

# Hierarchical Style-based Networks for Motion Synthesis

Jingwei Xu<sup>1\*</sup>, Huazhe Xu<sup>2\*</sup>, Bingbing Ni<sup>1\*\*</sup>, Xiaokang Yang<sup>1</sup>,  
Xiaolong Wang<sup>3</sup>, Trevor Darrell<sup>2</sup>

<sup>1</sup> Shanghai Jiao Tong University

<sup>2</sup> University of California, Berkeley

<sup>3</sup> University of California, San Diego

**Abstract.** Generating diverse and natural human motion is one of the long-standing goals for creating intelligent characters in the animated world. In this paper, we propose an unsupervised method for generating long-range, diverse and plausible behaviors to achieve a specific goal location. Our proposed method learns to model the motion of human by decomposing a long-range generation task in a hierarchical manner. Given the starting and ending states, a memory bank is used to retrieve motion references as source material for short-range clip generation. We first propose to explicitly disentangle the provided motion material into style and content counterparts via bi-linear transformation modelling, where diverse synthesis is achieved by free-form combination of these two components. The short-range clips are then connected to form a long-range motion sequence. Without ground truth annotation, we propose a parameterized bi-directional interpolation scheme to guarantee the physical validity and visual naturalness of generated results. On large-scale skeleton dataset, we show that the proposed method is able to synthesise long-range, diverse and plausible motion, which is also generalizable to unseen motion data during testing. Moreover, we demonstrate the generated sequences are useful as subgoals for actual physical execution in the animated world. Please refer to our project page for more synthesised results <sup>4</sup>.

**Keywords:** Long-range motion generation, Motion style transfer

## 1 Introduction

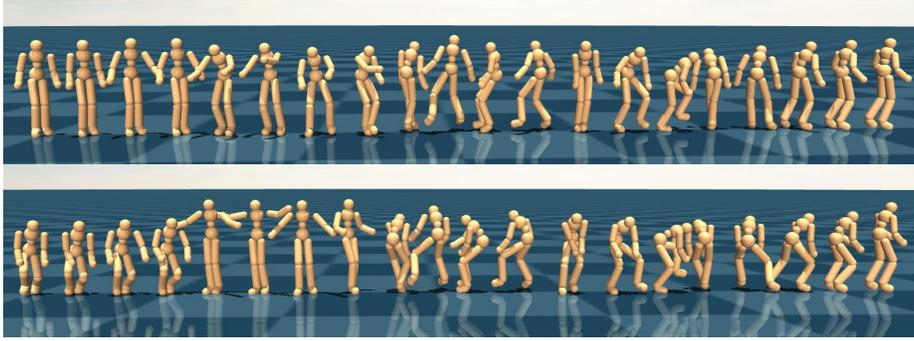
Human motion is naturally continuous in time and diverse between different individuals. The same action performed by different people can look very different, and even performed by the same actor twice could hardly be identical (Fig. 1). Capturing this diverse and stylized motion has been a long-standing demand in animation production and video games [10]. By generating this natural motion

---

\* denotes equal contribution.

\*\* corresponding author.

<sup>4</sup> <https://sites.google.com/view/hsnms/home>



**Fig. 1.** Visualization of motion sequences generated by proposed hierarchical style-based networks. The generation is achieved by transferring the “style” of multiple subsequences to new ones and then connect these clips smoothly in a temporal sequential manner. We present two synthesized motion sequences, which consist of the same motion “content”, but capture different motion “style”. The specific definition of “content” and “style” is described in the method part. The long-range and diverse motion generation is thus achieved by free-form composition of “content” and “style” and connecting short-range clips to a long-range one.

automatically, it provides useful tools for player customization of action skills in video games [10, 36]. However, synthesizing long-range and diverse motion is an extremely challenging task. On one hand, given the nature of multi-modal human behaviours, it is difficult to generate diverse motion without access to the distribution of motion states. On the other hand, it is common in sequential generative models that the error will accumulate through time, which restricts the maximum length of the generated motion sequence [6].

To generate natural and diverse motion, researchers have proposed statistical models based on optimization [34, 17, 2]. For example, Style Machine is a probabilistic generative model introduced in [2]. It is capable of generating motion with different styles (e.g., motion of novice ballet or modern dance of an expert) by optimizing with a cross-entropy framework. Motivated by this work, Motion texture [17] further proposed a two-level statistical model designed to capture diverse motion transitions, which achieves visually appealing state switching within seen sequence clips. However, these optimization-based methods are commonly restricted by the optimization complexity and can hardly be applied with large-scale dataset. Meanwhile, it is hard for these approaches to generalize to unseen distribution of data.

To adapt and generalize to large-scale data, deep neural networks are utilized in several recent works [11, 10, 36]. For example, Holden et al. [11] introduced to synthesize the motion sequence using deep networks by first taking the control signal as inputs, the outputs of the deep networks are then further edited via image-based style transfer techniques [7]. Although the adopted style transfer approach works well for transferring texture and color in the image domain, it

does not necessarily generalize to motion stylization. Without carefully designed training strategies for post-processing, unnatural pose and other artifacts, e.g., foot sliding, is commonly observed during motion synthesis [9].

In this work, we present hierarchical style-based networks, which leverage the large spectrum of human behaviours from unlabeled and unsegmented data to generate long-range, diverse and visually plausible motion sequences (as shown in Figure 1). Our framework is in a 2-level hierarchical structure: (i) Locally, our network first generates multiple short-range motion sequences independently; (ii) Globally, these short-range motion are then connected sequentially in time to a long-range motion sequence.

For obtaining diverse long-range motion sequences, we will first need to generate diverse short-range motion clips. To achieve this, we propose to disentangle the feature representation for short-range motion to two parts, representing the content and the style of the motion. The *content* refers to the specific action performed at each step, e.g., walking, running and dancing. The *style* refers to some pattern/property existing throughout the whole sequence, which keeps constant along with time. Take the walking motion for example: senior individuals and child walking sequences demonstrate two different styles for the same “walking” action. With this disentangled representation, we can obtain diverse motion from the free-form combination of motion content vectors and style vectors. However, labeling content and style is expensive and most of the time it is even hard to define the style semantically. Thus, instead of defining the style, we propose an unsupervised learning approach to automatically discover the disentanglement between content and style. Specifically, we utilize a bi-linear transformation to explicitly decompose the motion feature into two components, and the two features are then combined together to reconstruct the input motion in an auto-encoder structure. The key for disentanglement is that we enforce the intermediate style feature to be consistent in time.

