# SMooDi: Stylized Motion Diffusion Model

Lei Zhong<sup>1</sup>, Yiming Xie<sup>1</sup>, Varun Jampani<sup>2</sup>, Deqing Sun<sup>3</sup>, and Huaizu Jiang<sup>1</sup>

<sup>1</sup> Northeastern University <sup>2</sup> Stability AI <sup>3</sup> Google Research {le.zhong,xie.yim,h.jiang}@northeastern.edu varunjampani@gmail.com deqingsun@google.com

Abstract. We introduce a novel Stylized Motion Diffusion model, dubbed SMooDi, to generate stylized motion driven by content texts and style motion sequences. Unlike existing methods that either generate motion of various content or transfer style from one sequence to another, SMooDi can rapidly generate motion across a broad range of content and diverse styles. To this end, we tailor a pre-trained text-to-motion model for stylization. Specifically, we propose style guidance to ensure that the generated motion closely matches the reference style, alongside a lightweight style adaptor that directs the motion towards the desired style while ensuring realism. Experiments across various applications demonstrate that our proposed framework outperforms existing methods in stylized motion generation. Project Page: https://neu-vi.github.io/SMooDi/

Keywords: Motion synthesis  $\cdot$  Diffusion model  $\cdot$  Stylized motion

## 1 Introduction

We address the problem of generating stylized motion from a content text and a style motion sequence, as shown in Fig. 1. Human motion can typically be characterized by two components: content and style. Motion content represents the nature of a movement, such as walking and waving, and motion style reflects individual characteristics, such as personality traits (*e.g.*, old, childlike) and emotions (*e.g.*, happy, angry). Traditional pipelines create stylized motions via motion capture from actors, and are both labor-intensive and time-consuming. Therefore, decades of research have focused on developing automatic methods to assist stylized motion creation [1, 11, 19].

Motion style transfer [1, 45] is a practical and popular approach for the creation of stylized motion. It transfers the style from an existing style motion sequence to another existing content motion sequence. However, when a broad array of motion needs to be stylized, the pipeline may be inefficient – it would first require the collection of a large number of content motion sequences and then apply a motion style transfer method to process each sequence independently. Moreover, motion sequences are not always readily available, especially



Fig. 1: SMooDi can generate realistic, stylized human motions given a content text a style motion sequence. It also accepts a motion sequence as content input. Darker color indicates later frames in the sequence. To better showcase the stylized motion generation, we place the style label for the each of the style motion sequence. Note that such style labels are not used as model input and shown here for visualization purpose only. (Best viewed in color.)

for some customized content, such as running along a specific trajectory. They may still need to be created first by actors or animators for stylization.

Recent advances of human motion generation with diffusion models [10, 17] have shown impressive results of creating diverse and realistic human motions. But most efforts have concentrated on efficiently and accurately translating textual prompts into human motions, focusing on the *content* only [6,31,46]. Integrating the *style* condition to generate stylized motions remains under-explored.

Combining these two lines of research is a straightforward approach to tackle stylized motion generation, where a motion style transfer method [1,45] can be applied to each motion sequence generated by a text-driven motion diffusion model [6,31,46]. However, in addition to the aforementioned inefficiency issue, this approach has two more limitations. First, error may accumulate across the pipeline. As, motion style transfer methods are usually trained with high-quality real-world motion sequences, we empirically observe that their performance may significantly degrade for imperfect motions produced by text-to-motion techniques. Second, existing motion style transfer methods rely on specialized style datasets [1,27,52] with limited motion content, which restricts their applications to motion diffusion models.

In this paper, we present a novel stylized motion diffusion model, dubbed SMooDi, that customizes a pre-trained text-to-motion model for stylization. Built upon the pre-trained motion latent diffusion model (MLD) [6], SMooDi inherits MLD's ability to generate diverse motion content. At the same time, SMooDi can generate motions in a variety of styles according to different style reference conditions, as shown in Fig. 1. Our main novelty is the style modulation module, which consists of a style adaptor and a style guidance module. First of all, drawing inspiration from controllable image generation [61], the style adaptor is designed to predict residual features conditioned on style reference motion sequence within each attention layer of MLD. It is useful for incorporating the style condition while ensuring the realism of the generated motion. Second, we design both classifier-free and classifier-based style guidance to more precisely control the stylized motion generation. Specifically, the classifier-free style and content guidance are linearly combined, where we can easily strike a balance between preserving content and reflecting style within the generated motion. At the same time, we design an additional classifier-based style guidance mechanism. It is an analytic function quantifying the disparity between the generated motion and the style reference motion in a style-centric embedding space, whose gradients are subsequently employed to guide the generated motion closer to the intended style. Our style adaptor and guidance module are designed to be complementary, which lead to high-quality stylized motion generation. The two modules are jointly optimized in a feature space instead of sequence-wise separate stitching, thereby avoiding the error accumulation issue.

Although our approach is primarily designed for stylized motion generation driven by content text, we can utilize DDIM-Inversion [43] to identify the noisy latent corresponding to the content motion sequence. Following the same procedure as for text-driven content, SMooDi is capable of facilitating stylized motion generation based on content motion sequences. In other words, motion style transfer is a downstream application of our approach should it be desired in practice, *e.g.*, to stylize the already created motion sequences.

Experiments on the HumanML3D [16] and 100STYLE [27] datasets demonstrate that SMooDi surpasses other baseline models in generating stylized motion driven by content text, excelling in both content preservation and style reflection. More importantly, unlike previous methods that require individual fine-tuning for each style [12, 30, 54], SMooDi successfully integrates diverse content from the HumanML3D dataset and various styles from the 100STYLE dataset into a single model without requiring additional tuning during inference.

