
MOSPA: Human Motion Generation Driven by Spatial Audio

Shuyang Xu^{*,1}, Zhiyang Dou^{*,†,1}, Mingyi Shi¹, Liang Pan¹, Leo Ho¹, Jingbo Wang²,
Yuan Liu³, Cheng Lin⁴, Yuexin Ma⁵, Wenping Wang^{†,6}, Taku Komura^{†,1}

¹The University of Hong Kong ²Shanghai AI Lab
³The Hong Kong University of Science and Technology
⁴Macau University of Science and Technology
⁵ShanghaiTech University ⁶Texas A&M University




Figure 1: We introduce a novel human motion generation task centered on spatial audio-driven human motion synthesis. Top row: We curate a novel Spatial Audio-Driven Human Motion (SAM) dataset, including diverse spatial audio signals and high-quality 3D human motion pairs. Bottom row: We develop a generative framework for human *MOTION* generation driven by *SP*atial Audio (MOSPA) to produce high-quality, responsive human motion driven by spatial audio. We note that the motion generation results are both realistic and responsive, effectively capturing both the spatial and semantic features of spatial audio inputs.

Abstract

Enabling virtual humans to dynamically and realistically respond to diverse auditory stimuli remains a key challenge in character animation, demanding the integration of perceptual modeling and motion synthesis. Despite its significance, this task remains largely unexplored. Most previous works have primarily focused on mapping modalities like speech, audio, and music to generate human motion. As of yet, these models typically overlook the impact of spatial features encoded in spatial audio signals on human motion. To bridge this gap and enable high-quality modeling of human movements in response to spatial audio, we introduce the first comprehensive Spatial Audio-Driven Human Motion (SAM) dataset, which contains diverse and high-quality spatial audio and motion data. For benchmarking, we develop a simple yet effective diffusion-based generative framework for human *MOTION* generation driven by *SP*atial Audio, termed *MOSPA*, which faithfully captures the relationship between body motion and spatial audio through an effective fusion mechanism. Once trained, MOSPA can generate diverse realistic human

^{*}, [†] denote equal contributions and corresponding authors.

motions conditioned on varying spatial audio inputs. We perform a thorough investigation of the proposed dataset and conduct extensive experiments for benchmarking, where our method achieves state-of-the-art performance on this task. Our code and model are publicly available at our [website](#).

1 Introduction

Humans exhibit varying responses to different auditory inputs within a given space. For instance, when exposed to sharp, piercing sounds, individuals are likely to cover their ears and move away in the direction opposite to the sound source. Conversely, when the sound is soft and soothing, they may approach it out of curiosity or to investigate further. Therefore, generating realistic human motion for virtual characters to respond realistically to a variety of sounds in their environment is both a highly sought-after feature and is crucial for applications such as virtual reality, human-computer interaction, robotics, etc.

Unfortunately, while previous studies have extensively explored motion generation from action label [85, 29], text [93, 100, 77], music [78, 91, 46, 1], and speech [2, 97, 3, 102], human motion generation driven by *spatial audio* remains unexplored to the best of our knowledge. Unlike pure audio signals, e.g., music [70, 50, 78], speech [3, 97], the spatial audio signals not only does it encode semantics, but it also captures *spatial characteristics* that significantly influence body movements, requiring a specialized framework to accurately model motion responses to spatial audio stimuli.

To address this overlooked aspect, we propose to model the complex interactions between spatial audio inputs and human motion using a generative model. Since there is no such dataset tailored for this task, we first introduce the *SAM dataset* (Spatial Audio Motion dataset), which captures diverse human responses to various spatial audio conditions. This dataset is meticulously curated to include a wide range of spatial audio scenarios, enabling the study of motion conditioned on sound field variations. The *SAM dataset* has a total of more than 9 hours of motion, covering 27 common spatial audio scenarios and more than 70 audio clips. To ensure the diversity of the spatial audio, around 480 seconds of motion were captured for each audio clip at different positions in the character space. To ensure diverse motion responses to spatial audio, we introduce 20 distinct motion types (excluding motion genres) and 49 in total when including motion genres. We visualize samples from SAM in the top row of Fig. 1. See Appendix A for detailed statistics.

We further conduct benchmarking experiments on the proposed dataset, revealing the limitations of existing methods in this setting. To enable spatial audio-driven human motion generation, we introduce MOSPA, a simple yet effective framework tailored for this task. In real-world scenarios, human responses to sound are inherently influenced by spatial perception, intensity variations, directional cues, temporal dynamics, etc. Motivated by this, we generate motion by incorporating features extracted from the input spatial audio signals using [61]. Specifically, to capture intrinsic features across both temporal and spatial dimensions, we mainly utilize Mel-Frequency Cepstral Coefficients (MFCCs)[17] and Tempograms[28] to model the temporal characteristics of the audio. Additionally, we characterize the spatial audio by analyzing the root mean square (RMS) [61] energy, which quantifies signal intensity in audio processing.

These features enhance the effective modeling of the spatial and intensity variations of the spatial audio. To capture the distribution of spatial audio features and human motion dynamics effectively, we employ a diffusion-based generative model that ensures strong alignment between the two modalities—human motion and spatial audio signals. Leveraging diffusion models, MOSPA excels at modeling the complex interplay between spatial audio features and human motion. Besides, a residual feature fusion mechanism is employed to model the subtle influences of spatial audio on human movement.

Extensive evaluations on the SAM demonstrate that MOSPA achieves state-of-the-art performance on this task, outperforming existing baselines in generating realistic and diverse motion responses to spatial audio. Our contributions are summarized as follows:

- We introduce a novel task of spatial audio-conditioned motion generation and present the first comprehensive dataset SAM with over 9 hours of motion across diverse scenarios.
- We conduct extensive benchmarking and propose MOSPA, a diffusion-based generative framework tailored for modeling and generating diverse human motions from spatial audio.

- We achieve the SOTA performance on motion generation conditioned on spatial audio. Our dataset, code, and models will be publicly released for further research.

2 Related Work

Spatial Audio. Many studies have explored spatial audio modeling [98, 24, 99, 74, 37, 87, 44, 45]. For instance, [98] utilizes the natural synchronization between visual and audio modalities to learn models that jointly parse sounds and images without manual annotations. [24] leverages unlabeled audiovisual data to localize objects, such as moving vehicles, using only stereo sound at inference time. [99] reason about spatial sounds with large language models. Recently, spatial audio generation has been explored from text [74] and video [42]. [87] propose a method to model 3D spatial audio from body motion and speech. [37] presents a framework for spatial audio generation, capable of rendering 3D soundfields generated by human actions, including speech, footsteps, and hand-body interactions. Despite progress on spatial-audio tasks, generating human motion from spatial audio remains largely underexplored.

Conditional Motion Generation. Extensive efforts have been made into motion synthesis conditioning on user control signals [36, 72, 11, 79, 92], text [93, 100, 77, 30, 51, 53, 14], action [85, 29, 10, 39], music [78, 91, 46, 1], speech [2, 68, 38, 97, 3, 102], past trajectories [23, 9, 4, 90, 73], etc. We refer readers to [103] for a detailed survey of motion generation.