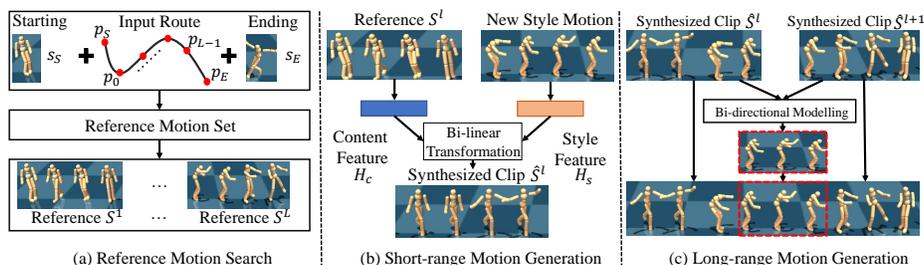
Given the synthesised short-range clips, we propose to connect them sequentially in time with motion interpolation between each two clips. Our network takes the starting and ending states from two different clips as inputs. To interpolate between these two states, we propose a parameterized representation which maintains the component scale of the generated sequences (e.g., bone lengths are fixed along time). Specifically, the representation is parameterized by a bi-directional LSTM model, which concurrently leverages the motion information of starting and ending states. In this way, we can obtain more plausible and visually appealing interpolation results.

Extensive experiments are conducted on large-scale human motion datasets [11]. We show that the proposed method outperforms existing motion generation baselines in terms of synthesis length, diversity and plausibility, e.g., being useful as sub-goals for actual physical execution in the animated world. We also demonstrate the proposed model is capable of synthesising novel motion based on unseen data without additional fine-tuning procedure, which indicates the generalization ability of our method.

## 2 Related Work

**Motion Interpolation.** Given start and end states, this task aims to synthesize intermediate states which smoothly translate between them [31]. For video interpolation [19, 18, 20, 22] where start and end states are two consecutive frames, the final result is expected to increase frame rate of original video to a higher value. Previous researches often utilize phase dynamics [20], flow based feature [19] and other motion information [22] to facilitate this task. Our work is different from this branch of work because there exists large motion gap between start and end states in our settings. Another branch of work is video completion [4, 18, 32]. It receives two *nonconsecutive* frames as input and aims to fill the motion gap between start and end states. [4] firstly attempts to solve this task and more specifically, propose to select out a rational path in the latent space with BFGS [3] algorithm. [18] incorporates the 3D convolution layers and LSTM network into a unified model, which tries to automatically find the optimal results for intermediate frames. Despite much progress has been made in this field, the high dimensional data (i.e., video frames) severely restricts video completion within *simple* and *seen* motion categories. However, we do not limit the start and end states belonging to the same sequence. Meanwhile, we expect the interpolated sequence as diverse as possible meanwhile with natural transition between synthesised states. This has not been deeply addressed in previous motion completion works [33, 15]. As a potential downstream application, our model could be used to construct motion planning [21] algorithm. Compared to goal-driven RL [14, 15], our model gets rid of requirements hard to achieve, i.e., known dynamics of agent, which is more general and applicable to more motion planning scenarios.

**Motion Synthesis in Computer Graphics.** In the context of computer graphics, there is a branch of researches [13, 23, 16, 28, 11] which also concentrate on motion generation, i.e., obtaining a continuous trajectory from a discrete set of poses. Our work shares similar target with this branch of work. However, we would like to emphasise that these works [28, 13, 23] are in parallel with ours and have completely different research routine on this task. More specifically, graphics methods [2, 17] focus on formulating an optimal and explicit statistic modelling framework, which is not easy scale-up to large datasets and unseen motion. For example, [2] introduce a fully data-driven method for articulated motion generation, which needs online optimization if given new demonstration data as inputs. The involved learning procedure is relatively more complex compared to ours and needs careful hyperparameter tuning. On the contrary, our model combines both advantages of deep model and purely data-driven method. Another related work [17] proposes a statistical model for approximation of original motion distribution, which is represented by a transition matrix indicating the possibility of state change. This method requires all motion data are available for generation, which is in turn restricted to seen data. Differently, our model could generalize to unseen data, which implies more practical value for downstream application.



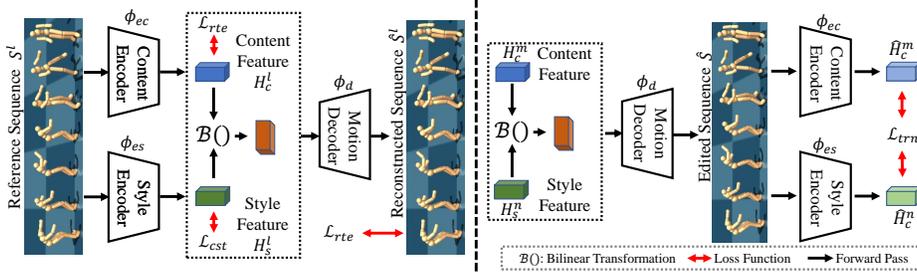
**Fig. 2.** Proposed framework for hierarchical style-based motion synthesis: (a) *Reference Motion Search*: We search for reference subsequences in the training set; (b) *Short-range Motion Generation*: a novel subsequence is then generated based on reference one with motion style transfer; (c) *Long-range Motion Generation*: all synthesized subsequences are connected together in temporal order with bi-directional modelling.

### 3 Method

In this work, we aim to synthesize visually natural, diverse and long-range motion sequence constrained by a pair of starting/ending states in a hierarchical manner. We denote the motion sequence with length of  $M$  as  $\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_i, \dots, \mathbf{s}_M]$ , where  $\mathbf{s}_i$  is the state at time stamp  $i$ . The starting and ending states are denoted as  $\mathbf{s}_S$  and  $\mathbf{s}_E$  respectively. Note that  $\mathbf{s} \in \mathbb{R}^{J \times 3}$ , where  $J$  refers to the number of joints, which is represented with the x-y-z Cartesian coordinate.

