

MoST: Motion Style Transformer between Diverse Action Contents

Boeun Kim^{1,2,3} Jungho Kim¹ Hyung Jin Chang³ Jin Young Choi²

¹Korea Electronics Technology Institute ²Seoul National University ³University of Birmingham

Abstract

While existing motion style transfer methods are effective between two motions with identical content, their performance significantly diminishes when transferring style between motions with different contents. This challenge lies in the lack of clear separation between content and style of a motion. To tackle this challenge, we propose a novel motion style transformer that effectively disentangles style from content and generates a plausible motion with transferred style from a source motion. Our distinctive approach to achieving the goal of disentanglement is twofold: (1) a new architecture for motion style transformer with ‘part-attentive style modulator across body parts’ and ‘Siamese encoders that encode style and content features separately’; (2) style disentanglement loss. Our method outperforms existing methods and demonstrates exceptionally high quality, particularly in motion pairs with different contents, without the need for heuristic post-processing. Codes are available at <https://github.com/Boeun-Kim/MoST>.

1. Introduction

Motion style transfer has received considerable attention for its potential to generate character animations reflecting human personalized characteristics in applications such as games, metaverse, and robotics. Our goal is to transfer stylistic characteristics from a source motion sequence (*style motion*) to a target motion sequence (*content motion*) without manually providing style labels. Performing this style transfer between motion sequences is a considerably more challenging task [1, 14, 20, 28] compared to generating motion when both content and style labels are manually provided [2, 25].

The most significant challenge is transfer failure, where the generated motion loses *content* of the *content motion* or does not reflect *style* of the *style motion*. Transfer failure often occurs when the *contents* of the *style motion* and *content motion* differ from each other, e.g. kick and punch. Certainly, its complexity arises from the significant variation in




Figure 1. Frequent failure cases in existing methods: (a) A result of MotionPuzzle [14] replicating *style motion*. (b) A result of Aberman *et al.* [1] showing complete failure with twisted motion. The character for the visualization is sourced from Mixamo [13].

style expression between different parts of the body for distinct *contents*. Despite the practical importance of transferring *styles* between different *content*, existing approaches have not primarily addressed these cases, nor have they been comprehensively evaluated, either qualitatively or quantitatively. Fig. 1 illustrates representative failure cases in existing methods. They fail to clearly disentangle *style* from source *content* and appropriately transfer it to the target *content*. Fig. 1 (a) illustrates the result of MotionPuzzle [14], where the generated motion, intended to be a jump, instead displays a kick. This method tends to replicate the motion pattern of the *style motion* rather than effectively transferring disentangled *style*. In Fig. 1 (b), Aberman *et al.* [1] utterly fails to generate a plausible motion. Aberman *et al.* recognized this issue, acknowledging that their method is less effective in motions other than walking and running.

The additional notable limitation found in recent studies [1, 14, 20, 28] is the requirement for heavy post-processing. These methods generate only the body pose of

each frame and cannot generate the global translation. Instead, they copy the global translation from the *content motion* in a post-processing step. However, this copying process degrades the quality of the generated motion, as highlighted by Park *et al.* [20] and Jang *et al.* [14] in their papers. Park *et al.* point out that inconsistencies may be caused between the generated pose and the copied global translation. Furthermore, methods of [1, 14, 20] heuristically grab the foot contact timings in the *content motion* and forcibly fix the generated motion to make the feet touch the ground at those specific timings. Without fixing, feet often penetrate the floor or float in the air in their generated motions.

To address the primary concern of transfer failure, we design a new framework called MoST and new loss functions for effectively transferring *style* between different *contents*. As *style* manifests in different patterns across various body parts based on *content*, MoST handles body part-specific style and content features. MoST comprises transformer-based Siamese motion encoders, a part-attentive style modulator (PSM), and a motion generator. In contrast to the previous methods [1, 14, 20] that build two separate networks for encoding content and style features, we introduce Siamese motion encoders capable of simultaneously extracting both features from an input motion. Subsequently, PSM modulates the raw style feature extracted from the *style motion* to align with the *content* of the *content motion* before being inserted into the generator. This modulation enables effective disentanglement of *style* from the *style motion* and its expression in the *content motion*. Furthermore, we propose new loss functions. The style disentanglement loss aims to distinctly separate *style* and *content*. This separation enhances the model’s robustness in transferring *style*, irrespective of the *content* of the *style motion*.

Our resulting motions are of high quality, eliminating the need for any heuristic post-processing. MoST achieves this by generating global translational motion along with pose sequences. In the proposed framework, the global translational feature and body part features exchange their information via attention, ensuring consistency between the generated poses and global translation. Moreover, to mitigate issues such as jittering and foot-skating, we introduce a physics-based loss. We evaluated the proposed model on two representative motion capture datasets [1, 30], verifying its performance both qualitatively and quantitatively. The main contributions of this study are summarized as follows:

- We design MoST, incorporating Siamese encoders and PSM, to effectively disentangle *style* from the source motion and align it with the target *content*.
- We propose novel loss functions to improve the model’s ability to disentangle *style* from *content* within the motion and generate plausible motion.
- MoST substantially outperforms existing methods, particularly when the two inputs have different con-

tents. Furthermore, our method achieves high fidelity in its output without the need for any post-processing.

2. Related Works

2.1. Motion and Image Style Transfer

Style transfer methods usually have been developed in the image domain first [4, 9, 12, 15, 18] and applied to the motion domain [1, 8, 10, 11, 20, 28]. In image style transfer, Gatys *et al.* [9] first obtained style and content features from a neural network rather than hand-crafted features. The style transfer model is optimized to yield features similar to both style and content features extracted from the given input images. Johnson *et al.* [15] proposed a real-time optimization framework by using the perceptual loss. The work of Holden *et al.* [10] was the first in motion style transfer, which adopts the optimization framework of Gatys *et al.* [9]. Holden *et al.* [11] and Du *et al.* [8] improved runtime and memory efficiency by replacing a training procedure similar to the approach of Johnson *et al.* [15].