Text-to-Motion. Text-to-motion generation has recently gained popularity as an intuitive and user-friendly approach to synthesizing diverse human body motions. Generative pre-trained transformer frameworks have been utilized for text-to-motion generation [93, 57]. Subsequently, various generation techniques have been explored, including diffusion models [77], latent diffusion models [12], autoregressive diffusion model [69, 11], denoising diffusion GANs [100], consistency models [16], and generative masked modeling [31]. Recent advancements include the integration of motion generation with large language models [41, 94] and investigations into the scaling laws for motion generation [58, 22]. Recently, controllable text-to-motion generation has gained attention, enabling motion synthesis conditioned on both text prompts and control signals, e.g., target control points [84, 79].

Music-to-Motion. Recent advancements in Music-to-Motion generation have been made [21, 70, 48, 78, 27, 88, 71]. DanceFormer [47] adopts a two-stage approach, generating key poses for beat synchronization followed by parametric motion curves for smooth, rhythm-aligned movements. Bailando [70] utilizes a VQ-VAE to encode motion features via a choreographic memory module. [1] introduces a diffusion-based probabilistic model for motion generation, using a Conformer-based architecture. EDGE [78] also applies a diffusion model for dance generation and editing. Furthermore, multimodal approaches incorporating language and music enhance generation quality [27, 88, 15].

Speech-to-Motion. We mainly review studies on audio-driven motion (gesture) generation [97, 13, 3, 2, 102]. Early works are mostly based on GAN models [26, 54, 65, 89], while the recent attempts are mainly based on the generative diffusion model [102]. For instance, [97] proposes a generative retrieval framework leveraging a large language model to efficiently retrieve semantically appropriate gesture candidates from a motion library in response to input speech. [2] introduces a co-speech gesture synthesis method by employing a segmentation pipeline for temporal alignment and disentangling speech-motion embeddings to capture both semantics and subtle variations.

While audio signals have been widely used in music- and speech-to-motion tasks, human motion synthesis driven by spatial audio remains largely unexplored. As a result, data-driven methods are highly constrained by limited paired data. The goal of this paper is to develop a comprehensive dataset and a novel approach for high-quality spatial audio-driven motion synthesis.

3 SAM Dataset

We first introduce the Spatial Audio-driven Motion (SAM) dataset designed for human motion synthesis conditioned on spatial audio. We focus on binaural audio, a common form of spatial audio that aligns with human and (most) animal perception and can be readily applied to robotic platforms. SAM consists of more than 9 hours of human motions with corresponding binaural audio, and more than 4M frames, covering 27 common spatial audio scenarios and 20 common reaction types in daily life without counting the motion genres. The majority of the audio clips are sourced from the AudioSet [25], while only a small portion is extracted from publicly available YouTube videos

Table 1: Statistics of the SAM dataset. The SAM dataset encompasses 27 common daily spatial audio scenarios, over 20 reaction types excluding the motion genres, and 49 reaction types (see details in Appendix A). The number of subjects covered in SAM is 12, where 5 of them are female and the remaining 7 are male. It is also the first dataset to incorporate *spatial* audio information, annotated with Sound Source Location (SSL). The total duration of the dataset exceeds 34K seconds.

Dataset	SSL	3DJoint _{pos/rot}	Model	Joints	Subjects	Seconds
Dance with Melody [75]	×	✓/×	-	21	-	5640
DanceNet [104]	×	✓/×	-	55	2	3472
AIST++ [48]	×	✓/✓	COCO/SMPL	17/24	30	18694
PopDanceSet [60]	×	✓/✓	COCO/SMPL	17/24	132	12819
FineDance [49]	×	✓/✓	SMPL & hand joints	52	27	52560
SAM (Ours)	✓	✓/✓	SMPL-X	55	12	34356





Figure 2: Visualization of samples from SAM with expected motions annotated. Red dots indicate the actor's trajectory, while the blue sphere represents the sound source. The SAM dataset ensures high diversity by encompassing a broad spectrum of audio types and varying sound source locations.

or through manual recording. More detailed information can be found in Tab. 1 and Appendix A. Visualization results are in Fig. 2.



(i) Mocap environment. (ii) The speakers.

Figure 3: Spatial audio-driven human motion data collection setup.




Figure 4: Statistics of action duration in the dataset.

Data Capture Settings. We utilize a Vicon motion capture system [56] to collect motion data and spatial audio signals. The motion capture is performed in a semi-open cage having a space of approximately $5\text{m} \times 10\text{m} \times 3\text{m}$, a structure covered by rope nets, with 28 mocap cameras mounted on the ropes and vertical supports recording at a frame rate of 120 Hz; See Fig. 3. The surrounding walls are standard painted concrete, resulting in a setting that resembles a typical indoor environment. In SAM, each audio clip is associated with 16 randomly sampled relative sound source locations, defined by combinations of different speakers and spatial positions relative to the subject. For each location, we capture three motion sequences corresponding to different reaction intensities: dull, neutral, and sensitive, resulting in a total of 48 motion sequences, each lasting 10 seconds. Fig. 4 shows the statistics of the approximate duration of the actions. The three motion genres define the varying degrees of responsiveness, decreasing from sensitive to dull. For instance, upon hearing an explosion, a dull individual might remain largely unreactive, whereas a sensitive one may immediately flee from the sound source. The total number of action types is 49. The percentage of action types covered within the dull, neutral, and sensitive motion genres are 28.57%, 34.69%, and 36.73%, respectively. To capture the binaural sound heard at the position of the actor, we employ two microphones to record the audio at the ear positions of the actors separately; See the inset. The two microphones are connected to a Deity PR-2 recorder [18] that has been synchronized with the Vicon mocap system in advance using a timecode with a frame rate of 30 FPS. With this setting, the stereo sound at the position of the actor can be recorded and has an accurate alignment with the corresponding motion.



Data Processing. All motion and audio clips are precisely aligned. The motions are re-targeted and converted from the original BVH format in Vicon to the SMPL-X [63] format. SMPL-X is a parametric 3D human body model that encompasses the body, hands, and face, comprising $N = 10,475$ vertices and $K = 55$ joints. Given shape parameters β and pose parameters θ , the SMPL-X model generates the corresponding body shape and pose through forward dynamics. We extract the locations of the sound sources in each motion clip. The sound source locations are then transformed into the local space of the character aligned with the SMPL-X local coordinate system (a.k.a the local frame).

4 Method

We introduce MOSPA, a diffusion-based probabilistic model that serves as a baseline for this novel task of spatial audio-driven human motion generation. First, we extract spatial audio features $\mathbf{a}$ using a feature extractor [61]. During motion generation, the extracted spatial audio feature $\mathbf{a}$ is combined with the sound source location $\mathbf{s}$ and the motion genre g as conditioning inputs. These inputs are passed to a denoiser $\mathcal{G}$, which is trained to reconstruct the original clean motion vector $\mathbf{x}_0$ by denoising the given noisy motion vector $\mathbf{x}_t$ at time step t . Mathematically, we have $\mathbf{x}_0 = \mathcal{G}(\mathbf{x}_t, t; \mathbf{a}, \mathbf{s}, g)$.

