Scalable Motion Style Transfer with Constrained Diffusion Generation

Wenjie Yin[1], Yi Yu[2*], Hang Yin[3], Danica Kragic[1], Mårten Björkman[1*]

[1]KTH Royal Institute of Technology, Sweden
[2]National Institute of Informatics, Japan
[3]University of Copenhagen, Denmark
{yinw, dani, celle}@kth.se, yiyu@nii.ac.jp, hayi@di.ku.dk

Abstract

Current training of motion style transfer systems relies on consistency losses across style domains to preserve contents, hindering its scalable application to a large number of domains and private data. Recent image transfer works show the potential of independent training on each domain by leveraging implicit bridging between diffusion models, with the content preservation, however, limited to simple data patterns. We address this by imposing biased sampling in backward diffusion while maintaining the domain independence in the training stage. We construct the bias from the source domain keyframes and apply them as the gradient of content constraints, yielding a framework with keyframe manifold constraint gradients (KMCGs). Our validation demonstrates the success of training separate models to transfer between as many as ten dance motion styles. Comprehensive experiments find a significant improvement in preserving motion contents in comparison to baseline and ablative diffusion-based style transfer models. In addition, we perform a human study for a subjective assessment of the quality of generated dance motions. The results validate the competitiveness of KMCGs.

Introduction

Human motion style is a complex and important aspect of human behavior, which can be perceived as motion characteristics that convey personality and temper or articulate socio-cultural factors. The ability to accurately capture and replicate motion style is crucial in numerous applications, such as video games, and choreography. Style transfer systems can streamline creation of various media, including images (Gatys, Ecker, and Bethge 2016), music (Brunner et al. 2018), and indeed human movements (Dong et al. 2020). For instance, in the realm of video games, unique modes of action correspond to various player operations and character states. Similarly, in the field of dance choreography, each dance genre has its own movement patterns, and the application of style transfer can assist choreographers in creating variations of particular movements.

Early methods for motion style transfer (Hsu, Pulli, and Popović 2005; Taylor and Hinton 2009; Smith et al. 2019)
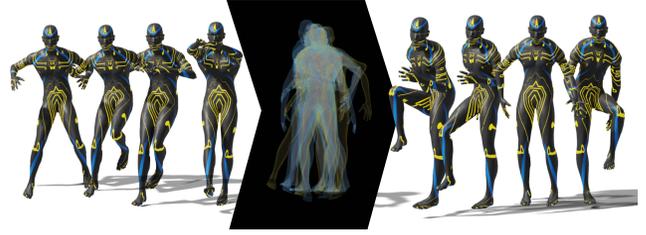
Figure 1: The proposed scalable motion style transfer system, with independently trained diffusion models. Given a source motion sequence, our system converts it to latent encoding, then constructs motion in the target style.

primarily relied on supervised learning, requiring paired and annotated data. However, obtaining matched human motion sequences is challenging and often involves a laborious preprocessing phase. Consequently, the predominant style transfer methods today are unpaired translations that do not demand direct mappings between motion samples. Common methods among these are the motion transfer systems based on CycleGAN (Dong et al. 2020; Yin et al. 2023b) for paired style domains and StarGAN (Chan, Irimia, and Ho 2020; Yin et al. 2023c) for a set of domains.

While these systems are capable of generating high-quality human motion, they exhibit considerable limitations when it comes to scalability towards new domains. Specifically, CycleGAN-based methods are trained on distinct pairs of motion styles, where each model is tailored for transfer between a specific pair of styles (Dong et al. 2020; Yin et al. 2023b), or establish a shared domain bridging multiple styles via StarGAN-based architecture (Chan, Irimia, and Ho 2020; Yin et al. 2023c).

The former implies training a quadratically increasing number of models, making it an impractical solution when transfer between a large number of style domains is considered. The latter faces finding such a shared domain, which is an enormously challenging task. Both methods also show limited incremental scalability and adaptability; a new style requires retraining of the models on the entire dataset. The requirement of simultanous access to the entire dataset highlights an additional drawback of these methods when parts of the data are sensitive due to privacy concerns, such as

in rehabilitation therapy. Developing methods that address these prominent challenges is thus critical for building practical motion style transfer systems.

In this paper, we propose a method based on Dual Diffusion Implicit Bridges (DDIBs) (Su et al. 2022) to mitigate both issues of system scalability and data privacy. DDIBs adopt the Schrödinger bridge perspective of diffusion models, showing certain information can be passed between diffusion latent spaces and used for image translation. Our system leverages this for a motion style transfer system with a two-step process. Given a source and a target model, the system initially utilizes the source model to obtain a latent encoding of motion at the final diffusion time step. This latent encoding is subsequently fed as the starting condition to the target model to generate the target motion (See Figure 1). The source and target domain models are completely decoupled, which allows for training the models separately. As a result, our system alleviates the need for a paired dataset that is typical for CycleGAN- and StarGAN-based methods, while upholding data privacy.

Our core technical contribution is improving content preservation of DDIBs in the motion style transfer domain. We find DDIBs struggle to retain the motion content faithfully when the analogy between source and target domains is low, consistent with the observations in image translation (Su et al. 2022). This is probably due to uncontrolled information encoding and decoding in independent diffusion processes. To this end, we propose a Keyframe Manifold Constraint Gradients (KMCGs) framework to improve content coherence of the target domain during inference. KMCGs uses keyframes from the source domain as context constraints and employ Manifold Constrained Gradient (Chung et al. 2022) to enforce these constraints during the second phase of DDIBs. Our experiments and quantative analysis on the 100STYLE (Mason, Starke, and Komura 2022) locomotion database, and the AIST++ (Tsuchida et al. 2019) dance database find KMCGs achieves successful style transfer and better content preservation, reflected on the cycle consistency property on these two motion datasets and two probabilistic divergence-based metrics. Additionally, our human study as a subjective evaluation find samples generated from KMCGs are preferred in general. Overall, these evaluations demonstrate that our system significantly outperforms baseline methods. A video of the summary and examples can be accessed via https://youtu.be/98wEWavjnxI.

