



Constructing Human Motion Manifold With Sequential Networks

Deok-Kyeong Jang and Sung-Hee Lee

Korea Advanced Institute of Science and Technology (KAIST), Daejeon, South Korea
{shadofex, sunghee.lee}@kaist.ac.kr

Abstract

This paper presents a novel recurrent neural network-based method to construct a latent motion manifold that can represent a wide range of human motions in a long sequence. We introduce several new components to increase the spatial and temporal coverage in motion space while retaining the details of motion capture data. These include new regularization terms for the motion manifold, combination of two complementary decoders for predicting joint rotations and joint velocities and the addition of the forward kinematics layer to consider both joint rotation and position errors. In addition, we propose a set of loss terms that improve the overall quality of the motion manifold from various aspects, such as the capability of reconstructing not only the motion but also the latent manifold vector, and the naturalness of the motion through adversarial loss. These components contribute to creating compact and versatile motion manifold that allows for creating new motions by performing random sampling and algebraic operations, such as interpolation and analogy, in the latent motion manifold.

Keywords: motion capture, animation, behavioural animation, animation

ACM CCS: • Computing methodologies → Dimensionality reduction and manifold learning; Neural networks; Motion processing

1. Introduction

Constructing a latent space for human motion is an important problem as it has a wide range of applications such as motion recognition, prediction, interpolation and synthesis. Ideal motion spaces should be compact, in the sense that random sampling in the space leads to plausible motions, and comprehensive so as to generate a wide range of human motions. In addition, locally linear arrangement of the semantically related hidden vectors would benefit motion synthesis, for example, by simple algebraic operations.

However, constructing a compact and versatile motion space and extracting valid motions from it remains a challenging problem because the body parts of human body are highly correlated in general actions and the joints are constrained to satisfy the bone lengths and the range of movement. The high dimensionality of the joint space adds additional difficulty to this problem.

In this paper, we present a novel framework to construct a latent motion manifold and to produce various human motions from the motion manifold. In order to embrace the temporal characteristic of human motion, our model is based on the sequence-to-sequence model. The unsupervised sequence-to-sequence models have been shown to be effective by previous studies on motion prediction [MBR17, PGA18]. Based on these studies, we develop several novel technical contributions to achieve a compact yet

versatile latent motion manifold and a motion generation method as follows.

First, our model is characterized by the combination of one encoder and two decoders. Given a motion manifold vector, one decoder learns to generate the joint rotation while the other learns to output joint rotation velocities. As will be discussed later, the joint rotation decoder has the advantage of reconstructing long-term motions better. In comparison, the joint velocity decoder has the advantage of improving the continuity of the motion. By complementing each other, our two decoder model shows a higher reconstruction accuracy than that of the single decoder model.

Second, unlike previous studies that deal with only either joint angles or joint positions, by adding a forward kinematics (FK) layer [VYCL18], our joint angle-based human representation achieves the advantage of satisfying bone-length constraints and simplifying joint limit representation. By additionally considering joint position computed by the FK layer while training, our method reduces the joint position error, which is visually more perceptible than the joint angle error.

Lastly, we introduce several loss functions, each of which contributes to enhancing the quality of the motion manifold in different aspects. A reconstruction loss reduces the difference between the reconstructed motion and the input motion and thus allows the