To summarize, our contributions are: (1) To our knowledge, SMooDi is the first approach that adapts a pre-trained text-to-motion model to generate diverse stylized motion. (2) We introduce a novel style modulation module that utilizes a stylized adaptor and a style classifier guidance to enable stylized motion generation while ensuring style reflection, content preservation, and realism. (3) Experiments demonstrate that SMooDi not only sets a new state of the art in stylized motion generation driven by content text but also achieves performance comparable to state-of-the-art methods in motion style transfer.

## 2 Related Work

### 2.1 Human Motion Generation

Human motion generation has attracted great attention [4, 5, 7, 9, 14, 16, 20, 32, 34–36, 47, 50, 51, 57, 58, 63]. Inspired by the impressive performance of diffusion models in image generation, a lot of works [6,8,13,18,22,23,25,29,32,33,39,41,46, 48,53,56,60,62,64] utilize diffusion models to generate human motion. MDM [46] facilitates high-quality generation and versatile conditioning, providing a solid baseline for novel motion generation tasks. MLD [6] minimizes computational overhead during both training and inference by establishing the diffusion process within the latent space. Driven by the efficacy of diffusion models for control and conditioning, several studies have leveraged pre-trained motion diffusion models to generate long-sequence motions [41], enable human-object interactions [29], and control the joint trajectory of generated motions [23,53]. However, there is no work exploring how to leverage pre-trained motion diffusion models to generate diverse stylized motion. While some studies [2,3,37] have enabled stylized motion generation in their diffusion pipeline, their methods are trained from scratch, and the supported styles are restricted by their motion content dataset. It is challenging for them to simultaneously support diverse motion content and style. In this work, we build upon a pre-trained motion diffusion model, MLD, and explore how to fine-tune it on a larger motion style dataset, 100STYLE, to learn diverse motion styles while retaining the ability to support motion generation across a wide range of content.

### 2.2 Motion Style Transfer

Recently, motion style transfer has seen quality enhancements through the adoption of various advanced neural architectures and generative models, such as graph neural networks [28], time-series models [27, 45], normalizing flows [49], and diffusion models [3,37]. Specifically, Aberman et al. [1] designed a two-branch generative adversarial network to disentangle motion style from content and facilitate their re-composition. Their approach effectively breaks the constraint of requiring a paired motion dataset. Motion Puzzle [19] realizes a framework that can control the style of individual body parts. Above methods extract both content and style features from the motion sequence. Moreover, Guo et al. [15] leverage the latent space of pre-trained motion models to enhance the extraction and infusion of motion content and style. However, a major limitation of these models is their reliance on specialized style datasets [1, 52] with limited motion content, which restricts their applications. In this work, we customize a pre-trained text-to-motion model for stylization, thus inheriting its ability to generate diverse motion content.

## 3 Stylized Motion Diffusion Model

In this section, we introduce our proposed SMooDi for incorporating style conditions from a style motion sequence into a content-oriented pre-trained motion



Fig. 2: Overview of SMooDi. Our model generates stylized human motions from content text and a style motion sequence. At the denoising step t, our model takes the content text  $\mathbf{c}$ , style motion  $\mathbf{s}$ , and noisy latent  $\mathbf{z}_t$  as input and predicts  $\epsilon_t$ , which is then transferred to  $\mathbf{z}_{t-1}$ . This denoising step is repeated T times to obtain the noise-free motion latent  $\mathbf{z}_0$ , which is fed into a motion decoder D to produce the stylized motion.

diffusion model (MLD [6]). Fig. 2 presents an overview of SMooDi. Following the setting in MLD [6], we place the diffusion process in the motion latent space. Let  $\epsilon_{\theta}$  denote the latent denoiser (a UNet parameterized by  $\theta$ ), and  $\{\boldsymbol{z}_t\}_{t=0}^T$  denote the sequence of noisy latents, where  $\boldsymbol{z}_T$  is a Gaussian noise. Given a content prompt  $\mathbf{c}$  and a style prompt  $\mathbf{s}$ , we define  $\epsilon_t = \epsilon_{\theta}(\boldsymbol{z}_t, t, \mathbf{c}, \mathbf{s})$  for the denoising at step t ( $0 < t \leq T$ ). A cleaner noisy latent  $\boldsymbol{z}_{t-1}$  can be obtained by subtracting  $\epsilon_t$  from  $\boldsymbol{z}_t$ . The denoising step is repeated T iterations until a clean latent  $\boldsymbol{z}_0$  is obtained. It can then be decoded by a motion decoder  $\mathbf{D}$  into a realistic motion sequence  $\boldsymbol{x} \in \mathbb{R}^{N \times H}$  that accurately reflects both the content and style conditions. Here, N represents the length of the motion sequence, and H is the dimension of human motion representations. We employ the same motion representations as in HumanML3D [16], where H = 263.

As shown in the Fig. 2, the content prompt is a text description, and the style prompt is provided by a reference style motion sequence  $\mathbf{s} \in \mathbb{R}^{N \times H}$ . In this section, we focus on using a text description as the content prompt  $\mathbf{c}$  to explain our proposed stylized motion diffusion model. By employing the DDIM-Inversion [43] to identify the noisy latent corresponding to a motion content sequence, we can effectively use motion sequences as content prompts to generate stylized motions. In other words, motion style transfer is a downstream application of our proposed approach.

Our proposed stylization module consists of a style adaptor and a style guidance module. We will explain them separately in the rest of this section.

### 3.1 Style Adaptor

Although LoRA has been successfully used to incorporate "style" into the models in image domain [21,42], they typically require training a separate LoRA for each style. In contrast, we focus on fine-tuning the model just once to adapt to



Fig. 3: Detailed illustration of our proposed style adaptor. The style adaptor is connected to the motion diffusion model via zero linear layer. The output of the style adaptor from each Transformer encoder is added to the motion diffusion model to steer the predicted noise towards the target style.

various motion styles, where adapting ControlNet [61] is more suitable. Therefore, we design a content-aware style adaptor based on ControlNet. This adaptor incorporates the motion style condition into the pre-trained MLD [6].