In summary, our contributions are:

- A motion style transfer system that generates motions ranging from fundamental human locomotion to sophisticated dance movements. Our system pioneers in terms of system scalability and data privacy, demonstrating efficient and independent training over ten styles.

- A technical method KMCGs that mitigates the content coherence issue of dual diffusion implicit bridges, which is found prominent in transferring complex motions.

- A comprehensive evaluation of the proposed style transfer system including both objective metrics and subjective human study, showing significant performances boost compared to baseline and ablative models.

Related Work

We now review relevant prior works, including diffusion-based motion synthesis, as well as motion style transfer.

Diffusion-based Motion Synthesis

In light of recent groundbreaking advances made possible by diffusion models (Ho, Jain, and Abbeel 2020; Song and Ermon 2019), there has been a surge of interest in extending these techniques to the 3D motion domain. The recent MotionDiffuse (Zhang et al. 2022a) is regarded as the pioneering diffusion-based framework for text-driven motion generation. Similar to MotionDiffuse, the concurrent MDM (Tevet et al. 2022) and FLAME (Kim, Kim, and Choi 2022) integrate diffusion models and pre-trained language models such as CLIP (Radford et al. 2021) and RoBERTa (Liu et al. 2019) for generating motion from natural language descriptions. In a parallel development, speech audio is considered a model input (Zhang et al. 2023; Alexanderson et al. 2022) for gesture synthesis. Compared with gesture synthesis, dance generation is perhaps a more complex and challenging, but relatively under-explored field. Alexanderson et al. (2022) pioneers diffusion models with Conformer (Zhang et al. 2022b) for music to dance generation. EDGE (Tseng, Castellon, and Liu 2022) and Magic (Li et al. 2022) are Transformer-based diffusion models. EDGE incorporates auxiliary losses to encourage physical realism. Our work explores diffusion-based motion transfer on both human locomotion and dance databases.

One line of study focuses on manipulating partial body movements and styles. The modeling is hence not only about what contents are expressed by gestures but also how they are executed. HumanMAC (Chen et al. 2023) achieves controllable prediction of any part of the body. FLAME (Kim, Kim, and Choi 2022) and MDM (Tevet et al. 2022) enable body part editing both frame-wise and joint-wise by adapting diffusion inpainting to motion data. Tevet et al. (2022) design a multi-task architecture to unify human motion synthesis and stylization. Alexanderson et al. (2022) control the style and strength of motion expression by guided diffusion (Dhariwal and Nichol 2021). Yin et al. (2023a) integrate multimodal transformer and autoregressive diffusion models for controllable motion generation and reconstruction. These works often require joint training of models for body parts and styles while this is not desired in our work for the application of dual diffusion implicit bridges (Su et al. 2022).

Motion Style Transfer

The field of computer animation has long grappled with the challenge of motion style transfer, a process that entails transferring animation from a source style to a target style while retaining key content aspects, such as structure, timing, spatial relationships, etc. Early research in motion style transfer has depended on handcrafted features (Amaya, Bruderlin, and Calvert 1996; Unuma, Anjyo, and Takeuchi 1995; Witkin and Popovic 1995). As style is an elusive concept to define precisely, contemporary studies tend to endorse data-driven methodologies for feature extraction.

Typical models utilized for style transfer include Cycle-GAN (Dong et al. 2020), AdaIN (Aberman et al. 2020), autoregressive flows (Wen et al. 2021), and etc. Additionally, some research concentrates on addressing real-time style transfer (Xia et al. 2015; Smith et al. 2019; Mason, Starke, and Komura 2022). Nonetheless, it is essential to highlight that these studies focus on relatively simplistic human motions, such as exercise and locomotion, where stylistic variation is restricted. CycleDance (Yin et al. 2023b) and StarDance (Yin et al. 2023c) address the transfer of dance movements that exhibit a substantial degree of complexity in terms of postures, transitions, rhythms, and artistic styles. However, CycleDance and StarDance suffer from severe drawbacks in rapid adaptation to alternative domains. As diffusion models advance, Alexanderson et al. (2022) demonstrate the ability to control the style and intensity of dance motion expression using a classifier-free guided diffusion model. Raab et al. (2023) propose the diffusion-based SinMDM to transfer the style of a given reference motion to learned motion motifs. Our proposed framework explores diffusion models with manifold constraints to facilitate the transfer of dance style. The process of style transfer hinges on the diffusion models independently trained within each respective domain. This method overcomes the adaptability limitations of CycleDance and StarDance. We further boost the performance of style transfer by imposing keyframe manifold constraints.

Method

In this section, we formulate the problem and provide preliminaries of diffusion models and Dual Diffusion Implicit Bridges (DDIBs). After this, the proposed system with Keyframe Manifold Constraint Gradients (KMCGs) is presented.
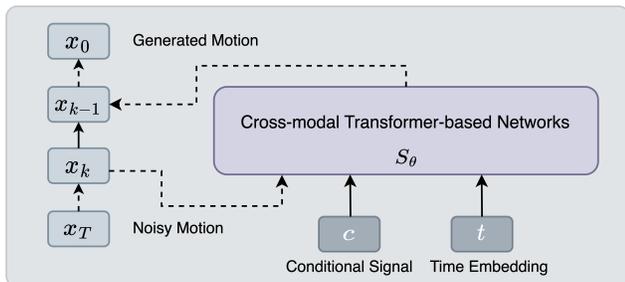

Figure 2: Architecture overview of diffusion models, conditional information acts as cross-attention context. The diffusion model takes noisy sequences and produces the estimated motion sequences.