Our motion synthesis framework contains 3 steps as shown in Fig. 2: (i) *Reference Motion Search*. Given the input starting/ending states, we define *Reference Motion Search* as first dividing the route to  $L$  segments by adding  $L - 1$  sub-goals in between on the ground; then, for a segment  $l$ , sampling 1 reference subsequences  $\mathbf{S}^l \in \mathbb{R}^{M \times J \times 3}$  from the dataset.<sup>5</sup> Each sub-goal is represented by a spatial point  $\mathbf{p}^l \in \mathbb{R}^2, l = 1, \dots, L - 1$ . We denote  $\mathbf{p}_S$  and  $\mathbf{p}_E$  as the projected locations of root joint of  $\mathbf{s}_S$  and  $\mathbf{s}_E$ . (ii) *Short-range Motion Generation*. Toward diverse motion generation, when generating subsequence  $\hat{\mathbf{S}}^l$ , we advocate a synthesis paradigm in the motion style transfer manner, which keeps the content identical to that of  $\mathbf{S}^l$  while changing the style based on another subsequence that is randomly sampled from the dataset. (iii) *Long-range Motion Generation*. Given the short-range motion subsequences, we connect each two consecutive subsequences by adding a transitional motion in between. We denote the transitional motion between  $\hat{\mathbf{S}}^l$  and  $\hat{\mathbf{S}}^{l+1}$  as  $\hat{\mathbf{S}}^{l,l+1} \in \mathbb{R}^{N \times J \times 3}$ , where  $N$  is the length of the sequence. The final long-range generated sequence is presented as:  $[\mathbf{s}_S, \hat{\mathbf{S}}^{0,1}, \hat{\mathbf{S}}^1, \dots, \hat{\mathbf{S}}^l, \hat{\mathbf{S}}^{l,l+1}, \hat{\mathbf{S}}^{l+1}, \dots, \hat{\mathbf{S}}^L, \hat{\mathbf{S}}^{L,L+1}, \mathbf{s}_E]$ .

<sup>5</sup> Please refer to the supplementary material for detailed description about the sampling procedure.



**Fig. 3.** Proposed model for short-range motion generation. Left part: We train an auto-encoder in self-supervised way, where bi-linear transformation is utilized for unsupervised content and style disentanglement. Right part: Training objective ( $\mathcal{L}_{trn}$ ) emphasising the consistency at feature level is incorporated into the proposed model. Note that the content and style features are from different subsequences if trained with  $\mathcal{L}_{trn}$ .

### 3.1 Short-range Motion Generation

In this section, we present the procedure of generating short-range motion clips. This part is formulated as a motion style transfer task, i.e., new motion clips are synthesised via altering the style of reference subsequence  $\mathbf{S}^l$  while keeping their motion contents unchanged. Recall that the *content* refers to the specific action performed at each step, e.g., walking, running and dancing, while the *style* refers to some pattern/property existing throughout the whole sequence, which keeps constant along with time. The free-form combination of content and style information is used for diverse synthesis of subsequences  $\hat{\mathbf{S}}^l$ . Without annotation of style, we first propose to learn corresponding representations in a disentangling manner.

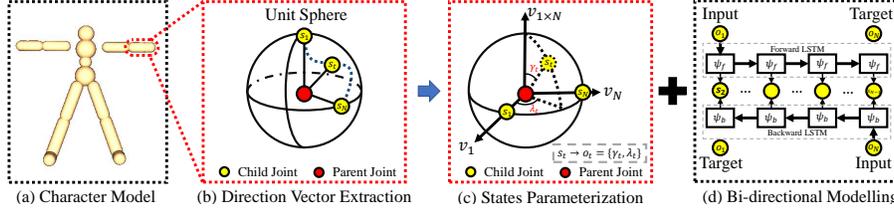
**Content and Style Disentanglement.** As shown in Fig. 3, two encoders,  $\phi_{ec}$  and  $\phi_{es}$ , are used to extract the content and style features as follows,

$$\mathbf{H}_c^l = \phi_{ec}(\mathbf{S}^l), \mathbf{H}_s^l = \phi_{es}(\mathbf{S}^l), \quad (1)$$

where  $\mathbf{S}^l$  is input subsequence,  $\mathbf{H}_c^l = \{\mathbf{h}_{c,i}^l\}_{i=1}^M$  and  $\mathbf{H}_s^l = \{\mathbf{h}_{s,i}^l\}_{i=1}^M$  are the content and style features respectively. Inspired by the success of content and style separation in character and image [29] fields, we propose to reconstruct the motion with a bi-linear transformation scheme:

$$\hat{\mathbf{s}}_i^l = \phi_d(\mathcal{B}(\mathbf{h}_{c,i}^l, \mathbf{h}_{s,i}^l)) \quad (2)$$

where  $\mathcal{B}(\cdot)$  is the bi-linear transformation [29], i.e.,  $\mathcal{B}(\mathbf{h}_{c,i}^l, \mathbf{h}_{s,i}^l) = \mathbf{h}_{c,i}^l \mathcal{W}(\mathbf{h}_{s,i}^l)^\top$ .  $\mathcal{W} \in \mathbb{R}^{C_{ctn} \times C_{out} \times C_{sty}}$  is a bi-linear transformation weight.  $\phi_d$  is the motion decoder, the subscript  $i$  indicates time step  $i$ . As a two-factor method, the bi-linear transformation possesses an elegant mathematical property, i.e., separability: their outputs are linear in either component when the other is kept unchanged. Facilitated by bi-linear transformation, the contribution of two components can



**Fig. 4.** Proposed model for long-range motion generation. Starting with skeleton model of human subject (leftmost), we extract the corresponding direction vector from 3D coordinate (second part), which are parameterized by the starting/ending states of one subsequence (third part). This leads to more compact solution space with learnable parameters  $(\gamma, \lambda)$ , which are generated by the proposed bi-directional model.

be effectively disentangled and fused into a representative latent feature that is naturally generalizable to unseen factor modes such as new contents [29].