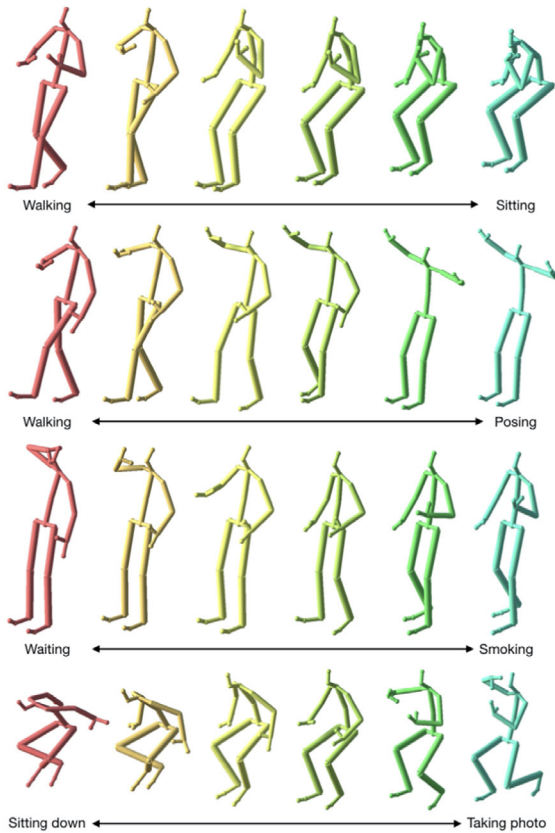


Figure 1: Examples of motion interpolation on the latent motion manifold generated by our method. The first and last columns are snapshots of two input motions, and the intermediate columns show the snapshots of four individual motions obtained by the linear interpolation on the motion manifold.

manifold to synthesize motion content and details observed in the training motion dataset. A regularizer loss improves the distribution quality of the motion manifold and thus enables random sampling and interpolation on the manifold. In addition, an adversarial loss increases the naturalness of the motions generated from the motion manifold.

In this paper, we show that, based on these technical contributions, our method allows for various practical applications such as random generation of motions, motion interpolation, motion denoising and motion analogy as will be shown in Section 5. The capability of our method is demonstrated by the comparison with other approaches, such as the seq2seq model [MBR17] and the convolution model [HSKJ15, HSK16].

The remaining part of this paper proceeds as follows: After reviewing previous studies related to our work in Section 2, we present our method and loss function in detail in Section 3. Sections 4 detail the data pre-processing and Section 5 reports a number of experiments performed to verify the effectiveness of our method. Section 6 discusses the limitations of our work, future research directions and concludes the paper. Our code and networks are available at https://github.com/DK-Jang/human_motion_manifold.

2. Related Work

Researchers have developed several methods to construct motion manifold to generate natural human motions, but compared with studies on manifold learning for other data such as image, research on motion data is scarce. Linear methods such as principal component analysis (PCA) can model human motion in only a local region. Chai *et al.* [CH05] apply local PCA to produce a motion manifold that includes a certain range of human motion, and apply it for synthesizing movements from low-dimensional inputs such as the position of end effectors. Lawrence [Law04] uses Gaussian Process Latent Variable Model (GPLVM) to find a low-dimensional latent space for high-dimensional motion data. Taylor *et al.* [THR07] propose a modified Restricted Boltzmann Machine that is able to deal with the temporal coherency of the motion data. Lee *et al.* [LWB*10] propose motion fields method, a novel representation of motion data, which allows for creating human motion responsive to arbitrary external disturbances. Recently, with the development of deep learning technology, a method of constructing a motion manifold by using Convolutional Neural Network (CNN)-based encoder was introduced by Holden *et al.* [HSKJ15, HSK16]. Butepage *et al.* [BBKK17] compare a number of deep learning frameworks for modelling human motion data.

Our method for constructing motion manifold is based on previous studies on sequence learning for motion to predict the joint position sequences of a 3D human body given past motions. Martinez *et al.* [MBR17] develop a novel sequence-to-sequence encoder-decoder model that predicts human motion given a short duration of past motion. The presented result is impressive but has a few limitations that sometimes implausible motions such as foot sliding are generated and the initial pose of the predicted motion is somewhat discontinuous from the input motion.

Pavlo *et al.* [PGA18] selectively use a joint rotation-based loss for short-term prediction and a joint position-based loss for long-term prediction. The latter includes FK to compute the joint positions. However, the basic sequence-to-sequence model can only predict short-term motions and has limitations in predicting non-trivial, long-term motions. In addition, a loss function that minimizes only the prediction error does not guarantee to construct compact and versatile motion manifold. Our method solves these problems by jointly considering joint rotation and position errors in the loss function and by adding regularization to the motion manifold.