4.1 Feature Representation

The two key vectors are the audio feature vector $\mathbf{a}$ and the motion vector $\mathbf{p}$. We carefully designed the structure of the two vectors in MOSPA.

Spatial Audio Feature Extraction. We first extract a range of audio features that capture intensity, temporal dynamics, and spatial characteristics. Inspired by [70], our feature set primarily includes *Mel-*

frequency cepstral coefficients (MFCCs), MFCC delta, constant- Q chromagram, short-time Fourier transform (STFT) of the chromagram, onset strength, tempogram, and beats [17, 67, 28, 61, 7, 20]. On top of these audio features, we additionally add the root mean square (RMS) energy E_{rms} of the audio [61], and the active frames F_{active} defined as $F_{active} = E_{rms} > 0.01$ to capture the distance information of the audio. The dimension of the audio feature vector for each ear is 1136. By concatenating the features from both ears, we obtain a combined feature vector $\mathbf{a}$ of dimension 2272. The detailed construction of the audio vector can be viewed in Appendix B.1.

Motion Representation. In this paper, we focus on body motion and leave the modeling of detailed finger movements to future work. Therefore, we exclude all the finger joints and retain only the first $J = 25$ body joints of the SMPL-X model [63]. In addition to the essential translation and joint rotations required for human pose representation, we introduce the residual feature fusion mechanism [30] to incorporate the global joint positions and the velocity of the joints to capture the nuanced difference in audio and further improve the accuracy of the generated samples. Each motion vector $\mathbf{x}$ is thus composed of the global positions $\mathbf{p} \in \mathbb{R}^{T \times (J \times 3)}$, the local rotations $\mathbf{r} \in \mathbb{R}^{T \times (J \times 6)}$ and the velocities $\mathbf{v} \in \mathbb{R}^{T \times (J \times 3)}$ of the joints (including the root), where $T = 240$ represents the number of frames in each motion sequence. The joint rotations are represented in the 6d format [101] to guarantee the continuity of the change ($\mathbf{x}_0 = (\mathbf{p}_0, \mathbf{r}_0, \mathbf{v}_0)$, $\mathbf{x} \in \mathbb{R}^{T \times (J \times 12)}$). The dimension of each motion vector is therefore 300.

4.2 Framework

Following [77, 11], the diffusion is modeled as a Markov chain process which progressively adds noise to clean motion vectors $\mathbf{x}_0$ in t time steps, i.e.

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\sqrt{\alpha_t} \mathbf{x}_{t-1}, (1 - \alpha_t) I) \quad (1)$$

where $\alpha_t \in (0, 1)$. The model then learns to gradually denoise a noisy motion vector $\mathbf{x}_t$ in t time steps, i.e. $p(\mathbf{x}_{t-1} | \mathbf{x}_t)$. We directly predict the clean sample $\hat{\mathbf{x}}_0$ in each diffusion step $\hat{\mathbf{x}}_0 = \mathcal{G}(\mathbf{x}_t, t; \mathbf{a}, \mathbf{s}, g)$, where $\mathbf{a}$ is the audio features, $\mathbf{s}$ is the sound source location and g is the motion genre. This strategy, employed by [77, 11, 66, 84], has been proved to be more efficient and accurate than predicting the noise ϵ_t , suggested by [35]. We employ an encoder-only Transformer to reverse the diffusion process and predict the clean samples. The timestep, motion, and conditioning signals are each projected into the same latent dimension using separate feed-forward networks. Random masks are applied to the audio features $\mathbf{a}$ and the sound source location (SSL) $\mathbf{s}$, after which all components are concatenated to form the complete token sequence $\mathbf{z}$. The tokens are positionally embedded afterward and input into a transformer to get the output $\hat{\mathbf{z}}$. The predicted clean sample is thus extracted from the last T tokens of $\hat{\mathbf{z}}$ by inputting it to another feed-forward network, where $T = 240$ is the length of the motion; see Fig. 5.

4.3 Loss Functions

We train MOSPA using the following loss functions. A simple mean squared error (MSE) loss is applied to the original clean sample and the predicted clean sample as the main objective: $\mathbb{E} \|\mathbf{x}_0 - \hat{\mathbf{x}}_0\|_2^2$. To guarantee the smooth variation on the predicted clean sample across frames, we also apply MSE loss to the rate of change of the vectors across frames: $\mathbb{E} \|\delta \hat{\mathbf{x}}_0 - \delta \mathbf{x}_0\|_2^2$. Combining the two simple losses we have $\mathcal{L}_{data} = \mathbb{E} \|\mathbf{x}_0 - \hat{\mathbf{x}}_0\|_2^2 + \mathbb{E} \|\delta \hat{\mathbf{x}}_0 - \delta \mathbf{x}_0\|_2^2$. Geometric losses, encompassing position loss and velocity loss, are also incorporated, as we rely solely on joint




Figure 5: The framework of MOSPA. We perform diffusion-based motion generation given spatial audio inputs. Specifically, Gaussian noise is added to the clean motion sample $\mathbf{x}_0$, generating a noisy motion vector $\mathbf{x}_t$, modeled as $q(\mathbf{x}_t | \mathbf{x}_{t-1})$. An encoder transformer then predicts the clean motion from the noisy motion $\mathbf{x}_t$, guided by extracted audio features $\mathbf{a}$, sound source location (SSL) $\mathbf{s}$, motion genre g , and timestep t .




Figure 6: Qualitative comparison of state-of-the-art methods for the spatial audio-to-motion task. We visualize motion results from five cases. MOSPA produces high-quality movements that closely correspond to the input spatial audio. We provide Expected Motion as a description for reference.

rotations and translations in motion vectors to represent poses: $\mathcal{L}_{geo} = \mathbb{E}\|FK(\hat{\mathbf{x}}_0) - FK(\mathbf{x}_0)\|_2^2 + \mathbb{E}\|\delta FK(\hat{\mathbf{x}}_0) - \delta FK(\mathbf{x}_0)\|_2^2$. Furthermore, foot sliding is prevented by introducing the foot contact loss $\mathcal{L}_{foot}$ that measures the inconsistency in the velocities of the foot joints between the ground truth and the predicted motions.

We also incorporate trajectory loss and joint rotation loss to underscore their importance in achieving the training objectives and accelerate the convergence of the model, defined as $\mathcal{L}_{traj} = \mathbb{E}\|\mathbf{traj}_0 - \hat{\mathbf{traj}}_0\|_2^2 + \mathbb{E}\|\delta \mathbf{traj}_0 - \delta \hat{\mathbf{traj}}_0\|_2^2$ and $\mathcal{L}_{rot} = \mathbb{E}\|\hat{\mathbf{r}}_0 - \mathbf{r}_0\|_2^2 + \mathbb{E}\|\delta \hat{\mathbf{r}}_0 - \delta \mathbf{r}_0\|_2^2$ respectively, where $\mathbf{traj}$ is the trajectory vector of the motion sequence and $\mathbf{r}$ is the joint rotations represented in the 6d format [101]. Given that trajectory and joint rotations are inherently encoded within the motion vectors, these supplementary losses represent an overlap with the existing loss terms, effectively amplifying the emphasis on trajectory and joint rotation accuracy through increased weighting. Empirically, we observe that this implementation accelerates model convergence and facilitates correct displacement direction generation in motion sequences. In sum, the total loss is given by:

$$\mathcal{L} = \lambda_{data}\mathcal{L}_{data} + \lambda_{geo}\mathcal{L}_{geo} + \lambda_{foot}\mathcal{L}_{foot} + \lambda_{traj}\mathcal{L}_{traj} + \lambda_{rot}\mathcal{L}_{rot} \quad (2)$$

All loss weights (λ) are initialized set to 1. At epoch 5,000 of the total 6,000 training epochs, λ_{traj} and λ_{rot} are increased to 3, thereby intensifying the emphasis on trajectory and rotation accuracy.