**Training Objective.** Encoder and decoder are jointly trained by a L2 reconstruction loss, i.e.,  $\mathcal{L}_{rec} = \|\mathbf{S}^l - \hat{\mathbf{S}}^l\|_2^2$ . To guarantee the style consistency of the reconstructed subsequence, an explicit L2 loss is incorporated into training procedure, i.e.,  $\mathcal{L}_{cst} = \|\hat{\mathbf{h}}_{s,i}^l - \hat{\mathbf{h}}_{s,j}^l\|_2^2$ , where  $i, j$  are randomly sampled two style features corresponding to time step  $i$  and  $j$  respectively. Meanwhile, the moving route  $\mathbf{c}^l$  (including root position and velocity, foot position and velocity [11]) is extracted as a part of content feature and trained with a L2 loss to prevent foot sliding, i.e.,  $\mathcal{L}_{rte} = \|\mathbf{c}^l - \hat{\mathbf{c}}^l\|_2^2$ , where  $\mathbf{c}^l$  is the input control signal. Note that the moving route is treated as a part of content feature  $\mathbf{h}_c^l$ , i.e.,  $\mathbf{c}^l \in \mathbf{h}_c^l$ .

At each iteration we sample a new batch of data for training, to further guarantee the consistency at feature level. Suppose a subsequence (denoted as  $\hat{\mathbf{S}}$ ) is generated with the content  $\mathbf{H}_c^m$  of  $\mathbf{S}^m$  and style  $\mathbf{H}_s^n$  of  $\mathbf{S}^n$ .  $\hat{\mathbf{S}}$  is subsequently fed into  $\phi_{ec}$  and  $\phi_{es}$  to obtain reconstructed features, i.e.,  $\hat{\mathbf{H}}_c^m$  and  $\hat{\mathbf{H}}_s^n$ . This part is trained with  $\mathcal{L}_{trn} = \|\mathbf{H}_c^m - \hat{\mathbf{H}}_c^m\|_2^2 + \|\mathbf{H}_s^n (\mathbf{H}_s^n)^\top - \hat{\mathbf{H}}_s^n (\hat{\mathbf{H}}_s^n)^\top\|_2^2$ . The latter part for style consistency is inspired from Gram Matrix utilized in image translation [7]. The high-order statistics of style feature has been proven to be more critical for translation in image domain. The final training objective is presented as follows,

$$\mathcal{L} = \mathcal{L}_{rec} + 0.01\mathcal{L}_{cst} + 0.5\mathcal{L}_{rte} + \mathcal{L}_{trn}. \quad (3)$$

Note that both the encoder and decoder are jointly trained under the supervision of  $\mathcal{L}_{trn}$ . We present the model architecture details in the supplementary material. Please refer to it.

### 3.2 Long-range Motion Generation

This part refers to connecting edited short-range subsequences  $(\hat{\mathbf{S}}^l, \hat{\mathbf{S}}^{l+1})$  into a long-range one with interpolated subsequence  $\hat{\mathbf{S}}^{l,l+1}$ . To achieve smooth and natural transition, we propose to parametrize the original 3D coordinate space

as shown in Fig. 4, which leads to a more compact output space consisting of parameterized representation rather than motion states. Details are shown as follows.

**Parameterized Representation.** Given the state  $\mathbf{s}_t$  at time step  $t$  we first obtain the direction vector  $\mathbf{v}_t(p, q) = \frac{\mathbf{s}_t(p) - \mathbf{s}_t(q)}{\|\mathbf{s}_t(p) - \mathbf{s}_t(q)\|_2}$ , where  $(p, q)$  corresponds to one child-parent pair of character joint according to the skeleton topology (Fig. 4(b)). The transition subsequence is generated in the direction vector space. Given two vectors  $(\mathbf{v}_1, \mathbf{v}_N)$  as bases (the length of transition is assumed to be  $N$ ), interpolation is essentially the combination of  $\mathbf{v}_1$  and  $\mathbf{v}_N$ . The interpolation procedure is inspired from Quaternion Slerp [27] but generalized to non-linear situation. As shown in Fig. 4(c), supposing  $\mathbf{v}_1$  and  $\mathbf{v}_N$  are two bases in 3D space, the third basis is obtained by outer product of them:  $\mathbf{v}_{1 \times N} = \mathbf{v}_1 \times \mathbf{v}_N$ . Therefore arbitrary direction vector  $\mathbf{v}_t$  could be represented in the parameterization manner as follows,

$$\mathbf{v}_t = \frac{\sin(1 - \gamma_t)A}{\sin A} \left( \frac{\sin(1 - \lambda_t)\Omega}{\sin \Omega} \mathbf{v}_1 + \frac{\sin(\lambda_t)\Omega}{\sin \Omega} \mathbf{v}_N \right) + \frac{\sin(\gamma_t A)}{\sin A} \mathbf{v}_{1 \times N}, \quad (4)$$

where  $\Omega$  is directed angle between  $\mathbf{v}_1$  and  $\mathbf{v}_N$ ,  $A = \pi/2$  because of outer product.  $\mathbf{v}_t$  is thus parameterized by  $\mathbf{v}_1$  and  $\mathbf{v}_N$  with  $(\gamma_t, \lambda_t)$ . Note that  $\gamma_t, \lambda_t$  are two learnable parameters in our work, which is modelled by a bi-directional LSTM [8] introduced in following paragraph.

Intuitively, the interpolation defined by Eqn. 4 is analogous to locating on the earth with the longitude and latitude. The Eqn. 4 conducts quaternion slerp twice, where the first one, inside the brackets, decides the ‘‘longitude’’ ( $\lambda$ ) and the second one, outside the brackets, decides the ‘‘latitude’’ ( $\gamma$ ). One one hand, the output dimension at each time stamp is reduced from 4J (quaternion) to 2J ( $\gamma, \lambda$ ), leading to a more compact solution space. On the other hand, outputs are naturally valid, i.e., unit vector required by direction vector, which avoids additional normalization procedure.