Recently, Huang *et al.* [12] introduced AdaIN, which simply replaces style statistics in certain layers during training. Inspired by the work of Huang *et al.* [12], Aberman *et al.* [1] proposed temporally invariant AdaIN which is applied in the temporal convolution layers. Kingma *et al.* [18, 19] first introduced a generative flow [7] into an image generation framework. Few years later, Wen *et al.* [28] adopted the generative flow model into motion style transfer, generating plausible motions with less jittering. However, the method is limited in capturing the global style feature owing to the extraction of the style feature from the small window. Furthermore, the model has large parameters and is extremely slower than other methods [1, 20]. Choi *et al.* [4] proposed an image-to-image translation method aimed at multiple domains. Park *et al.* [20] adopted this approach regarding the styles as domains. The style encoder successfully separates distinctive style features in the latent space due to the head layers, which are constructed separately for each style. However, a style label should be injected into the network along with the style motion sequence in this method. Furthermore, pose twisting and jittering are frequently shown in the generated motion. The recent methods of motion style transfer [1, 20, 28] have common shortcomings. The methods encounter difficulties in transferring styles between motions with different contents, and they require hard post-processing, such as fixing foot contact or copying global translation.

Most recently, StyleFormer [29] and StyTr² [6] have successfully applied the transformer architecture to image style transfer, generating stylized images through cross-attention between content and style features. Taking inspiration from these works, we introduce the transformer architecture to the motion encoder and generator. Motion requires more

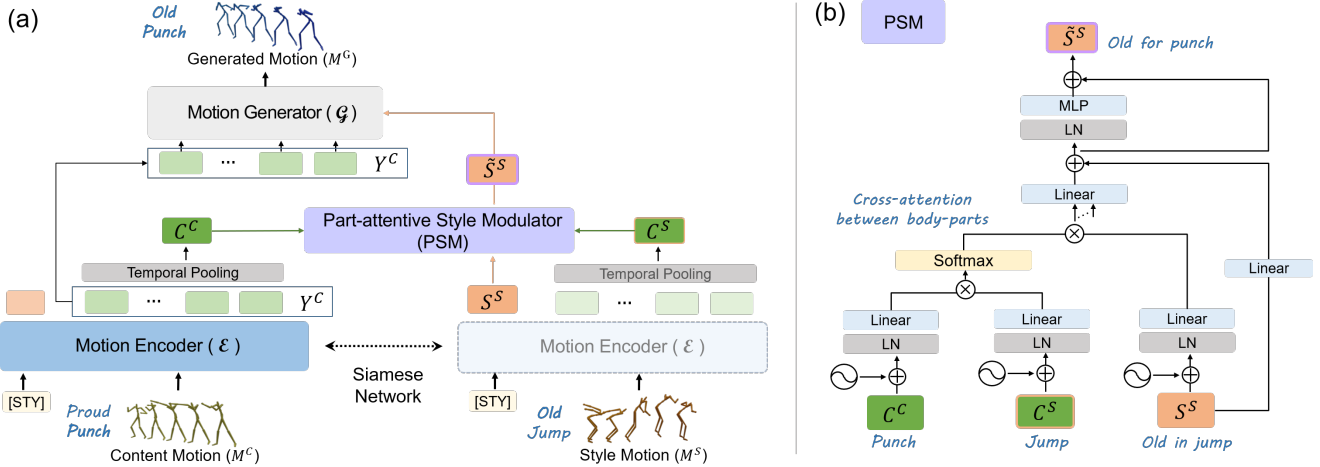


Figure 2. (a) Overall framework of MoST comprising Siamese motion encoders \mathcal{E} , motion generator \mathcal{G} , and part-attentive style modulator (PSM). PSM modulates style feature S^S under the condition of both contents of content motion and style motion, *i.e.*, C^C and C^S . \mathcal{G} generates final output motion with content dynamics feature Y^C and the modulated style feature \tilde{S}^S . (b) Detailed operations in PSM

precise modeling than images because even small temporal inconsistencies or incorrect joint movements are noticeable. Therefore, eliminating the original style and injecting a new style into the motion demands careful consideration. Our framework aims to disentangle content and style clearly and modulate style to be expressed effectively in the target motion, resulting in a plausible output motion.

2.2. Recent Studies on Motion Modeling

There are some methods that synthesize motion using given labels of style and content [2, 25]. Chang *et al.* introduced a diffusion-based model, and Tao *et al.* proposed an online synthesis method. Tang *et al.* introduced a method to stylize transitional motion between two frames [24]. Our goal is more complex compared to these methods, as we aim to transfer style between two input motions of diverse action types rather than using labels directly. Recently, there have been many studies that generate motion from a textual description [3, 5, 21, 26]. MotionCLIP [26] generates motions from arbitrary texts owing to the alignment of a human motion manifold and contrastive language-image pre-training (CLIP) [22] space. MLD [3] and MoFusion [5] generate realistic motions from long-sequence descriptions utilizing conditional diffusion models.

Similar to our task, MotionPuzzle [14] presents an innovative framework that transfers manually specified motion into a desired body part, allowing the combination of multiple target motions. However, the method has limitations in disentangling the style and content characteristics, resulting in the duplication of the *style motion* instead of transferring only the disentangled style. Nevertheless, MotionPuzzle provided essential inspiration for our work handling motion features by body parts for precise motion modeling.