Instead of learning to disentangle motion style from content from scratch on a large motion dataset, we redirect our focus towards capturing the motion style while ensuring the preservation of diverse motion content within the pre-trained MLD framework. Specifically, it consists of a trainable copy of the Transformer Encoder from the latent diffusion model in MLD. The architecture of the style adaptor is illustrated in Fig. 3. An independent style encoder is utilized to extract the style embedding from the style motion sequence  $\mathbf{s}$ . The style adaptor takes the same content prompt  $\mathbf{c}$ , the noised latent  $\mathbf{z}_t$  and timestep t as in MLD, and the extracted style embedding. Each Transformer layer in the original latent diffusion model and the style adaptor is connected via a linear layer, with both weight and bias initially set to zeros. As training progresses, the style adaptor learns the style constraints and gradually applies the learned feature corrections to the corresponding layers in the latent diffusion model, thereby implicitly steering the output towards the desired style.

### 3.2 Style Guidance

The style adaptor alone may not be sufficient to successfully incorporate the style condition. We further leverage both the classifier-free and classifier guidance to further enhance the stylization of a motion diffusion model. The combination of two types of guidance effectively ensures that generated motion meets multiple constraints while maintaining realism, complementing the style adaptor. **Classifier-free Style Guidance.** With the introduction of an extra style condition, we can divide the conditioned classifier-free guidance into two parts.

$$\epsilon_{\theta}(\boldsymbol{z}_{t}, t, \mathbf{c}, \mathbf{s}) = \epsilon_{\theta}(\boldsymbol{z}_{t}, t, \emptyset, \emptyset) + \underbrace{w_{c}(\epsilon_{\theta}(\boldsymbol{z}_{t}, t, \mathbf{c}, \emptyset) - \epsilon_{\theta}(\boldsymbol{z}_{t}, t, \emptyset, \emptyset))}_{\text{Classifier-free Content Guidance}} + \underbrace{w_{s}(\epsilon_{\theta}(\boldsymbol{z}_{t}, t, \mathbf{c}, \mathbf{s}) - \epsilon_{\theta}(\boldsymbol{z}_{t}, t, \mathbf{c}, \emptyset))}_{\text{Classifier-free Style Guidance}}, (1)$$

where  $w_c$  and  $w_s$  represent the strengths of the classifier-free guidance for the condition  $\mathbf{c}$  and  $\mathbf{s}$ , respectively. We slightly abuse the notations here by using  $\emptyset$  to denote a condition is not used. The classifier-free content guidance is the same as in MLD, which can facilitate the text-to-motion generation process in combination with the first term's unconditioned guidance. Our proposed classifier-free style guidance works in a similar way. Note that  $\epsilon_{\theta}(\mathbf{z}_t, t, \mathbf{c}, \mathbf{s})$  is MLD model with the style adaptor incorporated introduced in the previous section, which takes both a textual prompt and style motion sequence as input. By contrasting the text-driven denoising output with and without the style condition, it can highlight the effectiveness of the style input  $\mathbf{s}$  and facilitate the generation of stylized motion driven by content text. Our insight here is that by dividing the conditioned guidance into content and style components separately, we can easily strike a balance between preserving content and reflecting style within the generated motion.

To better understand the classifier-free content and style guidance, we visualize each of them through decoding denoised latent  $z_0$  into the motion space. As illustrated in Fig.4(a), the content guidance ensures the motion generation is faithful to the textual prompt, while the style guidance, as shown in Fig.4(b), emphasizes style-related characteristics in the output. Combining both forms of guidance results in a stylized motion that adheres to both content and style conditions, as illustrated in Fig. 4(c).

**Classifier-based Style Guidance.** To further improve the stylization of a motion diffusion model, we adopt the classifier guidance [10,59] to provide stronger guidance to the generated motion towards the desired style. The core of our classifier-based style guidance is a novel analytic function  $G(\mathbf{z}_t, t, \mathbf{s})$ , which calculates the  $L_1$  distance between the style embedding of the generated clean motion  $\hat{\mathbf{x}}_0$  at denoising step t and the reference style motions s. The function's gradient is utilized to steer the generated motion towards the desired style.

$$\epsilon_{\theta}(\boldsymbol{z}_t, t, \mathbf{c}, \mathbf{s}) = \epsilon_{\theta}(\boldsymbol{z}_t, t, \mathbf{c}, \mathbf{s}) + \tau \nabla_{\boldsymbol{z}_t} G(\boldsymbol{z}_t, t, \mathbf{s}),$$
  
$$G(\boldsymbol{z}_t, t, \mathbf{s}) = |f(\hat{\boldsymbol{x}}_0) - f(\mathbf{s})|,$$
(2)

where  $\tau$  adjusts the strength of reference-style guidance and f denotes the style feature extractor. The generated motion  $\hat{x}_0$  is obtained by first converting the denoising output latent  $z_t$  into the predicted clean latent as shown below:

$$\hat{\boldsymbol{z}}_0 = \frac{\boldsymbol{z}_t - \sqrt{1 - \alpha_t} \varepsilon_\theta(\boldsymbol{z}_t, t, \boldsymbol{c}, \boldsymbol{s})}{\sqrt{\alpha_t}},\tag{3}$$

where  $\alpha_t$  denotes the pre-defined noise scale in the forward process of the diffusion model. The predicted clean latent  $\hat{z}_0$  is then input into the motion decoder **D** to obtain the generated motion.

We obtain the style feature extractor by training a style classifier on the 100STYLE dataset [27] and removing its final layer. We refer the readers to more details in the supplementary materials. The training of the style feature extractor with ground-truth style labels for supervision enables it to effectively



**Fig. 4: Visual illustrations of the classifier-free and clasifier-based style guidance.** (a) and (b) respectively show the classifier-free content and style guidance; (c) displays the initial stylized motion resulting from the combination of (a) and (b); (d) illustrates the refined stylized motion modified by the classifier-based style guidance.

capture style-related features. Therefore, style classifier guidance can provide more guidance to the stylized motion generation.