In a broader perspective, our work is related with the studies on recognizing and generating human motion, which remains a challenging research topic due to the high dimensionality and dynamic nature of the human motion. Wu and Shao [WS14] propose a hierarchical dynamic framework that extracts top-level skeletal joint features and uses the learned representation to infer the probability of emissions to infer motion sequences. Du *et al.* [DWW15] and Wang *et al.* [WW17] use recurrent neural network (RNN) to model temporal motion sequences and propose hierarchical structure for action recognition. With regard to motion synthesis, Mittelman *et al.* [MKSL14] propose a new class of Recurrent Temporal Restricted Boltzmann Machine (RTRBM). The structured RTRBM explicitly graphs to model the dependency structure to improve the quality of motion synthesis. Fragkiadaki *et al.* [FLFM15] propose the encoder-recurrent-decoder (ERD) that combines representation

learning with learning temporal dynamics for recognition and prediction of human body pose in videos and motion capture. Jain *et al.* [JZSS16] propose structural RNN for combining the power of high-level spatio-temporal graphs.

3. Method

This section details our framework. After defining notations used in this paper, we explain the structure of the network and the design of the loss function for training.

3.1. Representation and notations

We denote the human motion set by \mathcal{Q} and corresponding random variable by \mathbf{Q} . A motion with a time range of $[t, t + \Delta t - 1]$ is written as $\mathbf{Q}_{t:(t+\Delta t-1)} = [\mathbf{q}_t, \dots, \mathbf{q}_{t+\Delta t-1}]$, where \mathbf{q}_t denotes the pose at time t . A pose is represented with a set of joint angles written in the exponential coordinates, that is, $\mathbf{q}_t = [q_{i,x}^t, q_{i,y}^t, q_{i,z}^t]_{i=1}^{n_{joint}}$, where $(q_{i,x}^t, q_{i,y}^t, q_{i,z}^t)$ are the three components of the exponential coordinates and n_{joint} is the number of joints. Therefore, the dimension of a human motion is $\mathcal{Q} \in \mathbb{R}^{\Delta t \times n_{joint} \times 3}$. Lastly, \mathbf{p}_t is the pose represented with the joint positions at time t corresponding to \mathbf{q}_t , and $\mathbf{P}_{t:(t+\Delta t-1)} = [\mathbf{p}_t, \dots, \mathbf{p}_{t+\Delta t-1}]$. \mathbf{P} is also a random variable of motion set \mathcal{Q} .

3.2. Motion manifold with sequential networks

We construct a motion manifold in an end-to-end unsupervised way using a network of sequential networks, with an objective to minimize the difference between the ground truth motion space distribution and the reconstructed motion space distribution extracted from the latent motion manifold. To this end, we develop a sequential model that consists of the RNN with gated recurrent unit (GRU). Our model has a sequence-to-sequence structure [MBR17], which is often used in machine translation. This RNN structure is effective for maintaining the temporal coherency in motion, and it is trained to generate a fixed length of motion (150 frames) in our study. As shown in Figure 2, our model includes the combination of one encoder and two decoders with a regularizer. The encoder takes the source motion as an input and maps it to the latent motion space. The regularizer encourages the encoded motion distribution to approximate some prior distribution. The two decoders are designed to map the latent motion space to joint angles and joint velocities, respectively. Details of our model are given next.

