similar inputs, resulting in freezing or drifting motion. On the other hand, probabilistic methods attempt to capture the distribution over possible motions for a given condition, allowing diverse and high-quality generation. In particular, Variational Autoencoder (VAE) [18, 47, 60, 61], or Diffusion Models [4, 39, 81, 83, 97] have been widely used. Although diffusion models have recently shown strong potential in generating high-quality motions [81], they generally operate on high-dimensional space

(same as the original data) [39, 81, 83], hence rendering them extremely challenging to use for scalable group generation. Meanwhile, standard VAE models often rely on a Gaussian latent space of fixed dimension for sampling a single latent vector to generate motions, making it difficult to adapt to a wide range of architectures. Besides, the fixed-size single latent space might also not be sufficient to represent the temporal dynamics of the whole motion sequence [7, 60].

Recent studies have demonstrated the significant advantages of learning motion features in the frequency domain. Specifically, [22, 74–76] have found that different motion skills can be effectively represented by several phase variables that can faithfully capture the spatial-temporal alignment of a wide range of movements. This can be used to combine with common deep motion synthesis networks to enhance its capability, facilitating natural, and stable motion generations by enforcing a unidirectional motion transition (i.e., no backward direction in time) to prevent the motion from being stuck temporally [74]. Unlike those prior works, where the phases are utilized as auxiliary signals to assist in the motion synthesis process [68, 74, 88], we propose to learn to generate the phases conditioned on the input audio signal in an end-to-end manner, eliminating the need of training extra components. Our generative phase manifold can be used to generate new motions efficiently, from which the latent motion curves are sampled in frequency domain and decoded into the original motion space by a learned decoder.

In this paper, our goal is to develop a scalable technique for group dance generation, a phase-based variational generative model for scalable group dance generation, namely Phase-conditioned Dance VAE (PDVAE). To our knowledge, PDVAE is the first method to represent the variational latent space using phase parameters in the frequency domain of the motion curves. Our method goes beyond the conventional VAE approach that typically relies on a single latent vector drawn from a Gaussian distribution, which is unable to adequately represent the temporal information of the motion sequence (e.g., the time dimension is squeezed out). Figure 1 illustrates an example of a dance motion sequence with 100 dancers generated by our model. In contrast to existing methods, our model can generate crowd-dance animations for an unlimited number of dancers without increasing computational burden (i.e., memory consumption remains constant), while still maintaining the performance and fidelity of the generated choreographies. To summarize, our key contributions are as follows:

- We introduce PDVAE, a phase-based variational generative model for scalable group dance generation. The method focuses on generating large-scale group dance under limited resources.
- To effectively learn the manifold that is group-consistent (i.e., dancers within a group lie upon the same manifold), we propose a group consistency loss that enforces the networks to encode the latent phase manifold to be identical for the same group given the input music.
- Extensive experiments along with thorough user study evaluations demonstrate the state-of-the-art performance of our model while achieving effective scalability.

2 Related Work

2.1 Music-driven Choreography

Crafting natural human choreography derived from music presents a multifaceted challenge [30, 62, 89]. Certain methodologies integrate constraints based on music-motion similarity matching to ensure coherence between the generated motion and the music [10, 33, 35, 43, 65, 69]. However, many of these approaches rely considerably on heuristic algorithms, stitching together pre-existing dance segments sourced from a limited music-dance database [10, 48]. While successful in generating extended and realistic dance sequences, these methods face limitations when endeavoring to create entirely novel dance fragments [57, 90].

In the recent period, advancements have been evident in the domain of converting music into dance movements through various techniques such as Convolutional Networks (CNN) [1, 8, 77, 91, 92, 99], Recurrent Networks (RNN) [3, 26, 77, 79, 87], Graph Neural Networks (GNN) [5, 12, 64, 95], Generative Adversarial Networks (GAN) [41, 77], or Transformer models [34, 44, 45, 49, 59, 71]. These methodologies typically depend on various inputs, including the present music and a concise history of previous dance movements, to forecast forthcoming sequences of human poses. This innovation introduces a music-text feature fusion module, engineered to amalgamate inputs into a motion decoder, thereby enabling the creation of dance sequences conditioned upon both musical and textual instructions [17].

While these techniques show promise in generating authentic and lifelike dance motions, they often struggle to synchronize movements seamlessly across multiple dancers [40]. Specifically, achieving coordination and harmony among dancers necessitates consideration of spatial and temporal relationships, their interactions, and the overall choreographic structure [84]. Consequently, further advancements in this domain are being pursued to tackle these complexities [25, 47, 59, 98]. For instance, Perez *et al.* [59] integrate a multimodal transformer encoder with a normalizing-flow-based decoder to estimate a probability distribution encompassing potential subsequent poses. Additionally, Feng *et al.* [11] enable long-term generation by imposing a motion repeat constraint to forecast future frames while considering historical motions. A recent study by Le *et al.* [39] investigates consistency and diversity factors between the generated motions of two or more dancers within a timeframe. However, these approaches are limited to a predefined number of dancers, heavily constrained by the maximum number of dancers in the dataset-provided videos.

2.2 Motion Manifold Learning

Motion manifold learning has attracted considerable attention in computer vision and artificial intelligence, with the primary goal of comprehending the fundamental structures inherent in human movement and dynamics [16, 24, 56, 82]. Its distinctive ability to generate human movement patterns presents numerous opportunities to comprehend intrinsic motion dynamics, manage nonlinear

relationships in motion data, and acquire contextual and hierarchical representations [29, 63, 88]. Consequently, various methodologies have emerged, each contributing distinct perspectives and techniques to advance the comprehension and synthesis of human motion [28, 38, 63, 68].

For instance, Holden *et al.* [23] generate character movements by mapping high-level parameters to the human motion manifold, allowing diverse movements without manual preprocessing, and enabling post-generation editing for natural, smooth motion sequences. MotionCLIP [80] introduces a 3D human motion auto-encoder aligned with the Contrastive Language-Image Pre-training (CLIP) model’s space, enabling semantic text-based motion generation, disentangled editing, and abstract language specification, leveraging CLIP’s rich semantic knowledge within the motion manifold. Recently, Sun *et al.* [78] utilize a VQ-VAE to acquire a low-dimensional manifold, effectively cleansing the motion sequences. Regarding the scalability of group dance generation, the concept of the motion manifold emerges as a potential solution for addressing the restricted number of dancers in the dataset. This direction enables the learning of a distribution that facilitates the extraction of dance motions, allowing for the presentation of cohesive group dance sequences.