4.4 Implementation Details

In MOSPA, the diffusion model is a transformer-based diffusion network [11, 77, 100]. The encoder transformer is configured with a latent dimension of 512, 8 heads, and 4 layers. We employ AdamW [55] as the optimizer with an initial value of 1×10^{-4} . The number of denoising steps used is 1000, and the noise schedule is cosine. The training phase concludes after 6,000 epochs. Exceeding these recommended epoch counts may degrade model quality due to overfitting. The

Table 2: Quantitative evaluation on the SAM, where MOSPA achieves higher alignment with the GT motion while maintaining high diversity, as reflected by the metrics. The error bar is the 95% confidence interval assuming normal distribution, and $\rightarrow$ means the closer to Real Motion the better.

Method	R-precision $\uparrow$			FID $\downarrow$	Diversity $\rightarrow$	APD $\rightarrow$
	Top1	Top2	Top3			
Real Motion	1.000 $\pm$ 0.000	1.000 $\pm$ 0.000	1.000 $\pm$ 0.000	0.001	23.616 $\pm$ 0.188	59.435
EDGE [78]	0.886 $\pm$ 0.005	0.960 $\pm$ 0.003	0.977 $\pm$ 0.002	13.993	23.099 $\pm$ 0.196	43.882
POPDG [60]	0.762 $\pm$ 0.006	0.886 $\pm$ 0.005	0.934 $\pm$ 0.003	20.967	22.536 $\pm$ 0.170	34.996
LODGE [50]	0.444 $\pm$ 0.006	0.594 $\pm$ 0.005	0.679 $\pm$ 0.004	102.289	21.101 $\pm$ 0.141	11.801
Bailando [70]	0.077 $\pm$ 0.003	0.134 $\pm$ 0.003	0.182 $\pm$ 0.004	168.396	17.347 $\pm$ 0.247	23.121
MOSPA	0.937$\pm$0.005	0.984$\pm$0.002	0.996$\pm$0.001	7.981	23.575$\pm$0.188	53.915

Table 3: Ablation study on MOSPA on the spatial audio-driven motion generation performance. The error bar is the 95% confidence interval assuming normal distribution, and $\rightarrow$ means the closer to real motions the better.

Latent Dim	Head Num	Diff Steps	Genre	R-precision $\uparrow$			FID $\downarrow$	Diversity $\rightarrow$	APD $\rightarrow$
				Top1	Top2	Top3			
Real Motion				1.000 $\pm$ 0.000	1.000 $\pm$ 0.000	1.000 $\pm$ 0.000	0.001	23.616 $\pm$ 0.188	59.435
512	8	1000	✓	0.937$\pm$0.005	0.984 $\pm$ 0.002	0.996 $\pm$ 0.001	7.981	23.575$\pm$0.188	53.915
256	8	1000	✓	0.891 $\pm$ 0.005	0.952 $\pm$ 0.002	0.971 $\pm$ 0.001	9.226	23.007 $\pm$ 0.198	55.175
512	4	1000	✓	0.923 $\pm$ 0.004	0.972 $\pm$ 0.002	0.986 $\pm$ 0.001	9.282	23.232 $\pm$ 0.170	56.572
512	8	100	✓	0.930 $\pm$ 0.004	0.980 $\pm$ 0.002	0.991 $\pm$ 0.001	8.456	23.351 $\pm$ 0.177	49.824
512	8	4	✓	0.934 $\pm$ 0.004	0.989$\pm$0.002	0.998$\pm$0.001	8.387	23.474 $\pm$ 0.192	49.507
512	8	1000	×	0.889 $\pm$ 0.005	0.958 $\pm$ 0.003	0.977 $\pm$ 0.002	10.930	23.150 $\pm$ 0.153	46.807

entire training process requires approximately 18 hours on a single RTX 4090 GPU with a batch size of 128.

5 Experiments

Experiment Setup. We use our SAM dataset to evaluate the spatial audio-driven motion generation task. As detailed in Sec. 3, it contains 9 hours of human motion with paired binaural audio and corresponding sound source locations, covering 27 common spatial audio scenarios and 20 common reaction types. The dataset is split into training, validation, and test sub-datasets at a common ratio of 8:1:1. Consequently, the training sub-dataset comprises 2,400 motion sequences, while the validation and test sub-datasets each contain approximately 300 motion sequences. To keep fair setting [10], the motions and the audio clips are both downsampled to the frame rate of 30 FPS. The character is rotated to face the negative y-axis and initially translated to the origin in the world space in all motion sequences, and the sound source locations (SSL) are transformed to the local space of the character in every single frame.

Baselines and Metrics. Our system is the first work to receive spatial audio as input to generate human motion results. To our best knowledge, as there is no other system achieving this, we made adaptations on other audio2motion methods, such as EDGE [78], POPDG [60], LODGE [50] and Bailando [70] by replacing their original audio input with our spatial audio feature as input. We evaluated four metrics, focusing on motion quality and diversity: 1) *R-precision*, *FID*, *Diversity*. These three metrics are calculated using the same setup proposed by [30]. Two bidirectional GRU are trained with a hidden size of 1024 for 1,500 epochs with a batch size of 64 to extract the audio features and the corresponding motion features, as suggested by [30]. Detailed implementation details of the feature extractor are provided in Appendix B.2. 2) *APD* [19, 33] is calculated by

$$APD(M) = \frac{1}{N(N-1)} \sum_{i=1}^N \sum_{\substack{j=1 \\ j \neq i}}^N \left(\sum_{t=1}^L \|s_t^i - s_t^j\|^2 \right)^{\frac{1}{2}},$$
 where $M = \{\mathbf{x}_i\}$ is the set of generated motion sequences, N is the number of motion sequences in the set M , L is the number of frames of each motion sequence, and $s_t^i \in \mathbf{x}_i$ is a state in the motion sequence $\mathbf{x}_i$.