3.2.1. Encoder

The encoder consists of a GRU and one linear layer, and Figure 2 shows the unrolled schematic diagram of the encoder. The Δt poses $[\mathbf{q}_t, \dots, \mathbf{q}_{t+\Delta t-1}]$ of a motion are input to the GRU sequentially. The GRU encodes the current frame while being conditioned by the previous frames with their hidden representation. Specifically, the pose \mathbf{q}_i in the i -th frame is encoded as follows:

$$h_i^{Enc} = \text{GRU}_{W_{Enc}}(h_{i-1}^{Enc}, \mathbf{q}_i), \quad (1)$$

where h_i is the hidden state at frame i , and $W_{Enc} \in \mathbb{R}^{3n_{joint} \times d_h}$ are the training parameters with d_h being the hidden dimension of the GRU. After the final pose of the input motion is read, one linear layer of

parameter $W_c \in \mathbb{R}^{d_h \times d_m}$ receives $h_{t+\Delta t-1}$ and compresses it to produce the d_m -dimensional code $Z \in \mathcal{Z}$ where \mathcal{Z} denotes the motion manifold. It is worth mentioning that this compression brings the benefit of denoising input data. Now the encoder mapping $Enc : \mathcal{Q} \rightarrow \mathcal{Z}$ is completed.

3.2.2. Latent motion manifold with the Wasserstein regularizer

We adopt the Wasserstein regularizer for matching the distribution $E_Z := \mathbb{E}_{P_Q}[E(Z | \mathbf{Q})]$ of the motion manifold to the desired prior distribution P_Z . Unlike the variational auto-encoder [RMW14], the sequential networks trained with the Wasserstein regularizer allows non-random encoders to deterministically map inputs to the latent codes, and thus it helps randomly sampled or interpolated points in the motion manifold correspond to plausible motions. Refer to [TBGS17] for more details about the Wasserstein regularizer.

3.2.3. Decoder with joint rotation and joint velocity

Our decoder model consists of two kinds: One decoder learns the joint rotation and the other learns joint rotational velocity as shown in Figure 2. Both decoders are based on the GRU while the connection structures of the two are different. Unlike the rotation decoder, the velocity decoder adds a residual connection between the input and the output to construct joint rotation. Each decoder then generates the reconstructed joint angle sequence in reverse temporal order as suggested by [SMS15]. The decoders are trained simultaneously with backpropagation.

This dual decoder model is based on the idea of [SMS15]. By combining the two decoders, we can alleviate the limitations of individual decoder models. The rotation decoder shows strength when reconstructing long-term motions because it learns joint angle itself. Conversely, it may cause pose discontinuity between frames. The velocity decoder has the advantage of reconstructing continuous human motion as it outputs difference between consecutive rotations, which is usually small and easier to learn. However, training velocities tends to be unstable in a long-term sequence because the longer the motion is, the more error is accumulated. As our two decoders have contrasting strengths and weaknesses, when combined, they complement each other in synergy.

Unlike previous studies about motion prediction, recognition and manifold [BBKK17, MBR17, HSKJ15, FLFM15, PGA18] in which either only the joint rotations or the joint positions are used, our model considers both the joint rotations and positions in the motion reconstruction loss term, L_R (see Equation (8)). Loss with joint angles has the advantage of preventing errors such as inconsistent bone length or deviation from human motion range, and thus learning with joint angle loss can generate plausible motions. However, rotation prediction is often paired with a loss that averages errors over joints by giving each joint the same weight. The ignorance of varying influence of different joints on the reconstructed motion can yield large errors in the important joints and degrade the quality of the generated poses.

The joint position loss minimizes the averaged position errors over 3D points, which better reflects perceptual differences between poses. To combine both joint rotations and positions in the motion