**Bi-directional Modelling.** As shown in Fig. 4(d), we utilize two LSTM [8] networks (forward and backward) to achieve natural and plausible interpolation in the parametrization space ( $\mathbf{o} = \{\lambda, \gamma\}$ ). The forward one (denoted as  $\psi_f$ ) predicts future states with  $\mathbf{o}_1$  as origin and  $\mathbf{o}_N$  as target. The counterpart (denoted as  $\psi_b$ ) takes the reverse direction, i.e.,  $\mathbf{o}_N$  as origin and  $\mathbf{o}_1$  as target. The whole modelling procedure is executed as follows,

$$\hat{\mathbf{o}}_{t+1, f} = \psi_f(\hat{\mathbf{o}}_{t, f}, \mathbf{o}_N), \hat{\mathbf{o}}_{t-1, b} = \psi_b(\hat{\mathbf{o}}_{t, b}, \mathbf{o}_1), \quad (5)$$

The complementary pair of forward/backward outputs, i.e.,  $(\hat{\mathbf{o}}_{t, f}, \hat{\mathbf{o}}_{N-t, b})$  are firstly converted back to the x-y-z coordinate space  $(\hat{\mathbf{s}}_{t, f}, \hat{\mathbf{s}}_{N-t, b})$  with forward kinematics [25], and then fused as follows,

$$\hat{\mathbf{s}}_t = \psi_{fse}(\hat{\mathbf{s}}_{t, f}, \hat{\mathbf{s}}_{N-t, b}), \mathcal{L}_{pos} = \|\mathbf{s}_t - \hat{\mathbf{s}}_t\|_2^2. \quad (6)$$

With supervision signal in the position space, the topology information of character skeleton is better utilized and motion generation artifacts, e.g. foot sliding, could be explicitly punished during training procedure.

**Adversarial Training: Generalizing beyond Training Sequence.** Above interpolation is trained with starting/ending pair belonging to the same sequence, i.e., with ground truth. However, proposed model should be tested with arbitrary state pairs for practical usage. It lacks ground truth signal for training under such condition. We thus utilize adversarial training to facilitate generalization ability of motion interpolation. More specifically, with  $\mathbf{s}_1, \mathbf{s}_N$  sampled from different sequences,  $\psi_{f,b}$  give prediction of intermediate states as described above, which are regarded as fake sequences. A motion discriminator  $D$  is further proposed for adversarial training as follows,

$$\mathcal{L}_D = \frac{1}{2}((1 - D(\mathbf{S}))^2 + D^2(G(\mathbf{s}_1, \mathbf{s}_N))), \mathcal{L}_G = (1 - D(G(\mathbf{s}_1, \mathbf{s}_N)))^2, \quad (7)$$

where  $G$  refers to proposed interpolation model in this section. The discriminator architecture follows the work of Chen et.al. [5], i.e., built based on residual block and designed for synthesising more realistic 3D poses.

### 3.3 Implementation Details

We implement the proposed framework with Pytorch [24] platform. Note that the long-range and short-range generation models are trained disjointly, but could be tested in the end-to-end manner without extra overhead. We empirically get suboptimal results with joint training, which is also time consuming to converge. The short-range part is trained w.r.t. Eqn. 3. As for the long-range generation part, the adversarial training scheme is adopted, i.e., alternatively supervised by  $\mathcal{L}_{pos} + 0.01\mathcal{L}_G$  and  $0.01\mathcal{L}_D$ . We adopt Adam [12] as the optimizer, where learning rate, learning decay and weight decay are set to  $5e^{-4}, 0.97, 1e^{-5}$  respectively. Note that learning decay is executed after the end of every epoch. Both parts are trained with 200 epoches or until converged.

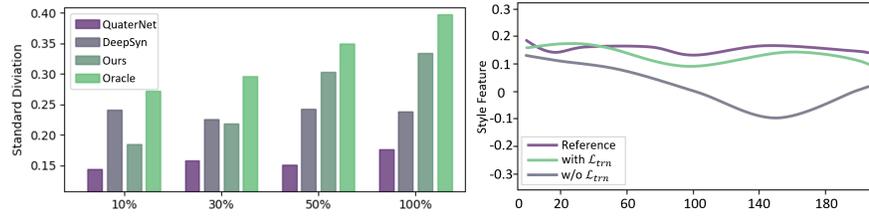
## 4 Experiments

### 4.1 Evaluation settings

**Datasets.** We use the CMU Mocap dataset <sup>6</sup> for training. Reference subsequences are obtained in a sliding window manner, where  $M = 120$  and  $N = 40$ . The dimension of state at a single time step ( $\mathbf{s}_t$ ) is 63, which is 21 joints with 3D coordinates. To demonstrate that our model could generate novel behaviour never seen during training, we keep a held-out reference set (denoted as  $\mathcal{D}_R$ ) from training data (denoted as  $\mathcal{D}_T$ ) for further testing. The held-out set provides reference subsequences during testing.

**Baselines.** Considering the hierarchical structure of our model, we present comparison experiments at both two levels (short-range and long-range). For short-range generation, we compare our model with two strong baselines: the work of Holden et.al. [11] (termed as DeepSyn in this paper) and QuaterNet [25].

<sup>6</sup> <http://mocap.cs.cmu.edu/>



**Fig. 5.** Motion diversity evaluation. X-axis is the proportion of data used for training. Y-axis is the deviation of joint coordinates. Our model is able to synthesize more diverse sequences than baseline models.

**Fig. 6.** Style feature visualization. We can see that facilitated by the loss  $\mathcal{L}_{trn}$ , the style feature of synthesized sequence is closer to the input one (reference) and kept relatively constant.

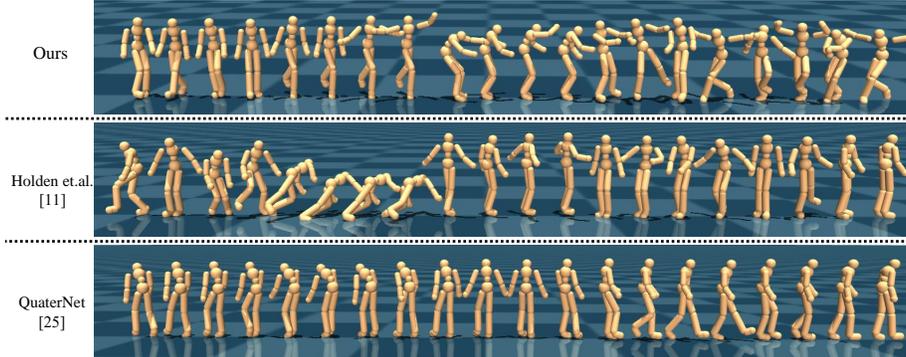
QuaterNet [25] is retrained with the same data used in this work. The input dimension is adjusted to match with our data. For long-range generation, we compare our model with temporal prediction baselines: HP-GAN [1] and MT-VAE [35]. For these two baselines, the ending states are feed as input with the default hyper-parameter setting, i.e., trained as goal-conditioned prediction models.

