REACTDANCE: HIERARCHICAL REPRESENTATION FOR HIGH-FIDELITY AND COHERENT LONG-FORM REACTIVE DANCE GENERATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Reactive dance generation (RDG), the task of generating a dance conditioned on a lead dancer's motion, holds significant promise for enhancing human-robot interaction and immersive digital entertainment. Despite progress in duet synchronization and motion-music alignment, two key challenges remain: generating fine-grained spatial interactions and ensuring long-term temporal coherence. In this work, we introduce **ReactDance**, a diffusion framework that operates on a novel hierarchical latent space to address these spatiotemporal challenges in RDG. First, for fine-grained spatial control and artistic expression, we propose Hierarchical Finite Scalar Quantization (**HFSQ**). This multi-scale motion representation effectively disentangles coarse body posture from subtle limb dynamics, enabling independent and detailed control over both aspects through a layered guidance mechanism. Second, to efficiently generate long sequences with high temporal coherence, we propose Blockwise Local Context (**BLC**), a non-autoregressive sampling strategy. Departing from slow, frame-by-frame generation, BLC partitions the sequence into blocks and synthesizes them in parallel via periodic causal masking and positional encodings. Coherence across these blocks is ensured by a dense sliding-window training approach that enriches the representation with local temporal context. Extensive experiments show that ReactDance substantially outperforms state-of-the-art methods in motion quality, long-term coherence, and sampling efficiency.

1 INTRODUCTION

Reactive Dance Generation (RDG) aims to synthesize lifelike and artistically coherent reactor motions that aligns with a lead dancer and accompanying music. This form of responsive interaction embodies a complex interplay of artistic expression and interpersonal dynamics that is a key challenge in fields like social robotics and human-agent interaction. The ability to generate such intricate, reactive behaviors holds significant potential for applications like virtual avatar animation for gaming and the metaverse (Hu, 2024; Zhang et al., 2023), as well as embodied AI and human-robot interaction (Granados et al., 2017; Chen et al., 2025).

Despite its promise, RDG faces two critical, unaddressed challenges: (1) modeling fine-grained spatial interactions and (2) maintaining long-term temporal coherence. While recent methods have advanced duet synchronization using tailored network architectures (Yao et al., 2023; Liang et al., 2024; Ghosh et al., 2024), physical constraints (Wang et al., 2024; Xu et al., 2024), or reinforcement learning (Siyao et al., 2024), they fail to resolve these core issues and fall short of generating truly authentic dance sequences. Their reliance on holistic, high-level constraints overlooks the subtle yet decisive local movements (*e.g.*, the whip-like 'boleo' in Tango), producing movements that are synchronized yet artistically sterile. Moreover, a fundamental difficulty for RDG is ensuring long-term temporal coherence. Models are typically trained on short clips, a practice driven by both the algorithmic complexity of learning long-term dependencies and practical computational constraints. This creates an inherent discrepancy between training and inference, leading to error accumulation that manifests as temporal drift and an eventual breakdown of synchronization.

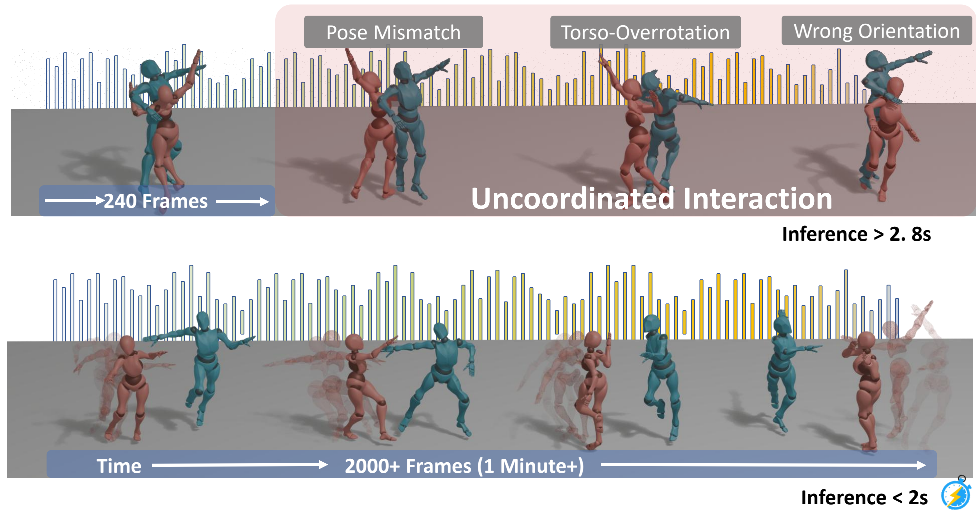

Figure 1: **ReactDance** generates high-fidelity, long-form reactive dance. Conditioned on a leader and music, our model uses a hierarchical representation to refine motion from coarse (translucent) to fine-grained (solid). Its parallel sampling mechanism produces coherent sequences exceeding 2000 frames within 2 seconds, avoiding the quality degradation and slow speed of autoregressive sampling.

To address these limitations, we draw inspiration from structural principles of dance theory and propose two core concepts for generating high-fidelity reactive motion:(1) *Hierarchical Motion Composition:* The principle that dance performance is inherently layered, with foundational full-body rhythms providing a coarse structure and fine-grained limb movements adding detailed semantics (Seham, 2016).(2) *Modular Temporal Coherence:* The observation that long-form choreography achieves consistency by composing shorter, coherent motion phrases and ensuring smooth transitions between them (Krug, 2022).

Grounded in these principles, we propose ReactDance, a two-stage diffusion framework that operationalizes this hierarchical and modular approach.

First, to achieve hierarchical refinement, we introduce a Hierarchical Finite Scalar Quantization (HFSQ) representation. Unlike traditional VQ-VAE models (Gong et al., 2023; Jiang et al., 2024) prone to codebook collapse, HFSQ combines the stability of FSQ (Mentzer et al., 2024) with a hierarchical structure (Yang et al., 2023) to build a robust, multi-scale latent space. This representation explicitly disentangles coarse global movements from fine-grained local details. Building on this, we develop Layer-Decoupled Classifier-Free Guidance (LDCFG), a technique providing independent, layered control over motion semantics at different granularities.

Second, to ensure long-term coherence, we introduce the Blockwise Local Context (BLC) sampling mechanism. BLC replaces fragile autoregressive sampling with efficient, parallel generation of all motion blocks. Long-term consistency is achieved by precisely aligning each block's temporal context at inference with the windows learned during our dense sliding-window (DSW) training. This alignment, enforced using periodic causal masking and positional encodings, effectively eliminates temporal drift. As a result, ReactDance efficiently generates coherent sequences exceeding 60 seconds within 2 seconds.

In summary, our main contributions are:

- We introduce ReactDance, a hierarchical framework for Reactive Dance Generation (RDG) that produces high-fidelity, coherent reactive dance sequences. It leverages a novel Hierarchical Finite Scalar Quantization (HFSQ) representation to progressively model the reactive performance from coarse body rhythms to fine-grained local movements, effectively addressing challenges in spatial interaction refinement.
- We propose Blockwise Local Context (BLC), a novel parallel sampling mechanism that enables coherent generation of long motion sequences within 2 seconds for over 2000 frames, building

on a temporally continuous latent space. By decomposing the generation process into blocks aligned with the temporal context of training, BLC mitigates error accumulation inherent in prior autoregressive methods, achieving state-of-the-art long-term stability and sampling efficiency.

- We propose Layer-Decoupled Classifier-Free Guidance (LDCFG), a novel guidance technique tailored for the HFSQ representation. LDCFG enables independent conditioning over different scales of the generated motion, offering finer-grained semantic manipulation of reactive performances compared to existing approaches.