Problem Formulation

Our study aims to develop a translation method across multiple motion domains, denoted as source domain $D^{(i)}$ and target domain $D^{(j)}$, without relying on jointly training over multiple data domains. In our scenario, we focus on conditional motion sequences. Given a sample sequence $x^{(i)}$ in
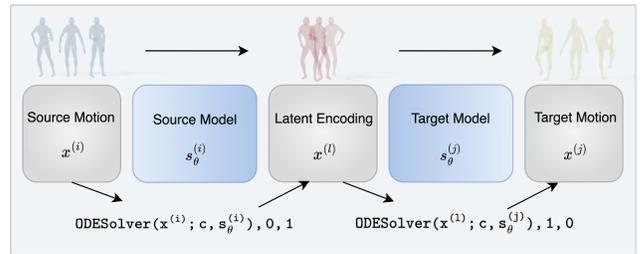

Figure 3: DDIBs leverage two models for style transfer. The source model converts the source sample from the source domain into the latent domain. The target model reverses the latent encoding to the target domain.

the source domain and a conditional signal c, our purpose is to transfer this sequence to the target domain, resulting in a sample sequence $x^{(j)}$ with the style of the target domain while striving to retain the content of the source domain.

Dual Diffusion Motion Transfer

To tackle the problem described above, we employ a strategy inspired by DDIB which use two separately trained Denoising Diffusion Implicit Models (DDIMs) (Song, Meng, and Ermon 2020) for image-to-image translation. With DDIMs, DDIBs achieve exact cycle consistency. We further boost the transfer performance by imposing keyframe context from the source motion as a manifold constraint.

Deep Diffusion Models. Continuous-time diffusion processes $\{x(t), t \in [0,1]\}$ of diffusion models are defined as forward Stochastic Differential Equations (SDEs) (Song et al. 2020),

$$d\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x},t)dt + g(t)d\boldsymbol{w}, \tag{1}$$

where $\boldsymbol{w}$ is the standard Wiener process running backward in time, $f(\boldsymbol{x},t)$ is the drift term, $g(t)$ is the scalar diffusion coefficient. The backward-time SDEs of Eqn. (1) are

$$d\boldsymbol{x} = \left[\boldsymbol{f}(\boldsymbol{x},t) - g^2 \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x})\right] dt + g(t)\,d\boldsymbol{w}, \tag{2}$$

where $\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x})$ is the score function of the noise perturbed data distribution at time t. The actual implementations of diffusion models are sample discrete times. The time horizon $t \in [0,1]$ is split up into T discretization segments as $\{x_k\}_{k=0}^{T}$, where k is an integer ranging from 0 to T. The forward and backward diffusion process can be defined as:

$$\boldsymbol{x}_k = a_k \boldsymbol{x}_0 + b_k z, \tag{3}$$

$$\boldsymbol{x}_{k-1} = f(\boldsymbol{x}_k, \boldsymbol{s}_\theta) + g(\boldsymbol{x}_k)z. \tag{4}$$

We base our diffusion model on the EDGE architecture (Tseng, Castellon, and Liu 2022) to generate human motion contingent upon conditional signals. This architecture utilizes a transformer-based diffusion model that accepts conditional feature vectors as input. It then generates corresponding motion sequences as depicted in Figure 2. The model incorporates a cross-attention mechanism, following (Saharia et al. 2022). We optimize the θ-parameterized score networks $\boldsymbol{s}_\theta$ with paired conditional signal c. The objective function is simplified as:

$$\mathbb{E}_{\boldsymbol{x},k} \|\boldsymbol{x} - \boldsymbol{s}_\theta(\boldsymbol{x}_k, k, \boldsymbol{c})\|_2^2. \tag{5}$$
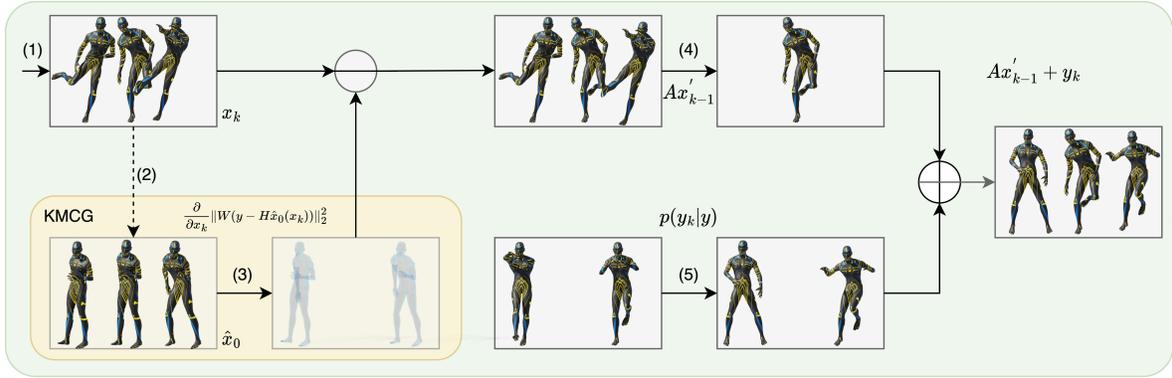
Figure 4: Visual schematic of the Keyframe Manifold Constraint Gradient (KMCGs) correction step. (1) unconditional reverse diffusion generates $\boldsymbol{x}_0$. (2) maps the noisy $\boldsymbol{x}_k$ to generate $\hat{\boldsymbol{x}}_0$. (3) the manifold constraint gradient $\frac{\partial}{\partial \boldsymbol{x}_k}\|\boldsymbol{W}(\boldsymbol{y}-\boldsymbol{H}\hat{\boldsymbol{x}}_0(\boldsymbol{x}_k))\|_2^2$ is applied to constraint on the context manifold. (4) Takes the orthogonal complement. (5) Samples from $p(\boldsymbol{y}_k\|\boldsymbol{y})$, then combines $\boldsymbol{A}\boldsymbol{x}'_{k-1}$ and $\boldsymbol{y}_k$ to further refine the motion transfer.