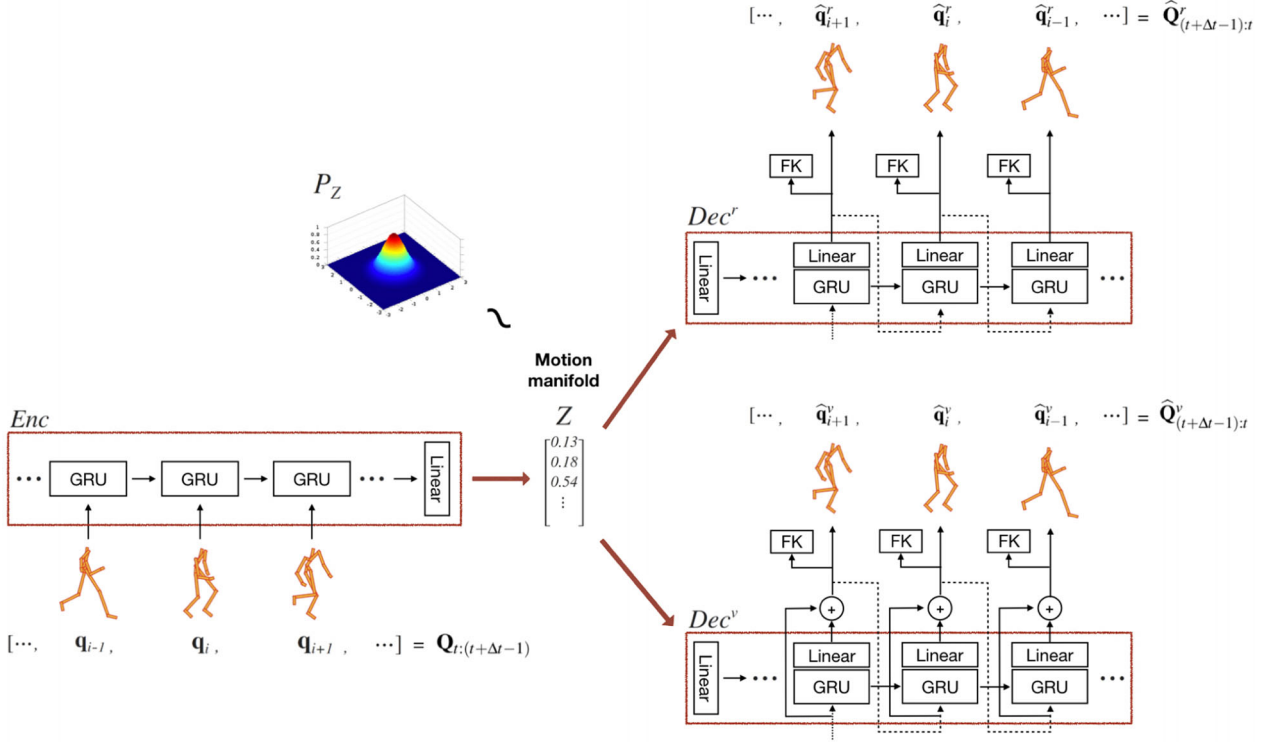


Figure 2: Structure of our sequential networks for constructing the motion manifold.

reconstruction loss L_R , we add an FK layer that computes the joint positions from the joint rotations. This allows for calculating the loss between the joint positions of the target motion and the reconstruction motion. The FK module is valid for network training because its output is differentiable with respect to joint rotation.

Finally, our method reconstructs the motion in the reverse order of the input sequence. Reversing the target sequence has an advantage in learning in that the first output frame of the decoder needs only to match the last frame input of the encoder, which allows for a continuous transition of hidden space vectors from the encoder to the decoders. Refer to [SMS15] for a theoretical background on this approach. Details of our decoder are explained next.

Joint rotation decoder. The unfolded schematic diagram of the joint rotation decoder is shown in the upper row in Figure 2. It first transforms an element of the motion manifold $z \in Z$ to a d_h -dimensional hidden space vector with a linear layer of parameter $W_e^r \in \mathbb{R}^{d_m \times d_h}$. Then, conditioned by the hidden space vector representing the future frames, the GRU and a linear layer outputs the reconstructed pose $\hat{\mathbf{q}}_i^r$ at the i -th frame given its next pose $\hat{\mathbf{q}}_{i+1}^r$:

$$h_i^{Dec^r} = \text{GRU}_{W_{Dec^r}}(h_{i+1}^{Dec^r}, \hat{\mathbf{q}}_{i+1}^r), \quad (2)$$

$$\hat{\mathbf{q}}_i^r = W_o^{rT} h_i^{Dec^r}, \quad (3)$$

where $W_{Dec^r} \in \mathbb{R}^{3n_{\text{joint}} \times d_h}$ is learning parameter of the GRU and $W_o^r \in \mathbb{R}^{d_h \times 3n_{\text{joint}}}$ is the parameter of the linear layer.