5.1 Comparisons

Qualitative Results. We demonstrate the qualitative comparison in Fig. 6. For the same input spatial audio, our methods show the superiority of producing high-quality and realistic response motion. Other methods often exhibit various limitations due to their unique model characteristics. EDGE [78] and POPDG [60] demonstrate relatively strong performance among the four baselines, sharing a diffusion-based foundation with MOSPA, despite differences in their encoding and decoding mechanisms. Their shortcomings in generated samples can primarily be attributed to model size and their strong focus on music-like audio. The bad performance of LODGE [50] is likely due to its specialization in long-term music-like audio, resulting in deficiencies when handling short-term audio information with abrupt feature changes. Similarly, Bailando [70] faces challenges in processing rapidly changing spatial audio. More critically, due to its separate training process for upper and lower body parts, Bailando occasionally produces distorted or disjointed motions when encountering sudden changes in spatial audio. Please watch our supplementary video for more results. Furthermore, we test MOSPA on out-of-distribution audio-source configurations. As shown in Fig. 8, it maintains motion quality and intent alignment, demonstrating robustness to unseen spatial setups.

Quantitative Results. The quantitative results are reported in Tab. 2. MOSPA achieves the best performance as shown by the lowest FID value and the highest R-precision values. Also, our generated motions exhibit the closest diversity and APD [19] values compared with the Real Motion, demonstrating the effectively balanced variation and precision. Bailando [70] has the worst performance among the four baselines in practice, as illustrated by the extremely high FID. The model possibly lacks the ability to perceive commonly heard sounds other than music and also the spatial information of the audio. Our method, overall speaking, still demonstrates competitive performance in spatial audio conditioned motion generation, which is proved by the low values in precision-related metrics and the high values in diversity-related metrics.

User Study. We conducted a user study with 25 participants to assess the perceptual quality of motion generation. Participants evaluated five models (MOSPA, EDGE, POPDG, LODGE, Bailando) alongside ground truth (GT), selecting the best motion for: **1) Human Intent Alignment:** Does the motion align with real-world intent? **2) Motion Quality:** Which has the highest movement quality? **3) GT Similarity:** Which best matches the GT motion? We provided GT motion and a textual description for reference. As shown in Fig. 7, MOSPA outperforms all baselines across all criteria, while LODGE and Bailando received the fewest selections, indicating limitations in generating realistic, semantically meaningful motions. See more details in Appendix C.




Figure 7: User study results. MOSPA outperforms other methods in intent alignment, motion quality, and similarity to ground truth. The bar chart shows the vote distribution across methods.

5.2 Ablation Study

We conducted ablation studies on the latent dimension, the number of attention heads, the diffusion step number, and the masking of motion genre, with results summarized in Tab. 3. All ablation experiments maintained consistent training epoch counts throughout.

Latent Dimension. The default latent dimension of MOSPA’s encoder transformer is 512. In our study, reducing it to 256 slightly increases the APD [19] value but also degrades the R-precision and the FID, leading to an overall decline in model performance, as seen in row 1 and row 2 in Tab. 3.

Number of Attention Heads. We reduced the number of attention heads in MOSPA’s encoder transformer from 8 to 4, observing degradation in almost all of the metrics except a slight improvement in APD [19]. This reduction compromises overall model performance without yielding significant improvements in training efficiency, as seen in row 1 and row 3 in Tab. 3.

Number of Diffusion Steps. We evaluated MOSPA with varying diffusion step numbers, reducing it from 1000 to 100 and further to 4, as detailed in rows 1, 4, and 5 of Tab. 3. Fewer steps slightly degrade the performance as shown by the increase in FID and degradation in diversity, thereby lowering the power limit of the model.




Figure 8: Test of MOSPA on out-of-distribution spatial audios. Descriptions of motions are provided for reference.




Figure 9: Spatial audio-driven physically simulated humanoid robot control based on [34]. Descriptions of expected motion are provided for reference.

Genre Masking. Masking motion genres leads to a degradation in model performance across all metrics, as demonstrated in row 1 and 6 of Tab. 3. Motion genres are required to provide a guidance for the model on the intensity of the expected motions.

We evaluate the contribution of the extracted audio features by conducting an ablation study on the effectiveness of MFCC [17] and tempogram [28] features. As shown in Tab. 4, improvements in FID and R-precision—two key metrics for assessing generative quality and correspondence—demonstrate their significance in model.

Table 4: Ablation study on the effect of MFCC [17] and tempogram [28] features.

MFCC	Tempogram	R-precision $\uparrow$			FID $\downarrow$
		Top-1	Top-2	Top-3	
$\checkmark$	$\checkmark$	0.937 ± 0.005	0.984 ± 0.002	0.996 ± 0.001	7.981
$\times$	$\checkmark$	0.907 ± 0.004	0.967 ± 0.002	0.983 ± 0.002	9.070
$\checkmark$	$\times$	0.917 ± 0.004	0.982 ± 0.002	0.994 ± 0.001	10.786

6 Conclusion

This introduces a novel task for enabling virtual humans to respond realistically to spatial auditory stimuli. We present a comprehensive SAM dataset, capturing human movement in response to spatial audio, and propose MOSPA, a diffusion-based generative model with an attention-based fusion mechanism. Once trained, MOSPA synthesizes diverse, high-quality motions that adapt to varying spatial audio inputs with binaural recording. Extensive evaluations show MOSPA achieves state-of-the-art performance on this task. *Limitations and Future Works.* Physical Correctness: While MOSPA generates diverse and semantically plausible motions, it lacks physical constraints, which may lead to physically implausible artifacts. Integrating physics-based control methods [19, 59, 76, 40, 95, 96, 62, 82] could improve motion realism and embodiment fidelity (see Fig. 9 for spatial audio-driven humanoid robot control). Body Modeling: This work focuses on body motion and omits finer-grained components such as hand gestures and facial expressions supported by SMPL-X [63]. Extending the model to full-body motion generation [57, 52, 86, 64, 5]—including hand motions—remains an important direction for future research. Scene Awareness: The current framework does not incorporate awareness of surrounding environments or physical scene geometry, limiting its ability to produce scene-consistent or contact-aware motions. Future extensions could integrate scene representations or affordance prediction [14, 80, 83, 8, 81] with spatial audio signals to enhance human motion generation.

Acknowledgements

This work is partly supported by the Innovation and Technology Commission of the HKSAR Government under the ITSP-Platform grant (Ref: ITS/335/23FP) and the InnoHK initiative (TransGP project). Part of the research was conducted in the JC STEM Lab of Robotics for Soft Materials, funded by The Hong Kong Jockey Club Charities Trust. We are grateful to Yiduo Hao, Chuan Guo, Chen Wang, Zitong Lan, and Peter Qingxuan Wu for their insightful discussions and constructive feedback, which have been helpful to the development of this research.