**Combination of the Two Style Guidance.** The classifier-free and classifierbased style guidance are designed to complement each other, each playing a vital role in accurately reflecting the target style in the generated motions. As illustrated in Fig. 4, the desired style motion is "arms open wide to the sides like an airplane". The classifier-free style guidance (Fig. 4(b)) can capture stylerelated characteristic in a reasonably accurate manner. When combined with the classifier-free content guidance, it depicts the desired style (Fig.4(c)). Refined further by the classifier-based style guidance, the stylization is more authentic, where the person's harms are more open (Fig.4(d)). In addition to such visual results, quantitative ablation studies also verify the effectiveness of our proposed both classifier-free and classifer-based style guidance.

At the same time, although classifier-based style guidance offers precise style control, its effectiveness may be compromised when the content text significantly diverges from locomotion-related movements. This is because the style feature function is trained solely on the 100STYLE dataset, which contains only such movements. Over-reliance on style classifier guidance risks producing motions that fail to execute the desired actions, leading to unrealistic and physically implausible movements. Therefore, we leverage a content-aware style adaptor that establishes the fundamental style direction, while style classifier guidance refines this base for a more precise outcome. The effectiveness of this design is verified in our ablation studies.

#### 3.3 Learning Scheme

Following [61], a straightforward approach to train SMooDi is to freeze the parameters of MLD and solely train the style adaptor on the 100STYLE dataset using the following loss function:

$$\mathcal{L}_{std} = \mathbb{E}_{\epsilon, \mathbf{z}} \left[ \left\| \epsilon_{\theta}(\mathbf{z}_t, t, \mathbf{c}, \mathbf{s}) - \epsilon \right\|_2^2 \right], \tag{4}$$

where  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$  represents the ground-truth noise added to  $z_0$ . In our experiments, however, we found that this loss function alone leads to an issue of

"content-forgetting", where the model progressively looses the MLD's ability to generate motions with diverse contents. To address this issue, we design a content prior preservation loss  $L_{pr}$ . Specifically, we randomly sample motions from the HumanML3D dataset to compute a prior preservation loss when fine-tuning SMooDi on the 100STYLE dataset.

$$\mathcal{L}_{pr} = \mathbb{E}_{\epsilon', \mathbf{z}'} \left[ \left\| \epsilon_{\theta}(\mathbf{z}'_{t}, t, \mathbf{c}', \mathbf{s}') - \epsilon' \right\|_{2}^{2} \right],$$
(5)

where  $\mathbf{z}'_t, \mathbf{c}'$  and  $\mathbf{s}'$  represents the motion latent, content prompt and style motion sequence derived from the HumanML3D dataset.  $\epsilon'$  is the noise map added to  $\mathbf{z}'_0$ . A similar solution is used in DreamBooth [40] to solve the "language drift" problem, where images generated from the frozen pretrained image generation model are utilized to enforce a class-prior preservation loss during model finetuning. Our content preservation loss can effectively mitigate content forgetting while learning diverse motion styles from the 100STYLE dataset.

To further encourage the style adaptor to focus on motion style, while also ensuring that the latent diffusion model in MLD handles motion content well, we introduce an additional cycle prior-preservation loss, inspired by [19,55]. Specifically, we start this process by randomly sampling content text and motion style sequences from both the 100STYLE and HumanML3D datasets simultaneously. Then, we intermix the content text and motion style from these sequences with each other. Finally, we repeat this process to reconstruct the original motion sequences. The formula is expressed as follows:

$$\mathcal{L}_{cyc} = \mathbb{E}_{\boldsymbol{z},\boldsymbol{z}',\boldsymbol{\epsilon},\boldsymbol{\epsilon}'} \left[ \left\| \boldsymbol{\epsilon}_{\theta}(\boldsymbol{z}_{t}^{sh}, t, \mathbf{c}, \boldsymbol{s}^{hs}) + \boldsymbol{\epsilon}_{\theta}(\boldsymbol{z}_{t}^{hs}, t, \mathbf{c}', \boldsymbol{s}^{sh}) - \boldsymbol{\epsilon} - \boldsymbol{\epsilon}' \right\|_{2}^{2} \right], \quad (6)$$

where  $s^{hs}$  denotes the motion sequence created by merging content from the HumanML3D dataset with style from the 100STYLE dataset. Similarly,  $s^{sh}$  represents the sequence where content is sourced from the 100STYLE dataset and style from the HumanML3D dataset. The noised latent codes  $z_t^{sh}$  and  $z_t^{hs}$  correspond to  $s^{sh}$  and  $s^{hs}$ , respectively. Essentially, the cycle prior-preservation loss exchanges diverse content and style between two datasets, encouraging the content text to remain invariant in the generated motion under forward and backward translation. Overall, the training loss function of our framework is defined as follows:

$$\mathcal{L}_{all} = \mathcal{L}_{std} + \lambda_{pr} \mathcal{L}_{pr} + \lambda_{cyc} \mathcal{L}_{cyc} \tag{7}$$

where  $\lambda_{pr}$  and  $\lambda_{cyc}$  are hyperparameters. We refer readers to the pseudocode and illustration for training in the supplementary material for more details.

## 4 Experiments

We conduct experiments on both stylized text2motion and motion style transfer to demonstrate the effectiveness of our framework. Both tasks use a motion sequence as a style prompt, with the primary difference being their content prompt input: the former utilizes text, while the latter relies on motion sequences.

**Table 1:** Comparison with baseline methods on stylized motion generation driven by content text, using a combination of the 100STYLE (providing style) and HumanML3D datasets (providing content).