## 4.2 Evaluation of Short-range Generation

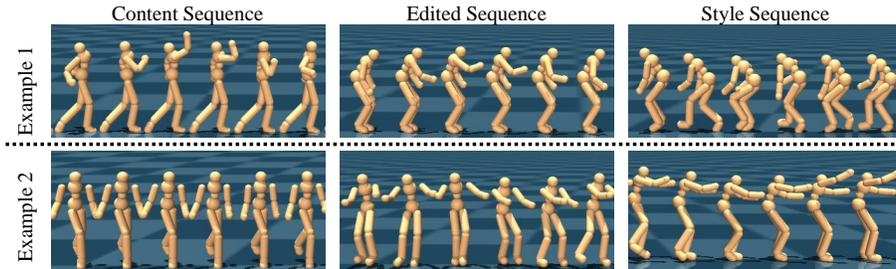
**Generation diversity evaluation.** We compare our model with DeepSyn [11] and QuaterNet [25] in evaluation of motion diversity, comparing the motion diversity under different percentages of data used for training. As illustrated in Fig. 5, 10%, 30%, 50% and 100% training data are used respectively. We calculate the averaged standard deviation of all joints with a higher value indicating more diverse. We can see that both DeepSyn [11] and QuaterNet [25] keep relative constant motion diversity which is comparable with training data used less than 30%, but largely inferior to the diversity of full training data. On the contrary, our model achieves higher motion diversity, which mainly benefits from the hierarchical style-based framework.

**Short-range generation visualization.** Fig. 7 shows the short-range results (from top to bottom generated motion: our model, DeepSyn [11] and QuaterNet [25]). We can observe that DeepSyn [11] synthesises a abnormal walking sequence with unnatural behavior (middle row in Fig. 7, fall-down pose during synthesis). QuaterNet [25] is able to generate a visually natural walking sequence but struggles to produce diverse motion behaviour (last row in Fig. 7, restricted within simple locomotion synthesis). Different from all these models, our hierarchical style-based model achieves natural state transition throughout the whole sequence (first row in Fig. 7), meanwhile provides natural and diverse motion (i.e., walking-turning-dancing) behaviours during generation.

**Style feature consistency evaluation.** As shown in Fig. 6, we plot one representative dimension of the learned style feature for a complete sequence. The purple/green/gray curves refer to the style feature of reference sequence,



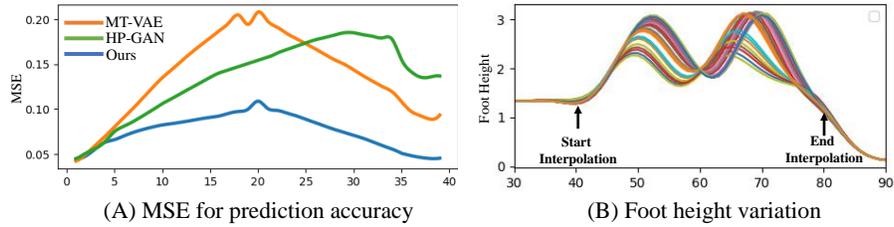
**Fig. 7.** Comparison of motion naturalness with DeepSyn [11] and QuaterNet [25].



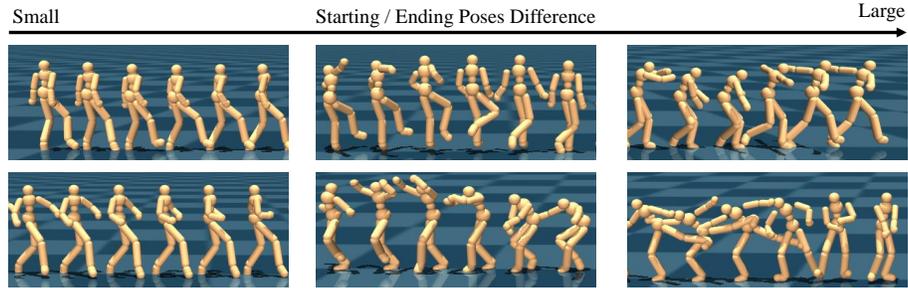
**Fig. 8.** Visualization of generated short-range subsequences. The edited sequence (middle column) well preserves the detailed motion from content sequence (e.g., the arm moving in the first row) and constant pattern from style sequence (e.g., the arm lifting in the second row).

synthesised sequence trained with and without  $\mathcal{L}_{trn}$  respectively. We can observe that training without the supervision of  $\mathcal{L}_{trn}$  leads to the transferred style feature (gray curve) drifts away from the reference one (purple curve). On the contrary, training with the supervision of  $\mathcal{L}_{trn}$  (green curve) effectively facilitates transferring the style pattern from reference sequence to the synthesised one (the green curve is close to the purple curve).

**Style-based generation evaluation.** As shown in Fig. 8, we provide two edited sequence examples (middle part) which possess the general motion style from one subsequence (right part), but detailed motion pattern from another one (left part). For the first sequence (top row in Fig. 8), the target style motion shows a walking sequence with back bent down (style), while the reference motion content (sampled from held-out set  $\mathcal{D}_R$ ) is a sequence with arm waving. We can notice that the synthesised motion (second column, top row) preserves the arm motion and learns the style pattern (back bent down) successfully. For the second sequence (bottom row in Fig. 8), the target style motion shows a walking



**Fig. 9.** Evaluation of long-range motion generation in terms of training MSE (A) and foot height variation (B). (A) shows that our model (blue line) achieves lower error than baseline models. (B) demonstrates the non-linearity of interpolation results.