References

- [1] Simon Alexanderson, Rajmund Nagy, Jonas Beskow, and Gustav Eje Henter. Listen, denoise, action! audio-driven motion synthesis with diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–20, 2023.
- [2] Tenglong Ao, Qingzhe Gao, Yuke Lou, Baoquan Chen, and Libin Liu. Rhythmic gesticulator: Rhythm-aware co-speech gesture synthesis with hierarchical neural embeddings. *ACM Transactions on Graphics (TOG)*, 41(6):1–19, 2022.
- [3] Tenglong Ao, Zeyi Zhang, and Libin Liu. Gesturediffuclip: Gesture diffusion model with clip latents. *ACM Transactions on Graphics (TOG)*, 42(4):1–18, 2023.
- [4] German Barquero, Sergio Escalera, and Cristina Palmero. Belfusion: Latent diffusion for behavior-driven human motion prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2317–2327, 2023.
- [5] Yuxuan Bian, Ailing Zeng, Xuan Ju, Xian Liu, Zhaoyang Zhang, Wei Liu, and Qiang Xu. Motioncraft: Crafting whole-body motion with plug-and-play multimodal controls. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1880–1888, 2025.
- [6] RB Blackman and JW Tukey. The measurement of power spectra dover publications. *Inc, New York*, 1958.
- [7] Sebastian Böck and Gerhard Widmer. Maximum filter vibrato suppression for onset detection. In *Proc. of the 16th Int. Conf. on Digital Audio Effects (DAFx), Maynooth, Ireland (Sept 2013)*, page 4. Citeseer, 2013.
- [8] Zhi Cen, Huaijin Pi, Sida Peng, Zehong Shen, Minghui Yang, Shuai Zhu, Hujun Bao, and Xiaowei Zhou. Generating human motion in 3d scenes from text descriptions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1855–1866, 2024.
- [9] Ling-Hao Chen, Jiawei Zhang, Yewen Li, Yiren Pang, Xiaobo Xia, and Tongliang Liu. Humanmac: Masked motion completion for human motion prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9544–9555, 2023.
- [10] Ling-Hao Chen, Shunlin Lu, Wenxun Dai, Zhiyang Dou, Xuan Ju, Jingbo Wang, Taku Komura, and Lei Zhang. Pay attention and move better: Harnessing attention for interactive motion generation and training-free editing. *arXiv preprint arXiv:2410.18977*, 2024.
- [11] Rui Chen, Mingyi Shi, Shaoli Huang, Ping Tan, Taku Komura, and Xuelin Chen. Taming diffusion probabilistic models for character control. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–10, 2024.
- [12] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18000–18010, 2023.
- [13] Qingrong Cheng, Xu Li, and Xinghui Fu. Siggature: Generalized co-speech gesture synthesis via semantic injection with large-scale pre-training diffusion models. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–11, 2024.
- [14] Peishan Cong, Ziyi Wang, Zhiyang Dou, Yiming Ren, Wei Yin, Kai Cheng, Yujing Sun, Xiaoxiao Long, Xinge Zhu, and Yuexin Ma. Laserhuman: language-guided scene-aware human motion generation in free environment. *arXiv preprint arXiv:2403.13307*, 2024.
- [15] Rishabh Dabral, Muhammad Hamza Mughal, Vladislav Golyanik, and Christian Theobalt. Mofusion: A framework for denoising-diffusion-based motion synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9760–9770, 2023.
- [16] Wenxun Dai, Ling-Hao Chen, Jingbo Wang, Jinpeng Liu, Bo Dai, and Yansong Tang. Motionlcm: Real-time controllable motion generation via latent consistency model. In *European Conference on Computer Vision*, pages 390–408. Springer, 2024.

- [17] Steven Davis and Paul Mermelstein. Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences. *IEEE transactions on acoustics, speech, and signal processing*, 28(4):357–366, 1980.
- [18] Deity Microphones. Pr-2. <https://deitymic.com/products/pr-2/>, 2025. Accessed: Feb 5, 2025.
- [19] Zhiyang Dou, Xuelin Chen, Qingnan Fan, Taku Komura, and Wenping Wang. C-ase: Learning conditional adversarial skill embeddings for physics-based characters. In *SIGGRAPH Asia 2023 Conference Papers*, pages 1–11, 2023.
- [20] Daniel PW Ellis. Beat tracking by dynamic programming. *Journal of New Music Research*, 36(1):51–60, 2007.
- [21] Di Fan, Lili Wan, Wanru Xu, and Shenghui Wang. A bi-directional attention guided cross-modal network for music based dance generation. *Computers and Electrical Engineering*, 103:108310, 2022.
- [22] Ke Fan, Shunlin Lu, Minyue Dai, Runyi Yu, Lixing Xiao, Zhiyang Dou, Juntong Dong, Lizhuang Ma, and Jingbo Wang. Go to zero: Towards zero-shot motion generation with million-scale data. *arXiv preprint arXiv:2507.07095*, 2025.
- [23] Yuming Feng, Zhiyang Dou, Ling-Hao Chen, Yuan Liu, Tianyu Li, Jingbo Wang, Zeyu Cao, Wenping Wang, Taku Komura, and Lingjie Liu. Motionwavelet: Human motion prediction via wavelet manifold learning. *arXiv preprint arXiv:2411.16964*, 2024.
- [24] Chuang Gan, Hang Zhao, Peihao Chen, David Cox, and Antonio Torralba. Self-supervised moving vehicle tracking with stereo sound. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7053–7062, 2019.
- [25] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 776–780. IEEE, 2017.
- [26] Shiry Ginosar, Amir Bar, Gefen Kohavi, Caroline Chan, Andrew Owens, and Jitendra Malik. Learning individual styles of conversational gesture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3497–3506, 2019.
- [27] 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*, pages 9942–9952, 2023.
- [28] Peter Grosche, Meinard Müller, and Frank Kurth. Cyclic tempogram—a mid-level tempo representation for musicsignals. In *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 5522–5525. IEEE, 2010.
- [29] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2021–2029, 2020.
- [30] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5152–5161, 2022.
- [31] Chuan Guo, Yuxuan Mu, Muhammad Gohar Javed, Sen Wang, and Li Cheng. Momask: Generative masked modeling of 3d human motions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1900–1910, 2024.
- [32] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06)*, pages 1735–1742. IEEE, 2006.
- [33] Mohamed Hassan, Duygu Ceylan, Ruben Villegas, Jun Saito, Jimei Yang, Yi Zhou, and Michael J Black. Stochastic scene-aware motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11374–11384, 2021.
- [34] Tairan He, Jiawei Gao, Wenli Xiao, Yuanhang Zhang, Zi Wang, Jiashun Wang, Zhengyi Luo, Guanqi He, Nikhil Sobanbab, Chaoyi Pan, et al. Asap: Aligning simulation and real-world physics for learning agile humanoid whole-body skills. *arXiv preprint arXiv:2502.01143*, 2025.
- [35] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [36] Daniel Holden, Taku Komura, and Jun Saito. Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017.
- [37] Chao Huang, Dejan Marković, Chenliang Xu, and Alexander Richard. Modeling and driving human body soundfields through acoustic primitives. In *European Conference on Computer Vision*, pages 1–17. Springer, 2024.