Algorithm 1: Motion style transfer with DDIBs

Input: source motion sample $\boldsymbol{x}^{(i)}\sim p_D^{(i)}$, conditional signal $\boldsymbol{c}$, source model $\boldsymbol{s}_\theta^{(i)}$, and target model $\boldsymbol{s}_\theta^{(j)}$
Transfer:
source to latent: $\boldsymbol{x}^{(l)} = \text{ODESolver}(\boldsymbol{x}^{(i)};\boldsymbol{c}, \boldsymbol{s}_\theta^{(i)}, 0, 1)$
latent to target: $\boldsymbol{x}^{(j)} = \text{ODESolver}(\boldsymbol{x}^{(l)};\boldsymbol{c}, \boldsymbol{s}_\theta^{(j)}, 1, 0)$
Output: $\boldsymbol{x}^{(j)}$, transferred motion in target domain.

Dual Diffusion Implicit Bridges. Unlike GAN-based approaches for style transfer (Yin et al. 2023b,c), which do not inherently ensure cycle consistency, the DDIBs (Su et al. 2022) leverage the connections between score-based generative models (Song et al. 2020) and the Schrödinger Bridges Problem (Chen, Georgiou, and Pavon 2016). To optimize cycle consistency across two domains, GAN-based methods require additional components during training, necessitating simultaneous access to multiple domains. This imposes constraints on both the system scalability and data privacy. In contrast, DDIBs establish deterministic bridges between distributions, operating as a form of entropy-regularized optimal transport. Consequently, cycle consistency is ensured up to the discretization errors of the ODE solvers, avoiding simultaneous access to multiple domains for consistency.

The DDIBs-based strategy is depicted in Figure 3 and Algorithm 1. The motion transfer process involves encoding source motion $\boldsymbol{x}^{(i)}$ and conditional signal $\boldsymbol{c}$ with the source diffusion model $\boldsymbol{s}_\theta^{(i)}$ to latent encoding $\boldsymbol{x}^{(l)}$ at the end time $t = 1$, and then decoding the latent encoding using the target model $\boldsymbol{s}_\theta^{(j)}$ to construct target motion $\boldsymbol{x}^{(j)}$ at time $t = 0$. Both steps are defined via ODEs. We then employ ODE solvers to solve the ODE and construct $\boldsymbol{x}^{(j)}$ at different times. In our experiment, we adopt DDIMs as the ODE solver. DDIMs generalize DDPMs via a class of non-Markovian diffusion processes that lead to the same training objective.

The trained diffusion models in DDIBs can be regarded as a summary of the domains of datasets. When translating, these models strive to generate motions in the target domain that is closest to the source motion in terms of optimal transport distances. This process is both a strength and a limitation of DDIBs. In image translation, when source and target domains share similarities, DDIBs typically succeed in identifying correct content. A similar phenomenon is observed in human locomotion domains, as patterns of locomotion across different domains are similar. DDIBs generally transfer correctly on human gaits. However, when datasets show less similarity (e.g., birds and dogs), DDIBs may struggle to produce translation results that accurately retain the postures. In human motion, dance movements present a greater complexity and lesser similarity. Consequently, when employing DDIBs for dance style transfer, the system sometimes fails to preserve the content of the movements while transferring the style.

Content Constrained Style Transfer. We propose incorporating additional corrective content to address the limitations of DDIBs discussed above:

$$\boldsymbol{y} = \boldsymbol{H}\boldsymbol{x}^i + \epsilon, \tag{6}$$

where $\boldsymbol{y}$ is the measured content for motion correction, and ϵ is the noise in the measurement. In our implementation, we extract keyframes from the source motion as corrective content to retain the postures from the source motion in the target motion. The keyframes are selected based on the acceleration of body joints.

The most straightforward method to incorporate additional information is explicitly adding these frames as part of model contexts. However, this approach has obvious drawbacks. It necessitates retraining the model, which negatively impacts the scalability of systems. Moreover, the inclusion of original motion frames as high-dimensional context can exacerbate the training stability. Nonetheless, we suggest this as an alternative way of enforcing domain information. A comparison to our proposed strategy can be found in the experimental evaluation section.

