

Unifying Human Motion Synthesis and Style Transfer with Denoising Diffusion Probabilistic Models

Ziyi Chang^a, Edmund J. C. Findlay, Haozheng Zhang and Hubert P. H. Shum^{bc}

Department of Computer Science, Durham University, Durham, UK
{ziyi.chang, edmund.findlay, haozheng.zhang, hubert.shum}@durham.ac.uk

Keywords: Diffusion Model, Animation Synthesis, Human Motion

Abstract: Generating realistic motions for digital humans is a core but challenging part of computer animations and games, as human motions are both diverse in content and rich in styles. While the latest deep learning approaches have made significant advancements in this domain, they mostly consider motion synthesis and style manipulation as two separate problems. This is mainly due to the challenge of learning both motion contents that account for the inter-class behaviour and styles that account for the intra-class behaviour effectively in a common representation. To tackle this challenge, we propose a denoising diffusion probabilistic model solution for styled motion synthesis. As diffusion models have a high capacity brought by the injection of stochasticity, we can represent both inter-class motion content and intra-class style behaviour in the same latent. This results in an integrated, end-to-end trained pipeline that facilitates the generation of optimal motion and exploration of content-style coupled latent space. To achieve high-quality results, we design a multi-task architecture of diffusion model that strategically generates aspects of human motions for local guidance. We also design adversarial and physical regulations for global guidance. We demonstrate superior performance with quantitative and qualitative results and validate the effectiveness of our multi-task architecture.

1 INTRODUCTION

The generation of realistic human motions has been an important but challenging task in computer graphics and vision. Compared to motion capture, generating motions allows obtaining data with affordable cost and abundant amount. Furthermore, the generated motions have a wide range of potential applications, e.g. making animations and games. As human motions are inherently diverse in both contents and styles, generating realistic motions is also challenging. For example, realistic human motions have a variety of contents, such as walking and running. Even within each content, motions vary from each other, e.g. strutting walking and depressed walking.

While recent years have witnessed some significant advances made by many deep learning approaches towards this domain, these approaches mainly consider motion synthesis and style manipulation as two separate tasks. Approaches to motion synthesis are usually concentrated on the generation of various contents (Holden et al., 2017; Mourot et al.,




Figure 1: Sampled results. (a) is walking. (b) is running. (c) is proud jumping. (d) is angry kicking.

2022), while the style transferring approaches focus on the manipulation of motion styles (Aberman et al., 2020; Ye et al., 2021). Although methods from the two separate tasks can be sequentially applied, the generated motions are likely to be better if the common space is modelled jointly.

The main challenge of generating realistic human motions is to effectively model the inter-class (i.e. content) and intra-class (i.e. style) behaviours in a

^a <https://orcid.org/0000-0003-0746-6826>

^b <https://orcid.org/0000-0001-5651-6039>

^c Corresponding author.

common representation. Human motions are composed of contents and styles (Aberman et al., 2020). The difference in contents dominates different inter-class behaviours, while the variation of styles accounts for the various intra-class motion behaviours. Modelling inter-class and intra-class behaviours of realistic human motions requires neural networks to learn a common representation in the latent space. The underlying distribution of such a latent space is greatly wider than that of modelling only one component, as both content and styles are inherently diverse. The increased diversity of data requires a higher capacity in neural networks to model the potential probabilistic distributions. However, traditional methods can only focus on learning the distribution of one component due to the limited capacity of their generative models (Bond-Taylor et al., 2021) as they suffer from sequential generation (Martinez et al., 2017), mode collapse (Dong et al., 2020), prior distribution assumption (Ling et al., 2020) or specially-designed architecture (Henter et al., 2020).

We propose a diffusion-based solution to unify human motion synthesis and style transfer. Because its mode coverage ability is wider than other previous neural networks, it is beneficial to model content and style in a common representation. Specifically, denoising diffusion probabilistic model (DDPM) is a diffusion-based generative model with high capacity due to the injection of stochasticity in the learning process, which is inspired by non-equilibrium thermodynamics in physics. To achieve high-quality motion synthesis and style-transfer results, we design a multi-task DDPM architecture that strategically models different aspects of realistic human motions, including joint angles, global movement trajectories, supporting foot patterns and physical regulations. Compared to the original DDPM (Ho et al., 2020), the multi-task architecture in our design increases the ability to model the data structure of human motions. Apart from only predicting the noise, our multi-task architecture further predicts other aspects of human motions with additional neural networks. To enhance the synthesized motions to be realistic globally, we also leverage adversarial training to coordinate different predictions from multiple tasks and ensure them to be harmonious with each other. In addition to adversarial regulation, we also leverage physical regulations to achieve global coherent motions.