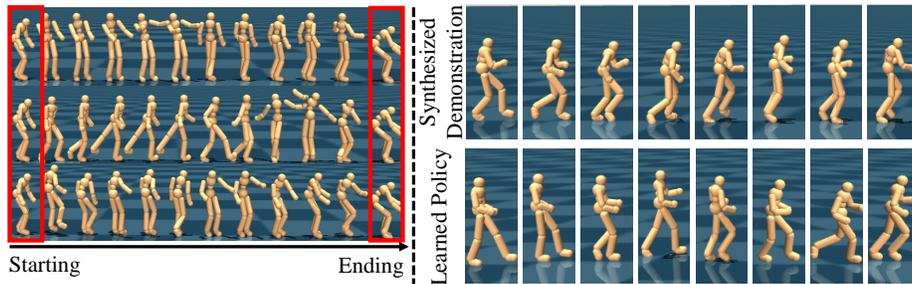
**Fig. 10.** States transition visualization for evaluation of smoothness.

sequence with arm lifted horizontally (style), while the reference motion content (sampled from training set  $\mathcal{D}_T$ ) is a regular walking sequence. The synthesised motion (second column, bottom row) fully captures the style of upper body meanwhile maintains the walking pattern from reference sequence.

### 4.3 Evaluation of Long-range Motion Generation

**Generation accuracy evaluation.** We compare with two goal-conditioned prediction models (HP-GAN [1] and MT-VAE [35]) evaluated by prediction accuracy. The training data is a subset chosen from original one, consisting of locomotion, punching, kicking and dancing sequences. The test data belongs to the same motion category of no overlap with training data. We calculate the MSE value (the lower the better) of predicted clips given starting/ending states from the same sequences. As shown in Fig. 9 (A), we can see that facilitated by the bi-directional modelling, our model could equally consider the contribution of starting/ending states, which leads to high transition accuracy (i.e., low MSE value) near the starting/ending states. Further boosted by the proposed parameterization method (i.e., more compact output space), the interpolation accuracy outperforms both HP-GAN [1] and MT-VAE [35] by a large margin.

**Interpolation sequence visualization.** Fig. 10 demonstrates long-range motion generation results given starting and ending states from different refer-



**Fig. 11.** Left: Motion generation given the same starting and ending states. Right: Long-range Motion generation with characteristic routes.

ence subsequences, respectively. Note that here the starting and ending states are sampled from held-out set ( $\mathcal{D}_R$ ) for evaluation. Synthesised motion clips of long-range part are termed as interpolation sequences for simplicity. Note that for all sequences shown in Fig. 10, the starting as well as ending states are from held-out reference set. We can observe that our model is able to generate smooth and natural transition when starting and ending states are similar (left part in Fig. 10). Moreover, when encountered large motion difference (right part in Fig. 10), e.g., from walking to dancing, turning back with a relatively large degree, our model still makes it to generate visually natural transition sequence.

**Non-trivial interpolation verification.** To evaluate whether our model learns non-linear interpolation between two sequences, we report the height variation of the right foot in a composed motion sequence. Meanwhile, we manually rotate the second sequence w.r.t. the final state of the first sequence to show that our model is robust to a large range of direction difference between two sequences. As shown in Fig. 9(B), we present multiple curves which correspond to different rotation angles. All recorded curves are highly non-linear but smooth between starting and ending points. Moreover, our model adaptively changes foot height with different rotation configurations, which indicates smooth and natural motion for better long-range motion generation.

**Visualization of final synthesized sequences.** With both short-range and long-range motion generation, we are able to synthesize final sequences. Recall that our model is constrained by given starting and ending states for motion interpolation. To this end, we present three synthesized sequences in Fig. 11, of which each result leverages three reference subsequences and four generated clips for motion generation. As shown in left part of Fig. 11, starting from the same state, we are able to generate long-range and visually natural motion boosted by the short-range and long-range motion generation. Meanwhile, we are able to generate diverse behaviour (shown as complex hand and arm motion) facilitated hierarchical style-based modelling.

**Reality Evaluation.** We evaluate the reality of the synthesised long-range sequence quantitatively. We train three classifiers with the ground truth as positive samples. Negative samples are generated by three models (i.e., DeepSyn [11],

**Table 1.** Classification Accuracy for evaluating the reality of synthesised sequences.

	DeepSyn [11]	QuaterNet [25]	Ours
DeepSyn [11]	—	37.6%	<b>66.3%</b>
QuaterNet [25]	41.9%	—	<b>75.3%</b>
Ours	28.3%	50.4%	—

QuaterNet [25], and Ours) respectively. As shown in Tab. 1, we report the proportion of generated sequences classified as positive ones. Row refers to the model used to train the classifier. Column refers to the model evaluated by the classifiers. Our model generates more realistic sequences than these two baselines (first two rows). The proportion of the other two models’ results classified as positive by the our model (third row) is lower.

**Distribution Similarity Evaluation.** Inspired by FVD score [30] we measure the distribution similarity between the original and generated sequences with motion features as inputs. The scores (the lower the better) are 281.5 (DeepSyn [11]), 341.1 (QuaterNet [25]), and 179.8 (ours), i.e., the distribution of our results is closer to that of the original data. Meanwhile, we compare the motion diversity of short-range sequences with those generated by the full model. Motion deviation (the higher the better) is reported, i.e., 0.226 (short-range sequence) and 0.347 (full sequence) respectively. We can see that the long-range generation model facilitates increasing the motion diversity by a large margin.

#### 4.4 Application

**Expert demonstration guidance for imitation learning.** Under a simulated environment with gravity constrain [26], unnatural motion violating physical law (e.g., severe joint twisting) is hard for a simulated agent to follow. To further show our model produces realistic motion, generated results are used for demonstration guidance of imitation learning. As shown in Fig. 11, the right part is demonstration synthesised by our model (top) while the bottom one is learned policy with [26]. We can see that the learned motion succeeds to follow the synthesised one.

## 5 Conclusion

In this work, we present hierarchical style-based networks to generate long-range, diverse and visually plausible motion sequences. Our framework is in a 2-level hierarchical structure involving short-range and long-range motion generation respectively. On large-scale skeleton dataset, we show that the proposed method is able to synthesise long-range, diverse and plausible motion, which is also generalizable to unseen motion data during testing.

**Acknowledgement.** This work was supported in part by BAIR and BDD. This work was also supported in part by grant No.18DZ1112300, No.61976137, No.61527804, No.U1611461, No.U19B2035, and No.2016YFB1001003.