Note that, as mentioned earlier, the decoder uses the reversed input motion as the target motion, so the reconstruction is performed

in the order of $\hat{\mathbf{Q}}_{(t+\Delta t-1):t}^r = [\hat{\mathbf{q}}_{t+\Delta t-1}^r, \dots, \hat{\mathbf{q}}_t^r]$. Unlike the encoder, the decoder uses the reconstructed result of the previous frame as the input [MBR17, LZX*17]. This is equivalent to the noise scheduling [BVJS15] without parameter tuning for long-term reconstruction, and it also helps prevent the overfitting. The initial input $\hat{\mathbf{q}}_{t+\Delta t}^r$ to the GRU is set zero because there is no reconstruction result of the previous frame. The reconstructed joint rotations are used to calculate the angle loss with respect to the target motion, and are also used to calculate the position $\hat{\mathbf{p}}_i^r$ through the FK layer.

$$\hat{\mathbf{p}}_i^r = \text{Forward Kinematics}(\hat{\mathbf{q}}_i^r). \quad (4)$$

After the last pose $\hat{\mathbf{q}}_t$ is generated, the joint decoder mapping $Dec^r: Z \rightarrow \mathcal{Q}$ is completed.

Joint velocity decoder. The joint velocity decoder has the similar structure to the joint rotation decoder. The main difference is that it has a residual connection to generate $\hat{\mathbf{q}}_i^v$.

$$h_i^{dec^v} = \text{GRU}_{W_{Dec^v}}(h_{i+1}^{Dec^v}, \hat{\mathbf{q}}_{i+1}^v), \quad (5)$$

$$\hat{\mathbf{q}}_i^v = W_o^v T h_i^{Dec^v} + \hat{\mathbf{q}}_{i+1}^v, \quad (6)$$

$$\hat{\mathbf{p}}_i^v = \text{Forward Kinematics}(\hat{\mathbf{q}}_i^v), \quad (7)$$

where $W_{Dec^v} \in \mathbb{R}^{3n_{\text{joint}} \times d_h}$ and W_o^v are the learning parameters. This residual network learns the difference between the current frame pose $\hat{\mathbf{q}}_i^v$ and the previous frame pose $\hat{\mathbf{q}}_{i+1}^v$. Therefore, the model predicts the angle difference or velocity and integrates it over time.