2 Related Works

2.1 Music-Driven Dance Generation

Music-driven dance generation has advanced from solo to multi-dancer interactions. Early music-driven dance generation relied on sequence-to-sequence models like LSTMs (Tang et al., 2018) and GANs (Li et al., 2022), but often suffered from rigid transitions and a lack of detail. While auto-regressive models (Li et al., 2021; Siyao et al., 2022) improved solo dance quality, they lacked mechanisms for duet interactions. Recent duet-specific models still fall short; approaches like DanY (Yao et al., 2023) and Duolando (Siyao et al., 2024) use single-scale representations, causing synchronization artifacts, while DuetGen's (Ghosh et al., 2025) coupled representation limits motion diversity. In contrast, we introduce a hierarchical representation that captures multi-scale spatial dynamics to generate high-fidelity, coherent reactive duets.

2.2 Hierarchical Motion Generation Models

The quality and controllability of generated motion are critically dependent on its underlying representation. Codebook-based autoregressive models (Gong et al., 2023; Jiang et al., 2024) quantize motion into discrete tokens but are prone to error accumulation and detail loss from VQ-VAE (Oord et al., 2017) quantization. While diffusion models (Qi et al., 2025; Liu et al., 2025) achieve state-of-the-art fidelity, they typically operate on flat representations, lacking a structure for multi-level control. Existing hierarchical models (Qi et al., 2023; Li et al., 2024) have focused on temporal, not spatial, refinement. Our work fills this gap by introducing a spatial hierarchy that enables both high-fidelity synthesis and multi-scale control over the generated motion.

2.3 Long Motion Sequence Generation

Generating coherent long motion sequences remains challenging. For text-to-motion, methods like MotionStreamer (Xiao et al., 2025) and priorMDM (Shafir et al., 2024a) use two-pass refinement to ensure smooth transitions between clips. The problem is harder in music-to-dance, which requires long-range musical coherence (Tsakalidis, 2020). Prior solutions are often inefficient, such as the iterative inpainting in EDGE (Tseng et al., 2023), or inflexible, like the primitive-based temporal hierarchy in Lodge (Li et al., 2024). In contrast, our approach enables efficient, parallel generation by training on overlapping windows, directly embedding temporal context into our representation and avoiding the costly overhead of previous methods.

3 Method

Our proposed framework, ReactDance, generates a coherent long reactive dance sequence in a two-stage process. First, an autoencoder with a Hierarchical Finite Scalar Quantizer (HFSQ) bottleneck learns to encode dance movements into a hierarchical latent representation. Second, a conditional diffusion model learns to generate HFSQ latent representations based on leader movements and background music, which are then decoded into the final motion.

3.1 Motion and Music Representation

Motion Representation.

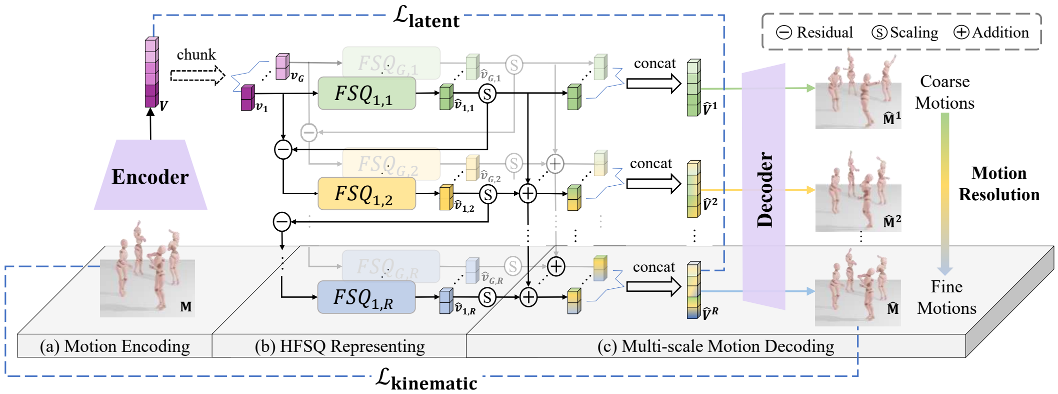

Figure 2: **Motion HFSQ overview.** The proposed motion HFSQ learns to progressively encode motion sequence into a hierarchical representation $\mathcal{V} = \{\hat{\boldsymbol{v}}_{g,r}\}_{g=1,r=1}^{G,R}$ and reconstructs motions via a grouped residual architecture. Building on FSQ, HFSQ eliminates codebook collapse while enhancing multi-scale motion expressiveness through grouped residual quantization, which integrates coarse-to-fine motion semantics.

We represent a dance sequence of N frames using a skeleton with $J = 22$ joints from the SMPL model (Loper et al., 2023). To better capture the dynamics of reactive dance, our approach utilizes an asymmetric representation for the leader and the reactor.

The leader's motion is modeled holistically as a single sequence of global joint positions, denoted as $\mathbf{M}_L \in \mathbb{R}^{N \times J \times 3}$. In contrast, the reactor's motion is decomposed into three independent components: (1) Upper-body motion ($\mathbf{M}_{Rup} \in \mathbb{R}^{N \times J_{up} \times 3}$): The localized positions of the J_{up} upper-body joints relative to the reactor's own root joint. (2) Lower-body motion ($\mathbf{M}_{Rdown} \in \mathbb{R}^{N \times J_{down} \times 3}$): The localized positions of the J_{down} lower-body joints relative to the reactor's own root joint. (3) Relative root distance ($\mathbf{M}_{tr} \in \mathbb{R}^{N \times 3}$): The displacement vector between the leader's and reactor's roots.

Music Representation. For the accompanying music, we use Librosa (Brian McFee et al., 2015) to extract a rich acoustic feature vector $\boldsymbol{c} \in \mathbb{R}^{54}$ for each corresponding motion frame. This vector comprises 20-dimensional MFCCs, their 20-dimensional deltas, 12-dimensional chroma features, and 1-dimensional onset strength and beat tracking signals.

Our goal is thus to generate the set of reactor motion components $\mathbf{M}_R = \{\mathbf{M}_{Rup}, \mathbf{M}_{Rdown}, \mathbf{M}_{tr}\}$ conditioned on the music $\boldsymbol{c}$ and the leader's motion $\mathbf{M}_L$. To process the complete reactor motion $\mathbf{M}_R$, each of its components is passed through a dedicated and structurally identical HFSQ stream. For clarity, the following sections will describe the core methodology for a single holistic stream.

3.2 Hierarchical Motion Representation via HFSQ

We learn a compressed and structured representation of motion using a hierarchical finite scalar quantizer, inspired by recent work in neural audio codecs (Yang et al., 2023; Mentzer et al., 2024).

Encoder and Quantizer. A temporal encoder, consisting of 1D convolutional layers, maps an input reactor motion sequence $\mathbf{M}_R \in \mathbb{R}^{N \times J_R}$ to a downsampled feature sequence $\boldsymbol{v} \in \mathbb{R}^{n \times d}$. This sequence is then processed by our HFSQ module, which operates in two stages:

- **Grouping:** The feature vector $\boldsymbol{v}$ is split into G parallel groups, $\boldsymbol{v} = [\boldsymbol{v}_1, \ldots, \boldsymbol{v}_G]$, where each group $\boldsymbol{v}_g \in \mathbb{R}^{n \times (d/G)}$.