3 Variational Phase Manifold Learning for Scalable Group Choreography

3.1 Task Definition

Given an input music sequence $\mathbf{a} = \{a_1, a_t, \dots, a_T\}$ with $t = \{1, \dots, T\}$ indicates the index of the music frames, our goal is to generate the group motion sequences of N arbitrary dancers: $\mathbf{x} = \{x_1^1, \dots, x_T^1; \dots; x_1^N, \dots, x_T^N\}$ where x_t^n is the pose of n -th dancer at frame t . We follow [19, 53] and represent dance as a sequence of poses in the 24-joint of the SMPL body model [54], using the 6D continuous rotation [96] for every joint, along with 3D joint positions and velocities. Additionally, the corresponding 3D root translation vectors are concatenated into the pose representations to involve the trajectory of motion. Previous group dance methods [40], which generate the whole group at once, cannot deal with the increasing number of dancers and can only create group sequences up to a pre-defined number of dancers, due to the vast complexity of the architecture. In contrast, we aim to generate group dance with an unlimited number of dancers.

3.2 Phase-conditioned Dance VAE

Our goal is to learn a continuous manifold such that the motion can be generated by sampling from this learned manifold. We assume that although different dancers within the same group may present visually distinctive movements, the properties of their motions, such as timing, periodicity, or temporal alignment are intrinsically similar. Drawing inspiration from [74], we aim to learn a generative phase representation for each group of dancer in order to synthesize their




Fig. 2: Overview of our Phase-conditioned Dance VAE (PDVAE) for scalable group dance generation. It consists of an Encoder, a Prior, and a Decoder network. During training, we encode the corresponding motion and music inputs into a latent phase manifold, which is variational and parameterized by the frequency domain parameters of periodic functions. The latent phases can be sampled from the manifold and then decoded back to the original data space to obtain new motions. The consistency loss $\mathcal{L}_{\text{csc}}$ is further imposed to constrain the manifold to be consistently unified for dancers that belong to the same group. At inference stage, only the Prior and the Decoder are used to synthesize group dances efficiently.

motion indefinitely. Our generative model is built upon the conditional Variational Autoencoder architecture [73], thanks to its diverse generation capability and fast sampling speed. However, instead of directly encoding the data into a Gaussian latent distribution as in common VAE approaches [15, 41, 47, 60, 61], we model the latent variational distribution by the phase parameters extracted from the latent motion curve, which we call variational phase manifold. The latent phase manifold is well-structured and can well describe key characteristics of motion (such as its timing, local periodicity, and transition), which benefits learning motion features [74].

The overview of our Phase-conditioned Dance VAE is illustrated in Figure 2. Specifically, the model contains three main networks: an encoder $\mathcal{E}$ to capture the approximate posterior distribution conditioned on both motion and music $q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{a})$, a prior network $\mathcal{P}$ to learn the conditional prior given only the music $p_\theta(\mathbf{z}|\mathbf{a})$, and a decoder $\mathcal{D}$ to learn to reconstruct the data from the latent distribution $p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{a})$. The new motion is generated by sampling the frequency-domain parameters predicted by the prior network, which is then passed through the decoder network to reconstruct the motion in the original data space. Furthermore, we adopt Transformer-based architecture [85] in each network to effectively capture long-range dependencies and holistic context of the dance sequence.

Encoder. The encoder $\mathcal{E}$ is expected to take both the motion and music feature sequence as input, and produce a distribution over possible latent variables capturing the cross-modal relationship between them. To transform the joint input space into a learned phase manifold, we adopt the Transformer decoder architecture where the Cross-Attention mechanism [85] is utilized to learn the relationship between the motion and the music. Accordingly, the output of the encoder is a batch of latent curves (i.e., the activation sequences per channel) that can particularly capture different spatial and temporal aspects of the motion sequence. However, instead of training the model to directly reconstruct the input motion from the extracted latent curves, we further enforce each channel of the latent space to have a periodic functional form (i.e., sinusoidal). This enables us to effectively learn a compact parameterization for each latent channel from a small set of parameters in the frequency domain.

Generative Variational Phase Manifold. Here we focus on learning the periodicity and non-linear temporal alignment of the motion in the latent space. In particular, given the output latent curves from the encoder $\mathbf{L} = \mathcal{E}(x, a) \in \mathbb{R}^{D \times T}$ with D is the number of desired phase channels to be extracted from the motion, we parameterize each latent curve in $\mathbf{L}$ using a sinusoidal function with amplitude ($\mathbf{A}$), frequency ($\mathbf{F}$), offset ($\mathbf{B}$) and phase shift ($\mathbf{S}$) parameters [74]. To allow for variational phase manifold learning, we opt to predict two sets of parameters $\mu_{\mathcal{E}} = \{\mu^A; \mu^F; \mu^B; \mu^S\}$ and $\sigma_{\mathcal{E}} = \{\sigma^A; \sigma^F; \sigma^B; \sigma^S\}$, which corresponds to the mean and variance of $\mathbb{R}^{4D}$ dimensional Gaussian distribution:

$$q_{\phi}(\mathbf{z}|\mathbf{x}, \mathbf{a}) = \mathcal{N}(\mathbf{z}; \mu_{\mathcal{E}}, \sigma_{\mathcal{E}}) \quad (1)$$