To avoid extra training, we use an additional corrective term by manifold constraints (Chung et al. 2022). We take the keyframes from the source motion sample as a constraint, which leads to the defined keyframe manifold constrained gradients (KMCGs). This term can synergistically operate with previous solvers to steer the reverse diffusion process closer to the context manifold. The reverse diffusion process in equation 4 can be replaced by:

$$\boldsymbol{x}_{k-1}' = \boldsymbol{f}(\boldsymbol{x}_k, s_\theta) - \alpha \frac{\partial}{\partial \boldsymbol{x}_k} \|(\boldsymbol{y} - \boldsymbol{H}\hat{\boldsymbol{x}}_0(\boldsymbol{x}_k))\|_2^2 + g(\boldsymbol{x}_k)\boldsymbol{z}, \tag{7}$$

$$\boldsymbol{x}_{k-1} = \boldsymbol{A}\boldsymbol{x}_{k-1}' + \boldsymbol{b}_k, \tag{8}$$

where $\hat{\boldsymbol{x}}_0$ is estimated $\boldsymbol{x}_0$. α depends on the noise covariance. Other parameters are defined by $\boldsymbol{A} = \boldsymbol{I} - \boldsymbol{H}^T\boldsymbol{H}$, $\boldsymbol{b}_k = \boldsymbol{H}^T\boldsymbol{y}$. Here we omit the target domain index j in the reverse process for the notation brevity.

Although keyframes in the target domain are not available in style translation problem formulation, we found that the keyframe context in the source motion as a manifold constraint can also significantly improve the performance of dance style transfer. We illustrate our scheme in Figure 4. The gradient term $\frac{\partial}{\partial \boldsymbol{x}_k}\|\boldsymbol{W}(\boldsymbol{y} - \boldsymbol{H}\hat{\boldsymbol{x}}_0(\boldsymbol{x}_k))\|_2^2$ incorporates the information of $\boldsymbol{y}$ so the gradient of corrective term stays on the context manifold.

Experiments

To validate the capabilities of our motion style transfer system, we are eager to answer the questions: 1) can the system achieve scalable training? 2) can KMCGs enhance content preservation? 3) will the introduction of KMCGs compromise the strength of style transfer? In what follows, we present both objective and subjective evaluations of the system on domain standard human motion datasets with results favouring the proposed system and KMCGs.

Experimental Setting and Datasets

Experimental Setting. We compare our proposed system with KMCGs (DDIBs-gradient) to the baseline DDIBs-vanilla. For a fair comparison, we also add another ablation setting, DDIBs-explicit, which also has a direct access to source domain keyframes but incoporated through cross-attention. To remain indpendent training, DDIBs-explicit trains each style model with its own keyframes. To evaluate the system performance, we quantitatively assess the transfer strength and content preservation based on two probabilistic divergence-based metrics. The concrete metrics are presented below. Qualitative user studies are also important in the evaluation of generative models as well. As a complementary, we performed a subjective evaluation in the form of online survey to evaluate human perceived quality in terms of naturalness and content preservation.

Datasets. We evaluate our system on the 100STYLE (Mason, Starke, and Komura 2022) locomotion database and the AIST++ (Tsuchida et al. 2019) dance database. We downsample both motion datasets to 30 fps and use 150-frame clips for experiments.

For the 100STYLE dataset, we employ forward walking movements that encompass 100 diverse styles, such as angry, neutral, stiff, etc. We transform the motion context into a form that includes 24 body joints with 3D rotations (72-dimensional vector), supplemented with frame-wise delta translation and delta rotation of the root joint (3-dimensional feature) as control signals.

The AIST++ dataset includes ten dance genres: break, pop, locking, waacking, middle hip-hop, LA hip-hop, house, krump, street-jazz, and ballet-jazz dance. We used the same pose representation as in (Tseng, Castellon, and Liu 2022) that represent dance as sequences of poses in 24-joint SMPL format (Loper et al. 2015) using the 6-DOF rotation representation and 4-dimensional binary foot contact label (151-dimensional vector). The music features are extracted with a frozen Jukebox model (Dhariwal et al. 2020), resulting in a 4800-dimensional feature.

Evaluations

Cycle Consistency. A desirable feature of the motion style transfer system is the cycle consistency property, which means transforming a motion sequence from the source domain to the target domain, and then back to the source, will recover the original data point in the source domain.
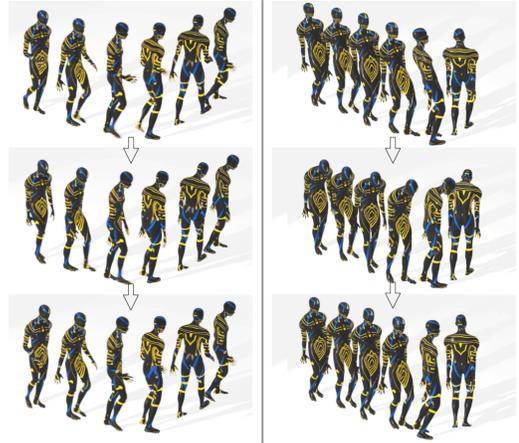

Figure 5: Cycle consistency: Translating the locomotion from the source style to the target style, and then back to the source style. Left: 'proud' to 'neutral' to 'proud'. Right: 'leanback' to 'depressed' to 'leanback'.

Figure 5 illustrates the cycle consistency property guaranteed by DDIBs-based systems. It concerns the locomotion dataset. Starting from the source domain, DDIBs first obtain the latent encodings and construct the motion in the target domain. Next, DDIBs do translations in the reverse direction, transforming the motion back to the latent and the source domain. After this round trip, motions are approximately mapped to their original patterns. This cycle consistency can also be observed in the dance dataset as shown in Figure 6. DDIBs-based transfer systems restore almost the exact same motion patterns as the original ones, which implies that noised latent encoding indeed carried informa-
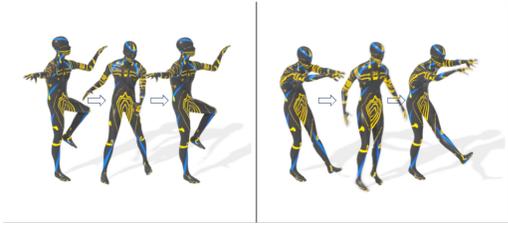
Figure 6: Cycle consistency: Translating the dance motion from the source style to the target style, and then back to the source style. Left: 'breaking' to 'hip-hop' to 'breaking'. Right: 'jazz-ballet' to 'waacking' to 'jazz-ballet'.