3. Method

3.1. Overall Framework

The proposed framework aims to transfer the *style* from a given *style motion* M^S to a given *content motion* M^C to generated $M^G = \text{MoST}(M^C, M^S)$. As depicted in Fig. 2 (a), our framework comprises Siamese motion encoders \mathcal{E} , motion generator \mathcal{G} , and a part-attentive style modulator (PSM). \mathcal{E} is designed to encode both a content feature and a style feature simultaneously from a single motion. We utilize \mathcal{E} to encode both M^C and M^S as $\mathcal{E}(M^C) = \{S^C, Y^C\}$, $\mathcal{E}(M^S) = \{S^S, Y^S\}$, where $Y^{(\cdot)}$ denotes frame-level content dynamics features and $S^{(\cdot)}$ denotes a body part-specific style feature. The content feature of M^S , which is unobtainable by the style encoders in previous methods, is utilized as additional information for subsequent steps. Subsequently, S^S is modulated through the PSM as $\text{PSM}(S^S|C^S, C^C) = \tilde{S}^S$, where C^C and C^S are obtained by temporally pooling Y^C and Y^S . Finally, \mathcal{G} produces the generated motion M^G from M^C based on the modulated style feature \tilde{S}^S as $\mathcal{G}(Y^C, \tilde{S}^S) = M^G$.

3.2. Motion Representation

The motion sequence is defined by $M = \{m_t^j, m_t^{\text{root}}, v_t | j = 1, \dots, J; t = 1, \dots, T\}$, where J is the number of body joints and T is the maximum length of the model input. $\{m_t^j \in \mathbb{R}^7 | j = 1, \dots, J\}$ denotes local motion, where m_t^j represents the j -th joint vector in the t -th frame. $m_t^{\text{root}} \in \mathbb{R}^7$ and $v_t \in \mathbb{R}^4$ denote global translation and global velocity, respectively. Note that M^C , M^S , and M^G have the same form as M . m_t^j , m_t^{root} , and v_t contain both positional and angular representations. More details are available in the supplementary material.

We generate motion embeddings that are the input tensors of \mathcal{E} . Inspired by the method of [14], we grouped the joints into P body parts, accounting for the body structure. Moreover, to consider global translational motion, which existing methods overlook, we acquire a separate embedding of global translation in addition to body part embeddings. i -th body part embedding at t -th frame, is written as

$$\bar{p}_{(t,i)} = \text{FC}(p_{(t,i)}) \in \mathbb{R}^d, p_{(t,i)} = \parallel_j^{J^i} m_t^j \in \mathbb{R}^{7N_{J^i}}, \quad (1)$$

where $p_{(t,i)}$ is a concatenated vector of local joint vectors in the i -th body part. J^i is the set of joints within the i -th part, and N_{J^i} denotes the number of joints in J^i . d denotes the feature dimension of each body part embedding. $\text{FC}(\cdot)$ and \parallel indicate the fully-connected layer and the concatenation operation, respectively. The global translation embedding \bar{g}_t is obtained as

$$\begin{aligned} \bar{g}_t &= [\bar{m}_t^{\text{root}}; \bar{v}_t] \in \mathbb{R}^d, \\ \bar{m}_t^{\text{root}} &= \text{FC}(m_t^{\text{root}}) \in \mathbb{R}^{d/2}, \quad \bar{v}_t = \text{FC}(v_t) \in \mathbb{R}^{d/2}. \end{aligned} \quad (2)$$

Subsequently, the final embedding tensor X is expressed as

$$\begin{aligned} X &= [X_1, \dots, X_t, \dots, X_T]^\top \in \mathbb{R}^{T \times (P+1) \times d}, \quad (3) \\ X_t &= [\bar{p}_{(t,1)}, \dots, \bar{p}_{(t,i)}, \dots, \bar{p}_{(t,P)}, \bar{g}_t]^\top \in \mathbb{R}^{(P+1) \times d}. \end{aligned}$$

3.3. Siamese Motion Encoders

The frameworks of existing methods [1, 14, 20] consist of separate style encoders and content encoders. However, we introduce a Siamese architecture for both encoders to eliminate redundancy. Furthermore, we design the encoder to extract both style and content features from each motion, which are utilized in the next step. In contrast to the style encoders of Park *et al.* [20] and Wen *et al.* [28], which can only extract style from a short-range time window, our model extracts a global style feature across the entire motion sequence.

Motion encoder, $\mathcal{E}(M^{(\cdot)}) = \{S^{(\cdot)}, Y^{(\cdot)}\}$, comprises N stacked transformer blocks, and each block contains body part attention and temporal attention modules. We adopt the transformer architecture from the method of [16], but we employ part attention instead of joint attention. Additionally, we introduce a style token to aggregate the style features across an entire motion sequence. Let us denote Z^{n-1} as the input motion tensor of the n -th block. The input tensor of the first block is represented as

$$Z^0 = [X_{\text{sty}}, X_1, \dots, X_T]^\top \in \mathbb{R}^{(T+1) \times (P+1) \times d}, \quad (4)$$

where $X_{\text{sty}} \in \mathbb{R}^{(P+1) \times d}$ is a trainable style token. A positional embedding tensor is defined as $E = [E_1, \dots, E_t, \dots, E_{(T+1)}]^\top \in \mathbb{R}^{(T+1) \times (P+1) \times d}$ [16, 27]. In each block, the part attention module learns dependencies between body parts at the t -th frame as

$$\bar{Z}_t^{n-1} = \text{MHA}(\text{LN}(Z_t^{n-1} + E_t)) + (Z_t^{n-1} + E_t), \quad (5)$$

where $Z_t^{n-1} \in \mathbb{R}^{(P+1) \times d}$ is a sliced vector of Z^{n-1} . $\text{LN}(\cdot)$ denotes the layer normalization operator and $\text{MHA}(\cdot)$ denotes the multi-head attention operator. Subsequently, the temporal attention module learns the dependencies between frames as

$$Z^n = \text{MHA}(\text{LN}(\bar{Z}^{n-1} + E)) + (\bar{Z}^{n-1} + E), \quad (6)$$

where $\bar{Z}^{n-1} = [\bar{z}_{\text{sty}}^{n-1}, \dots, \bar{z}_T^{n-1}]^\top \in \mathbb{R}^{(T+1) \times (P+1) \times d}$ and $\bar{z}_t^{n-1} \in \mathbb{R}^{(P+1) \times d}$ is a vectorized form of \bar{Z}_t^{n-1} . We extract the style feature matrix from the output of the $(N-1)$ -th block as

$$S^{(\cdot)} = Z_{\text{sty}}^{N-1} \in \mathbb{R}^{(P+1) \times d}. \quad (7)$$