We demonstrate the superior performance of our proposed solution by both quantitative and qualitative experiments on the dataset proposed by (Xia et al., 2015) and validate the effectiveness of our design by ablation study. In addition, the Fréchet inception dis-

tance (FID) (Heusel et al., 2017) has been reported to measure the distribution difference between the original and generated motion datasets. Our method achieves the lowest FID, signifying superior performance. We also provide several synthesized results by figures to show the quality of our method. We also validate different components of our multi-task architecture by ablation study.

An existing work presents preliminary results of synthesizing human motions by diffusion models (Findlay et al., 2022) where a barebone prototype to show the potential feasibility that walking motion with different styles can be synthesised with diffusion models. However, (Findlay et al., 2022) suffers from low motion quality in non-walking motion due to the simple network design. In this paper, we present an improved solution that is capable of generating high-quality motion of different contents with styles. Furthermore, we showcase the improvement over (Findlay et al., 2022) in our experiments. This paper presents the following contributions:

- We present a single-stage pipeline unifying human motion synthesis and style transfer for high-quality motion creation. The source code is open on <https://github.com/mrzzzy2021/StyledMotionSynthesis>
- To effectively represent the coupled representation of both inter-class motion contents and intra-class motion styles in a common latent space, we propose a denoising diffusion probabilistic model solution that has a large learning capacity for modelling the diverse data structure.
- To generate high-quality results, we propose a multi-task network architecture that leverages both local guidance, including joint angles, movement trajectories and supporting foot patterns, and global guidance, including physical and adversarial regulations.

2 RELATED WORK

Our work is mainly related to two research fields, neural motion generation and generative models. First, we will review the recent advances in neural motion generation, mainly including human motion synthesis and motion style transfer. In the generative models part, we will mainly discuss diffusion models.

2.1 Neural Motion Generation

Most motion synthesis works are prediction-based. Given a past pose or partial sequence, a model will

predict future poses. Recent advances in deep generative models such as Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), long short-term memory (LSTM) models and Flow-based models have seen a strong performance with this approach (Mourot et al., 2022).

MoGlow (Henter et al., 2020) used an LSTM-based normalising flow model inspired by recent applications of such models to video predicting. When provided with a context of previous poses and a control signal containing the character’s future angular velocity, the model can generate the next pose in the sequence. Characters can accurately follow various trajectories for long and short-term motions by using this approach. However, this method can only generate the motions one frame at a time, which is not ideal for real-time applications such as computer graphics. The generated motions are diverse but often suffer inaccuracies, such as foot sliding and joint crossover. Wen et al. also show that this model architecture performs strongly for unsupervised motion style transfer. Although the results of flow-based models are promising, they are often parameter inefficient compared to other generative models such as GANs.

Improving on the control signals from MoGlow (Henter et al., 2020) for human locomotion, Ling et al. propose Motion VAEs (Ling et al., 2020). Using a two-stage approach, the authors first train an encoder and mixture-of-experts decoder on the pose sequences dataset of locomotion data. They then train a reinforcement learning algorithm to predict the latent space input to be passed to the decoder based on the previous pose and a control signal. The final method is able to control a character using a joystick as well as by providing a trajectory or a target. Adding noise during the second stage training process allows for diversity when generating controlled inputs. However, the algorithm demonstrates bias since the generated data tend to have right-side behaviour whilst the training data contains equal amounts of right and left-handed data.

The action-conditioned motion generation is typically considered a more difficult task than motion prediction-based synthesis since less input information is supplied to the model. The Text2Action (Ahn et al., 2018) uses a recurrent neural network (RNN) encoder to take a sentence as input and output an embedding. This embedding is then used to condition a generator RNN which will output a motion based on the original sentence. A separate recurrent discriminator is also trained to minimise perceptual differences between generated motion and the training data. The system struggled to achieve sufficient quality motions when trained end-to-end and relied on

a multistage training procedure. ACTOR (Petrovich et al., 2021) builds on Text2Action by introducing a transformer-based VAE approach. Different from (Ahn et al., 2018), this method does not require multistage training and the training is more stable as it is not adversarial. The transformer will generate the appropriate pose sequence when supplied with the desired action. This model significantly outperforms their baseline gated recurrent unit (GRU) model on multiple datasets. Due to the autoregressive nature of the decoder, the method can be adapted to generate actions of varying lengths.

2.2 Diffusion Models

Diffusion models are firstly proposed by (Sohl-Dickstein et al., 2015) for data modelling. The diffusion phenomenon is driven by a stochastic dynamic system in physics and destroys the original data structure. Aiming to model the diffusion and its reverse process, (Sohl-Dickstein et al., 2015) proposes diffusion models and uses a neural network to recover the data distribution. (Sohl-Dickstein et al., 2015) uses the variational lower bound to optimize the negative log-likelihood of diffusion models.