To do so, we first apply differentiable Fast Fourier Transform (FFT) to each channel of the latent curve $\mathbf{L}$ and create the zero-indexed matrix of Fourier coefficients as $\mathbf{c} = FFT(\mathbf{L})$ with $\mathbf{c} \in \mathbb{C}^{D \times K+1}$, $K = \lfloor \frac{T}{2} \rfloor$. Correspondingly, we compute the per channel power spectrum $\mathbf{p} \in \mathbb{R}^{D \times K+1}$ as $\mathbf{p}_{i,j} = \frac{2}{N} |\mathbf{c}_{i,j}|^2$, where i is the channel index and j is the index for the frequency bands. Correspondingly, the distributional mean parameters of the periodic sinusoidal function are then calculated as follows:

$$\mu_i^A = \sqrt{\frac{2}{T} \sum_{j=1}^K \mathbf{p}_{i,j}}, \quad \mu_i^F = \frac{\sum_{j=1}^K \mathbf{f}_j \cdot \mathbf{p}_{i,j}}{\sum_{j=1}^K \mathbf{p}_{i,j}}, \quad \mu_i^B = \frac{\mathbf{c}_{i,0}}{T}, \quad (2)$$

where $\mathbf{f} = (0, \frac{1}{T}, \dots, \frac{K}{T})$ is the frequencies vector. At the same time, the phase shift $\mathbf{S}$ is predicted using a fully-connected (FC) layer with two arctan activation as:

$$(s_y, s_x) = FC(\mathbf{L}_i), \quad \mu_i^S = \arctan(s_y, s_x), \quad (3)$$

To predict the distributional variance of the phase amplitude and phase shift parameters $\{\sigma^A, \sigma^S\}$, We additionally apply a separate two-layer MLP network over each channel of the latent curves, similar to Equation 3. The variational latent phase parameters are sampled by utilizing parameterization trick [36],

i.e., $\mathbf{A} \sim \mathcal{N}(\mu^A, \sigma^A)$ and $\mathbf{S} \sim \mathcal{N}(\mu^S, \sigma^S)$. In our experiments, we find that sampling the phase frequency $\mathbf{F}$ and offset $\mathbf{B}$ often produce unstable and non-coherent group movements. This might be because the frequency amplitudes of the dancers within the same group are likely to associate with the rhythmic pattern of the musical beats while the offsets capture their alignment, thereby should be consistent with each other. Therefore, we treat those parameters as deterministic by constraining their variance to zero.

Finally, the sampled set of phase parameters $\mathbf{z} = \{\mathbf{A}; \mathbf{F}; \mathbf{B}; \mathbf{S}\}$ are used to reconstruct a parametric latent space consisting of multiple periodic curves to represent each intrinsic property of the motion by:

$$\hat{\mathbf{L}} = \mathbf{A} \cdot \sin(2\pi \cdot (\mathbf{F} \cdot \mathcal{T} - \mathbf{S})) + \mathbf{B} \quad (4)$$

where $\mathcal{T}$ is a known time window series obtained by evenly spacing the timesteps from 0 to T . Intuitively this curve construction procedure can be viewed as a "quantization" layer to enforce the network to learn to represent the motion features in the frequency domain, which is useful in representing different aspects of human motion such as their timing and periodicity. In the last step, a decoder is utilized to reconstruct the original motion signals from the set of parametric latent curves.

Decoder. To decode the latent space into the original motion space, previous works [60, 61] have to use a sinusoidal positional encoding sequence with duration T as the proxy input to the sequence decoder. This is because their latent space is only formed by single latent vectors following a Gaussian distribution, which cannot span the time dimension. However, we observe that it usually results in unstable and inconsistent movements, as the proxy sequence is generic and usually contains less meaningful information for the decoder. Meanwhile, our method does not suffer from this problem as our latent space is built on multiple curves that can represent the motion information through time, thanks to the phase parameters. Subsequently, our decoder $\mathcal{D}$ is based on Transformer decoder architecture that takes the constructed parametric latent curve, as well as the music features as inputs, to reconstruct the corresponding dance motions. Here, we also utilize the cross-attention model [85] where we consider the sequence of and music features as key and value along with the sampled latent curves as the query. The output of the decoder is a sequence of T vectors in $\mathbb{R}^D$, which is then projected back to the original motion dimensions through a linear layer, to obtain the reconstructed outputs $\hat{\mathbf{x}} = p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{a})$. We additionally employ a global trajectory predictor to predict the global translation of the root joint based on the generated local motions [19, 46], in order to avoid intersection problems between dancers. We provide details of our global predictor in the supplemental.

Prior Network. Since the ground-truth motion is generally inaccessible at test time (i.e., we only have access to the music), we also need to learn a prior $\mathcal{P}$ to match the posterior distribution of motion from which the latent phase can be sampled. Specifically, We follow the procedure similar to Section 3.2 (Equation 1-4) to predict the Gaussian distribution conditioned on the music

sequence $\mathbf{a}$, which is then used for sampling the latent phases:

$$p_\theta(\mathbf{z}|\mathbf{a}) = \mathcal{N}(\mathbf{z}; \mu_{\mathcal{P}}, \sigma_{\mathcal{P}}) \quad (5)$$

where a Transformer encoder is used to encode the input conditioning music sequence and predict the corresponding $\mu_{\mathcal{P}}$ and $\sigma_{\mathcal{P}}$. We implement the prior network similarly to the encoder network, however, we use self-attention mechanism [85] to capture the global music context. Learning the conditional prior is crucial for the conditional VAE to generalize to diverse types of music and motion. Intuitively speaking, each latent variable $\mathbf{z}$ is expected to represent possible dance motions $\mathbf{x}$ conforming to the music context $\mathbf{a}$. Therefore, the prior should be able to encode different latent distributions given different musics.

3.3 Training

During training, we consider the following variational lower bound [73] to mainly train our dance generation VAE model:

$$\log p_\theta(\mathbf{x}|\mathbf{a}) \geq \mathbb{E}_{q_\phi} [\log p_\theta(\mathbf{x}|\mathbf{z}, \mathbf{a})] - D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x}, \mathbf{a}) \| p_\theta(\mathbf{z}|\mathbf{a})) \quad (6)$$

In practice, we apply the conditional VAE loss as similar to [61], which is defined as the weighted sum $\mathcal{L}_{\text{cvae}} = \mathcal{L}_{\text{rec}} + \lambda_{\text{KL}} \mathcal{L}_{\text{KL}}$. In particular, the reconstruction term $\mathcal{L}_{\text{rec}}$ measures the motion reconstruction error given the decoder output (via a smooth-L1 loss). The KL divergence term $\mathcal{L}_{\text{KL}}$ compares the divergence D_{KL} between the posterior and the prior distribution to enforce them to be close to each other.