Table 1: Cycle consistency of motion style transfer on locomotion and dance dataset. The numbers are the averaged L_2 distances between the original motion sequences and the motion after cycle translation. The motion data are standardized to have unit variance.

Method/dataset	100STYLE	AIST++
DDIBs-vanilla	0.0198 ± 0.0083	0.0264 ± 0.0035
DDIBs-gradient	0.0192 ± 0.0079	0.0232 ± 0.0032

tion that may recover content. Table 1 reports quantitative results on cycle-consistent translation among both locomotion and dance styles, by DDIB-vanilla and DDIB-gradient. The reported values are negligibly small and endorse the cycle-consistent property in the human motion domain.

Transfer Performance. The primary objective of this system is to transfer the motion style from a specified source domain to a particular target domain. We conduct assessments from both objective and subjective perspectives to ensure a comprehensive evaluation of the complex motion patterns common in dance. For the objective evaluation, we analyze the transfer strength and content preservation by 90 dance sequences for each style. Style transfer was performed on each ablated system and among all possible style pairs. Two metrics based on Fréchet distance are adopted (Yin et al. 2023c), which is computed by Equation 9:

$$\text{FID} = \|\mu_r - \mu_g\|_2^2 + \text{Tr}(\Sigma_r + \Sigma_g - 2\sqrt{\Sigma_r \Sigma_g}), \quad (9)$$

where (μ_r, Σ_r) and (μ_g, Σ_g) are, respectively, the mean and the covariance matrix of the real and generated dance movement distribution.

Table 2: Quantitative objective evaluation: Fréchet Motion distance (FMD) and Fréchet Pose distance (FPD) are employed to evaluate the performance of the proposed transfer system and the baseline system.

Method	FPD↓	FMD↓
StarDance	0.4138	0.1816
DDIBs-vanilla	0.2904	0.1295
DDIBs-explicit	0.1748	0.1313
DDIBs-gradient	0.1208	0.1214

The Fréchet motion distance (FMD) quantifies transfer strength, which measures the extent that the motion is transferred from the source domain to the target domain. FMD
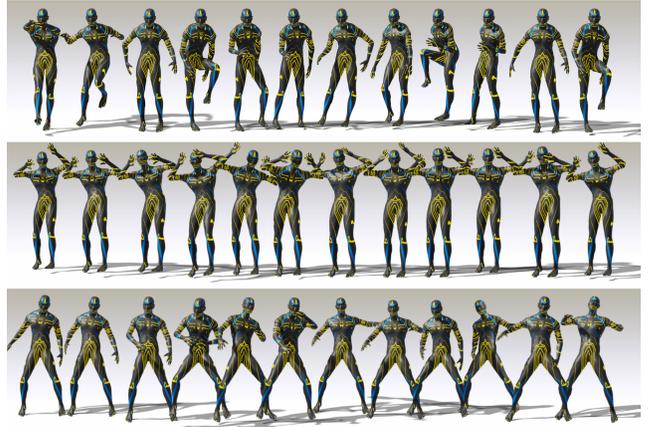

Figure 7: An example of transferring locking dance sequences (top) to krump dance using DDIBs-vanilla (middle) and DDIBs-gradient (bottom).

calculates the distance between the true and generated dance motion distributions. The body joint acceleration and velocity are adopted as style-correlated features. The Fréchet pose distance (FPD) evaluates content preservation, which measures how well the salient poses of the source motion sequence are preserved after the transfer. The salient poses are detected by local maxima in joint acceleration and are normalized with respect to the hip-centric origin.

The quantitative results for the proposed DDIBs-gradient and baselines across various dance style pairings are presented in Table 2. For StarDance, establishing a shared domain that bridges multiple styles presents a significant challenge. Other methods exhibit comparable performance on transfer strength, shows the introduced constraints do not compromise style transfer. However, in terms of content preservation, the baseline system DDIBs-vanilla encounters difficulties, as indicated by the higher FPD value relative to DDIBs-explicit and DDIBs-gradient. The improvement in performance can be attributed to the explicit imposition of context information or through manifold constraint gradients. An example of synthesized motion sequences showcasing the transfer of dance style from locking to krump is provided in Figure 7. The middle sequence is generated by DDIBs-vanilla, whereas the bottom sequence is created by DDIBs-gradient, which integrates keyframe context through manifold constraint gradients. By comparing the poses in each column, the DDIBs-gradient achieves a higher similarity in body postures to the source gesture in terms of body orientations and limb/body shapes, thus preserving more content. Interestingly, DDIBs-gradient outperforms DDIBs-explicit. One possible explanation is that DDIBs-gradient gets to impose source domain related information across multiple intermediate diffusion steps. DDIBs-gradient is even more favoured for the possibility of directly using pre-trained models without requiring additional training and hyperparameter searching.

Scalability and Data Privacy We further compare our proposed method with CycleDance concerning the scalabil-

Table 3: Comparison of scalability between CycleDance and DDIBs-based methods.

Dataset	Adult2Child	AIST+	100STYLE	PowerWash
Number of Styles	2	10	100	11080
Number of Models				
CycleDance	2	90	9900	122901720
DDIBs-based	2	10	100	11080

ity. Table 3 demonstrates that as the diversity of styles increases, our method scales the number of models linearly. In contrast, CycleDance surges the number of models quadratically, which would strain resources prohibitively. Consequently, CycleDance is appropriate for scenarios with a limited range of styles. As the scope of styles expands, our approach, characterized by its improved scalability, becomes increasingly advantageous. In instances where data sensitivity is a priority due to privacy issues, such as with the Emopain@Home (Olugbade et al. 2023) dataset in rehabilitation therapy or the PowerWash (Vuorre et al. 2023) dataset in player behavior analysis. Our method secures data privacy by decoupling the training procedure.