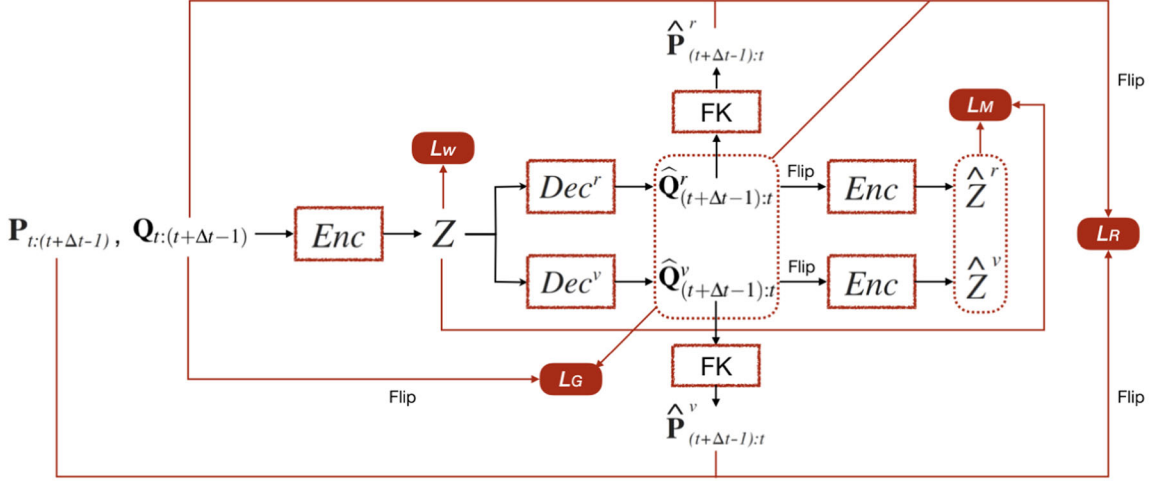


Figure 3: Each loss term is evaluated from the data processed in the network pipeline shown with black arrows. Red arrows indicate the data used for the individual loss terms.

After the last pose is generated, the joint velocity decoder mapping $Dec^v : Z \rightarrow Q$ is completed.

3.3. Training the motion manifold

We model a number of loss functions, each of which contributes to enhancing the quality of the motion generated from the motion manifold from different perspectives. To reduce the reconstruction loss, we employ two kinds of loss functions: motion reconstruction loss L_R that encourages a motion to be reconstructed after going through the encoder and decoder, and manifold reconstruction loss L_M that helps a latent vector be reconstructed after going through the decoder and encoder. In addition, we include Wasserstein loss L_W that penalizes the discrepancy between P_Z and the distribution E_Z induced by the encoder, and an adversarial loss L_G to achieve more natural motions from the motion manifold. Figure 3 shows overview of our loss functions.

Motion reconstruction loss. The motion reconstruction loss penalizes the difference between the motion and the reconstructed motion, which is obtained by encoding the motion followed by decoding it. Specifically, we measure the discrepancy of both the joint rotation angle \mathbf{q} and the joint position \mathbf{p} as follows:

$$L_R = L_{ang} + w_p L_{pos}, \quad (8)$$

$$L_{ang} = \sum_i^{n_{joint}} \|\hat{\mathbf{q}}_i^r - \mathbf{q}_i\| + \|\hat{\mathbf{q}}_i^v - \mathbf{q}_i\|, \quad (9)$$

$$L_{pos} = \sum_i^{n_{joint}} \|\hat{\mathbf{p}}_i^r - \mathbf{p}_i\| + \|\hat{\mathbf{p}}_i^v - \mathbf{p}_i\|, \quad (10)$$

where $\|\cdot\|$ is the Euclidean norm and w_p ($= 5$ in our experiment) is the weight of the position error.

Manifold reconstruction loss. A latent code sampled from the latent distribution should be reconstructed after decoding and

encoding. Manifold reconstruction loss encourages this reciprocal mapping between the motions and the manifold space. To this end, we apply L_1 loss similar to [LTH*18]. We draw a motion manifold vector Z from the encoded motion sequences and reconstruct it with $\hat{Z}^r = Enc(Dec^r(Z))$ and $\hat{Z}^v = Enc(Dec^v(Z))$, where $Z = Enc(\mathbf{Q}_{t:(t+\Delta t-1)})$.

$$L_M = \|\hat{Z}^r - Z\|_1 + \|\hat{Z}^v - Z\|_1. \quad (11)$$

Wasserstein regularizer loss. In order to make the manifold space have a particular desired prior distribution so that we can efficiently sample from the distribution, we use the Wasserstein regularizer that penalizes deviation of the distribution E_Z of the latent manifold from the desired prior distribution P_Z .

$$L_W = \text{MMD}_k(P_Z, E_Z), \quad (12)$$

where $P_Z(Z) = \mathcal{N}(Z; \mathbf{0}, \sigma_z^2 \cdot \mathbf{I}_d)$ is modelled as the multivariate normal distribution with σ_z^2 being decided through validation. We use the maximum mean discrepancy MMD_k to measure the divergence between two distributions with the inverse multi-quadratics kernel $k(x, y) = C/(C + \|x - y\|_2^2)$ with $C = 2Z_{\text{dim}}\sigma_z^2$. We set $\sigma_z^2 = 1$ and the dimension of motion manifold space $Z_{\text{dim}} = 64$.