The denoising diffusion probabilistic model (DDPM) is firstly proposed by (Ho et al., 2020). Instead of using variational lower bound, (Ho et al., 2020) simplifies the loss function by predicting the added noise. As reported by (Ho et al., 2020), predicting the added noise achieves better results than predicting the intermediate data. Many improvements have also been proposed to DDPM. (Nichol and Dhariwal, 2021) proposes to predict the covariance matrix during the reverse process and implements a cosine noise schedule instead of a linear schedule. (Dhariwal and Nichol, 2021) proposes a guidance-based diffusion model where auxiliary classifiers have been pre-trained to provide guidance to the generation. (Song et al., 2020) proposes a denoising diffusion implicit model (DDIM) to accelerate the sampling speed of DDPM. (Kawar et al., 2022) proposes a denoising diffusion restoration model (DDRM) to specialize DDPM in image restoration problems.

Hybrid approaches which combine a DDPM with another type of generative model (e.g., VAEs or GANs) have shown early promise in addressing the issues with diffusion models. (Vahdat et al., 2021) propose to train a VAE and then train a diffusion model in the latent space. This method surpasses all previous diffusion models when trained on image data and only requires around 100 sampling steps due to the reduced complexity of the latent space compared to the original data. However, training both models is

a multistage process as the VAE needs to be trained separately and then trained with the diffusion model. It results in more computational resources than a standard diffusion model to achieve these results. Denoising diffusion GANs (Xiao et al., 2021) use a reformulation of the diffusion process where the original data is estimated instead of the noise added. As a result, the model is less prone to mode collapse than GAN models and is also more stable to train. However, the quality of the generated samples is still slightly lower than those generated by previous diffusion models.

Most recently, diffusion models have achieved impressive results for text-to-image generation. GLIDE (Nichol et al., 2021) uses an encoding transformer to input a caption to a diffusion model. Imagen (Saharia et al., 2022) proposes to leverage generic large language models that are pretrained on text-only corpora for text-to-image translation. DALLE-2 (Ramesh et al., 2022) proposes a two-stage model for text-guided image generation. Their two stages include a prior that generates a CLIP image embedding given a text caption, and a decoder that generates an image conditioned on the image embedding. Such models for text-to-image transition usually are large. For example, (Nichol et al., 2021) has 3.5 billion parameters in size and is extremely computationally intensive to train. However, when combined with another 1.5 billion parameters, the upsampling diffusion model can generate high-resolution samples when provided with a text input.

Diffusion models have excelled at a range of tasks on many different data domains. They offer many benefits over other model classes, providing both high-quality and highly diverse samples at the cost of sampling time. However, despite their success, they have not yet been applied to motion data.

3 METHODOLOGY

Generating realistic human motions is challenging due to the diversity of the common representation for the coupled content-style information. Motions are composed of inter-class contents and intra-class styles (Aberman et al., 2020). Previous methods do not have sufficient learning capacity for modelling the coupled representation in a common latent space, such that they learn content and style modelling separately as motion synthesis and style transfer. However, separating as two tasks usually results in sub-optimal motions and under-explored joint distribution space.

To tackle the challenge, we propose our denoising diffusion probabilistic model (DDPM) solution for styled motion synthesis as shown in Figure 2. DDPM

is a diffusion-based generative model that injects and removes Gaussian noise progressively. The introduction of stochasticity in the learning process brings a larger capacity to the neural network. Compared to previous methods, DDPM is a more suitable candidate for styled motion synthesis as it does not suffer from the lack of generation diversity. Our proposed DDPM solution facilitates an integrated, end-to-end framework with increased learning capacity. The motion creation benefits more from our single-stage pipeline when compared to previously two-stage solutions. The integration enables direct optimization over the joint common representation space, and the underlying manifold is better explored to obtain better performance.

The training process of our proposed pipeline shown in Figure 2 is designed to be a multi-task architecture. The pipeline is optimized over a set of predictions on different motion aspects, including joint angles, foot patterns, global movements and physical regulations. Additionally, a discriminator is applied for adversarial optimization over our pipeline.

We unify human motion synthesis and style transfer into our proposed end-to-end pipeline. The framework leverages DDPM to increase the learning capacity to model the diversity in motion data. The architecture of our proposed pipeline is designed to be multi-task with a set of corresponding auxiliary losses for a high-quality generation.

3.1 Problem Formulation

Generating realistic human motions is an important but challenging task. Modelling human motions requires the neural network to learn the inter-class motion contents and the intra-class motion styles simultaneously. The joint distribution of contents and styles describes the underlying motion manifold in the latent space. However, the great diversity of both human contents and styles demands a high learning capacity of neural networks to maintain generation diversity, leading to a separate treatment in previous studies.