The conditional VAE objective above is calculated for each dancer separately and cannot capture the correlation between dancers within a group. Therefore, it is essential to impose consistency among dancers and avoid strange effects such as unsynchronized dance. To this end, we propose a group consistency loss by enforcing the latent phase manifold to be similar for the same group, given the input music. Specifically, we first calculate the phase manifold features based on the frequency domain parameters as follows:

$$\mathbf{P}_{2i-1} = \mathbf{A}_i \sin(2\pi \cdot \mathbf{S}_i), \quad \mathbf{P}_{2i} = \mathbf{A}_i \cos(2\pi \cdot \mathbf{S}_i) \quad (7)$$

where $\mathbf{P} \in \mathbb{R}^{2D}$ is the phase manifold vector that encodes the spatial-temporal information of the motion state. The phase feature may look similar to the positional encodings of transformers [85] in the sense that both use multi-resolution sinusoidal functions. However, the phase feature is much richer in terms of representation capacity since it learns to embed the spatial (body joints) and temporal (positions in time) information of the motion curves, whereas the positional encodings only encode the position of words. Finally, our consistency objective is to constrain phase manifold between dancers within a group to be as close as possible to each other, which is formulated as:

$$\mathcal{L}_{\text{csc}} = D_{\text{KL}}(q_\phi(\mathbf{z}|\mathbf{x}^m, \mathbf{a}) \| (q_\phi(\mathbf{z}|\mathbf{x}^n, \mathbf{a}))) + \|\mathbf{P}^m - \mathbf{P}^n\|_2^2 \quad (8)$$

where the first term encourages the network to map different dancers belonging to the same group ($\mathbf{x}^m$ and $\mathbf{x}^n$) into the same set of distributional phase parameters while the second term penalizes the discrepancy in their corresponding phase manifolds. In general, this loss is applied to ensure every dancer is embedded into a single unified manifold that can effectively represent their corresponding group. To summarize, our total training loss is defined as the combination of the VAE loss and the consistency loss $\mathcal{L} = \mathcal{L}_{\text{vae}} + \lambda_{\text{csc}} \mathcal{L}_{\text{csc}}$.

For testing, we can efficiently generate motions for an unlimited number of dancers by indefinitely drawing samples from the learned continuous group-consistent phase manifold. It is noteworthy that for inference, we only need to process the prior network once to obtain the latent distribution. To generate a new motion, we can sample from this latent (Gaussian) distribution and use the decoder to decode it back to the motion space. This approach is much more efficient and has significantly higher scalability than previous approaches [39, 40] that is limited by the number of dancers processed simultaneously by the entire network.

4 Experiments

4.1 Implementation Details

Our model was trained on 4 NVIDIA V100 GPUs using Adam optimizer with a fixed learning rate of 10^{-4} and a mini-batch size of 32 per GPU. For training losses, the weights are empirically set to $\lambda_{\text{KL}} = 5 \times 10^{-4}$ and $\lambda_{\text{csc}} = 10^{-4}$, respectively. The Transformer encoders and decoders consist of 5 layers for both encoder, decoder, and prior Network with 512-dimensional hidden units. Meanwhile, the number of latent phase channels is set to 256. To further capture the periodic nature of the phase feature, we also use Siren activation following the initialization scheme as in [70]. This can effectively model the periodicity inherent in the motion data, and thus can benefit motion synthesis [68].

4.2 Experimental Settings

Dataset In our experiments, we utilize the AIOZ-GDance [40] and AIST-M [90] datasets. AIOZ-GDance is the largest music-driven dataset focusing on group dance, encompassing paired music and 3D group motions extracted from in-the-wild videos through a semi-automatic process. This dataset spans 7 dance styles and 16 music genres. For consistency, we adhere to the training and testing split outlined in [40] during our experiments.

Evaluation Protocol We employ several metrics to assess the quality of individual dance motions, including Frechet Inception Distance (FID) [20, 49], Motion-Music Consistency (MMC) [49], and Generation Diversity (GenDiv) [26, 41, 49], along with the Physical Foot Contact score (PFC) [83]. Specifically, the FID score gauges the realism of individual dance movements concerning the

ground-truth dance. MMC assesses the matching similarity between motion and music beats, reflecting how well-generated dances synchronize with the music’s rhythm. GenDiv is computed as the average pairwise distance of kinetic features among motions [58]. PFC evaluates the physical plausibility of foot movements by determining the agreement between the acceleration of the character’s center of mass and the foot’s velocity.

In assessing the quality of group dance, we adopt three metrics outlined in [40]: Group Motion Realism (GMR), Group Motion Correlation (GMC), and Trajectory Intersection Frequency (TIF). Broadly, GMR gauges the realism of generated group motions in comparison to ground-truth data, employing Frechet Inception Distance on extracted group motion features. GMC evaluates the synchronization among dancers within the generated group by computing their cross-correlation. TIF quantifies the frequency of collisions among the generated dancers during their dance movements.

Baselines Our method is subjected to comparison with various recent techniques in music-driven dance generation, namely FACT [49], Transflower [59], and EDGE [83]. These approaches are adapted for benchmarking within the context of group dance generation, as outlined in [40], considering that their original designs were tailored for single-dance scenarios. Additionally, our evaluation includes a comparison with GDanceR [40], GCD [39], and DanY [90]. All of the mentioned works are specifically designed for the generation of group choreography.

Table 1: Performance comparison.