- [38] Yinghao Huang, Leo Ho, Dafei Qin, Mingyi Shi, and Taku Komura. Interact: Capture and modelling of realistic, expressive and interactive activities between two persons in daily scenarios. *arXiv preprint arXiv:2405.11690*, 2024.
- [39] Yiming Huang, Weilin Wan, Yue Yang, Chris Callison-Burch, Mark Yatskar, and Lingjie Liu. Como: Controllable motion generation through language guided pose code editing. In *European Conference on Computer Vision*, pages 180–196. Springer, 2024.
- [40] Yiming Huang, Zhiyang Dou, and Lingjie Liu. Modskill: Physical character skill modularization. *arXiv preprint arXiv:2502.14140*, 2025.
- [41] 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*, 36:20067–20079, 2023.
- [42] Jaeyeon Kim, Heeseung Yun, and Gunhee Kim. Visage: Video-to-spatial audio generation. *arXiv preprint arXiv:2506.12199*, 2025.
- [43] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [44] Zitong Lan, Chenhao Zheng, Zhiwei Zheng, and Mingmin Zhao. Acoustic volume rendering for neural impulse response fields. *arXiv preprint arXiv:2411.06307*, 2024.
- [45] Zitong Lan, Yiduo Hao, and Mingmin Zhao. Resounding acoustic fields with reciprocity, 2025.
- [46] Nhat Le, Tuong Do, Khoa Do, Hien Nguyen, Erman Tjiputra, Quang D Tran, and Anh Nguyen. Controllable group choreography using contrastive diffusion. *ACM Transactions on Graphics (TOG)*, 42(6): 1–14, 2023.
- [47] 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*, pages 1272–1279, 2022.
- [48] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13401–13412, 2021.
- [49] Ronghui Li, Junfan Zhao, Yachao Zhang, Mingyang Su, Zeping Ren, Han Zhang, Yansong Tang, and Xiu Li. Finedance: A fine-grained choreography dataset for 3d full body dance generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10234–10243, 2023.
- [50] 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*, pages 1524–1534, 2024.
- [51] Zhouyingcheng Liao, Mingyuan Zhang, Wenjia Wang, Lei Yang, and Taku Komura. Rmd: A simple baseline for more general human motion generation via training-free retrieval-augmented motion diffuse. *arXiv preprint arXiv:2412.04343*, 2024.
- [52] Jing Lin, Ailing Zeng, Shunlin Lu, Yuanhao Cai, Ruimao Zhang, Haoqian Wang, and Lei Zhang. Motion-x: A large-scale 3d expressive whole-body human motion dataset. *Advances in Neural Information Processing Systems*, 36:25268–25280, 2023.
- [53] Jinpeng Liu, Wenxun Dai, Chunyu Wang, Yiji Cheng, Yansong Tang, and Xin Tong. Plan, posture and go: Towards open-world text-to-motion generation. *arXiv preprint arXiv:2312.14828*, 2023.
- [54] Xian Liu, Qianyi Wu, Hang Zhou, Yinghao Xu, Rui Qian, Xinyi Lin, Xiaowei Zhou, Wayne Wu, Bo Dai, and Bolei Zhou. Learning hierarchical cross-modal association for co-speech gesture generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10462–10472, 2022.
- [55] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [56] Vicon Motion Systems Ltd. Vicon motion capture systems, 2025. Accessed: Feb 5, 2025.
- [57] Shunlin Lu, Ling-Hao Chen, Ailing Zeng, Jing Lin, Ruimao Zhang, Lei Zhang, and Heung-Yeung Shum. Humantomato: Text-aligned whole-body motion generation. *arXiv preprint arXiv:2310.12978*, 2023.
- [58] Shunlin Lu, Jingbo Wang, Zeyu Lu, Ling-Hao Chen, Wenxun Dai, Junting Dong, Zhiyang Dou, Bo Dai, and Ruimao Zhang. Scamo: Exploring the scaling law in autoregressive motion generation model. *arXiv preprint arXiv:2412.14559*, 2024.
- [59] Zhengyi Luo, Jinkun Cao, Kris Kitani, Weipeng Xu, et al. Perpetual humanoid control for real-time simulated avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10895–10904, 2023.

- [60] Zhenye Luo, Min Ren, Xuecai Hu, Yongzhen Huang, and Li Yao. Popdg: Popular 3d dance generation with popdanceset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26984–26993, 2024.
- [61] Brian McFee, Colin Raffel, Dawen Liang, Daniel PW Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. librosa: Audio and music signal analysis in python. *SciPy*, 2015:18–24, 2015.
- [62] Liang Pan, Zeshi Yang, Zhiyang Dou, Wenjia Wang, Buzhen Huang, Bo Dai, Taku Komura, and Jingbo Wang. Tokenhsi: Unified synthesis of physical human-scene interactions through task tokenization. *arXiv preprint arXiv:2503.19901*, 2025.
- [63] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10975–10985, 2019.
- [64] Huaijin Pi, Zhi Cen, Zhiyang Dou, and Taku Komura. Coda: Coordinated diffusion noise optimization for whole-body manipulation of articulated objects. *arXiv preprint arXiv:2505.21437*, 2025.
- [65] Shenhan Qian, Zhi Tu, Yihao Zhi, Wen Liu, and Shenghua Gao. Speech drives templates: Co-speech gesture synthesis with learned templates. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11077–11086, 2021.
- [66] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- [67] Christian Schölkhuber and Anssi Klapuri. Constant-q transform toolbox for music processing. In *7th sound and music computing conference, Barcelona, Spain*, pages 3–64. SMC, 2010.
- [68] Mingyi Shi, Dafei Qin, Leo Ho, Zhouyingcheng Liao, Yinghao Huang, Junichi Yamagishi, and Taku Komura. It takes two: Real-time co-speech two-person’s interaction generation via reactive auto-regressive diffusion model. *arXiv preprint arXiv:2412.02419*, 2024.
- [69] Yi Shi, Jingbo Wang, Xuekun Jiang, Bingkun Lin, Bo Dai, and Xue Bin Peng. Interactive character control with auto-regressive motion diffusion models. *ACM Transactions on Graphics (TOG)*, 43(4): 1–14, 2024.
- [70] Li Siyao, Weijiang Yu, Tianpei Gu, Chunze Lin, Quan Wang, Chen Qian, Chen Change Loy, and Ziwei Liu. Bailando: 3d dance generation by actor-critic gpt with choreographic memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11050–11059, 2022.
- [71] 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. *arXiv preprint arXiv:2403.18811*, 2024.
- [72] Sebastian Starke, Ian Mason, and Taku Komura. Deepphase: Periodic autoencoders for learning motion phase manifolds. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022.
- [73] Jiarui Sun and Girish Chowdhary. Comusion: Towards consistent stochastic human motion prediction via motion diffusion. In *European Conference on Computer Vision*, pages 18–36. Springer, 2024.
- [74] Peiwen Sun, Sitong Cheng, Xiangtai Li, Zhen Ye, Huadai Liu, Honggang Zhang, Wei Xue, and Yike Guo. Both ears wide open: Towards language-driven spatial audio generation. *arXiv preprint arXiv:2410.10676*, 2024.
- [75] Taoran Tang, Jia Jia, and Hanyang Mao. Dance with melody: An lstm-autoencoder approach to music-oriented dance synthesis. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 1598–1606, 2018.
- [76] Chen Tessler, Yunrong Guo, Ofir Nabati, Gal Chechik, and Xue Bin Peng. Maskedmimic: Unified physics-based character control through masked motion inpainting. *ACM Transactions on Graphics (TOG)*, 43(6):1–21, 2024.
- [77] 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*, 2023.
- [78] Jonathan Tseng, Rodrigo Castellon, and Karen Liu. Edge: Editable dance generation from music. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 448–458, 2023.
- [79] Weilin Wan, Zhiyang Dou, Taku Komura, Wenping Wang, Dinesh Jayaraman, and Lingjie Liu. Tlcontrol: Trajectory and language control for human motion synthesis. *arXiv preprint arXiv:2311.17135*, 2023.
- [80] Jingbo Wang, Sijie Yan, Bo Dai, and Dahua Lin. Scene-aware generative network for human motion synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12206–12215, 2021.