- **Cascaded Residual Quantization:** Each group $\boldsymbol{v}_g$ is quantized sequentially over R residual stages. Let the initial residual be $\boldsymbol{e}_{g,1} = \boldsymbol{v}_g$. For each stage $r \in [1, \ldots, R]$, we compute:

$$\boldsymbol{z}_{g,r} = \mathrm{FSQ}_r(\boldsymbol{e}_{g,r}); \quad \hat{\boldsymbol{v}}_{g,r} = \mathrm{Dequantize}(\boldsymbol{z}_{g,r}); \quad \boldsymbol{e}_{g,r+1} = \boldsymbol{e}_{g,r} - s_r \cdot \hat{\boldsymbol{v}}_{g,r}. \tag{1}$$

The core output is the set of continuous reconstructions $\hat{\boldsymbol{v}}_{g,r}$ from each hierarchical level. We define the complete hierarchical continuous representation as the set of these vectors, denoted by $\mathcal{V}$:

$$\mathcal{V} = \{\hat{\boldsymbol{v}}_{g,r}\}_{g=1,r=1}^{G,R}. \tag{2}$$

As we will detail, this structured set $\mathcal{V}$ serves as the target space for our diffusion model.

Decoder and Training Objective. The decoder reconstructs the motion from the dequantized features. The final feature vector $\hat{\boldsymbol{v}}$ is obtained by summing the reconstructions within each group and concatenating them: $\hat{\boldsymbol{v}}_g = \sum_{r=1}^{R} s_r \cdot \hat{\boldsymbol{v}}_{g,r}$ and $\hat{\boldsymbol{v}} = [\hat{\boldsymbol{v}}_1, \ldots, \hat{\boldsymbol{v}}_G]$. The HFSQ autoencoder is trained to minimize a combination of a latent loss and physical plausibility losses on the final decoded motion $\hat{\mathbf{M}}_R = \mathrm{Dec}(\hat{\boldsymbol{v}})$:

$$\mathcal{L}_{\text{HFSQ}} = \lambda_{kin}\mathcal{L}_{\text{kinematic}} + \lambda_{lat}\mathcal{L}_{latent}, \quad \text{where} \tag{3}$$

$$\mathcal{L}_{\text{kinematic}} = \sum_{k\in\{\text{pos,vel,acc}\}} \lambda_{Pk} \cdot \|\mathbf{M}_R^{(k)} - \hat{\mathbf{M}}_R^{(k)}\|_1, \tag{4}$$

$$\mathcal{L}_{latent} = \|\boldsymbol{v} - \mathrm{sg}(\hat{\boldsymbol{v}})\|_2^2 + \beta\|\,\mathrm{sg}(\boldsymbol{v}) - \hat{\boldsymbol{v}}\|_2^2. \tag{5}$$

The $\mathbf{M}^{(k)}$ denotes the k-th order temporal derivative (position, velocity, acceleration) of the ground-truth motion, and $\hat{\mathbf{M}}^{(k)}$ denotes the reconstructed counterparts. This objective ensures that the final reconstruction is both physically realistic and faithful to the encoder's output.

Latent Robustness via Progressive Masking.

To construct a structured HFSQ latent space, we corrupt the hierarchical representation $\mathcal{V}$ (Eq. 2) using two complementary techniques. First, residual masking promotes inter-layer independence by stochastically occluding higher-level FSQ layers while perturbing the base layer with scaled Gaussian noise. Second, code masking improves the decoder's robustness by randomly masking dimensions within individual FSQ codes. Together, these operations regularize the latent space, improving the decoder's tolerance to imperfect inputs and thereby enhancing the final motion quality.

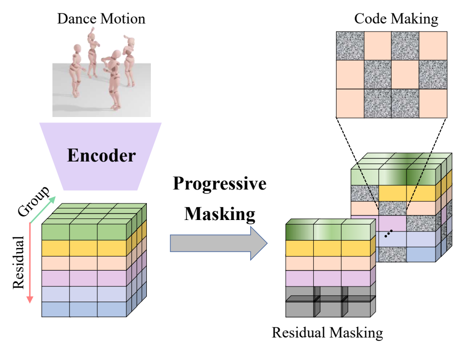

Figure 3: **Progressive masking overview.**

3.3 Latent Diffusion on Hierarchical Features

Generative Target. Unlike methods that diffuse on discrete codes or the final aggregated feature, our diffusion model operates on the hierarchical continuous representation $\mathcal{V}$ (Eq. 2). This strategy enables the model to learn the coarse-to-fine structure of motion within the continuous space defined by our HFSQ module. The final motion is then robustly decoded from this generated representation.

Diffusion Framework. We model the generation process using a DDPM (Ho et al., 2020). The forward process gradually adds Gaussian noise to the HFSQ latents $\boldsymbol{x}_0 = \mathcal{V}$ over T timesteps. The reverse process trains a Transformer-based network $\mathcal{G}_\theta$ to denoise the corrupted latents $\boldsymbol{x}_t$ at timestep t, conditioned on music c and leader motion $\mathbf{M}_L$. To leverage the hierarchical structure of $\mathcal{V}$, we apply independent weights to the loss for each residual scale. The latent diffusion loss is:

$$\mathcal{L}_{\text{simple}} = \sum_{r=1}^{R} \lambda^r \cdot \mathbb{E}_{\boldsymbol{x}_0,t,\epsilon}\left[\|\mathcal{V}^r - \mathcal{G}_\theta(\boldsymbol{x}_t, t, \boldsymbol{c}, \mathbf{M}_L)^r\|_2^2\right], \tag{6}$$

where $\mathcal{V}^r = \{\hat{\boldsymbol{v}}_{g,r}\}_{g=1}^{G}$ represents the set of feature vectors at the r-th level, and $\mathcal{G}_\theta(\cdot)^r$ is the corresponding portion of the network's output. The weights λ^r balance the contribution of each residual level, allowing the model to focus on coarse-to-fine details during generation. Moreover, we apply kinematic loss $\mathcal{L}_{kinematic}$ (Eq. 4), foot contact loss $\mathcal{L}_{fc}$, leader-reactor relative orientation loss $\mathcal{L}_{ro}$ (Liang et al., 2024) to enforce physical plausibility and interaction coherence. The complete loss is defined in Eq. 7, with λ_{kin}, λ_{fc} and λ_{ro} as balancing weights:

$$\mathcal{L}_{\text{diffusion}} = \mathcal{L}_{\text{simple}} + \lambda_{\text{kin}}\mathcal{L}_{\text{kinematic}} + \lambda_{\text{fc}}\mathcal{L}_{\text{fc}} + \lambda_{\text{ro}}\mathcal{L}_{\text{ro}}. \tag{7}$$

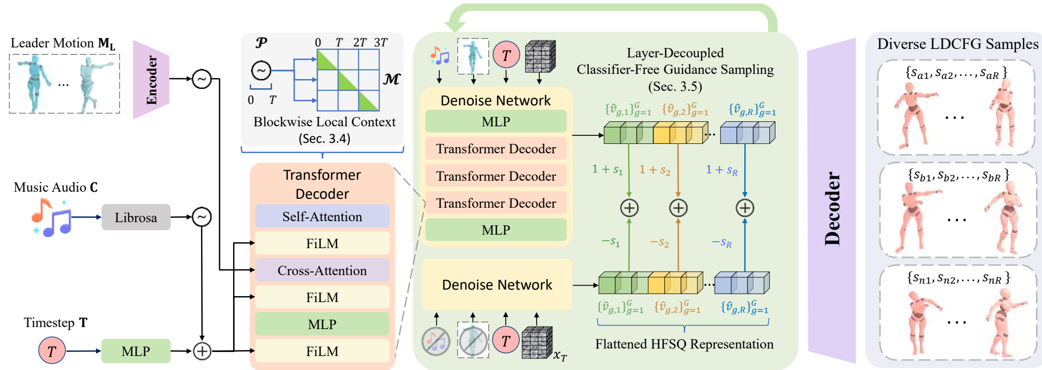

Figure 4: **ReactDance Pipeline Overview.** Our ReactDance generates long, high-fidelity reactive dance sequences conditioned on leader motion and music. The core is a diffusion model that learns to denoise hierarchical HFSQ latents. Leader motion is injected via cross-attention, while music features are fused using a FiLM layer. For coherent generation of long sequences, our Blockwise Local Context (BLC) sampling strategy partitions the timeline into parallel blocks with aligned temporal contexts. Within each denoising step, Layer-Decoupled Classifier-Free Guidance (LDCFG) provides fine-grained control by applying independent guidance weights to each HFSQ scale. Finally, the denoised latents are decoded into a dance sequence coherently aligned with the input conditions.