## References

1. Barsoum, E., Kender, J., Liu, Z.: HP-GAN: probabilistic 3d human motion prediction via GAN. In: CVPR Workshops. pp. 1418–1427 (2018)
2. Brand, M., Hertzmann, A.: Style machines. In: SIGGRAPH. pp. 183–192 (2000)
3. Byrd, R.H., Lu, P., Nocedal, J., Zhu, C.: A limited memory algorithm for bound constrained optimization. *SIAM J. Scientific Computing* **16**(5), 1190–1208 (1995)
4. Cai, H., Bai, C., Tai, Y., Tang, C.: Deep video generation, prediction and completion of human action sequences. In: ECCV. pp. 374–390 (2018)
5. Chen, C., Tyagi, A., Agrawal, A., Drover, D., MV, R., Stojanov, S., Rehg, J.M.: Unsupervised 3d pose estimation with geometric self-supervision. In: CVPR. pp. 5714–5724 (2019)
6. Finn, C., Goodfellow, I.J., Levine, S.: Unsupervised learning for physical interaction through video prediction. In: NIPS. pp. 64–72 (2016)
7. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: CVPR. pp. 2414–2423 (2016)
8. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Computation* **9**(8), 1735–1780 (1997)
9. Holden, D., Habibie, I., Kusajima, I., Komura, T.: Fast neural style transfer for motion data. *IEEE Computer Graphics and Applications* **37**(4), 42–49 (2017)
10. Holden, D., Komura, T., Saito, J.: Phase-functioned neural networks for character control. *ACM Trans. Graph.* **36**(4), 42:1–42:13 (2017)
11. Holden, D., Saito, J., Komura, T.: A deep learning framework for character motion synthesis and editing. *ACM Trans. Graph.* **35**(4), 138:1–138:11 (2016)
12. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
13. Kovar, L., Gleicher, M.: Flexible automatic motion blending with registration curves. In: SIGGRAPH/Eurographics Symposium. pp. 214–224 (2003)
14. Kulkarni, T.D., Narasimhan, K., Saeedi, A., Tenenbaum, J.: Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. In: NIPS. pp. 3675–3683 (2016)
15. Lee, Y., Sun, S., Somasundaram, S., Hu, E.S., Lim, J.J.: Composing complex skills by learning transition policies. In: ICLR (2019)
16. Levine, S., Wang, J.M., Haraux, A., Popovic, Z., Koltun, V.: Continuous character control with low-dimensional embeddings. *ACM Trans. Graph.* **31**(4), 28:1–28:10 (2012)
17. Li, Y., Wang, T., Shum, H.: Motion texture: a two-level statistical model for character motion synthesis. *ACM Trans. Graph.* **21**(3), 465–472 (2002)
18. Li, Y., Roblek, D., Tagliasacchi, M.: From here to there: Video inbetweening using direct 3d convolutions. *CoRR* **abs/1905.10240** (2019)
19. Liu, Z., Yeh, R.A., Tang, X., Liu, Y., Agarwala, A.: Video frame synthesis using deep voxel flow. In: ICCV. pp. 4473–4481 (2017)
20. Meyer, S., Wang, O., Zimmer, H., Grosse, M., Sorkine-Hornung, A.: Phase-based frame interpolation for video. In: CVPR. pp. 1410–1418 (2015)
21. Myers, D.R.: *Robot Motion: Planning and Control* edited by michael brady M.I.T. press, cambridge ma, usa, 1983 (£33.95). *Robotica* **1**(2), 109 (1983)
22. Niklaus, S., Mai, L., Liu, F.: Video frame interpolation via adaptive separable convolution. In: ICCV. pp. 261–270 (2017)
23. Park, S.I., Shin, H.J., Shin, S.Y.: On-line locomotion generation based on motion blending. In: SIGGRAPH/Eurographics Symposium. pp. 105–111 (2002)

24. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: NeurIPS. pp. 8024–8035 (2019)
25. Pavlo, D., Feichtenhofer, C., Auli, M.: Modeling human motion with quaternion-based neural networks. In: IJCV (2019)
26. Peng, X.B., Abbeel, P., Levine, S., van de Panne, M.: Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics (TOG)* **37**(4), 143 (2018)
27. Shoemake, K.: Animating rotation with quaternion curves. In: SIGGRAPH. pp. 245–254 (1985)
28. Tan, C.I., Tai, W.: Characteristics preserving racer animation: a data-driven race path synthesis in formation space. *Journal of Visualization and Computer Animation* **23**(3-4), 215–223 (2012)
29. Tenenbaum, J.B., Freeman, W.T.: Separating style and content. In: NeurIPS. pp. 662–668 (1996)
30. Unterthiner, T., van Steenkiste, S., Kurach, K., Marinier, R., Michalski, M., Gelly, S.: Towards accurate generative models of video: A new metric & challenges. *CoRR* **abs/1812.01717** (2018), <http://arxiv.org/abs/1812.01717>
31. Urtasun, R., Fleet, D.J., Geiger, A., Popovic, J., Darrell, T., Lawrence, N.D.: Topologically-constrained latent variable models. In: ICML. pp. 1080–1087 (2008)
32. Wexler, Y., Shechtman, E., Irani, M.: Space-time completion of video. *IEEE Trans. Pattern Anal. Mach. Intell.* **29**(3), 463–476 (2007)
33. Xia, G., Sun, H., Liu, Q., Hang, R.: Learning-based sphere nonlinear interpolation for motion synthesis. *IEEE Trans. Industrial Informatics* **15**(5), 2927–2937 (2019). <https://doi.org/10.1109/TII.2019.2894113>, <http://doi.org/10.1109/TII.2019.2894113>
34. Xia, S., Wang, C., Chai, J., Hodgins, J.K.: Realtime style transfer for unlabeled heterogeneous human motion. *ACM Trans. Graph.* **34**(4), 119:1–119:10 (2015)
35. Yan, X., Rastogi, A., Villegas, R., Sunkavalli, K., Shechtman, E., Hadap, S., Yumer, E., Lee, H.: MT-VAE: learning motion transformations to generate multimodal human dynamics. In: ECCV. pp. 276–293 (2018)
36. Zhang, H., Starke, S., Komura, T., Saito, J.: Mode-adaptive neural networks for quadruped motion control. *ACM Trans. Graph.* **37**(4)