Method	FID↓F	Foot skating	MM Dist $\downarrow$	R-precision ↑	Diversity –	$\rightarrow$ SRA $\uparrow$
		ratio $\downarrow$		(Top-3)		
Ours	1.609	0.124	4.477	0.571	9.235	72.418
MLD+Motion Puzzle [19]	6.127	0.185	6.467	0.290	6.4762	63.769
MLD+Aberman et al. [1]	3.309	0.347	5.983	0.406	8.816	54.367
ChatGPT+MLD	0.614	0.131	4.313	0.605	8.836	4.819



Fig. 5: Qualitative comparisons of our approach and baseline methods on two stylized motion generation task.

**Datasets.** We utilize the HumanML3D dataset [16] as our motion content dataset and the 100STYLE dataset [27] as our motion style dataset. The HumanML3D dataset is the largest motion capture dataset, featuring text annotations and comprising 14,646 motions and 44,970 motion annotations. Following the processing approach outlined in [16], we preprocess the HumanML3D dataset to obtain consistent motion representations. On the other hand, the 100STYLE dataset [27], being the largest motion style dataset, comprises up to 1,125 minutes of motion sequences, showcasing a wide array of 100 diverse locomotion styles. Due to differences in skeletons between the 100STYLE dataset and HumanML3D, we retarget the motions from 100STYLE to match the HumanML3D (SMPL-H) skeleton. Following this alignment, we apply the same processing steps as used for the HumanML3D dataset to preprocess the 100STYLE dataset. Moreover, as the 100STYLE dataset lacks text descriptions, we leverage MotionGPT [20] to generate pseudo text descriptions for the motion sequences in the 100STYLE dataset.

**Evaluation metrics** are designed to assess three dimensions: Content Preservation, Style Reflection, and Realism. For content preservation and style reflection assessment, we employ metrics consistent with those used in [6]: motion-retrieval precision (R precision), Multi-modal Distance (MM Dist), Diversity, and Frechet Inception Distance (FID). Additionally, recognizing the common foot skating issues in kinetics-based motion generation methods, we incorporate the foot skating ratio metric proposed by [23] into our motion quality evaluation. For style reflection, we employ Style Recognition Accuracy (SRA) [19]. During evaluation, we randomly select a content text from the HumanML3D dataset and a motion style sequence from the 100STYLE dataset to generate the stylized motion. We then use a pre-trained style classifier to compute the SRA for the generated motion. It's noteworthy that some motion style labels in the 100STYLE dataset, like 'kick' and 'jump,' inherently convey motion content, which may conflict with the content text in HumanML3D dataset. To address this, we categorize the motion style labels into server groups following the approach by Kim et al. [24], Specifically excluding the 'ACT' group ensures that only motion style labels not conflicting with motion content are considered when computing the SRA metric. Further details about the evaluation metrics are provided in the supplementary material.

**Baselines.** For motion style transfer task, we compare our methods with two state-of-the-art methods, namely Motion Puzzle [19] and Aberman et al. [1]. To ensure a fair comparison, we train the compared methods under the same settings as ours, using a combined dataset comprising HumanML3D and 100STYLE. Due to the constraints of the multi-class discriminator in Aberman et al., which requires style labels, we adopt the training method outlined in Motion Puzzle to eliminate the need for style labels. For stylized text2motion task, we compare our method against baselines capable of generating stylized motion from content text and style motion sequences. The straightforward baselines involve applying motion style transfer methods to the motion sequences generated by the text2motion model. To align with our approach that uses a pre-trained motion diffusion model, the text2motion models in the baselines for stylized motion generation select MLD, and the motion style transfer methods are consistent with those used in the motion style transfer task. For the stylized text2motion task, we compare our method against two kinds of baselines capable of generating stylized motion from content text and style motion sequences. The first kind of baseline involves applying motion style transfer methods to the motion sequences generated by the text2motion model. To align with our approach that uses a pretrained motion diffusion model, the text2motion models select MLD [6], and the motion style transfer methods are consistent with those used in the motion style transfer task. The second kind of baseline involves using ChatGPT to merge style labels from 100STYLE with text from HumanML3D into a sentence. For example, given the content text 'a person walks.' and the style label 'old,' we obtain 'an elderly person walks.' This merged sentence is then fed to MLD.

#### 4.1 Comparison to Baseline Methods

**Quantitative and Qualitative** For the task of stylized text2motion, Table 1 reports the comparisons of our method with the three baseline methods.

As shown in the  $3^{rd}$  row of Table 1, ChatGPT+MLD only achieves around 5.29% in terms of SRA, indicating that MLD cannot enable stylized generation from text alone, even though it contains style descriptions. Notably, our method outperforms the two baselines that combine MLD with motion style transfer methods in all metrics.

Table 2: Comparison with baseline methods on motion style transfer.

(a) Evaluation on HumanML3D dataset

(b) Evaluation on Xia dataset

Method	Foot skatin ratio ↓	g FID $\downarrow$	$SRA\uparrow(\%)$	Method	Foot skating ratio↓	FID↓	$SRA\uparrow(\%)$	$CRA\uparrow(\%)$
Ours	<b>0.095</b>	<b>1.582</b>	65.147	Ours	0.0317	<b>4.663</b>	61.111	<b>45.555</b>
Motion Puzzle [19]	0.197	6.871	67.233	Motion Puzzle [19]	0.0316	5.360	67.778	25.556
(Aberman et al [1])	0.338	3.892	61.006	(Aberman et al [1])	<b>0.0260</b>	5.681	56.667	34.444

Specifically, our method performs better than MLD+Motion Puzzle in the SRA metric by 13.56% and significantly outperforms MLD+Aberman et al. in the FID metric by **51.38%** and **0.64%** in the R-precision metric. The first row of Fig. 5 validates our observation, where the motion generated by our method performs better in adhering to both content and style constraints than baseline methods. In contrast, MLP+Aberman et al. can successfully perform the action but fail to reflect the motion style in Fig.5(b), while Motion Puzzle can accurately reflect the motion style but struggles to effectively perform the action in Fig.5(a).