3.4 Efficient Long Sequence Generation

Generating long, coherent dance sequences is a critical challenge, as the desired inference length K often far exceeds the model's training window W. A common approach is to apply diffusion models autoregressively—sequentially generating and stitching together short clips. However, this method has significant drawbacks: (1) slow inference due to its serial nature, (2) error accumulation, where early mistakes degrade the quality of later parts, and (3) blending artifacts at the overlaps between clips. To overcome these issues, we introduce the Blockwise Local Context (BLC), a non-autoregressive sampling strategy that generates all parts of a long sequence in parallel while ensuring both local and global coherence. BLC's design rests on two principles: intra-block consistency and inter-block continuity.

Intra-block consistency ensures that each block is generated with high fidelity. We achieve this by aligning the temporal context of each inference block with that of the training window, using periodic applications of causal masking and positional encoding. Specifically, we apply a periodic causal attention mask $\mathcal{M}$ across the entire sequence of length K. This mask is a block-diagonal matrix(Fig. 4) that enforces causality strictly within each non-overlapping block of size W. This structure prevents information from leaking between blocks, thereby halting error propagation. Complementing the mask, we introduce a phase-aligned periodic positional encoding $\mathcal{P}_i$ that resets temporal phase for each block:

$$\mathcal{P}_i = \sin\left(\frac{\pi(i \bmod W)}{W}\right) \oplus \cos\left(\frac{\pi(i \bmod W)}{W}\right), \tag{8}$$

where i is the global temporal position. This strategy makes every block's positional context appear identical to the fixed-length window the model saw during training, ensuring it generalizes effectively regardless of a block's actual position in the long sequence.

Inter-block continuity ensures seamless transitions between these parallel blocks. This is achieved by training the entire framework with a dense sliding window approach, with $m \ll T$ where m is the sliding stride. By exposing the model to a vast number of overlapping subsequences during training, it implicitly learns how to naturally begin and end motion phrases from any temporal point. During inference, this learned behavior ensures that the overlapping regions between blocks are generated with smooth, natural transitions, creating a coherent long-form performance without costly sequential generation or explicit blending.

3.5 Layer-Decoupled Classifier-Free Guidance

Classifier-free guidance (CFG) is a powerful technique for controlling the trade-off between fidelity and diversity in generative models (Ho & Salimans, 2021). However, in motion generation, conventional CFG applies a single guidance weight across the entire motion representation (Tevet et al., 2023; Shafir et al., 2024b). This is a significant limitation for complex, hierarchical tasks like reactive dance, as it cannot distinguish between different motion semantics encoded at different scales (*e.g.*, overall body posture vs. fine-grained limb articulation).

To address this, we propose Layer-Decoupled Classifier-Free Guidance (LDCFG), a method that enables fine-grained, multi-scale control over the generation process. As shown in Fig. 4, LDCFG assigns an independent guidance strength to each scale of our HFSQ representation:

$$\hat{\boldsymbol{x}}_0^r = (1 + s_r)\mathcal{G}_\theta(\boldsymbol{x}_t^r, t, \boldsymbol{c}, \mathbf{M}_L) - s_r\mathcal{G}_\theta(\boldsymbol{x}_t^r, t, \emptyset, \emptyset), \tag{9}$$

where s_r is the guidance strength for the r-th residual scale.

4 Experiment

4.1 Datasets and Experimental Setting

Dataset. We conduct experiments on the DD100 dataset (Siyao et al., 2024), which comprises 1.95 hours of paired dance motion and music data spanning 10 musical genres. We adopt the original dataset's training/test splits. All motion sequences are split into 8-second clips (30 frames per second) with a sliding window stride of 4.

Evaluation Metrics. We adopt a comprehensive set of quantitative metrics to evaluate the generated motion from four key perspectives: (1) Motion Quality. To assess the realism of individual motions, we measure distributional similarity and reconstruction accuracy. Distributional metrics include Fréchet Inception Distance (FID) and Diversity (Div) on both kinematic ($\text{FID}_k, \text{Div}_k$) (Onuma et al., 2008) and graphical ($\text{FID}_g, \text{Div}_g$) (Müller et al., 2005) features. Reconstruction accuracy against the ground truth is measured by Mean Per-Joint Position Error (MPJPE) for static poses and Mean Per-Joint Velocity Error (MPJVE) for dynamics (Ghosh et al., 2024). (2) Interaction Quality. To evaluate leader-reactor coordination, we compute FID and Diversity on cross-distance features ($\text{FID}_{cd}, \text{Div}_{cd}$), which describe the evolving spatial relationships between the dancers' key joints. (3) Music-Dance Alignment. We use two metrics: the Beat Align Score (BAS) (Siyao et al., 2022) to measure synchronization between motion peaks and music beats, and the Beat Echo Degree (BED) (Siyao et al., 2024) to quantify rhythmic consistency between the dancers. (4) Generation Efficiency. We report the Average Inference Time per Sequence (AITS) in seconds.

4.2 Comparison with Baseline Methods

Baselines. We compare our model against several state-of-the-art (SOTA) methods. For a fair comparison, all baselines were adapted for the reactive dance generation (RDG) task and retrained from scratch on the DD100 dataset using the SMPL body model. We integrated the leader's motion as a direct condition for several models: the text-driven GestureLSM (Liu et al., 2025) was modified by replacing its text encoder with a leader motion encoder, while for the music-driven EDGE (Tseng et al., 2023), we added the features from a leader motion encoder to its existing music features. Other baselines required more targeted modifications. The text-to-interaction model InterGen (Liang et al., 2024) was adapted by replacing its text encoder with a music encoder and changing its weight-shared network to be independent for each dancer. For the group dance model TCDiff (Dai et al., 2025), we altered its navigator module to predict the reactor's trajectory from the leader's motion. Finally, Duolando (Siyao et al., 2024), an existing RDG model, was retrained on our dataset to ensure body model consistency (SMPL vs. its native SMPL-X).

Analysis. As reported in Table 1, our method, ReactDance, demonstrates superior performance across the majority of metrics. Note that these results are computed on the entire test set, which features long sequences with an average length of 2066 frames. This highlights our model's strong generalization capability and consistency in long-term generation. The significant lead in realism ($\text{FID}_k, \text{FID}_g$) and accuracy (MPJPE, MPJVE) stems from our effective HFSQ representation, which enables the diffusion model to learn motion hierarchically from coarse postural dynamics to fine-grained

Table 1: Comparison with state-of-the-art methods on the test dataset of DD100 dataset. Symbols ↑, ↓, and → indicate the higher, lower and closer to Ground Truth are better. **Bold** and <u>underline</u> indicate the best and second best results. A dotted line separates methods adopted from single-person motion generation (below) and multi-person counterparts (above).