We propose a new problem called styled motion synthesis from an integration view and a framework, as shown in Figure 3. Previous studies have been conducted in developing neural networks for human motion synthesis or motion transfer, separately. The lack of unification formulation brings incoordination across neural networks in practice, undermining the quality of generated results. Instead of optimizing separate neural networks for human motion synthesis and motion style transfer, the new problem concerns the integration of the two tasks. Styled motion synthesis demands a neural network to generate motions

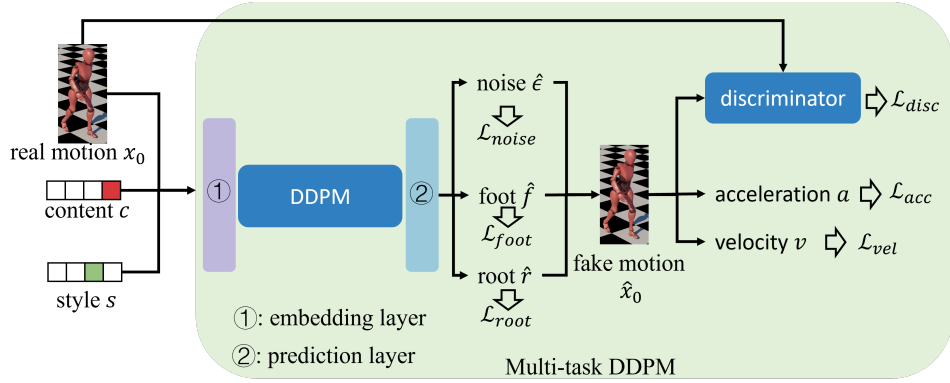


Figure 2: An overview of our proposed framework.

with styles from noise distribution that is easy to sample from in an end-to-end manner.




Figure 3: Our proposed end-to-end framework for the styled motion synthesis task.

To formally propose this problem, we provide a formulation. Let x and \hat{x} be our original motion data of all frames and generated motion data of all frames, respectively. The distributions of the original motion and generated motion are denoted as p_{ori} and p_{gen} . The neural network $F(\cdot)$ in styled motion synthesis starts from a noise ϵ that is sampled from a distribution p_{noise} and generates \hat{x} that follows the distribution of original data x . Styled motion synthesis is formulated as follows:

$$\hat{x} = F(\epsilon; \theta) \sim p_{gen} \approx p_{ori}, \quad (1)$$

where the noise is sampled from the noise distribution $\epsilon \sim p_{noise}$ and θ denotes the parameters to be optimized in the neural network $F(\cdot)$.

3.2 DDPM Backbone

We propose to use the denoising diffusion probabilistic model since it has the potential to possess enough capacity to model diversity compared with existing works. Realistic human motions are inherently diverse as motions differ from each other in terms of the inter-class contents and the intra-class styles. Modelling motions requires a larger capacity of networks in data modelling. Previous methods such as (Aberman et al., 2020) only focus on modelling one component as their networks are limited in capacity. This

limitation usually leads to a lack of diversity in their generated results.

Compared with the existing methods, the denoising diffusion probabilistic model (DDPM) (Sohl-Dickstein et al., 2015; Ho et al., 2020) has an increased learning capacity due to the stochasticity introduced in the neural network. DDPM is essentially designed to formulate the generation process as a stochastic system. Both the injection and noise removal are designed to be tractable for changing probabilistic distributions. In addition, the stochastic characteristic in DDPM extends the exploration scope of the latent space, facilitating the network to model data diversity.

The training process of DDPM is designed to gradually inject (the first line in Eq. 2) and remove (the second line in Eq. 2) a Gaussian noise $\epsilon \sim \mathcal{N}(0, I)$ to the input data $x_0 = x$ within t diffusion steps. The optimization (see the third line in Eq. 2) is based on minimizing the prediction of what noise has been added. The whole training process is formulated as follows:

$$\begin{cases} p_t(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\sigma}_t}x_0, (1 - \bar{\sigma}_t)I), \\ p_t(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \frac{1}{\sqrt{1-\sigma_t}}(x_t - \frac{\sigma_t}{\sqrt{1-\bar{\sigma}_t}}\hat{\epsilon}), \sigma_t I), \\ \mathcal{L} = \mathbb{E}_{t, x_0, \epsilon} \|\epsilon - \hat{\epsilon}\|_2^2, \end{cases} \quad (2)$$

where $\sigma_t \in (0, 1)$ is the noise schedule, $\bar{\sigma}_t = \prod_{i=1}^t (1 - \sigma_i)$, $\hat{\epsilon}$ is the prediction of the neural network and \mathcal{L} is the loss function of DDPM in (Ho et al., 2020). When training is completed, the generation process starts from an isotropic Gaussian noise $x_T \sim \mathcal{N}(0, I)$ and iteratively compute the previous state x_{t-1} over T diffusion steps, which is formulated as:

$$x_{t-1} = \frac{1}{\sqrt{1-\sigma_t}}(x_t - \frac{\sigma_t}{\sqrt{1-\bar{\sigma}_t}}\hat{\epsilon}) + \sigma_t z, \quad (3)$$

where $z \sim \mathcal{N}(0, I)$ is a Gaussian noise.