To obtain a content feature, $\hat{Z}^{N-1} = [Z_1^{N-1}, \dots, Z_T^{N-1}]^\top$ go through the last transformer block. In the last block, instance normalization (IN) is applied to remove style characteristics from the motion features [1, 12, 20] as

$$\check{Z}_{t,i}^{N-1} = \frac{\hat{Z}_{t,i}^{N-1} - \mu(\hat{Z}^{N-1})}{\sigma(\hat{Z}^{N-1})}, \quad (8)$$

where $\mu(\cdot)$ and $\sigma(\cdot)$ are operations that yield d -dimensional channel-wise mean and variance vectors from a tensor. \check{Z}^{N-1} passes through body part attention and temporal attention same as Z^{n-1} in Eqs.(5) and (6). Subsequently, the frame-level content dynamics feature is encoded as

$$Y^{(\cdot)} = [\check{Z}_1^N, \dots, \check{Z}_T^N]^\top \in \mathbb{R}^{T \times (P+1) \times d}. \quad (9)$$

3.4. Part-Attentive Style Modulator

Part-attentive style modulator (PSM) modulates the style feature S^S , which originates from M^S , to be more effectively expressed in M^C . The detailed operation is illustrated in Fig. 2 (b). Initially, S^S is expressed within the context of C^S , and our objective is to derive \tilde{S}^S , suitable for integration into C^C . We initiate the process by computing cross-attention between two content features, C^S and C^C , where $C^{(\cdot)} \in \mathbb{R}^{(P+1) \times d}$ is obtained by temporally pooling $Y^{(\cdot)}$. This cross-attention identifies how **style** should be transmitted from a specific body part of C^S to a corresponding body part of C^C , *e.g.* the legs in jump to the arm in punch. Subsequently, \tilde{S}^S is generated by linear combinations of part-specific style features of $S^S \in \mathbb{R}^{(P+1) \times d}$ according to the cross-attention. Consequently, PSM prevents the transmission of motion from an undesired body part of M^S , *i.e.* it only transfers **style** well-disentangled from **content** in M^S . As a result, the generated motion preserves the distinct **content** of M^C instead of confused combinations of motions from M^S and M^C . The cross-attention mechanism is applied as

$$\begin{aligned} \text{cross-MHA}(\bar{C}^C, \bar{C}^S, \bar{S}^S) &= (\parallel_{i=1}^h H_i) W_H, \quad (10) \\ H_i &= \text{Attn}(\text{LN}(\bar{C}^C) W_i^Q, \text{LN}(\bar{C}^S) W_i^K, \text{LN}(\bar{S}^S) W_i^V), \\ \bar{C}^C &= C^C + E_p, \quad \bar{C}^S = C^S + E_p, \quad \bar{S}^S = S^S + E_p, \end{aligned}$$

where $\text{Attn}(Q, K, V)$ indicates $\text{softmax}(QK^\top/\sqrt{d})V$. $W^Q, W^K, W^V \in \mathbb{R}^{d \times d'}$ and $W_H \in \mathbb{R}^{d' \times d}$ are weight matrices for projection. d' and h denote the projection dimension and the number of heads, respectively. As the index indicating the body part can serve as crucial information, we incorporate positional embedding $E_p \in \mathbb{R}^{(P+1) \times d}$ for each input. Finally, \tilde{S}^S is calculated as

$$\check{S}^S = \text{FC}(S^S) + \text{cross-MHA}(\bar{C}^C, \bar{C}^S, \bar{S}^S), \quad (11)$$

$$\tilde{S}^S = \check{S}^S + \text{MLP}(\text{LN}(\check{S}^S)), \quad (12)$$

where $\text{MLP}(\cdot)$ denotes multi-layer perceptron.

3.5. Motion Generator

The motion generator, denoted as $\mathcal{G}(Y^C, \tilde{S}^S) = M^G$, is composed of N stacked transformer blocks, similar to \mathcal{E} . Y^C is introduced as an input tensor for the first transformer block, denoted as $U^0 = Y^C$, where U^{n-1} represents the input tensor of the n -th block. Since Y^C encapsulates extensive information regarding the frame-level dynamics of each body part, **content** can be reliably conveyed to M^G . To incorporate \tilde{S}^S , we employ AdaIN [12], a method proven to be effective in previous works [1, 20]. In each transformer block, AdaIN is applied as

$$\hat{U}_{t,i}^{n-1} = \gamma \left(\frac{U_{t,i}^{n-1} - \mu(U^{n-1})}{\sigma(U^{n-1})} \right) + \beta, \quad (13)$$

$$\gamma = \text{FC}(\tilde{S}^S) \in \mathbb{R}^d, \quad \beta = \text{FC}(\tilde{S}^S) \in \mathbb{R}^d,$$

where β and γ are parameters modifying the channel-wise mean and variance, respectively. \hat{U}^{n-1} passes through part attention and temporal attention consistent with Z^{n-1} in Eqs.(5) and (6). The local motion, global translation, and global velocity information of the generated motion are reconstructed by the output feature tensor of the final transformer block, \hat{U}^N , as

$$\{\hat{m}_t^j \in \mathbb{R}^7 | j \in J^i\} = \text{FC}(\hat{U}_{t,i}^N), \quad (i \leq P), \quad (14)$$

$$\hat{m}_t^{\text{root}} = \text{FC}(\hat{U}_{t,(P+1)}^N) \in \mathbb{R}^7, \quad (15)$$

$$\hat{v}_t = \text{FC}(\hat{U}_{t,(P+1)}^N) \in \mathbb{R}^4. \quad (16)$$

Finally, the generated motion becomes

$$M_G = \{\hat{m}_t^j, \hat{m}_t^{\text{root}}, \hat{v}_t | j = 1, \dots, J; t = 1, \dots, T\}.$$