Dataset	Method	FID↓	MMC↑	GenDiv↑	PFC↓	GMR↓	GMC↑	TIF↓
AIOZ-GDANCE [40]	FACT [49]	56.20	0.222	8.64	3.52	101.52	62.68	0.321
	Transflower [59]	37.73	0.217	8.74	3.07	81.17	60.78	0.332
	EDGE [83]	31.40	0.264	9.57	2.63	63.35	61.72	0.356
	GDANCER [40]	43.90	0.250	9.23	3.05	51.27	79.01	0.217
	GCD [39]	31.16	0.261	10.87	2.53	31.47	80.97	0.167
	PDVAE (ours)	31.01	0.271	10.98	2.33	30.08	84.52	0.163
AIST-M [90]	GDANCER [40]	52.90	0.222	6.52	1.93	65.13	60.56	0.121
	GCD [39]	35.36	0.245	10.97	1.52	42.52	72.15	0.083
	DanY [90]	40.25	0.240	11.40	1.65	50.29	63.53	0.137
	PDVAE (ours)	31.49	0.257	11.81	1.42	41.24	78.64	0.076

4.3 Experimental Results

Quality Comparison Table 1 presents a comparison among the baselines FACT [49], Transflower [59], EDGE [83], GDanceR [40], GCD [39], and our proposed manifold-based method. The results clearly demonstrate that our model almost outperforms all baselines across all evaluations on two datasets AIOZ-GDANCE [40] and AIST-M [90]. We observe that recent diffusion-based dance generation models such as EDGE or GCD yield competitive performance on both single-dance metrics (FID, MMC, GenDiv, and PFC) and group dance

metrics (GMR, GMC, and TIF). However, due to limitations in their training procedures, they still struggle when dealing with generating multiple dancing motions in the context of a high quantity of dancers, as shown by their low performance compared to our proposal. This result implies that our method successfully preserves the performance of dancing motions when the number of dancers is increased. Besides, Figure 3 also shows that our proposed methods outperform other state-of-the-art methods such as GDanceR or GCD in dealing with monotonous, repetitive, sinking, and overlapping dance motions.




Fig. 3: Visualization of a dancing sample between different methods. GDanceR displays monotonous, repetitive, or sinking dance motions. GCD exhibits more divergence in dance motions, yet dancers may intersect since their optimization does not address this issue explicitly. Blue boxes mark these issues. In contrast, our manifold-based solution ensures the divergence of dancing motions, while the phase motion path demonstrates its effectiveness in addressing floating and crossing issues in group dances.

Table 2: Performance of group dance generation methods when we increase the number of generated dancers. The experiments are done with common consumer GPUs with 4GB memory. (N/A means models could not run due to inadequate memory footprint)

#Generated Dancers	Method	FID↓	MMC↑	GenDiv↑	GMR↓	GMC↑	TIF↓
5	GDanceR [40]	44.19	0.249	8.99	55.05	78.72	0.218
	GCD [39]	35.08	0.264	9.92	38.43	81.44	0.168
	PDVAE (ours)	31.35	0.268	10.05	32.58	84.56	0.161
10	GDanceR [40]	57.93	0.225	7.63	73.32	72.79	0.386
	GCD [39]	N/A (Memory)	N/A (Memory)	N/A (Memory)	N/A (Memory)	N/A (Memory)	N/A (Memory)
	PDVAE (ours)	32.19	0.269	8.99	34.32	86.96	0.193
100	GDanceR [40]	N/A (Memory)	N/A (Memory)	N/A (Memory)	N/A (Memory)	N/A (Memory)	N/A (Memory)
	GCD [39]	N/A (Memory)	N/A (Memory)	N/A (Memory)	N/A (Memory)	N/A (Memory)	N/A (Memory)
	PDVAE (ours)	30.97	0.268	8.76	38.13	85.73	0.222

Scalable Generation Analysis Table 2 illustrates the performance of different group dance generation methods (GCD, GDancer, and Ours) when generating

dance movements with an increasing number of dancers. When the number of dancers is increased to 10, GCD appears to run out of memory, which is also observed in GDanceR when the number of dancers increases to 100. Regardless of the number of dancers, our method consistently achieves stable and competitive results. This implies that our proposed method successfully addresses the scalability issue in group dance generation without compromising the overall performance of each individual dance motion. Figure 4 illustrates the memory consumption to generate dance motions in groups for each method. Noticeably, our proposal still achieves the highest performance while consuming immensely fewer resources required for generating group dance motions (See Figure 5 for illustrations). This, again, indicates that our method successfully breaks the barrier of limited generated dancers by using the manifold.

Ablation Analysis Table 3 presents the performance improvements achieved through the integration of consistency loss and phase manifold. Additionally, we showcase the effectiveness of our proposed approach across three different backbones: Transformer [49, 85], LSTM [21], and CNN [99]. Evaluation metrics including FID, GMR, and GMC are utilized. The results indicate that the absence of consistency loss leads to an increase in GMR and a decrease in GMC, suggesting a significant enhancement in the realism and correlation of group dance motions facilitated by the inclusion of the proposed objective. Meanwhile, with out the phase manifold, the model exhibits remarkably higher scores in both the FID and GMR metrics, suggesting that phase manifold can effectively improve the distinction in dance motions while maintaining the realism of group dances, even when the number of dancers in a group is large. In comparing three backbones—Transformer, LSTM, and CNN—we have observed that the chosen Transformer backbone achieved the best results compared to LSTM or CNN.

Method	FID↓	GMR↓	GMC↑
Ours	31.01	30.08	84.52
Ours w/o Consistency Loss	35.35	57.63	66.72
Ours w/o Phase Manifold	41.78	45.32	77.93
Ours w LSTM	41.29	47.47	71.82
Ours w CNN	36.99	44.94	75.77

Table 3: Module contribution and loss analysis.




Fig. 4: Memory usage vs. number of dancers in different dance generators.

4.4 User Study

User studies are vital for evaluating generative models, as user perception is pivotal for downstream applications; thus, we conducted two studies with around 70 participants each, diverse in background, with experience in music and dance, aged between 20 to 40, consisting of approximately 47% females and 53% males, to assess the effectiveness of our approach in group choreography generation.

In the user study, we aim to assess the realism of sample outputs with more and more dancers. Specifically, participants assign scores ranging from 0 to 10 to evaluate the realism of each dance clip with 2 to 10 dancers. Figure 6 shows that, across all methods, the more the number of dancers is increased, the lower the realism is found. However, the drop in realism of our proposed method is the least compared to GCD and GDanceR. The results indicate our method’s good performance compared to other baselines when the number of dancers increases.