3.3 Multi-task DDPM for Styled Motion Synthesis

3.3.1 Local Guidance

To achieve high-quality generated results, we propose to leverage a multi-task DDPM architecture for styled human motion synthesis. Directly applying DDPM to motion synthesis does not produce high-quality results due to the lack of consideration of other inherent aspects of motions, such as foot patterns.

Our proposed multi-task design considers the optimization of trajectory movements, foot patterns and joint angle predictions, while the original DDPM only minimizes the difference between the predicted noise and the ground truth noise. We provide DDPM with extra guidance from other aspects of motion by the proposed multi-task design. With such guidance, the learnt distribution of coupled contents and styles is close to the manifold of realistic motions in the common latent space.

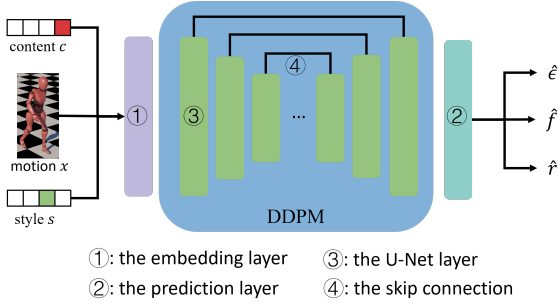


Figure 4: Our proposed multi-task DDPM pipeline for styled motion synthesis.

As shown in Figure 4, our multi-task DDPM takes a motion clip $x_0 = x$, the content c and the style s as inputs. We represent the motion x with joint angles. Both the content c and the style s are converted into one-hot embeddings with embedding layers before they are used in DDPM. We apply the U-Net with attention and skip connections in our DDPM to model the manifold conformed by motions x , the inter-class contents c and the intra-class styles s .

Our multi-task DDPM produces estimation on the noise $\hat{\epsilon}$, global movements \hat{r} and supporting foot patterns \hat{f} based on the three inputs. For supporting foot patterns, we use binary values to indicate a foot contacts to the floor. We leverage the noise prediction $\hat{\epsilon}$ instead of directly predicting latent joint angles \hat{x} for joint angles because DDPM shows better performance by predicting noise $\hat{\epsilon}$ than latent intermediate result \hat{x} (Ho et al., 2020). To optimize the three pre-

dictions, we formulate loss functions as:

$$\mathcal{L}_{noise} = \mathbb{E}_{x_0, t, s, c} \|\epsilon - \hat{\epsilon}\|_2^2, \quad (4)$$

$$\mathcal{L}_{foot} = \mathbb{E}_{x_0, t, s, c} \|f - \hat{f}\|_2^2, \quad (5)$$

$$\mathcal{L}_{root} = \mathbb{E}_{x_0, t, s, c} \|r - \hat{r}\|_2^2. \quad (6)$$

3.3.2 Global Guidance

Apart from separately providing guidance on different aspects, we also propose to apply global guidance for the conformity of all previous aspects. We propose to apply physical regulations and a discriminator based on a reconstruction formulation in DDPM to enhance the generated motions to be realistic.

For our global guidance, we propose to utilize a reconstruction formulation (Eq. 8), which is derived from the arbitrary query property (Eq. 7) in DDPM. The process of gradually adding Gaussian noise to the input x_0 as described in the first line of Eq. 2 allows the arbitrary query on intermediate state x_t at any diffusion step t , which is written as:

$$x_t = \sqrt{\bar{\sigma}_t}x_0 + \sqrt{1 - \bar{\sigma}_t}\epsilon_t. \quad (7)$$

The arbitrary query property is denoted by the arrow from x_0 to x_t in Figure 5. Our neural network makes




Figure 5: Our proposed multi-task conditional DDPM pipeline for styled motion synthesis.

predictions on the added noise ϵ_t , which enables us to estimate x_0 reversely. Therefore, based on the prediction $\hat{\epsilon}$, the reconstruction estimation \hat{x}_0 is written as:

$$\hat{x}_0 = \frac{x_t - \sqrt{1 - \bar{\sigma}_t}\hat{\epsilon}}{\sqrt{\bar{\sigma}_t}}, \quad (8)$$

which is denoted as the arrow from $\hat{\epsilon}_t$ to \hat{x}_0 in Figure 5. By the reconstruction formulation, we inversely obtain an estimation of motions \hat{x}_0 at any diffusion step t .