3.6. Loss

We introduce the *style disentanglement loss*, aiming to effectively separate **style** from **content**. It increases the robustness of our model in generating a well-stylized motion regardless of **content** of M^S . It is written as

$$L_D = \mathbb{E}_{M^C, M_a^S, M_b^S \sim \mathbb{M}} \left\| \text{MoST}(M^C, M_a^S) - \text{MoST}(M^C, M_b^S) \right\|_2, \quad (17)$$

where \mathbb{M} denotes the training set. We denote $\|\{\cdot\}\|_2$ as a sum of l_2 -norms of all vectors in set $\{\cdot\}$ for simplicity. L_D minimizes the discrepancy between the generated motions stylized by M_a^S and M_b^S , where M_a^S and M_b^S denote two *style motions* that have identical style labels but different content labels. L_D induces the model to clearly remove **content** from M^S , independent of the specific **content** present in M^S , thereby avoiding the blending of **content** into M^G .

Furthermore, we introduce the *physics-based loss* to mitigate pose jittering and improve foot-contact stability. It comprises regularization terms of velocity, acceleration, and foot contact as

$$L_{\text{phy}} = \lambda_{\text{vel}} R_{\text{vel}} + \lambda_{\text{acc}} R_{\text{acc}} + \lambda_{\text{foot}} R_{\text{foot}}, \quad (18)$$

where λ s denote weights of each regularization term. R_{vel} and R_{acc} are adopted by the method of [23] and [14]. Assuming that $\text{MoST}(M^G, M^C)$ should be identical to M^C , it should have the same foot contact timing as M^C . Therefore, we introduce foot contact regularization as

$$R_{\text{foot}} = \sum_j^{J_{\text{feet}}} \left(\frac{1}{N_{T_j}} \sum_t^{T_j} \|\dot{\hat{m}}_t^j\|_2 \right), \quad (19)$$

where $\dot{\hat{m}}_t^j = \hat{m}_{t+1}^j - \hat{m}_t^j$. J_{feet} includes the toe and heel joints of both feet and T_j is a set of frames where the foot joint j contacts the floor in M^C . N_{T_j} indicates the number of frames in T_j . Additionally, commonly used loss functions from existing methods [1, 14, 20] are utilized: adversarial (L_{adv}), reconstruction (L_{recon}), and cycle consistency (L_{cyc}) losses. The final loss is calculated as

$$L = \lambda_D L_D + \lambda_{\text{phy}} L_{\text{phy}} + L_{\text{pre}}, \quad (20)$$

$$L_{\text{pre}} = \lambda_{\text{adv}} L_{\text{adv}} + \lambda_{\text{recon}} L_{\text{recon}} + \lambda_{\text{cyc}} L_{\text{cyc}}.$$

4. Experiments

Implementation Details.

We evaluate our model using two different motion datasets captured by Xia *et al.* [30] (Xia dataset) and Aberman *et al.* [1] (BFA dataset) following [1], [20], and [28]. In both datasets, the skeleton has 31 joints, and we select 21 joints among them for training following [1, 20]. Xia dataset consists of eight motion styles and six motion contents. BFA dataset contains 16 styles and has long-length motion sequences not separated by the types of contents. The motion sequences are down-sampled into 60fps which is half of the original frame rate. The sequence is augmented by applying a random crop during training. \mathcal{E} and \mathcal{G} comprise three transformer blocks. We construct a discriminator \mathcal{D} as a single-block transformer for L_{adv} . T is set to 200. We set $\lambda_{\text{adv}}, \lambda_D, \lambda_{\text{vel}}, \lambda_{\text{foot}} = 1$, $\lambda_{\text{recon}}, \lambda_{\text{cyc}} = 3$, and $\lambda_{\text{acc}} = 0.1$. Adam optimizer [17] is used for training, with a learning rate of e^{-5} for \mathcal{E} , \mathcal{G} , and e^{-6} for \mathcal{D} . Our model




Figure 3. Description of evaluation metrics, using easy-to-recognize label notations. Note that our model uses only motion data

Table 1. Motion style transfer results evaluated in content consistency (CC) and style consistency++ (SC++) on Xia dataset [30]. The performances are reported in both cases: when M^C and M^S have identical content labels and different content labels

Methods	CC ↓			SC++ ↓		
	average	same cnt.	diff. cnt.	average	same cnt.	diff. cnt.
MotionPuzzle [14]	51.4	8.6	66.9	76.0	58.0	82.5
Aberman <i>et al.</i> [1]	46.0	42.0	47.0	189.7	220.8	178.4
Park <i>et al.</i> [20]	38.4	22.6	44.1	65.7	59.6	67.9
Wen <i>et al.</i> [28]	18.5	7.8	22.3	80.8	76.2	82.4
MoST	8.5	8.0	8.7	63.0	62.2	63.2

is trained for 300K iterations with a batch size of 8.

Evaluation Metrics.

Motion style transfer methods lack a standardized evaluation metric, hence, we adopt content consistency (CC) and style consistency (SC) proposed by Wen *et al.* [28]. Furthermore, we extend SC to SC++ to evaluate all the test motion pairs including the cases not covered by SC. Fig. 3 explains the metrics through straightforward examples. CC measures how effectively the **content** of M^C is retained in M^G . CC calculates the l_2 -distance between M^G and M^C and averages it across test pairs. In contrast, SC evaluates how well the **style** of M^S is reflected in M^G . SC calculates the similarity between M^G and M^S ; however, a problem arises when the **content** of M^S differs from that of M^C and M^G , as shown in ‘case 2’ of Fig. 3. We cannot find the reference motion to compare among the input motions. Therefore, Wen *et al.* only evaluate cases where the **content** labels of M^C and M^S are identical, *i.e.* ‘case 1’. To go beyond this limitation and evaluate all motion pairs, in SC++, M^G is compared to the pseudo ground truth motions in the training set, which share the identical style and content categories as those of M^G . Formulas are provided in the supplementary material. A crucial factor to note is that a style transition is deemed successful when both CC and SC++ are low, as these metrics exhibit a trade-off relationship. For instance, if a model outputs unmodified M^C as-is to generate M^G , SC++ will yield a high value, although CC may be low.