For the task of motion style transfer, since it does not take content text as input, text-motion related metrics such as MM Dist, R-precision, and Diversity are not applicable and thus are not reported. Part (a) of Table 2 presents a comparison between our method and the two baseline methods, using the HumanML3D dataset as the motion content source and drawing motion styles from the 100STYLE dataset. Our method delivers competitive results in the SRA metric and excels in the FID and foot skating ratio metrics. Specifically, we see a substantial **59.35%** improvement in the FID metric over Aberman et al. [1]. To more effectively compare the generalizability of different methods, we conduct experiments on the Xia dataset [52], a small, specific motion style dataset that was unseen by our and the baseline models during training. Because motion content labels are present in the Xia dataset, we report the Content Recognition Accuracy (CRA). Part (b) of Table 2 showcases the results. Our method maintains competitive performance in the SRA metric, with only a marginal 9.83% decrease in SRA compared to Motion Puzzle. On the contrary, our method exhibits a significant 32.26% increase in the CAR metric relative to Aberman et al., and a notable 78.26% enhancement over Motion Puzzle. Our method achieves a better balance between style reflection and content preservation. The second row of Fig. 5 validates this observation. It is worth noting that, unlike other motion style transfer methods, our method does not incorporate objectives for enabling stylized motion generation using motion content sequences. Through simple DDIM-Inversion and without any additional optimization or regularization, our method achieves performance comparable to existing motion style transfer methods.

**User Study.** Due to the highly subjective nature of stylized motion, we conduct User studies using pairwise comparisons to further evaluate our proposed method in the tasks of stylized motion generation and motion style transfer. We

Method	FID↓	Foot skating ratio↓	MM Dist $\downarrow$	$\begin{array}{c} \text{R-precision} \uparrow \\ \text{(Top-3)} \end{array}$	$Diversity \rightarrow$	$SRA(\%)\uparrow$
Ours (on all)	1.609	0.124	4.477	0.571	9.235	72.418
$w/o L_{cyc}$	2.046	0.136	4.465	0.569	8.869	64.866
$w/o L_{pr} + L_{cyc}$	5.996	0.166	6.098	0.335	7.456	81.841
w/o classifier-based	1.050	0.111	4.085	0.630	9.445	20.245
w/o adaptor	2.984	0.123	4.526	0.550	8.372	69.952

Table 3: Ablation Studies on HumanML3D Content and 100STYLE Styles.



Fig. 6: Visual comparisons of the ablation designs and our full model.

recruited 22 human subjects to participate in the study. In each test, participants are presented with two 4-second video clips synthesized by our method and one comparison method. They are then required to select their preferred clip while considering *Realism*, *Style Reflection*, and *Content Preservation* dimensions, respectively. As shown in Fig. 7, our method receives more user appreciation compared to two baselines across three dimensions in two tasks. Further user study details are provided in supplementary.

### 4.2 Ablation studies

To validate the effectiveness of our framework's design choices, we have conducted several ablation studies: the first assesses the impact of each loss function term, while the second evaluates the influence of the style adaptor and style guidance during sampling.

Loss Components.Firstly, we exclude the *cycle-prior* term in the loss function, denoted as  $w/o \ L_{cyc}$ . Comparing the results in the 1<sup>st</sup> and 2<sup>nd</sup> rows in Table 3, we observe that our full model outperforms in all content preservation and style reflection metrics. The motion generated by our approach can still perform the content adhering to the text description but performs worse in accurately reflecting the motion style, as reflected by the arms not being fully extended horizontally.

Since the cycle prior-preservation term is built upon the prior-preservation term, it is meaningless to exclude  $L_{pr}$  while retaining  $L_{cyc}$ . Therefore, we further exclude both the *prior-preservation* and *cycle-prior* term, denoted as  $w/o L_{pr} + L_{cyc}$  in the 3<sup>rd</sup> row. By comparing the results in the 2<sup>nd</sup> and 3<sup>rd</sup> rows of Table 3, we notice that while the number of SRAs is higher in the third row, other metrics show a significant decline. Specifically, in terms of the FID metric, performance deteriorates by more than **229%**. Indeed, without  $L_{pr}$ , the



Fig. 7: User Study on two stylized motion generation tasks.

model tends to lose the ability to translate content text into corresponding motion, a phenomenon named 'content-forgetting' as described in Sec. 3.3. Fig. 6(a) validates our observation, showing that the content in the generated motion significantly deviates from the text descriptions and closely resembles the style motion sequence.

Style Adaptor and Style Guidance. Initially, we compare our model to a variant without the classifier-based guidance, w/o classifier-based, to demonstrate its effectiveness. The 5<sup>th</sup> row of Table 3 presents the results. Consistent with the findings in [19], a reasonable trade-off between content preservation and style reflection is observed. Although classifier-based style guidance may slightly affect the content preservation metrics, it significantly boosts the model's performance in the SRA metric, yielding an impressive 208% improvement.Fig. 6(d) demonstrates that, without classifier-based style guidance, the generated motions can reflect the motion style, yet they still fall short of fully achieving the target style. Classifier-based style guidance can effectively bridge this gap.

Subsequently, as shown in Table 3, we evaluate a variant without the style adaptor, denoted as w/o adaptor (the last row). In cases where the SRA values are close, the style adaptor improved the FID metric by about 80.46%. Fig. 6(c) shows that the generated motion can greatly perform the 'walk' action while successfully reflecting the style, but fails to perform the 'sit' action. This indicates that the effectiveness of classifier-based style guidance diminishes when the content text deviates from locomotion-related movements. Relying solely on it may even adversely affect action performance.