We propose to fulfil the physical regulations based on the estimated motion \hat{x} . These physical regulations enforce the physical properties in real world, encouraging natural and coherent motions to be generated by our network. Specifically, the state of joints in two

consecutive frames is approximately the same as humans move consecutively in real world. Furthermore, the velocity of joints between two consecutive frames in realistic human motions does not change abruptly in real world following the inertia constraint. Based on the two observations, we propose to regulate the velocity and acceleration in our estimated motion \hat{x} to encourage natural and coherent results. We use the rotation of joints, denoted as j_t , to calculate the velocity and acceleration. Therefore, we propose the following two loss functions:

$$\mathcal{L}_{vel} = \mathbb{E}_{t, j \in \hat{x}} \|j_t - j_{t-1}\|_2^2, \quad (9)$$

$$\mathcal{L}_{acc} = \mathbb{E}_{t, j \in \hat{x}} \|j_t - 2j_{t-1} + j_{t-2}\|_2^2. \quad (10)$$

Apart from physical regulations, we also propose an adversarial design to enhance the consonance between these tasks. We design a discriminator to provide the global harmony guidance. As we predict the joint angles \hat{x} , global movements \hat{r} and supporting foot patterns \hat{f} separately in local guidance, the proposed discriminator examines whether the three estimated values are harmoniously combined into a natural motion clip. Following the adversarial training paradigm, our discriminator is trained to distinguish the generated motion clip and the ground truth motion clip. The adversarial loss is written as:

$$\mathcal{L}_{disc} = \|D(x_0, r, f) - 1\|_2^2 + \|D(\hat{x}_0, \hat{r}, \hat{f}) - 0\|_2^2. \quad (11)$$

Overall, our multi-task DDPM is optimized by a set of respective loss functions on multiple tasks. The overall loss is written as:

$$\begin{aligned} \mathcal{L} = & \lambda_1 \mathcal{L}_{noise} + \lambda_2 \mathcal{L}_{foot} + \lambda_3 \mathcal{L}_{root} \\ & + \lambda_4 \mathcal{L}_{disc} + \lambda_5 \mathcal{L}_{vel} + \lambda_6 \mathcal{L}_{acc}, \end{aligned} \quad (12)$$

where $\lambda_1 = \lambda_2 = \lambda_3 = \lambda_4 = 1$ and $\lambda_5 = \lambda_6 = 0.01$.

4 EXPERIMENTAL RESULTS

We train our method on a public motion dataset (Xia et al., 2015) and evaluate the performance quantitatively and qualitatively. We compared our method with the initial DDPM method (Ho et al., 2020) and the existing DDPM-based method (Findlay et al., 2022) for evaluating performance quantitatively. In addition, we provide several generated motions for qualitative evaluation on generating motions with different contents and styles. To validate the effectiveness of each component in our proposed multi-task architecture design, we conduct a comprehensive ablation study.

4.1 Implementation Details

We trained all models on an NVIDIA RTX 3090 GPU and used 32-bit floating-point arithmetic. The training process can be completed within approximately one day. The generation process for each instance takes around 20 seconds. The hyperparameters in our experiments are shown in Table 1.

Table 1: Hyperparameters

Learning Rate	0.0002
Discriminator Learning Rate	0.0001
Adam β_1	0.9
Adam β_2	0.999
Adam ϵ	$1e^{-8}$
Batch Size	128
number of timesteps T	1000
EMA decay rate m	0.9999

We train our model with a publicly available dataset proposed by (Xia et al., 2015). The dataset contains six content classes and eight different styles for each type of motion. This dataset offers a wide range of styles and contents, allowing us to assess our model’s diversity by combining different styles and contents during motion generation. Additionally, we apply a Gaussian filter and inverse kinematics as the post-processing.

4.2 Quantitative Comparison

We choose to evaluate our models using the Fréchet inception distance (FID) (Heusel et al., 2017). This metric compares the distribution of motions generated by our diffusion model with the distribution of motions in our dataset. Given a set of generated motions from our model and a set of real motions from our dataset, we calculate a multidimensional Gaussian distribution from the features of a neural network for these datasets, which are represented as $\mathcal{N}(\mu_g, \Sigma_g)$ for our generated data and $\mathcal{N}(\mu_r, \Sigma_r)$ for our real data. We use the penultimate features from our pre-trained classifier for content motions. The FID metric is calculated by:

$$FID = \|\mu_r - \mu_g\|_2^2 + \text{tr}(\Sigma_g + \Sigma_r - 2(\Sigma_r \Sigma_g)^{\frac{1}{2}}). \quad (13)$$

We compare our method with an existing DDPM-based method (Findlay et al., 2022), and the DDPM baseline (Ho et al., 2020) on the dataset proposed by (Xia et al., 2015). All three methods have been trained using the same hyperparameter setting as shown in Table 1. The FID score is computed by the same




Figure 6: Generated motions with different contents. (a) is walking. (b) is running. (c) is jumping. (d) is kicking.

amount of generated data with the same pre-trained classifier. The results have been reported in Table 2. As shown in Table 2, our method achieves the

Table 2: Quantitative comparison.

Model	FID (\downarrow)
(Findlay et al., 2022)	158.47
(Ho et al., 2020)	198.67
Ours	56.73

best FID score compared with previous studies. This means the distribution of the generated data by our method has less difference from the ground truth dataset. In addition, we observe that the low performance in (Ho et al., 2020) is mainly due to the lack of constraints on motion-specific aspects. Although (Findlay et al., 2022) adds the adversarial training, it still does not perform well because it suffers from low motion quality in non-walking motion caused by its simple network design. Based on the observations above, our method achieves the best generative result with our proposed multi-task DDPM.