4.1. Comparison with State-of-the-Art Methods

We compare MoST to the state-of-the-art methods in motion style transfer, including the method of Aberman *et al.* [1], Park *et al.* [20], Wen *et al.* [28] and Jang *et al.* (MotionPuzzle) [14]. All the provided post-processing were employed in existing methods [1, 14, 20, 28], including copying global translation, removing foot skating, and correcting global velocity. In contrast, our result uses the raw output motion as-is. Note that the method of Park *et al.* [20] utilizes a style label in addition to the *style motion* during both

training and inference times, while other methods do not.

As indicated in Table 1, our method achieves the lowest values in both CC and SC++. Specifically, MoST exhibits a significant drop in CC. This suggests that the well-disentangled **style** is injected into the appropriate body part without compromising the **content** of M^C . In particular, MoST demonstrates strong performance in both scenarios, where the input motions have the same **content** or different **contents**. In contrast, other methods exhibit a significant performance degradation when handling input motions with different **contents**. MotionPuzzle tends to generate motion by replicating M^S . The high CC value highlights this limitation, indicating that **content** of M^C is not preserved. For the same reason, MotionPuzzle shows a low SC++ when the inputs have the same **contents** but a significantly high SC++ when they are different. The method of Aberman *et al.* [1] yields high CC and SC++ scores, aligned with qualitative results that indicate style transfer failures in numerous cases. The method of Park *et al.* [20] yields a satisfactory result in SC++; however, there is a significant disparity in CC compared to our method. In their results, twisted motion is often observed in certain body parts. Conversely, the method proposed by Wen *et al.* [28] demonstrates satisfactory results in CC but exhibits high SC++. Their generated motions are usually smooth and realistic; however, they do not effectively reflect the style.

We present the qualitative results in Fig. 4. We also have attached a video for better visualization in the supplementary material. Despite not employing heuristic post-processing, our generated motions are both plausible and well-stylized. They exhibit clear body movements corresponding to the desired **content**, while **style** is accurately expressed in the appropriate body parts. In contrast, the methods of Aberman *et al.*, Park *et al.*, and MotionPuzzle tend to fail in preserving the **content**, resulting in awkward motion where the **contents** of the two inputs are blended in different body parts. Moreover, it is difficult to discern the **style** expression in their results. While the method of Wen *et al.* produces plausible motions, it fails to effectively swap **styles**. MoST exhibits favorable results for the unseen style dur-