Fig. 5: Visualization of scalable dancers.




Fig. 6: Realism between different methods when number of dancers is varied.

5 Discussion

While our approach leverages the VAE as a primary solution for generating a manifold, it is important to acknowledge certain inherent limitations associated with this choice. One notable challenge is the susceptibility to issues such as posterior collapse and unstable sampling within the VAE framework. These challenges can result in generated group dance motions that may not consistently meet performance expectations.

One specific manifestation of this limitation is the potential for false decoding when sampling points that lie too far from the learned distribution. This scenario can lead to unexpected rotations or disruptions in the physics of the generated content. The impact of this problem becomes evident in instances where the generated samples deviate significantly from the anticipated distribution, introducing inaccuracies and distortions.

To address these challenges, we recognize the need for ongoing efforts to mitigate the effects of posterior collapse and unstable sampling. While the problem is acknowledged, our approach incorporates measures to limit its impact. Future research directions could explore alternative generative models or additional techniques to enhance the robustness and reliability of the generated results in the face of these identified limitations.

References

1. Ahn, H., Kim, J., Kim, K., Oh, S.: Generative autoregressive networks for 3d dancing move synthesis from music. *IEEE Robotics and Automation Letters (RA-L)* (2020)
2. Alaoui, S.F., Henry, C., Jacquemin, C.: Physical modelling for interactive installations and the performing arts. *International Journal of Performance Arts and Digital Media* (2014)
3. Alemi, O., François, J., Pasquier, P.: Groovenet: Real-time music-driven dance movement generation using artificial neural networks. *Networks* (2017)
4. Alexanderson, S., Nagy, R., Beskow, J., Henter, G.E.: Listen, denoise, action! audio-driven motion synthesis with diffusion models. *ACM Transactions on Graphics (TOG)* (2023)
5. Au, H.Y., Chen, J., Jiang, J., Guo, Y.: Choreograph: Music-conditioned automatic dance choreography over a style and tempo consistent dynamic graph. In: *ACM International Conference on Multimedia* (2022)
6. Bisig, D.: Generative dance-a taxonomy and survey. In: *International Conference on Movement and Computing* (2022)
7. Cervantes, P., Sekikawa, Y., Sato, I., Shinoda, K.: Implicit neural representations for variable length human motion generation. In: *ECCV* (2022)
8. Chan, C., Ginosar, S., Zhou, T., Efros, A.A.: Everybody dance now. In: *ICCV* (2019)
9. Chen, K., Tan, Z., Lei, J., Zhang, S.H., Guo, Y.C., Zhang, W., Hu, S.M.: Choreomaster: choreography-oriented music-driven dance synthesis. *ACM Transactions on Graphics (TOG)* (2021)
10. Fan, R., Xu, S., Geng, W.: Example-based automatic music-driven conventional dance motion synthesis. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* (2011)
11. Feng, B., Ao, T., Liu, Z., Ju, W., Liu, L., Zhang, M.: Robust dancer: Long-term 3d dance synthesis using unpaired data. *arXiv* (2023)
12. 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)
13. Fink, B., Bläsing, B., Ravignani, A., Shackelford, T.K.: Evolution and functions of human dance. *Evolution and Human Behavior* (2021)
14. Gärtner, E., Metz, L., Andriluka, M., Freeman, C.D., Sminchisescu, C.: Transformer-based learned optimization. In: *CVPR* (2023)
15. Ghorbani, S., Ferstl, Y., Holden, D., Troje, N.F., Carbonneau, M.A.: Zeroeggs: Zero-shot example-based gesture generation from speech. In: *Computer Graphics Forum* (2023)
16. Ghosh, A., Cheema, N., Oguz, C., Theobalt, C., Slusallek, P.: Synthesis of compositional animations from textual descriptions. In: *ICCV* (2021)
17. Gong, K., Lian, D., Chang, H., Guo, C., Zuo, X., Jiang, Z., Wang, X.: Tm2d: Bimodality driven 3d dance generation via music-text integration. In: *ICCV* (2023)
18. Guo, C., Zou, S., Zuo, X., Wang, S., Ji, W., Li, X., Cheng, L.: Generating diverse and natural 3d human motions from text. In: *CVPR* (2022)
19. He, C., Saito, J., Zachary, J., Rushmeier, H., Zhou, Y.: Nemf: Neural motion fields for kinematic animation. *NeurIPS* **35**, 4244–4256 (2022)
20. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: *NeurIPS* (2017)

21. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Computation* (1997)
22. Holden, D., Komura, T., Saito, J.: Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)* **36**(4), 1–13 (2017)
23. Holden, D., Saito, J., Komura, T.: A deep learning framework for character motion synthesis and editing. *ACM Transactions on Graphics (TOG)* (2016)
24. Holden, D., Saito, J., Komura, T., Joyce, T.: Learning motion manifolds with convolutional autoencoders. In: *SIGGRAPH Asia 2015 technical briefs* (2015)
25. Hong, F., Zhang, M., Pan, L., Cai, Z., Yang, L., Liu, Z.: Avatarclip: zero-shot text-driven generation and animation of 3d avatars. *ACM Transactions on Graphics (TOG)* (2022)
26. Huang, R., Hu, H., Wu, W., Sawada, K., Zhang, M.: Dance revolution: Long sequence dance generation with music via curriculum learning. *CoRR* (2020)
27. Huang, Y., Zhang, J., Liu, S., Bao, Q., Zeng, D., Chen, Z., Liu, W.: Genre-conditioned long-term 3d dance generation driven by music. In: *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 4858–4862. IEEE (2022)
28. Jiang, C., Cornman, A., Park, C., Sapp, B., Zhou, Y., Anguelov, D., et al.: Motiondiffuser: Controllable multi-agent motion prediction using diffusion. In: *CVPR* (2023)
29. Jiang, Y., Won, J., Ye, Y., Liu, C.K.: Drop: Dynamics responses from human motion prior and projective dynamics. In: *SIGGRAPH Asia 2023 Conference papers* (2023)
30. Joshi, M., Chakrabarty, S.: An extensive review of computational dance automation techniques and applications. *Proceedings of the Royal Society A* (2021)
31. Karunratanakul, K., Preechakul, K., Suwajanakorn, S., Tang, S.: Guided motion diffusion for controllable human motion synthesis. In: *ICCV* (2023)
32. Kico, I., Grammalidis, N., Christidis, Y., Liarokapis, F.: Digitization and visualization of folk dances in cultural heritage: A review. *Inventions* (2018)
33. Kim, J.W., Fouad, H., Hahn, J.K.: Making them dance. In: *AAAI Fall Symposium: Aurally Informed Performance* (2006)
34. Kim, J., Oh, H., Kim, S., Tong, H., Lee, S.: A brand new dance partner: Music-conditioned pluralistic dancing controlled by multiple dance genres. In: *CVPR* (2022)
35. Kim, T.h., Park, S.I., Shin, S.Y.: Rhythmic-motion synthesis based on motion-beat analysis. *ACM Transactions on Graphics (TOG)* (2003)
36. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: *ICLR* (2014)
37. Kolotouros, N., Alldieck, T., Zanfir, A., Bazavan, E., Fieraru, M., Sminchisescu, C.: Dreamhuman: Animatable 3d avatars from text. In: *NeurIPS* (2024)
38. Lannan, N., Zhou, L., Fan, G.: Human motion enhancement via tobit kalman filter-assisted autoencoder. *IEEE Access* **10**, 29233–29251 (2022)
39. Le, N., Do, T., Do, K., Nguyen, H., Tjiputra, E., Tran, Q.D., Nguyen, A.: Controllable group choreography using contrastive diffusion. *ACM Transactions on Graphics (TOG)* (2023)
40. Le, N., Pham, T., Do, T., Tjiputra, E., Tran, Q.D., Nguyen, A.: Music-driven group choreography. In: *CVPR* (2023)
41. 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)
42. Lee, L., Lin, Z., Hu, R., Gong, Z., Kumar, A., Li, T., Li, S., Hui, P.: When creators meet the metaverse: A survey on computational arts. *CoRR* (2021)

43. Lee, M., Lee, K., Park, J.: Music similarity-based approach to generating dance motion sequence. *Multimedia tools and applications* (2013)
44. Li, B., Zhao, Y., Sheng, L.: Dancenet3d: Music based dance generation with parametric motion transformer. In: *AAAI* (2022)
45. Li, B., Zhao, Y., Zhelun, S., Sheng, L.: Danceformer: Music conditioned 3d dance generation with parametric motion transformer. In: *Proceedings of the AAAI Conference on Artificial Intelligence* (2022)
46. Li, J., Villegas, R., Ceylan, D., Yang, J., Kuang, Z., Li, H., Zhao, Y.: Task-generic hierarchical human motion prior using vaes. In: *International Conference on 3D Vision (3DV)* (2021)
47. Li, J., Kang, D., Pei, W., Zhe, X., Zhang, Y., He, Z., Bao, L.: Audio2gestures: Generating diverse gestures from speech audio with conditional variational autoencoders. In: *ICCV* (2021)
48. Li, R., Zhao, J., Zhang, Y., Su, M., Ren, Z., Zhang, H., Tang, Y., Li, X.: Finedance: A fine-grained choreography dataset for 3d full body dance generation. In: *ICCV* (2023)
49. Li, R., Yang, S., Ross, D.A., Kanazawa, A.: Ai choreographer: Music conditioned 3d dance generation with aist++. In: *ICCV* (2021)
50. Li, W., Dai, B., Zhou, Z., Yao, Q., Wang, B.: Controlling character motions without observable driving source. In: *WACV* (2024)
51. Liang, H., Zhang, W., Li, W., Yu, J., Xu, L.: Intergen: Diffusion-based multi-human motion generation under complex interactions. *arXiv preprint arXiv:2304.05684* (2023)
52. Lin, J., Zeng, A., Lu, S., Cai, Y., Zhang, R., Wang, H., Zhang, L.: Motion-x: A large-scale 3d expressive whole-body human motion dataset. *NeurIPS* (2024)
53. Ling, H.Y., Zinno, F., Cheng, G., Van De Panne, M.: Character controllers using motion vaes. *ACM Transactions on Graphics (TOG)* **39**(4), 40–1 (2020)
54. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics (TOG)* (2015)
55. Martinez, J., Black, M.J., Romero, J.: On human motion prediction using recurrent neural networks. In: *CVPR* (2017)
56. Mo, C.A., Hu, K., Long, C., Wang, Z.: Continuous intermediate token learning with implicit motion manifold for keyframe based motion interpolation. In: *CVPR* (2023)
57. Ofli, F., Erzin, E., Yemez, Y., Tekalp, A.M.: Learn2dance: Learning statistical music-to-dance mappings for choreography synthesis. *IEEE Transactions on Multimedia (TMM)* (2011)
58. Onuma, K., Faloutsos, C., Hodgins, J.K.: Fmdistance: A fast and effective distance function for motion capture data. In: *Eurographics* (2008)
59. Perez, G.V., Beskow, J., Henter, G., Holzapfel, A., Oudeyer, P.Y., Alexanderson, S.: Transflower: probabilistic autoregressive dance generation with multimodal attention. *ACM Transactions on Graphics (TOG)* (2021)
60. Petrovich, M., Black, M.J., Varol, G.: Action-conditioned 3d human motion synthesis with transformer vae. In: *ICCV* (2021)
61. Petrovich, M., Black, M.J., Varol, G.: Temos: Generating diverse human motions from textual descriptions. In: *ECCV* (2022)
62. Qi, Q., Zhuo, L., Zhang, A., Liao, Y., Fang, F., Liu, S., Yan, S.: Diffdance: Cascaded human motion diffusion model for dance generation. In: *Proceedings of the 31st ACM International Conference on Multimedia*. pp. 1374–1382 (2023)
63. Raab, S., Leibovitch, I., Li, P., Aberman, K., Sorkine-Hornung, O., Cohen-Or, D.: Modi: Unconditional motion synthesis from diverse data. In: *CVPR* (2023)