4.3 Qualitative evaluations

Apart from the quantitative comparison, we also provide several qualitative evaluations. After training is completed, we use our model to generate motion samples. The number of total frames is kept the same as that of the training dataset, which is 32 frames for each motion clip. Figure 6 has shown the generated motions with different contents from our model. Our results show that our method is capable of generating different motions with different contents.

In Figure 7, we have shown several sampled instances of styled walking motions. Our method produces different walking motions with styles, e.g. the old walking motion in Figure 7. As our method is

based on stochastic generation, the capacity of our model is adequate for modelling various styles. Additionally, we also provide styled motions with other contents, such as running, in Figure 8.

4.4 Ablation Study

We perform an ablation study to evaluate the effectiveness of multi-task architecture design. We use the same dataset proposed by (Xia et al., 2015) and keep the same hyperparameter settings as shown in Table 1. The FID score has been reported on the same amount of generated data and the same classifier.

Table 3: Ablation study on multi-task architecture.

Model	FID (\downarrow)
w/o foot loss	74.68
w/o root loss	118.20
w/o physical loss	139.44
w/o discriminator	106.29
Ours - full	56.73

As shown in Table 3, our method achieves the best FID score among different ablations. From the results of the ablation study, we validate the effectiveness of our multi-task DDPM architecture. We observe a performance drop on the FID score by removing the discriminator. In our design, the discriminator harmonizes the separately estimated components for local guidance. The adversarial loss encourages to generate natural and coherent motions. If we remove the root prediction, the performance also drops a lot. As previously discussed, the trajectory of movements, which are represented by root, are mostly represented in the inter-class motion content. This result shows that human motions are heavily dependent on the trajectory of movements. The physical loss, including velocity and acceleration, provides another important motion information, which is represented by the intra-class styles. Within the same content, the difference between the actions mainly happens in the joint movement where velocity and acceleration can describe the intra-class motion styles. Our method without the foot loss still does not perform well. This is because estimating supporting foot patterns helps our neural network generate natural and coherent motions. Therefore, our ablation validates that each component in our multi-task DDPM architecture is effective towards generating high-quality motions.




Figure 7: Walking motion with styles. (a) is angry walking. (b) is sexy walking. (c) is proud walking. (d) is strutting walking. (e) is neutral walking. (f) is depressed walking. (g) is childlike walking. (h) is old walking.

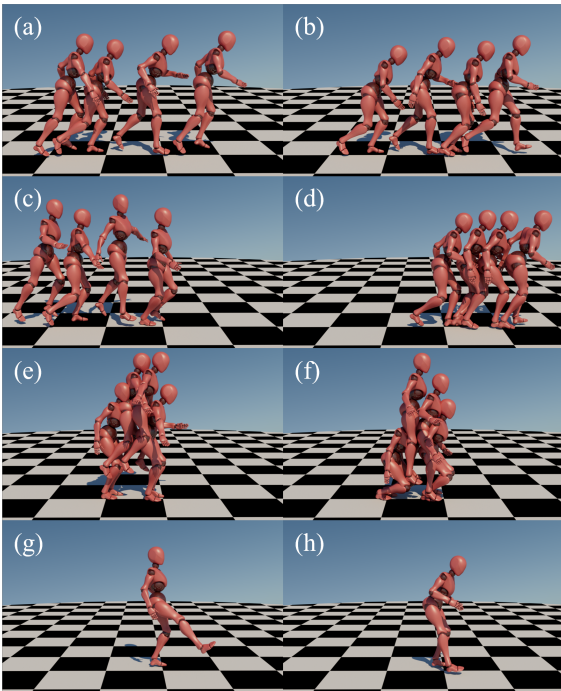


Figure 8: We also provide generated styled motions with other contents. (a) is angry running. (b) is depressed running. (c) is strutting running. (d) is old running. (e) is sexy jumping. (f) is proud jumping. (g) is angry kicking. (h) is old kicking.

5 CONCLUSION

Stylized motion synthesis is a critical and challenging problem with broad applications in many areas. We propose an end-to-end framework for stylized motion synthesis where we successfully integrate two separate tasks, i.e. human motion synthesis and motion style transfer, into one pipeline. We represent our motions as the inter-class content component and the intra-class style component coupled in a common la-

tent space. Such integration brings benefits more in optimal motion generation and thorough exploration of the coupled content-style latent space than previously two-staged frameworks. As motions contain a vast span of potential instances, we propose to leverage diffusion models in our framework to capture the diversity of motion data. We take DDPM for our implementation while our framework is generally compatible with other diffusion models. Our DDPM is designed to be a multi-task architecture for stylized motion synthesis. The multi-task system is optimized by the local guidance including joint angles, trajectory of movements, supporting foot patterns and the global guidance including physical regulations and a discriminator based on the reconstruction formulation to encourage the generation of natural and coherent motions. Both qualitative and quantitative experimental results have shown the superior performance of our method. Our ablation study validates the effectiveness of our proposed multi-task architecture.