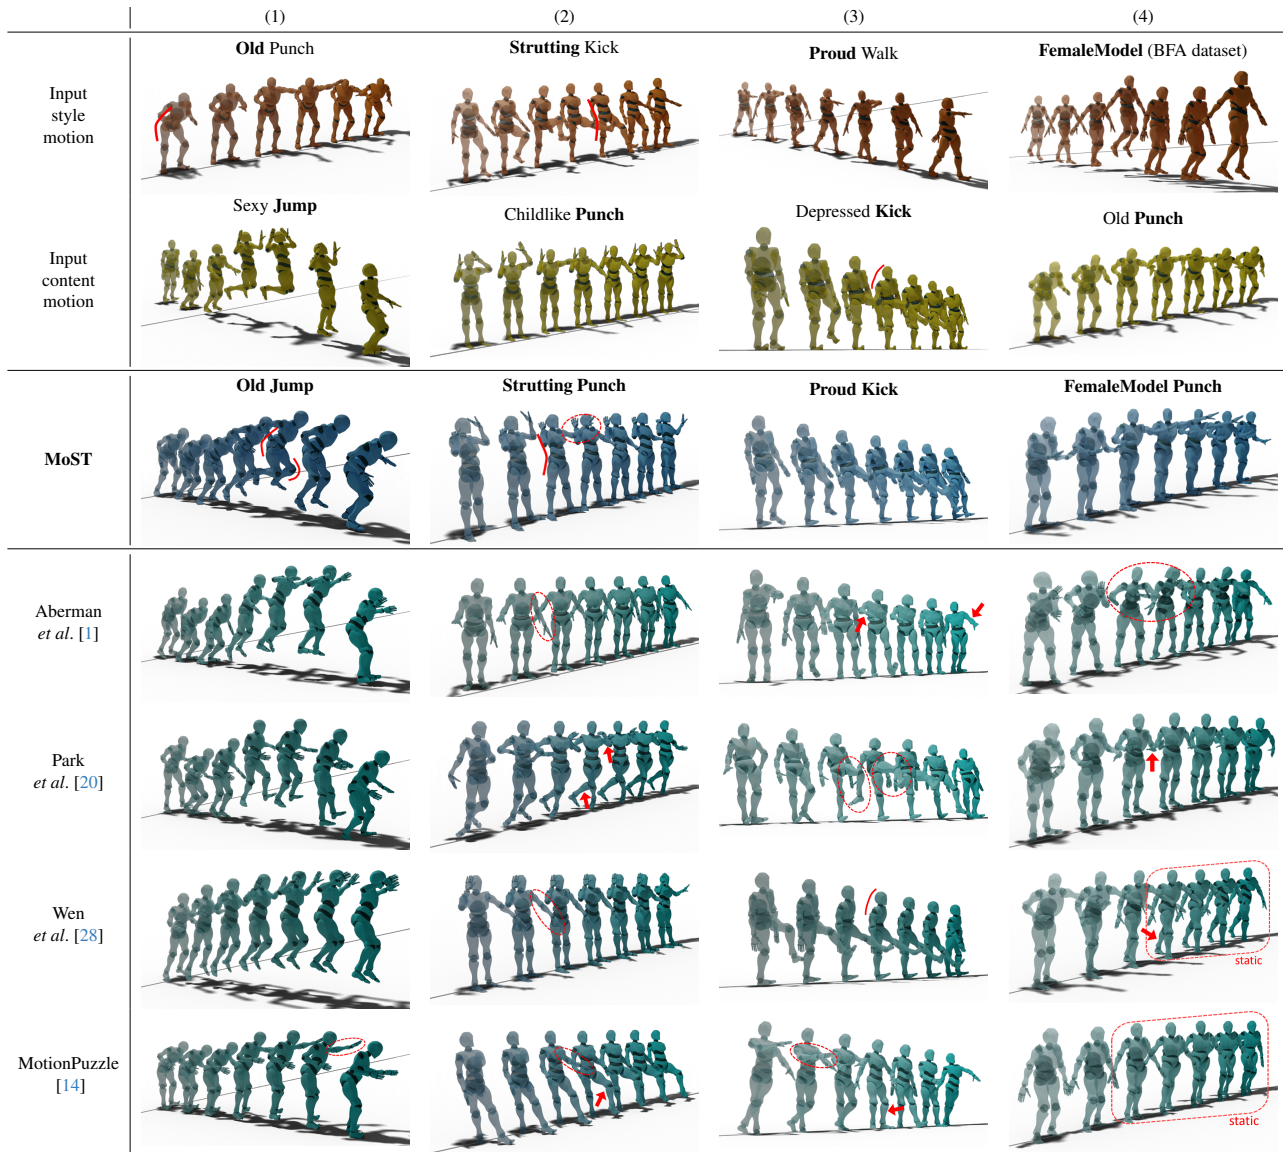


Figure 4. Qualitative results in Xia [30] and BFA [1] datasets. Please refer to the red indications. (1) Our method better reflects the style of *old* in comparison to other existing methods, accurately representing both the bent upper body and leg. (2) Other methods fail to preserve the *content* of *punch*, instead, they result in peculiar leg movements or body twists. On the other hand, our result accurately depicts *strutting punch*, where the upper body leans backward. (3) The results of [1] and [14] do not exhibit a *kick*, instead, their arm moves. [20] yields twisted leg movements. (4) Unlike our method, others fail to preserve the content of *punch*, resulting in vibrations in static poses or twists

ing training, which is the style in BFA dataset [1]. However, other methods fail to preserve *content* as shown in Fig. 4 (4)

4.2. Self Study

In Table 2, we evaluate the impact of the proposed loss function and PSM. When PSM is applied, CC significantly decreases, confirming that the style modulation prevents the generation of movement in undesired body parts. L_D substantially reduces both CC and SC^{++} , highlighting the importance of disentangling *style* and *content*. The model

achieves the best performance when both L_D and PSM are employed. The effectiveness of L_D is further demonstrated in the visualization of the style feature space. Fig. 5 (a) illustrates the space of \tilde{S}^S before applying L_D , where the style features have been mixed across certain categories. Fig. 5 (b) shows that L_D enforces complete separation among different styles. Furthermore, the transition from Fig. 5 (c) to (b) shows that the raw style features become distinctive through PSM. Fig. 6 shows how PSM modulates S^S according to different M^C . Fig. 6 (a) illustrates the space of S^S ,

Table 2. Ablation study for verifying the proposed L_D , and PSM on Xia dataset [30]. L_{pre} and L_{phy} are applied by default.

PSM	L_D	CC ↓	SC++ ↓
		37.4	69.5
✓		19.7	66.1
	✓	9.3	63.2
✓	✓	8.5	63.0

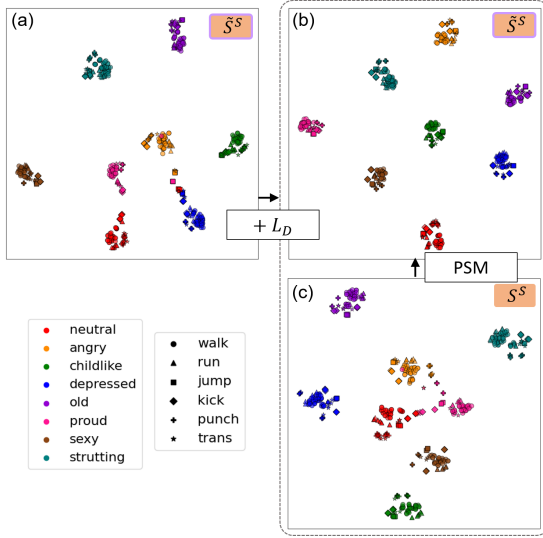


Figure 5. (a-b) Visualization of the modulated style feature (\tilde{S}^S) space of MoST in different loss settings. L_{pre} and L_{phy} are applied by default in (a). L_D is additionally introduced in (b). All training and testing data are used as *style motion*, and a single data point in the test set is employed for *content motion*. The spaces are projected in 2D through t-SNE. The samples are visualized with different shapes according to their content labels and different colors according to their style labels. (c) Space of S^S before PSM. All loss functions are applied




Figure 6. Visualization of S^S and \tilde{S}^S spaces projected through t-SNE for all $56 (M^C) \times 56 (M^S)$ motion pairs in the test set.

where each point represents an overlap of every $56 M^C$ instances, as S^S is yielded independently of M^C . After PSM, the data points are spread according to the paired M^C in the \tilde{S}^S space.




Figure 7. Cross-attention maps of PSM. The highest column (↓) indicates the body part in the *style motion* from which the style originates. The highest element (★) pinpoints the body part in the *content motion* that will receive the style. Traj. refers to global translation. The symbols ↓ and ★ are indicators related to traj.