	Solo Metrics							Interactive Metrics			
Method	FID$_k$(↓)	FID$_g$(↓)	Div$_k$(→)	Div$_g$(→)	MPJPE(↓)	MPJVE(↓)	BAS(→)	FID$_{cd}$(↓)	Div$_{cd}$(→)	BED(→)	AITS(↓)
Ground Truth	-	-	10.86	7.82	-	-	0.1791	-	12.53	0.5308	-
GestureLSM	<u>14.65</u>	<u>34.23</u>	12.20	9.00	<u>171.37</u>	19.01	**0.1734**	40.09	9.45	0.1903	3.13
EDGE	68.68	113.25	5.62	10.05	235.52	20.76	<u>0.1854</u>	1523.27	**12.37**	0.1904	<u>2.91</u>
TCDiff	105.97	251.35	14.23	24.54	182.64	22.91	0.2270	1472.00	<u>11.31</u>	0.2304	3.34
InterGen	35.89	47.28	12.24	9.13	210.66	18.69	0.1634	176.19	18.10	0.2746	5.22
Duolando	27.68	35.01	<u>10.95</u>	<u>8.70</u>	174.54	<u>18.72</u>	0.2086	<u>17.49</u>	14.73	<u>0.3285</u>	4.41
ReactDance	**5.57**	**7.63**	**10.82**	**7.76**	**132.99**	**15.68**	0.2031	**14.17**	10.58	**0.3863**	**1.75**

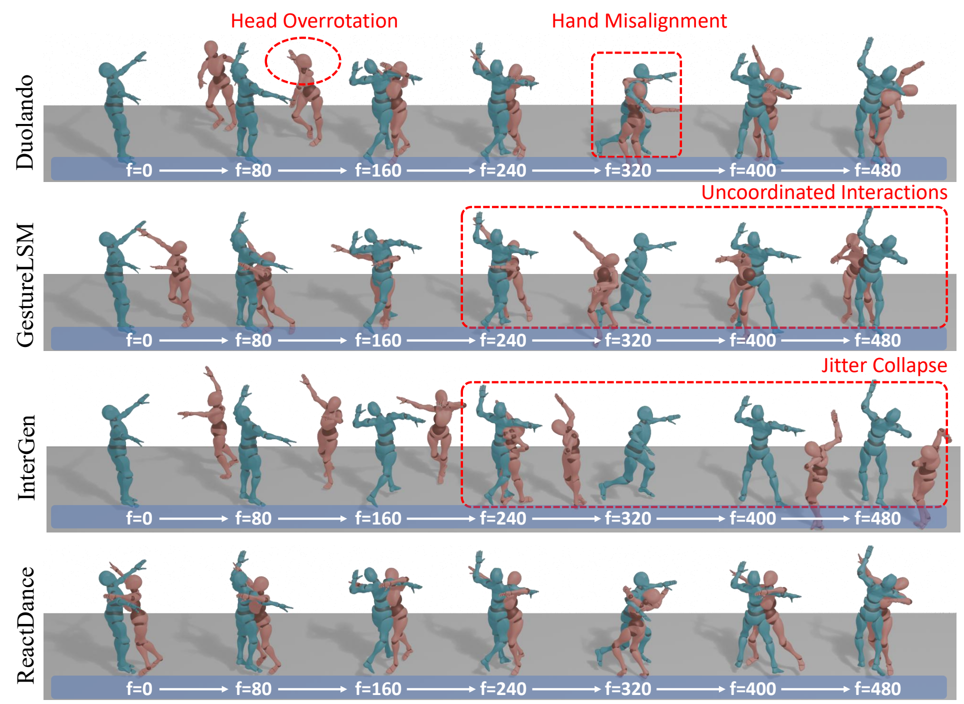

Figure 5: **Qualitative comparison of reactive dance generation.** Given the same leader motion, Duolando produces unnatural head rotations and GestureLSM shows uncoordinated interactions. InterGen's motion collapses into unrealistic jitter when generalizing beyond its training horizon. In contrast, our model generates a fluid and coherent reactive motion.

articulations. This high-fidelity representation directly enhances duet coherence, where ReactDance significantly outperforms all baselines (*e.g.*, achieving a FID$_{cd}$ of 14.17 and a BED of 0.3863). Furthermore, ReactDance also achieves the most efficient generation with an AITS of 1.75s due to the parallel generation capability of our BLC sampling strategy. In contrast, competing methods exhibit key limitations. Architectures with a two-stage design like GestureLSM and Duolando, while competitive in solo motion quality, are constrained by their representations. GestureLSM's use of a single-scale latent space for generation limits its ability to capture hierarchical interactions (evidenced by a high FID$_{cd}$ of 40.09), while Duolando's VQ-VAE is less effective at fine-grained reconstruction. Other models like EDGE and TCDiff perform poorly overall, as they learn directly on the high-frequency raw motion space and lack explicit mechanisms for interaction modeling.

User Study. We conducted a user study to evaluate the perceptual quality of long-form (> 60 seconds) reactive dance sequences, involving 20 participants. Each participant completed 15 paired comparisons, rating videos generated by our method against baseline methods on a 5-point Likert scale across three criteria: Motion Naturalness, Music Alignment, and Interaction Coordination. As detailed in the appendix (Fig. 9), our method consistently outperformed all baselines across each criterion, demonstrating a strong and consistent participant preference. These results highlight our method's capability to produce high-quality, coherent interactions over extended durations.

4.3 Ablation Study

Effects of HFSQ and Progressive Masking. We ablate our Hierarchical Finite Scalar Quantizer (HFSQ) and Progressive Masking (PM) strategy in Table 2. First, replacing HFSQ with a comparable RVQ-VAE significantly degrades performance (FID_g increases from 7.63 to 36.73; MPJPE increases from 132.99 to 139.42), underscoring the superior motion representation capability of HFSQ. Second, removing the PM strategy trades a marginal MPJPE improvement for a substantial drop in realism (FID_g increases to 10.46), demonstrating that PM is crucial for synthesizing high-fidelity motion. Additional results are available in the supplementary materials.

Table 2: Ablation study of HFSQ and Progressive Masking (PM).

	Solo Metrics			
Method	$FID_g(\downarrow)$	$Div_g(\rightarrow)$	MPJPE($\downarrow$)	BAS($\rightarrow$)
Ground Truth	–	7.82	–	0.1791
w/o HFSQ	36.73	8.82	139.42	0.2040
w/o PM	10.46	7.65	**132.19**	**0.2003**
ReactDance (full)	**7.63**	**7.76**	132.99	0.2031

Effects of Blockwise Local Context Sampling. We ablate the two core components of our Blockwise Local Context (BLC) sampling to validate its contributions to duet coherence and sampling efficiency, with results presented in Table 3. We analyze the dense sliding window component by increasing the training stride from 4 (our full model) to 16 and 64. We observe a significant degradation in coherence, as evidenced by the substantial increase in FID_{cd} and the decline in BED. Moreover, we replace our periodic positional encoding (PPE) and causal attention masking (PCAM) with the latent stitching operation (Tseng et al., 2023). This change harms both sampling efficiency (AITS increases from 1.75 to 2.03) and generation quality (FID_{cd} increases from 14.17 to 18.59). These results confirm a dense training stride is crucial for coherence, and our periodic sampling mechanism is key to achieving both high quality and efficiency.

Table 3: Ablation of Dense Sliding Window (DSW) training and periodic sampling strategies (PPE & PCAM).

	Duet Metrics			
Method	$FID_{cd}(\downarrow)$	$Div_{cd}(\rightarrow)$	BED($\rightarrow$)	AITS($\downarrow$)
Ground Truth	–	12.53	0.5308	–
s=16	22.69	10.06	0.3065	–
s=64	39.50	12.08	0.2840	–
w/o PPE, PCAM	18.59	11.17	0.3218	2.03
ReactDance (full)	**14.17**	10.58	**0.3863**	**1.75**

Effect of Layer-Decoupled Guidance. Our diffusion model operates on hierarchical HFSQ latents, enabling us to apply independent classifier-free guidance weights s_r of Eq. (9) to each motion scale. This provides fine-grained control over the fidelity-diversity trade-off across different semantic levels. For instance, a user can enforce strong postural alignment with high guidance on coarse scales while encouraging creative limb articulation with lower guidance on fine scales. A detailed analysis demonstrating this flexibility is available in the supplementary materials.