Our proposed pipeline is generally compatible with various diffusion-based generative models. In addition, our pipeline has the potential to model other motions (e.g., animal motions). However, as the motions become increasingly diverse, we may need new generative models with higher learning capacity to model the coupled content-style representations.

There are also a few potential future directions towards the stylized human motion synthesis with diffusion models. One future work is that we will try our pipeline with a larger set of solutions containing also non-diffusion solutions and also try longer clip synthesis.

A future direction could be the stability and controllability of the generation process via a stochastic system. Although the capacity of models has been greatly increased, the generation process is inherently trained under heavy uncertainty brought by noise in-

jection. The stochastic property makes the generation difficult to achieve higher controllability and stability. The stochastic system such as DDPM requires further mathematical theories on stochasticity to achieve higher stability and better controllability.




Figure 9: Punching is a challenging case.

Another potential direction in future is the control of limb ends, e.g. feet and hands. Figure 9 shows a challenging case. In human motions, the ends of limbs usually have tiny but complicated movements that DDPM is hard to model. Although we estimate foot patterns and trajectory of movements, the control of limb ends still has the potential to be improved by several ways such as physics-based guidance.

REFERENCES

- Aberman, K., Weng, Y., Lischinski, D., Cohen-Or, D., and Chen, B. (2020). Unpaired motion style transfer from video to animation. *ACM Transactions on Graphics (TOG)*, 39(4):64–1.
- Ahn, H., Ha, T., Choi, Y., Yoo, H., and Oh, S. (2018). Text2action: Generative adversarial synthesis from language to action. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5915–5920. IEEE.
- Bond-Taylor, S., Leach, A., Long, Y., and Willcocks, C. G. (2021). Deep generative modelling: A comparative review of vaes, gans, normalizing flows, energy-based and autoregressive models. *arXiv preprint arXiv:2103.04922*.
- Dhariwal, P. and Nichol, A. (2021). Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794.
- Dong, Y., Aristidou, A., Shamir, A., Mahler, M., and Jain, E. (2020). Adult2child: Motion style transfer using cyclegans. In *Motion, interaction and games*, pages 1–11.
- Findlay, E. J., Zhang, H., Chang, Z., and Shum, H. P. (2022). Denoising diffusion probabilistic models for styled walking synthesis. *arXiv preprint arXiv:2209.14828*.
- Henter, G. E., Alexanderson, S., and Beskow, J. (2020). Moglow: Probabilistic and controllable motion synthesis using normalising flows. *ACM Transactions on Graphics (TOG)*, 39(6):1–14.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. (2017). Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J., Jain, A., and Abbeel, P. (2020). Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851.
- Holden, D., Komura, T., and Saito, J. (2017). Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)*, 36(4):1–13.
- Kawar, B., Elad, M., Ermon, S., and Song, J. (2022). Denoising diffusion restoration models. *arXiv preprint arXiv:2201.11793*.
- Ling, H. Y., Zinno, F., Cheng, G., and Van De Panne, M. (2020). Character controllers using motion vaes. *ACM Transactions on Graphics (TOG)*, 39(4):40–1.
- Martinez, J., Black, M. J., and Romero, J. (2017). On human motion prediction using recurrent neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2891–2900.
- Mourot, L., Hoyet, L., Le Clerc, F., Schnitzler, F., and Hellier, P. (2022). A survey on deep learning for skeleton-based human animation. In *Computer Graphics Forum*, volume 41, pages 122–157. Wiley Online Library.
- Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., and Chen, M. (2021). Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*.
- Nichol, A. Q. and Dhariwal, P. (2021). Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR.
- Petrovich, M., Black, M. J., and Varol, G. (2021). Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10985–10995.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. (2022). Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*.
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S. K. S., Ayan, B. K., Mahdavi, S. S., Lopes, R. G., et al. (2022). Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*.
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. (2015). Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR.
- Song, J., Meng, C., and Ermon, S. (2020). Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Vahdat, A., Kreis, K., and Kautz, J. (2021). Score-based generative modeling in latent space. *Advances in*

Neural Information Processing Systems, 34:11287–11302.

Xia, S., Wang, C., Chai, J., and Hodgins, J. (2015). Realtime style transfer for unlabeled heterogeneous human motion. *ACM Transactions on Graphics (TOG)*, 34(4):1–10.

Xiao, Z., Kreis, K., and Vahdat, A. (2021). Tackling the generative learning trilemma with denoising diffusion gans. *arXiv preprint arXiv:2112.07804*.

Ye, Z., Wu, H., and Jia, J. (2021). Human motion modeling with deep learning: A survey. *AI Open*.