In Fig. 7, we have visualized cross-attention maps of PSM. PSM identifies the activated body parts and transfers style from/to those parts. Specifically, in (a), style of the arm in the ‘Punch’ motion is transferred to the leg in the ‘Kick,’ as indicated by the attention map. In (b) and (c), high attention values are observed for the left leg in the ‘Kick’ and both legs in the ‘Walk,’ respectively. Notably, PSM incorporates global translation alongside body parts. The symbol ★ in (b) represents the application of style to traj, while ↓ in (c) represents the origin of style from traj.

5. Conclusions

MoST is designed to effectively disentangle style and content of input motions and transfer styles between them. The proposed loss functions successfully train MoST to generate well-stylized motion without compromising content. Our method outperforms existing methods significantly, especially in motion pairs with different contents.

Limitations and Future Works. Some samples may exhibit the foot contact problem, as the physics-based loss does not forcibly eliminate it. Testing time optimization could be applied for the perfect removal. While the transformer architecture offers advantages in handling sequences of diverse lengths, it requires specifying the maximum length of the model. We plan to expand the model for few-shot learning, to handle small datasets due to the high cost of motion data acquisition.

Acknowledgement

This work was supported by IITP grant funded by MSIT [No.RS-2022-00164860] and [No.2021-0-01343].

References

- [1] Kfir Aberman, Yijia Weng, Dani Lischinski, Daniel Cohen-Or, and Baoquan Chen. Unpaired motion style transfer from video to animation. *ACM Transactions on Graphics (TOG)*, 39(4):64–1, 2020. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [3](#)
- [2] Ziyi Chang, Edmund JC Findlay, Haozheng Zhang, and Hubert PH Shum. Unifying human motion synthesis and style transfer with denoising diffusion probabilistic models. *arXiv preprint arXiv:2212.08526*, 2022. [1](#), [3](#)
- [3] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18000–18010, 2023. [3](#)
- [4] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8188–8197, 2020. [2](#)
- [5] Rishabh Dabral, Muhammad Hamza Mughal, Vladislav Golyanik, and Christian Theobalt. Mofusion: A framework for denoising-diffusion-based motion synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9760–9770, 2023. [3](#)
- [6] Yingying Deng, Fan Tang, Weiming Dong, Chongyang Ma, Xingjia Pan, Lei Wang, and Changsheng Xu. Stytr2: Image style transfer with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11326–11336, 2022. [2](#)
- [7] Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014. [2](#)
- [8] Han Du, Erik Herrmann, Janis Sprenger, Noshaba Cheema, Somayeh Hosseini, Klaus Fischer, and Philipp Slusallek. Stylistic locomotion modeling with conditional variational autoencoder. In *Eurographics (Short Papers)*, pages 9–12, 2019. [2](#)
- [9] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. [2](#)
- [10] Daniel Holden, Jun Saito, and Taku Komura. A deep learning framework for character motion synthesis and editing. *ACM Transactions on Graphics (TOG)*, 35(4):1–11, 2016. [2](#)
- [11] Daniel Holden, Taku Komura, and Jun Saito. Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017. [2](#)
- [12] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017. [2](#), [4](#), [5](#)
- [13] Mixamo Inc. Mixamo. <https://www.mixamo.com>. [1](#)
- [14] Deok-Kyeong Jang, Soomin Park, and Sung-Hee Lee. Motion puzzle: Arbitrary motion style transfer by body part. *ACM Transactions on Graphics (TOG)*, 41(3):1–16, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [15] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 694–711. Springer, 2016. [2](#)
- [16] Boeun Kim, Hyung Jin Chang, Jungho Kim, and Jin Young Choi. Global-local motion transformer for unsupervised skeleton-based action learning. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IV*, pages 209–225. Springer, 2022. [4](#)
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [5](#)
- [18] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018. [2](#)
- [19] Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. Improved variational inference with inverse autoregressive flow. *Advances in neural information processing systems*, 29, 2016. [2](#)
- [20] Soomin Park, Deok-Kyeong Jang, and Sung-Hee Lee. Diverse motion stylization for multiple style domains via spatial-temporal graph-based generative model. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 4(3):1–17, 2021. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [3](#)
- [21] Mathis Petrovich, Michael J Black, and Gül Varol. Temos: Generating diverse human motions from textual descriptions. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, pages 480–497. Springer, 2022. [3](#)
- [22] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [3](#)
- [23] Davis Remppe, Leonidas J Guibas, Aaron Hertzmann, Bryan Russell, Ruben Villegas, and Jimei Yang. Contact and human dynamics from monocular video. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 71–87. Springer, 2020. [5](#), [2](#)
- [24] Xiangjun Tang, Linjun Wu, He Wang, Bo Hu, Xu Gong, Yuchen Liao, Songnan Li, Qilong Kou, and Xiaogang Jin. Rsm: Real-time stylized motion transition for characters. *arXiv preprint arXiv:2306.11970*, 2023. [3](#)
- [25] Tianxin Tao, Xiaohang Zhan, Zhongquan Chen, and Michiel van de Panne. Style-erd: responsive and coherent online motion style transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6593–6603, 2022. [1](#), [3](#)
- [26] Guy Tevet, Brian Gordon, Amir Hertz, Amit H Bermano, and Daniel Cohen-Or. Motionclip: Exposing human motion generation to clip space. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, pages 358–374. Springer, 2022. [3](#)
- [27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia

- Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 4
- [28] Yu-Hui Wen, Zhipeng Yang, Hongbo Fu, Lin Gao, Yanan Sun, and Yong-Jin Liu. Autoregressive stylized motion synthesis with generative flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13612–13621, 2021. 1, 2, 4, 5, 6, 7
- [29] Xiaolei Wu, Zhihao Hu, Lu Sheng, and Dong Xu. Style-former: Real-time arbitrary style transfer via parametric style composition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14618–14627, 2021. 2
- [30] Shihong Xia, Congyi Wang, Jinxiang Chai, and Jessica Hodgins. Realtime style transfer for unlabeled heterogeneous human motion. *ACM Transactions on Graphics (TOG)*, 34(4):1–10, 2015. 2, 5, 6, 7, 8, 4