5 Conclusion

We present ReactDance, a diffusion framework for generating high-fidelity and coherent long-form reactive dance. Our method is built upon a hierarchical latent representation that captures multi-scale spatiotemporal dynamics. This representation further enables two key contributions: an efficient parallel sampling strategy for long sequence generation and a multi-level conditioning mechanism that offers fine-grained artistic control. ReactDance offers a comprehensive system for generating dynamic body motion under complex conditions, providing a versatile tool for music-driven character animation and digital choreography.

Limitation. Despite the effectiveness of our HFSQ representation, it lacks explicit semantic meaning across scales, which reduces control interpretability. Additionally, finger motions are omitted mainly due to the noisy hand data of DD100 dataset, which are crucial for nuanced gestures. Future work will focus on imbuing our representation with semantic meaning for more intuitive control and incorporating detailed hand articulations to broaden the model's expressive range.

References

Brian McFee, Colin Raffel, Dawen Liang, Daniel P.W. Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. Librosa: Audio and Music Signal Analysis in Python. In *Proceedings of the 14th Python in Science Conference*, pp. 18–24, 2015.

Haoran Chen, Yiteng Xu, Yiming Ren, Yaoqin Ye, Xinran Li, Ning Ding, Peishan Cong, Ziyi Wang, Bushi Liu, Yuhan Chen, et al. Symbiosim: Human-in-the-loop simulation platform for bidirectional continuing learning in human-robot interaction. *arXiv preprint arXiv:2502.07358*, 2025.

Xiangning Chen, Chen Liang, Da Huang, Esteban Real, Kaiyuan Wang, Hieu Pham, Xuanyi Dong, Thang Luong, Cho-Jui Hsieh, Yifeng Lu, et al. Symbolic discovery of optimization algorithms. *Advances in neural information processing systems*, 36:49205–49233, 2023.

Yuqin Dai, Wanlu Zhu, Ronghui Li, Zeping Ren, Xiangzheng Zhou, Jixuan Ying, Jun Li, and Jian Yang. Harmonious music-driven group choreography with trajectory-controllable diffusion. *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 2645–2653, Apr. 2025.

Anindita Ghosh, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Philipp Slusallek. Remos: 3d motion-conditioned reaction synthesis for two-person interactions. In *European Conference on Computer Vision*, pp. 418–437. Springer, 2024.

Anindita Ghosh, Bing Zhou, Rishabh Dabral, Jian Wang, Vladislav Golyanik, Christian Theobalt, Philipp Slusallek, and Chuan Guo. Duetgen: Music driven two-person dance generation via hierarchical masked modeling. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pp. 1–11, 2025.

Kehong Gong, Dongze Lian, Heng Chang, Chuan Guo, Zihang Jiang, Xinxin Zuo, Michael Bi Mi, and Xinchao Wang. TM2D: Bimodality driven 3d dance generation via music-text integration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9908–9918. IEEE, 2023.

Diego Felipe Paez Granados, Breno A Yamamoto, Hiroko Kamide, Jun Kinugawa, and Kazuhiro Kosuge. Dance teaching by a robot: Combining cognitive and physical human–robot interaction for supporting the skill learning process. *IEEE Robotics and Automation Letters*, 2(3):1452–1459, 2017.

Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, pp. 12. Curran Associates Inc., 2020.

Li Hu. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8153–8163, June 2024.

Biao Jiang, Xin Chen, Wen Liu, Jingyi Yu, Gang Yu, and Tao Chen. Motiongpt: Human motion as a foreign language. *Advances in Neural Information Processing Systems*, 36, 2024.

Maximilian Krug. Temporal procedures of mutual alignment and synchronization in collaborative meaning-making activities in a dance rehearsal. *Frontiers in Communication*, 7:957894, 2022.

Buyu Li, Yongchi Zhao, Shi Zhelun, and Lu Sheng. Danceformer: Music conditioned 3d dance generation with parametric motion transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 1272–1279. AAAI, 2022.

Ronghui Li, Yuxiang Zhang, Yachao Zhang, Hongwen Zhang, Jie Guo, Yan Zhang, Yebin Liu, and Xiu Li. Lodge: A coarse to fine diffusion network for long dance generation guided by the characteristic dance primitives. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1524–1534. IEEE, 2024.

Ruilong Li, Shan Yang, David A. Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13381–13392. IEEE, 2021.

Han Liang, Wenqian Zhang, Wenxuan Li, Jingyi Yu, and Lan Xu. Intergen: Diffusion-based multi-human motion generation under complex interactions. *International Journal of Computer Vision*, pp. 1–21, 2024.

Pinxin Liu, Luchuan Song, Junhua Huang, and Chenliang Xu. GestureLSM: Latent Shortcut based Co-Speech Gesture Generation with Spatial-Temporal Modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025.

Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pp. 851–866. 2023.

Fabian Mentzer, David Minnen, Eirikur Agustsson, and Michael Tschannen. Finite scalar quantization: VQ-VAE made simple. In *Proceedings of the International Conference on Learning Representations*. OpenReview.net, 2024.

Meinard Müller, Tido Röder, and Michael Clausen. Efficient content-based retrieval of motion capture data. *ACM Trans. Graph.*, 24(3):677–685, July 2005.

Kensuke Onuma, Christos Faloutsos, and Jessica K. Hodgins. Fmdistance: A fast and effective distance function for motion capture data. In Katerina Mania and Eric R. Reinhard (eds.), *Eurographics*, pp. 83–86. Eurographics Association, 2008.

Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Adv. Neural Inform. Process. Syst.*, 30, 2017.

Qiaosong Qi, Le Zhuo, Aixi Zhang, Yue Liao, Fei Fang, Si Liu, and Shuicheng Yan. Diffdance: Cascaded human motion diffusion model for dance generation. In *Proceedings of the 31st ACM International Conference on Multimedia*, pp. 1374–1382. ACM, 2023.

Xingqun Qi, Yatian Wang, Hengyuan Zhang, Jiahao Pan, Wei Xue, Shanghang Zhang, Wenhan Luo, Qifeng Liu, and Yike Guo. Co3gesture: Towards coherent concurrent co-speech 3d gesture generation with interactive diffusion. In *ICLR*, 2025.

Jenny Seham. Public scholarship in dance: Teaching, choreography, research, service, and assessment for community engagement. *Journal of Dance Education*, 16(4):159–159, 2016.

Yoni Shafir, Guy Tevet, Roy Kapon, and Amit Haim Bermano. Human motion diffusion as a generative prior. In *The Twelfth International Conference on Learning Representations*, 2024a.

Yoni Shafir, Guy Tevet, Roy Kapon, and Amit Haim Bermano. Human motion diffusion as a generative prior. In *The Twelfth International Conference on Learning Representations*. OpenReview.net, 2024b.

Li Siyao, Weijiang Yu, Tianpei Gu, Chunze Lin, Quan Wang, Chen Qian, Chen Change Loy, and Ziwei Liu. Bailando: 3D dance generation by actor-critic GPT with choreographic memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11040–11049. IEEE, 2022.

Li Siyao, Tianpei Gu, Zhitao Yang, Zhengyu Lin, Ziwei Liu, Henghui Ding, Lei Yang, and Chen Change Loy. Duolando: Follower GPT with off-policy reinforcement learning for dance accompaniment. In *Proceedings of the International Conference on Learning Representations*. OpenReview.net, 2024.

Taoran Tang, Jia Jia, and Hanyang Mao. Dance with melody: An lstm-autoencoder approach to music-oriented dance synthesis. In *Proceedings of the 26th ACM international conference on Multimedia*, pp. 1598–1606. ACM, 2018.

Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*. OpenReview.net, 2023.

Konstantin Tsakalidis. *Choreography: craft and vision: developing and structuring dance for Solo, Duet and groups*. Stage Verlag, 2020.

Jonathan Tseng, Rodrigo Castellon, and C. Karen Liu. EDGE: Editable dance generation from music. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 448–458. IEEE, 2023.

Zhenzhi Wang, Jingbo Wang, Yixuan Li, Dahua Lin, and Bo Dai. Intercontrol: Zero-shot human interaction generation by controlling every joint. In *NeurIPS*, 2024.

Lixing Xiao, Shunlin Lu, Huaijin Pi, Ke Fan, Liang Pan, Yueer Zhou, Ziyong Feng, Xiaowei Zhou, Sida Peng, and Jingbo Wang. Motionstreamer: Streaming motion generation via diffusion-based autoregressive model in causal latent space. 2025.

Liang Xu, Yizhou Zhou, Yichao Yan, Xin Jin, Wenhan Zhu, Fengyun Rao, Xiaokang Yang, and Wenjun Zeng. Regennet: Towards human action-reaction synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1759–1769. IEEE, 2024.

Dongchao Yang, Songxiang Liu, Rongjie Huang, Jinchuan Tian, Chao Weng, and Yuexian Zou. Hifi-codec: Group-residual vector quantization for high fidelity audio codec. *arXiv preprint arXiv:2305.02765*, 2023.

Siyue Yao, Mingjie Sun, Bingliang Li, Fengyu Yang, Junle Wang, and Ruimao Zhang. Dance with you: The diversity controllable dancer generation via diffusion models. In *Proceedings of the ACM International Conference on Multimedia*, pp. 8504–8514. ACM, 2023.

Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 3836–3847, 2023.

6 APPENDIX

In the appendix, we present:

- Section 6.1: The Use of Large Language Models (LLMs).
- Section 6.2: Training Hyper-parameters.
- Section 6.3: More Results of Ablation on HFSQ Representation and Progressive Masking.
- Section 6.4: More Results of Layer-Decoupled Guidance.
- Section 6.5: User Study Results.

6.1 THE USE OF LARGE LANGUAGE MODELS

In the preparation of this manuscript and its accompanying code, we utilized Large Language Models (LLMs) as assistive tools. We wish to clarify that the core scientific contributions—including the problem formulation, the design of our proposed method, and the experimental analysis—are entirely the work of the authors. The use of LLMs was confined to the following supporting tasks:

- **Manuscript Preparation:** LLMs were employed for language refinement. This included improving grammatical correctness, rephrasing sentences for better clarity and flow, and ensuring stylistic consistency throughout the paper. All final text was carefully reviewed and edited by the authors to ensure it accurately reflects our research and contributions.
- **Software Development:** LLMs assisted in writing auxiliary and boilerplate code. Specific use cases included generating scripts for data preprocessing and visualization, creating utility functions, and debugging isolated code segments. The core algorithmic implementation of our proposed method was developed exclusively by the authors.

In all instances, the LLM served as a productivity and quality-enhancement tool. The authors maintain full intellectual responsibility and ownership for all conceptual and experimental aspects of this work.

6.2 TRAINING HYPER-PARAMETERS

Our framework is implemented in PyTorch and trained on a single NVIDIA RTX L40 GPU. The motion HFSQ is trained for 200 epochs using the Lion optimizer (Chen et al., 2023) with initial learning rate $3e^{-5}$ and a stepwise learning rate scheduler decays the rate by 0.998 per epoch. For diffusion training, the denoise network is trained for 1500 epochs using the Adamw optimizer with initial learning rate $2e^{-4}$ and a stepwise learning rate scheduler decays the rate by 0.998 per epoch. Batch size is set to 256 across all stages.

6.3 MORE RESULTS OF ABLATION ON HFSQ REPRESENTATION AND PROGRESSIVE MASKING

The visualization in Fig. 6 supports the quantitative findings in the main paper. The ablated model using an RVQ-VAE often generates poses that are individually plausible but contextually incorrect within the duet. This leads to clear visual artifacts such as inaccurate relative positioning and distance between the dancers, as well as failures in fine-grained hand-to-hand or hand-to-body contact. This demonstrates that HFSQ's hierarchical structure is crucial for effectively representing the complex spatial dynamics of dance interactions.

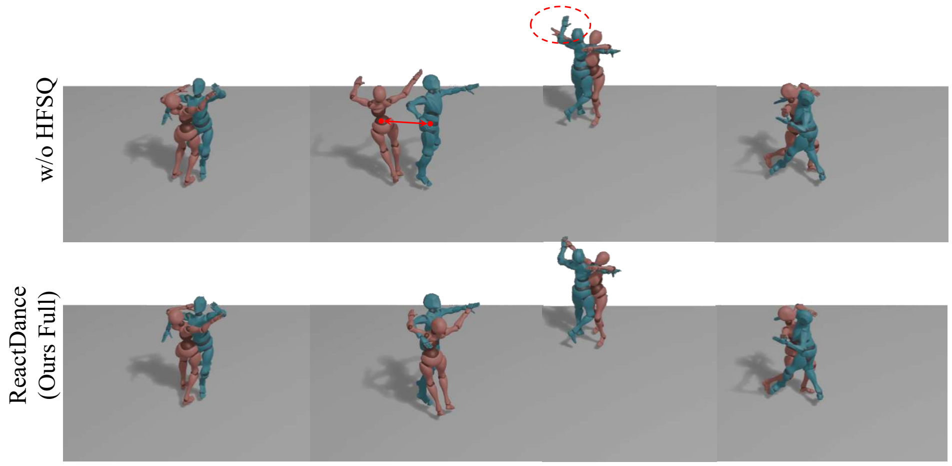

Figure 6: **Qualitative comparison for the HFSQ ablation.** While our full model maintains correct spatial relationships, the model trained without HFSQ (using an RVQ-VAE) fails to capture precise inter-personal dynamics, resulting in incorrect relative distancing and mismatched hand interactions (highlighted in red circles).

As shown in Fig. 7, removing the PM strategy introduces severe interaction artifacts. This aligns with the quantitative trade-off noted in the main paper, where the ablated model achieves a slightly lower MPJPE but a much worse FID_g. Without the regularization provided by PM, the model may reconstruct individual poses more closely but fails to learn robust interaction context. This leads to fundamental errors in character orientation relative to the leader and results in physically implausible outcomes, such as body part interpenetration and collisions. This highlights that the PM strategy is essential for learning the physical constraints and relational logic of reactive dance.

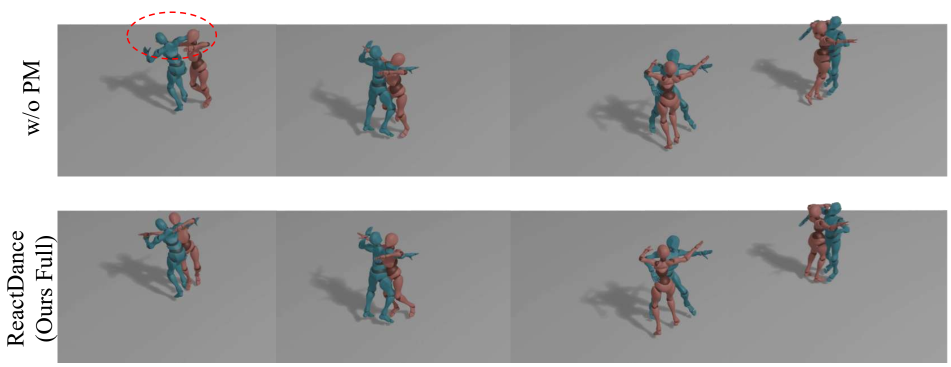

Figure 7: **Qualitative comparison for the Progressive Masking (PM) ablation.** Our full model generates a coherent and physically plausible interaction. In contrast, removing the PM strategy leads to critical errors in interaction logic, such as the reactor adopting an incorrect orientation relative to the leader, resulting in interpenetration artifacts (highlighted).