## 5 Conclusion

In this work, we introduce the Stylized Motion Diffusion Model, a novel approach that leverages a pre-trained motion diffusion model to facilitate stylized motion generation driven by content text. By integrating a style adaptor and style classifier guidance, our method is capable of producing realistic human motions that accurately reflect both the content text descriptions and the desired motion style from motion sequences. Through detailed ablation studies, we have demonstrated the effectiveness of each component in our framework.

## References

- 1. Aberman, K., Weng, Y., Lischinski, D., Cohen-Or, D., Chen, B.: Unpaired motion style transfer from video to animation. TOG (2020)
- Alexanderson, S., Nagy, R., Beskow, J., Henter, G.E.: Listen, denoise, action! audiodriven motion synthesis with diffusion models. TOG (2023)
- 3. Ao, T., Zhang, Z., Liu, L.: Gesturediffuclip: Gesture diffusion model with clip latents. TOG (2023)
- Cen, Z., Pi, H., Peng, S., Shen, Z., Yang, M., Zhu, S., Bao, H., Zhou, X.: Generating human motion in 3d scenes from text descriptions. In: CVPR (2024)
- Chen, L.H., Lu, S., Zeng, A., Zhang, H., Wang, B., Zhang, R., Zhang, L.: Motionllm: Understanding human behaviors from human motions and videos. ArXiv (2024)
- Chen, X., Jiang, B., Liu, W., Huang, Z., Fu, B., Chen, T., Yu, G.: Executing your commands via motion diffusion in latent space. In: CVPR (2023)
- 7. Cohan, S., Tevet, G., Reda, D., Peng, X.B., van de Panne, M.: Flexible motion in-betweening with diffusion models. ArXiv (2024)
- Dabral, R., Mughal, M.H., Golyanik, V., Theobalt, C.: Mofusion: A framework for denoising-diffusion-based motion synthesis. In: CVPR (2023)
- 9. Dai, W., Chen, L.H., Wang, J., Liu, J., Dai, B., Tang, Y.: Motionlcm: Real-time controllable motion generation via latent consistency model. ArXiv (2024)
- Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. NeurIPS (2021)
- Du, H., Herrmann, E., Sprenger, J., Fischer, K., Slusallek, P.: Stylistic locomotion modeling and synthesis using variational generative models. In: Proceedings of the 12th ACM SIGGRAPH Conference on Motion, Interaction and Games (2019)
- Everaert, M.N., Bocchio, M., Arpa, S., Süsstrunk, S., Achanta, R.: Diffusion in style. In: ICCV (2023)
- Ghosh, A., Dabral, R., Golyanik, V., Theobalt, C., Slusallek, P.: Remos: Reactive 3d motion synthesis for two-person interactions. In: ArXiv (2023)
- Guo, C., Mu, Y., Javed, M.G., Wang, S., Cheng, L.: Momask: Generative masked modeling of 3d human motions. In: CVPR (2024)
- Guo, C., Mu, Y., Zuo, X., Dai, P., Yan, Y., Lu, J., Cheng, L.: Generative human motion stylization in latent space. ArXiv (2024)
- Guo, C., Zou, S., Zuo, X., Wang, S., Ji, W., Li, X., Cheng, L.: Generating diverse and natural 3d human motions from text. In: CVPR (2022)
- 17. Ho, J., Salimans, T.: Classifier-free diffusion guidance. ArXiv (2022)
- Huang, S., Wang, Z., Li, P., Jia, B., Liu, T., Zhu, Y., Liang, W., Zhu, S.C.: Diffusion-based generation, optimization, and planning in 3d scenes. In: CVPR (2023)
- Jang, D.K., Park, S., Lee, S.H.: Motion puzzle: Arbitrary motion style transfer by body part. TOG (2022)
- 20. Jiang, B., Chen, X., Liu, W., Yu, J., Yu, G., Chen, T.: Motiongpt: Human motion as a foreign language. ArXiv (2023)
- Jones, M., Wang, S.Y., Kumari, N., Bau, D., Zhu, J.Y.: Customizing text-to-image models with a single image pair. ArXiv (2024)
- Karunratanakul, K., Preechakul, K., Aksan, E., Beeler, T., Suwajanakorn, S., Tang, S.: Optimizing diffusion noise can serve as universal motion priors. In: Arxiv (2023)