User Study. To obtain a more comprehensive evaluation of our system and the baseline system, we conducted a user study in addition to the objective evaluation, asking participants to rate three aspects: motion naturalness, transfer strength, and content preservation. An online survey was performed to evaluate the transfer tasks. In this study, 21 participants between were recruited. Participants were between 24 and 35 years of age, 71.4% male and 28.6% female. We blend videos for each source and target dance sequence. During the survey, participants were presented with a source dance video clip followed by a generated target dance clip. To avoid potential order effects, the order of the target dance clips was randomly shuffled. Each target dance clip was generated either from the DDIBs-vanilla or DDIBs-gradient system. The participants were allowed to view the clips multiple times.

For motion naturalness and transfer strength, the participants were asked to rate how much they agree with the statement that the motion is natural or transferred to the target style, using a 1-5 Likert scale, where "1" means strongly disagree, and "5" means strongly agree. Figure 8 indicates that the performance on these two aspects is similar. We use the two one-sided tests to determine if the means of the two systems evaluations are equivalent. The equivalence margin was set at $\delta = 0.5$. The differences are 0.2 and 0.35, which fall within the specified equivalence margin. The results showed there is no statistically significant difference between the means for the motion naturalness and transfer strength of DDIBs-vanilla and the DDIBs-gradient, which indicates that the constraints from source motion content will not lead to a decline in the motion naturalness and style transfer. For content preservation, the participants were asked to identify the features they believe are preserved between the source and the target motion, including orientations, limbs shape, body shape, and rhythmic patterns (Newlove and Dalby 2004). Figure 9 presents the overall statistics for the four aspects. To assess the statistical sig-
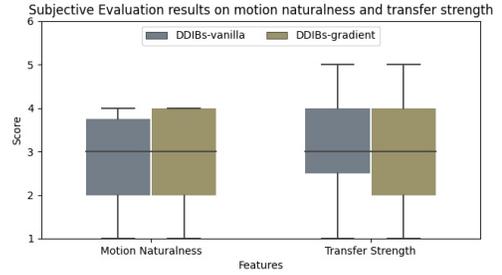

Figure 8: Subjective evaluation results in motion naturalness and transfer strength. Equivalence tests were performed.
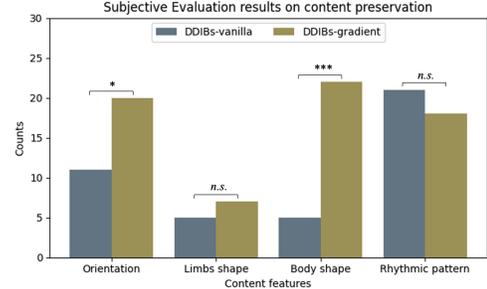

Figure 9: Subjective evaluation results in content preservation. Statistical significance was determined using the McNemar test (∗ means $p < 0.1$, ∗ ∗ ∗ means $p < 0.001$, and $n.s.$ means $p > 0.1$).

nificance of the differences between the DDIBs-vanilla and DDIBs-gradient systems, we conducted a McNemar test. The test showed no significant differences for 'limb shape' ($p = 0.25$), and 'rhythmic' ($p = 0.47$). However, there was a statistically significant difference for 'orientations' ($p = 0.07$) and 'body shape' ($p = 0.0000961$). Both systems scored high on 'rhythmic', attributable to the DDIBs-based transfer system, which naturally preserved the rhythmic patterns due to its grounding in pre-trained music-conditioned diffusion models. In terms of preserving body trunk shape, the DDIBs-gradient outperformed DDIBs-vanilla because the posture information was provided as a constraint. As limb movements are more complex, participants reported that this content was not well preserved in either system.

Conclusions

This study tackles the challenging task of motion style transfer. We propose a dual diffusion-based method with keyframe manifold constrain gradients. Our solution first converts the source motion into a latent encoding by pre-trained source model, and then produces motion in the target domain by the pretrained target model. This framework addresses the scalability and data privacy issues associated with GAN-based motion style transfer systems. We improve the transfer performance on content preservation by guiding the transfer process with keyframe manifold constraint gradients. Extensive evaluations demonstrate the efficacy and superior performance of the proposed method.

Acknowledgments

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.

Alexanderson, S.; Nagy, R.; Beskow, J.; and Henter, G. E. 2022. Listen, denoise, action! Audio-driven motion synthesis with diffusion models. *arXiv preprint arXiv:2211.09707*.

Amaya, K.; Bruderlin, A.; and Calvert, T. 1996. Emotion from motion. In *Graphics interface*, volume 96, 222–229. Toronto, Canada.

Brunner, G.; Wang, Y.; Wattenhofer, R.; and Zhao, S. 2018. Symbolic music genre transfer with cyclegan. In *2018 ieee 30th international conference on tools with artificial intelligence (ictai)*, 786–793. IEEE.

Chan, J. C.; Irimia, A.-S.; and Ho, E. 2020. Emotion Transfer for 3D Hand Motion using StarGAN.

Chen, L.-H.; Zhang, J.; Li, Y.; Pang, Y.; Xia, X.; and Liu, T. 2023. HumanMAC: Masked Motion Completion for Human Motion Prediction. *arXiv preprint arXiv:2302.03665*.