64. Ren, X., Li, H., Huang, Z., Chen, Q.: Self-supervised dance video synthesis conditioned on music. In: ACMMM (2020)
65. Safonova, A., Hodgins, J.K.: Construction and optimal search of interpolated motion graphs. In: SIGGRAPH (2007)
66. Shafir, Y., Tevet, G., Kapon, R., Bermano, A.H.: Human motion diffusion as a generative prior. arXiv (2023)
67. Shi, J.: Application of 3d computer aided system in dance creation and learning. In: International Conference on Machine Learning and Big Data Analytics for IoT Security and Privacy (2021)
68. Shi, M., Starke, S., Ye, Y., Komura, T., Won, J.: Phasemp: Robust 3d pose estimation via phase-conditioned human motion prior. In: ICCV (2023)
69. Shiratori, T., Nakazawa, A., Ikeuchi, K.: Dancing-to-music character animation. *Computer Graphics Forum* (2006)
70. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. In: NeurIPS (2020)
71. Siyao, L., Yu, W., Gu, T., Lin, C., Wang, Q., Qian, C., Loy, C.C., Liu, Z.: Bailando: 3d dance generation by actor-critic gpt with choreographic memory. In: CVPR (2022)
72. Soga, A., Umino, B., Longstaff, J.S.: Automatic composition of ballet sequences using a 3d motion archive. In: 1st South-Eastern European Digitization Initiative Conference (2005)
73. Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. In: NeurIPS (2015)
74. Starke, S., Mason, I., Komura, T.: Deepphase: Periodic autoencoders for learning motion phase manifolds. *ACM Transactions on Graphics (TOG)* (2022)
75. Starke, S., Zhang, H., Komura, T., Saito, J.: Neural state machine for character-scene interactions. *ACM Transactions on Graphics (TOG)* **38**(6), 209–1 (2019)
76. Starke, S., Zhao, Y., Komura, T., Zaman, K.: Local motion phases for learning multi-contact character movements. *ACM Transactions on Graphics (TOG)* **39**(4), 54–1 (2020)
77. Sun, G., Wong, Y., Cheng, Z., Kankanhalli, M.S., Geng, W., Li, X.: Deepdance: Music-to-dance motion choreography with adversarial learning. *IEEE Transactions on Multimedia (TMM)* (2020)
78. Sun, J., Wang, C., Hu, H., Lai, H., Jin, Z., Hu, J.F.: You never stop dancing: Non-freezing dance generation via bank-constrained manifold projection. In: NeurIPS (2022)
79. Tang, T., Jia, J., Mao, H.: Dance with melody: An lstm-autoencoder approach to music-oriented dance synthesis. In: ACMMM (2018)
80. Tevet, G., Gordon, B., Hertz, A., Bermano, A.H., Cohen-Or, D.: Motionclip: Exposing human motion generation to clip space. In: European Conference on Computer Vision (2022)
81. Tevet, G., Raab, S., Gordon, B., Shafir, Y., Cohen-Or, D., Bermano, A.H.: Human motion diffusion model. In: ICLR (2023)
82. Tiwari, G., Antić, D., Lenssen, J.E., Sarafianos, N., Tung, T., Pons-Moll, G.: Pose-ndf: Modeling human pose manifolds with neural distance fields. In: ECCV (2022)
83. Tseng, J., Castellon, R., Liu, C.K.: Edge: Editable dance generation from music. In: CVPR (2023)
84. 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: ISMIR (2019)

85. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017)
86. Wang, Z., Jia, J., Wu, H., Xing, J., Cai, J., Meng, F., Chen, G., Wang, Y.: Group-dancer: Music to multi-people dance synthesis with style collaboration. In: ACM International Conference on Multimedia (2022)
87. Yalta, N., Watanabe, S., Nakadai, K., Ogata, T.: Weakly-supervised deep recurrent neural networks for basic dance step generation. In: IJCNN (2019)
88. Yang, S., Wu, Z., Li, M., Zhang, Z., Hao, L., Bao, W., Zhuang, H.: Qpgesture: Quantization-based and phase-guided motion matching for natural speech-driven gesture generation. In: CVPR (2023)
89. Yang, Z., Wen, Y.H., Chen, S.Y., Liu, X., Gao, Y., Liu, Y.J., Gao, L., Fu, H.: Keyframe control of music-driven 3d dance generation. *IEEE Transactions on Visualization and Computer Graphics* (2023)
90. Yao, S., Sun, M., Li, B., Yang, F., Wang, J., Zhang, R.: Dance with you: The diversity controllable dancer generation via diffusion models. In: Proceedings of the 31st ACM International Conference on Multimedia. pp. 8504–8514 (2023)
91. Ye, Z., Wu, H., Jia, J., Bu, Y., Chen, W., Meng, F., Wang, Y.: Choreonet: Towards music to dance synthesis with choreographic action unit. In: ACMMM (2020)
92. Yin, W., Yin, H., Baraka, K., Kragic, D., Björkman, M.: Dance style transfer with cross-modal transformer. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 5058–5067 (2023)
93. Yu, Z., Cheng, W., Liu, X., Wu, W., Lin, K.Y.: Monohuman: Animatable human neural field from monocular video. In: CVPR (2023)
94. Zhang, X., Yang, S., Xu, Y., Zhang, W., Gao, L.: Mining and applying composition knowledge of dance moves for style-concentrated dance generation. In: AAAI (2023)
95. Zhou, L., Luo, Y.: A spatio-temporal learning for music conditioned dance generation. In: International Conference on Multimodal Interaction (2022)
96. Zhou, Y., Barnes, C., Lu, J., Yang, J., Li, H.: On the continuity of rotation representations in neural networks. In: CVPR (2019)
97. Zhou, Z., Wang, B.: Ude: A unified driving engine for human motion generation. In: CVPR (2023)
98. Zhu, Y., Olszewski, K., Wu, Y., Achlioptas, P., Chai, M., Yan, Y., Tulyakov, S.: Quantized gan for complex music generation from dance videos. In: ECCV (2022)
99. Zhuang, W., Wang, C., Chai, J., Wang, Y., Shao, M., Xia, S.: Music2dance: Dancenet for music-driven dance generation. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* (2022)