- 16 Zhong et al.
- 23. Karunratanakul, K., Preechakul, K., Suwajanakorn, S., Tang, S.: Gmd: Controllable human motion synthesis via guided diffusion models. In: ICCV (2023)
- Kim, H.J., Lee, S.H.: Perceptual characteristics by motion style category. In: Eurographics (Short Papers) (2019)
- Kulkarni, N., Rempe, D., Genova, K., Kundu, A., Johnson, J., Fouhey, D., Guibas, L.: Nifty: Neural object interaction fields for guided human motion synthesis. ArXiv (2023)
- 26. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. ArXiv (2017)
- 27. Mason, I., Starke, S., Komura, T.: Real-time style modelling of human locomotion via feature-wise transformations and local motion phases. Proceedings of the ACM on Computer Graphics and Interactive Techniques (2022)
- Park, S., Jang, D.K., Lee, S.H.: Diverse motion stylization for multiple style domains via spatial-temporal graph-based generative model. Proceedings of the ACM on Computer Graphics and Interactive Techniques (2021)
- 29. Peng, X., Xie, Y., Wu, Z., Jampani, V., Sun, D., Jiang, H.: Hoi-diff: Text-driven synthesis of 3d human-object interactions using diffusion models. ArXiv (2023)
- Peng, X.B., Ma, Z., Abbeel, P., Levine, S., Kanazawa, A.: Amp: Adversarial motion priors for stylized physics-based character control. TOG (2021)
- Petrovich, M., Black, M.J., Varol, G.: Temos: Generating diverse human motions from textual descriptions. In: ECCV (2022)
- Petrovich, M., Litany, O., Iqbal, U., Black, M.J., Varol, G., Bin Peng, X., Rempe, D.: Multi-track timeline control for text-driven 3d human motion generation. In: CVPR (2024)
- Pi, H., Peng, S., Yang, M., Zhou, X., Bao, H.: Hierarchical generation of humanobject interactions with diffusion probabilistic models. In: ICCV (2023)
- 34. Pinyoanuntapong, E., Saleem, M.U., Wang, P., Lee, M., Das, S., Chen, C.: Bamm: Bidirectional autoregressive motion model. ArXiv (2024)
- Pinyoanuntapong, E., Wang, P., Lee, M., Chen, C.: Mmm: Generative masked motion model. In: CVPR (2024)
- Raab, S., Gat, I., Sala, N., Tevet, G., Shalev-Arkushin, R., Fried, O., Bermano, A.H., Cohen-Or, D.: Monkey see, monkey do: Harnessing self-attention in motion diffusion for zero-shot motion transfer. ArXiv (2024)
- Raab, S., Leibovitch, I., Tevet, G., Arar, M., Bermano, A.H., Cohen-Or, D.: Single motion diffusion. ArXiv (2023)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
- Rempe, D., Luo, Z., Peng, X.B., Yuan, Y., Kitani, K., Kreis, K., Fidler, S., Litany, O.: Trace and pace: Controllable pedestrian animation via guided trajectory diffusion. In: CVPR (2023)
- Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In: CVPR (2023)
- 41. Shafir, Y., Tevet, G., Kapon, R., Bermano, A.H.: Human motion diffusion as a generative prior. ArXiv (2023)
- 42. Shah, V., Ruiz, N., Cole, F., Lu, E., Lazebnik, S., Li, Y., Jampani, V.: Ziplora: Any subject in any style by effectively merging loras. ArXiv (2023)
- 43. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. ArXiv (2020)
- 44. Song, Y., Dhariwal, P., Chen, M., Sutskever, I.: Consistency models. ArXiv (2023)
- 45. Tao, T., Zhan, X., Chen, Z., van de Panne, M.: Style-erd: Responsive and coherent online motion style transfer. Arxiv (2022)

17

- 46. Tevet, G., Raab, S., Gordon, B., Shafir, Y., Cohen-or, D., Bermano, A.H.: Human motion diffusion model. In: ICLR (2023)
- 47. Wan, W., Dou, Z., Komura, T., Wang, W., Jayaraman, D., Liu, L.: Tlcontrol: Trajectory and language control for human motion synthesis. ArXiv (2023)
- Wang, Z., Wang, J., Lin, D., Dai, B.: Intercontrol: Generate human motion interactions by controlling every joint. ArXiv (2023)
- 49. Wen, Y.H., Yang, Z., Fu, H., Gao, L., Sun, Y., Liu, Y.J.: Autoregressive stylized motion synthesis with generative flow. In: CVPR (2021)
- Wu, Q., Zhao, Y., Wang, Y., Tai, Y.W., Tang, C.K.: Motionllm: Multimodal motion-language learning with large language models. ArXiv (2024)
- 51. Wu, Q., Shi, Y., Huang, X., Yu, J., Xu, L., Wang, J.: Thor: Text to human-object interaction diffusion via relation intervention. ArXiv (2024)
- 52. Xia, S., Wang, C., Chai, J., Hodgins, J.: Realtime style transfer for unlabeled heterogeneous human motion. TOG (2015)
- 53. Xie, Y., Jampani, V., Zhong, L., Sun, D., Jiang, H.: Omnicontrol: Control any joint at any time for human motion generation. In: ICLR (2024)
- Xu, P., Xie, K., Andrews, S., Kry, P.G., Neff, M., McGuire, M., Karamouzas, I., Zordan, V.: Adaptnet: Policy adaptation for physics-based character control. TOG (2023)
- 55. Xu, S., Ma, Z., Huang, Y., Lee, H., Chai, J.: Cyclenet: Rethinking cycle consistency in text-guided diffusion for image manipulation. NeurIPS (2024)
- Xu, S., Li, Z., Wang, Y.X., Gui, L.Y.: Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In: ICCV (2023)
- 57. Xu, S., Wang, Z., Wang, Y.X., Gui, L.Y.: Interdreamer: Zero-shot text to 3d dynamic human-object interaction. ArXiv (2024)
- 58. Yi, H., Thies, J., Black, M.J., Peng, X.B., Rempe, D.: Generating human interaction motions in scenes with text control. ArXiv (2024)
- Yu, J., Wang, Y., Zhao, C., Ghanem, B., Zhang, J.: Freedom: Training-free energyguided conditional diffusion model. ArXiv (2023)
- Yuan, Y., Song, J., Iqbal, U., Vahdat, A., Kautz, J.: Physdiff: Physics-guided human motion diffusion model. In: ICCV (2023)
- Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: ICCV (2023)
- Zhang, M., Cai, Z., Pan, L., Hong, F., Guo, X., Yang, L., Liu, Z.: Motiondiffuse: Text-driven human motion generation with diffusion model. PAMI (2024)
- Zhang, Y., Huang, D., Liu, B., Tang, S., Lu, Y., Chen, L., Bai, L., Chu, Q., Yu, N., Ouyang, W.: Motiongpt: Finetuned llms are general-purpose motion generators. ArXiv (2023)
- Zhou, W., Dou, Z., Cao, Z., Liao, Z., Wang, J., Wang, W., Liu, Y., Komura, T., Wang, W., Liu, L.: Emdm: Efficient motion diffusion model for fast, high-quality motion generation. Arxiv (2023)
- Zhou, Y., Barnes, C., Lu, J., Yang, J., Li, H.: On the continuity of rotation representations in neural networks. In: CVPR (2019)