Chen, Y.; Georgiou, T. T.; and Pavon, M. 2016. On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint. *Journal of Optimization Theory and Applications*, 169: 671–691.

Chung, H.; Sim, B.; Ryu, D.; and Ye, J. C. 2022. Improving diffusion models for inverse problems using manifold constraints. *Advances in Neural Information Processing Systems*, 35: 25683–25696.

Dhariwal, P.; Jun, H.; Payne, C.; Kim, J. W.; Radford, A.; and Sutskever, I. 2020. Jukebox: A generative model for music. *arXiv preprint arXiv:2005.00341*.

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*, 1–11.

Gatys, L. A.; Ecker, A. S.; and Bethge, M. 2016. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2414–2423.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.

Hsu, E.; Pulli, K.; and Popović, J. 2005. Style translation for human motion. In *ACM SIGGRAPH 2005 Papers*, 1082–1089.

Kim, J.; Kim, J.; and Choi, S. 2022. Flame: Free-form language-based motion synthesis & editing. *arXiv preprint arXiv:2209.00349*.

Li, R.; Zhao, J.; Zhang, Y.; Su, M.; Ren, Z.; Zhang, H.; and Li, X. 2022. Magic: Multi Art Genre Intelligent Choreography Dataset and Network for 3D Dance Generation. *arXiv preprint arXiv:2212.03741*.

Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Loper, M.; Mahmood, N.; Romero, J.; Pons-Moll, G.; and Black, M. J. 2015. SMPL: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6): 1–16.

Mason, I.; Starke, S.; and Komura, T. 2022. Real-Time Style Modelling of Human Locomotion via Feature-Wise Transformations and Local Motion Phases. *arXiv preprint arXiv:2201.04439*.

Newlove, J.; and Dalby, J. 2004. *Laban for all*. Taylor & Francis US.

Olugbade, T.; Buono, R.; Potapov, K.; Bujorianu, A.; Williams, A.; de Ossorno Garcia, S.; Gold, N.; Holloway, C.; and Berthouze, N. 2023. The EmoPain@ Home Dataset: Capturing Pain Level and Activity Recognition for People with Chronic Pain in Their Homes.

Raab, S.; Leibovitch, I.; Tevet, G.; Arar, M.; Bermano, A. H.; and Cohen-Or, D. 2023. Single Motion Diffusion. *arXiv preprint arXiv:2302.05905*.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763. PMLR.

Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35: 36479–36494.

Smith, H. J.; Cao, C.; Neff, M.; and Wang, Y. 2019. Efficient neural networks for real-time motion style transfer. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 2(2): 1–17.

Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.

Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32.

Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.

Su, X.; Song, J.; Meng, C.; and Ermon, S. 2022. Dual diffusion implicit bridges for image-to-image translation. *arXiv preprint arXiv:2203.08382*.

Taylor, G. W.; and Hinton, G. E. 2009. Factored conditional restricted Boltzmann machines for modeling motion style. In *Proceedings of the 26th annual international conference on machine learning*, 1025–1032.

Tevet, G.; Raab, S.; Gordon, B.; Shafir, Y.; Cohen-Or, D.; and Bermano, A. H. 2022. Human Motion Diffusion Model. *arXiv*.

Tseng, J.; Castellon, R.; and Liu, C. K. 2022. EDGE: Editable Dance Generation From Music. *arXiv preprint arXiv:2211.10658*.

Tsuchida, S.; Fukayama, S.; Hamasaki, M.; and Goto, M. 2019. AIST Dance Video Database: Multi-Genre, Multi-Dancer, and Multi-Camera Database for Dance Information Processing. In *ISMIR*, volume 1, 6.

Unuma, M.; Anjyo, K.; and Takeuchi, R. 1995. Fourier principles for emotion-based human figure animation. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, 91–96.

Vuorre, M.; Magnusson, K.; Johannes, N.; Butlin, J.; and Przybylski, A. K. 2023. An intensive longitudinal dataset of in-game player behaviour and well-being in PowerWash Simulator. *Scientific Data*, 10(1): 622.

Wen, Y.-H.; Yang, Z.; Fu, H.; Gao, L.; Sun, Y.; and Liu, Y.-J. 2021. Autoregressive Stylized Motion Synthesis with Generative Flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13612–13621.

Witkin, A.; and Popovic, Z. 1995. Motion warping. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, 105–108.

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.

Yin, W.; Tu, R.; Yin, H.; Kragic, D.; Kjellström, H.; and Björkman, M. 2023a. Controllable Motion Synthesis and Reconstruction with Autoregressive Diffusion Models. *arXiv preprint arXiv:2304.04681*.

Yin, W.; Yin, H.; Baraka, K.; Kragic, D.; and Björkman, M. 2023b. Dance style transfer with cross-modal transformer. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 5058–5067.

Yin, W.; Yin, H.; Baraka, K.; Kragic, D.; and Björkman, M. 2023c. Multimodal dance style transfer. *Machine Vision and Applications*, 34(4): 1–14.

Zhang, F.; Ji, N.; Gao, F.; and Li, Y. 2023. DiffMotion: Speech-Driven Gesture Synthesis Using Denoising Diffusion Model. *arXiv preprint arXiv:2301.10047*.

Zhang, M.; Cai, Z.; Pan, L.; Hong, F.; Guo, X.; Yang, L.; and Liu, Z. 2022a. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv*.

Zhang, M.; Liu, C.; Chen, Y.; Lei, Z.; and Wang, M. 2022b. Music-to-Dance Generation with Multiple Conformer. In *Proceedings of the 2022 International Conference on Multimedia Retrieval*, 34–38.