6.4 MORE DETAILS ABOUT LAYER-DECOUPLED GUIDANCE

Our Layer-Decoupled Classifier-Free Guidance (LDCFG) is designed to overcome the limitation of a single, global guidance weight used in conventional CFG. By assigning independent guidance strengths s_r (Eq. 9) to each hierarchical layer of our HFSQ representation, LDCFG provides fine-grained control over the fidelity-diversity trade-off at different semantic levels.

We conducted a quantitative analysis by sampling with different LDCFG weight pairs $S = [s_{\text{coarse}}, s_{\text{fine}}]$, where s_{coarse} and s_{fine} control the coarse- and fine-grained motion scales, respectively. The results are presented in Table 4. The data reveals a clear trade-off between solo motion fidelity

Table 4: Solo and Duet metrics with diverse Layer-Decoupled Classifier-Free Guidance (LDCFG) weights. Best and second-best results are in **bold** and underline, respectively.

LDCFG Weight	Solo Metrics							Interactive Metrics		
	FID$_k$(↓)	FID$_g$(↓)	Div$_k$(→)	Div$_g$(→)	MPJPE(↓)	MPJVE(↓)	BAS(→)	FID$_{cd}$(↓)	Div$_{cd}$(→)	BED(→)
Ground Truth	–	–	10.86	7.82	–	–	0.1791	–	12.53	0.5308
$S = [1.0, 1.0]$	8.09	10.37	10.74	7.69	155.30	17.59	0.2198	18.15	9.82	0.3470
$S = [1.0, 1.2]$	6.17	8.10	10.79	7.74	138.68	16.29	0.2154	13.66	10.55	0.3709
$S = [1.0, 1.5]$	5.77	7.70	10.80	7.75	134.78	15.90	0.2091	**13.40**	**10.65**	0.3835
$S = [1.2, 1.2]$	**5.57**	7.63	10.82	7.76	132.99	15.68	0.2031	14.17	10.58	0.3863
$S = [1.2, 1.5]$	5.62	**7.29**	10.83	7.79	**132.20**	**15.43**	0.1978	18.57	10.14	0.3901
$S = [1.5, 1.5]$	6.07	7.62	**10.85**	**7.81**	135.27	15.62	**0.1961**	24.56	9.68	0.3925

and interaction quality. For instance, increasing the guidance weights (e.g., $S = [1.2, 1.5]$) improves reconstruction metrics like MPJPE and MPJVE, but at the cost of harming the interaction quality (FID$_{cd}$ increases significantly to 18.57). Conversely, high guidance on both scales ($S = [1.5, 1.5]$) improves diversity (Div$_k$, Div$_g$) but leads to the worst interaction scores (FID$_{cd}$ of 24.56). Our chosen setting, $S = [1.2, 1.2]$ (highlighted), provides an optimal balance, achieving the best FID$_k$ and strong reconstruction quality while maintaining competitive interactive metrics. This demonstrates that LDCFG enables a nuanced tuning of realism against interaction coherence.

Furthermore, because our representation disentangles body components, we can apply LDCFG for targeted artistic control. Fig. 8 visualizes this capability, showcasing independent guidance applied to the upper body, lower body, and their combination to create distinct stylistic motion variations while maintaining plausible interactions.

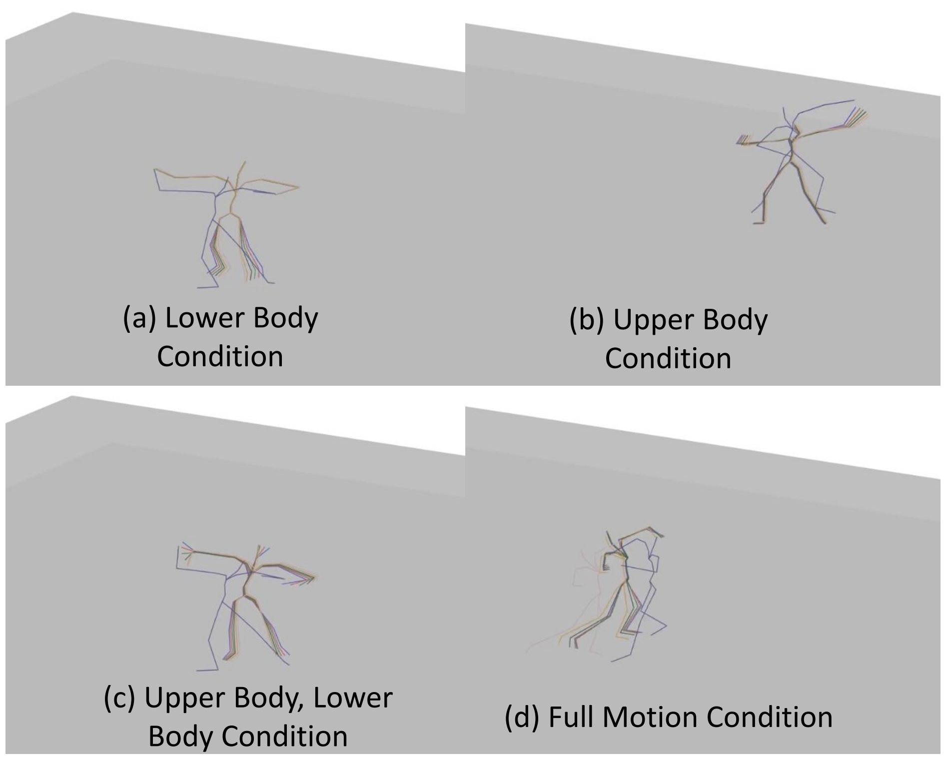

Figure 8: **Qualitative visualization of LDCFG's body-part control.** We apply decoupled guidance with varying strengths to different body regions. The blue robot is the leader. The other colored robots are generated by our model, each with a different LDCFG setting applied to the (a) lower body, (b) upper body, (c) their combination, or (d) full body global motion demonstrating fine-grained artistic control.

6.5 User Study Results

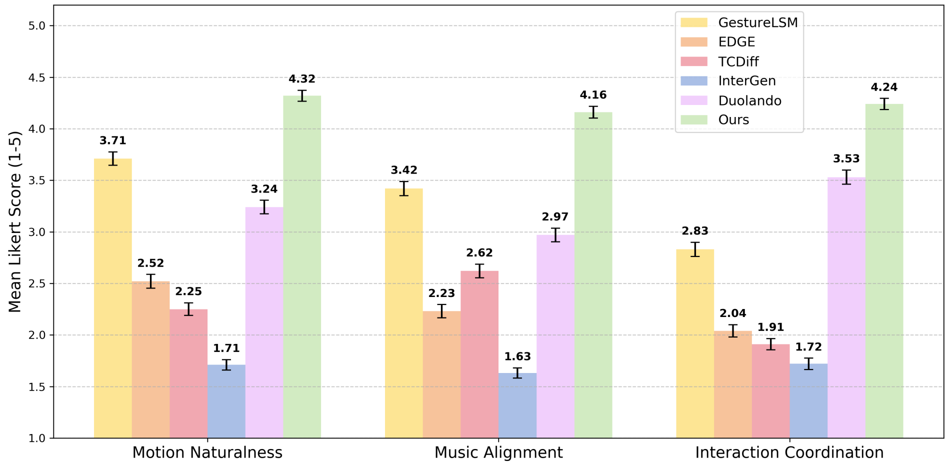

Figure 9: **User study results.** ReactDance consistently outperforms all baselines across all criterions.