- [81] Jingbo Wang, Yu Rong, Jingyuan Liu, Sijie Yan, Dahua Lin, and Bo Dai. Towards diverse and natural scene-aware 3d human motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20460–20469, 2022.
- [82] Wenjia Wang, Liang Pan, Zhiyang Dou, Jidong Mei, Zhouyingcheng Liao, Yuke Lou, Yifan Wu, Lei Yang, Jingbo Wang, and Taku Komura. Sims: Simulating stylized human-scene interactions with retrieval-augmented script generation. *arXiv preprint arXiv:2411.19921*, 2024.
- [83] Zan Wang, Yixin Chen, Baoxiong Jia, Puhao Li, Jinlu Zhang, Jingze Zhang, Tengyu Liu, Yixin Zhu, Wei Liang, and Siyuan Huang. Move as you say interact as you can: Language-guided human motion generation with scene affordance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 433–444, 2024.
- [84] Yiming Xie, Varun Jampani, Lei Zhong, Deqing Sun, and Huaizu Jiang. Omnicontrol: Control any joint at any time for human motion generation. *arXiv preprint arXiv:2310.08580*, 2023.
- [85] Liang Xu, Ziyang Song, Dongliang Wang, Jing Su, Zhicheng Fang, Chenjing Ding, Weihao Gan, Yichao Yan, Xin Jin, Xiaokang Yang, et al. Actformer: A gan-based transformer towards general action-conditioned 3d human motion generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2228–2238, 2023.
- [86] Sirui Xu, Hung Yu Ling, Yu-Xiong Wang, and Liang-Yan Gui. Intermimic: Towards universal whole-body control for physics-based human-object interactions. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 12266–12277, 2025.
- [87] Xudong Xu, Dejan Markovic, Jacob Sandakly, Todd Keebler, Steven Krenn, and Alexander Richard. Sounding bodies: modeling 3d spatial sound of humans using body pose and audio. *Advances in Neural Information Processing Systems*, 36:44740–44752, 2023.
- [88] Han Yang, Kun Su, Yutong Zhang, Jiaben Chen, Kaizhi Qian, Gaowen Liu, and Chuang Gan. Unimomo: Unified text, music and motion generation. *arXiv preprint arXiv:2410.04534*, 2024.
- [89] Youngwoo Yoon, Bok Cha, Joo-Haeng Lee, Minsu Jang, Jaeyeon Lee, Jaehong Kim, and Geehyuk Lee. Speech gesture generation from the trimodal context of text, audio, and speaker identity. *ACM Transactions on Graphics (TOG)*, 39(6):1–16, 2020.
- [90] Ye Yuan and Kris Kitani. Diverse trajectory forecasting with determinantal point processes. *arXiv preprint arXiv:1907.04967*, 2019.
- [91] Canyu Zhang, Youbao Tang, Ning Zhang, Rueli-Sung Lin, Mei Han, Jing Xiao, and Song Wang. Bidirectional autoregressive diffusion model for dance generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 687–696, 2024.
- [92] He Zhang, Sebastian Starke, Taku Komura, and Jun Saito. Mode-adaptive neural networks for quadruped motion control. *ACM Transactions on Graphics (TOG)*, 37(4):1–11, 2018.
- [93] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Yong Zhang, Hongwei Zhao, Hongtao Lu, Xi Shen, and Ying Shan. Generating human motion from textual descriptions with discrete representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14730–14740, 2023.
- [94] Mingyuan Zhang, Daisheng Jin, Chenyang Gu, Fangzhou Hong, Zhongang Cai, Jingfang Huang, Chongzhi Zhang, Xinying Guo, Lei Yang, Ying He, et al. Large motion model for unified multi-modal motion generation. In *European Conference on Computer Vision*, pages 397–421. Springer, 2024.
- [95] Yufei Zhang, Jeffrey O Kephart, Zijun Cui, and Qiang Ji. Physpt: Physics-aware pretrained transformer for estimating human dynamics from monocular videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2305–2317, 2024.
- [96] Yufei Zhang, Jeffrey O Kephart, and Qiang Ji. Incorporating physics principles for precise human motion prediction. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6164–6174, 2024.
- [97] Zeyi Zhang, Tenglong Ao, Yuyao Zhang, Qingzhe Gao, Chuan Lin, Baoquan Chen, and Libin Liu. Semantic gesticulator: Semantics-aware co-speech gesture synthesis. *ACM Transactions on Graphics (TOG)*, 43(4):1–17, 2024.
- [98] Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba. The sound of pixels. In *Proceedings of the European conference on computer vision (ECCV)*, pages 570–586, 2018.
- [99] Zhisheng Zheng, Puyuan Peng, Ziyang Ma, Xie Chen, Eunsol Choi, and David Harwath. Bat: Learning to reason about spatial sounds with large language models. *arXiv preprint arXiv:2402.01591*, 2024.

- [100] Wenyang Zhou, Zhiyang Dou, Zeyu Cao, Zhouyingcheng Liao, Jingbo Wang, Wenjia Wang, Yuan Liu, Taku Komura, Wenping Wang, and Lingjie Liu. Emdm: Efficient motion diffusion model for fast and high-quality motion generation. In *European Conference on Computer Vision*, pages 18–38. Springer, 2025.
- [101] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5745–5753, 2019.
- [102] Lingting Zhu, Xian Liu, Xuanyu Liu, Rui Qian, Ziwei Liu, and Lequan Yu. Taming diffusion models for audio-driven co-speech gesture generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10544–10553, 2023.
- [103] Wentao Zhu, Xiaoxuan Ma, Dongwoo Ro, Hai Ci, Jinlu Zhang, Jiaxin Shi, Feng Gao, Qi Tian, and Yizhou Wang. Human motion generation: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [104] Wenlin Zhuang, Congyi Wang, Jinxiang Chai, Yangang Wang, Ming Shao, and Siyu Xia. Music2dance: Dancenet for music-driven dance generation. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)*, 18(2):1–21, 2022.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Claimed in abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Stated in Sec. 6

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [No]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Stated in Sec. 5

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The dataset and the codes will be released after acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Stated in Sec. 5

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Stated in Sec. 5

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Stated in Sec. 5, with one GPU (Nvidia RTX 4090 GPU).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Yes.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: Stated in Sec. 6.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to

generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We follow the original license of previous work.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [Yes]

Justification: Stated in Sec. 5.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [Yes]

Justification: Research with human subjects has been known and approved by the authors' university and supervisor.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [Yes]

Justification: Editing (e.g., grammar, spelling, word choice)

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.