Adversarial loss. Finally, we employ the least squares generative adversarial network (LSGAN) to match the distribution of generated motion to the real motion data distribution, that is, to promote motions generated by our model to be indistinguishable from real motions.

$$L_D = \frac{1}{2} \sum_{\hat{\mathbf{Q}}_{t:(t+\Delta t-1)}} [D(\hat{\mathbf{Q}}_{t:(t+\Delta t-1)}) - 0]^2 + \frac{1}{2} \sum_{\mathbf{Q}_{t:(t+\Delta t-1)}} [D(\mathbf{Q}_{t:(t+\Delta t-1)}) - 1]^2 \quad (13)$$

$$L_G = \frac{1}{2} \sum_{\hat{\mathbf{Q}}_{t:(t+\Delta t-1)}} [D(\hat{\mathbf{Q}}_{t:(t+\Delta t-1)}) - 1]^2, \quad (14)$$

where the discriminator D tries to distinguish between the reconstructed motions and the real motions. The discriminator is then used to help our decoder generate realistic motions.

Total loss. We jointly train the encoder, joint rotation decoder, joint velocity decoder and discriminator to optimize the total objective function, which is a weighted sum of the reconstruction loss, Wasserstein regularizer loss and adversarial loss. The total objective function of manifold network is:

$$\min_{Enc, Dec^r, Dec^v} L(Enc, Dec^r, Dec^v) \quad (15)$$

$$= L_R + \lambda_M L_M + \lambda_W L_W + \lambda_G L_G$$

and the discriminator loss is:

$$\min_D L(D) = \lambda_G L_D, \quad (16)$$

where weighting parameters λ_M , λ_W and λ_G are 0.001, 0.1, and 0.001 determined through validation.

4. Data Pre-Processing

We tested our method with H3.6M dataset. Every motion in the dataset has the same skeletal structure. All the poses are represented with the position and orientation of the root and the joint rotations expressed with the exponential coordinates. For the training, motion clips of 150 frames are randomly selected from the input motion sequence and used to learn a motion manifold. The root position in the transverse plane is removed and other data are normalized for better performance. We will explain how motion dataset is processed.

H3.6M dataset. H3.6M dataset [IPOS14] consists of 15 activities such as walking, smoking, discussion, taking pictures and phoning performed by seven subjects. We reduce 32 joints in the original data to 17 joints by removing redundant joints as done by [MBR17], and configured all data to have a frame rate of 25 Hz. Therefore, 150 frames motion applied to our model cover 6 s. The activities of subject S5 were used as the test data and those of the remaining subjects S1, S6, S7, S8, S9 and S11 were used as the training data. Some motion data contain noises such as joint popping, but were used without noise removal.

5. Experimental Results

We perform several experiments to evaluate the performance of our method. First, we compare the reconstruction accuracy of the proposed model with its own variations with some components ablated as well as the sequence-to-sequence model proposed by [MBR17]. Next, we test random sampling, motion interpolation via motion manifold, and motion denoising, followed by an experiment for motion analogies. For these tests, we use the joint rotation decoder to generate motions. We qualitatively compare the result of motion interpolation and motion analogies with that of [HSKJ15].¹ All experiments were conducted with test sets not included in the training set.

¹[HSKJ15] is not compared with ours with respect to the reconstruction quality as it deals only with joint positions and not joint angles.

The supplemental video shows the resulting motions from the experiments.

5.1. Motion and manifold reconstruction

We assess the accuracy of the reconstructed motion $\hat{\mathbf{Q}}$ with respect to the input motion \mathbf{Q} , as well as the accuracy of the reconstructed motion manifold vector $\hat{\mathbf{z}}$ with respect to the motion manifold vector \mathbf{z} obtained by encoding a motion. The results are provided in Table 1. Generally, the reconstruction accuracy and the data generation quality of a manifold conflict with each other to some degree. As our purpose is to achieve a motion manifold that supports not only the motion reconstruction but also motion generation, it is important to strike a balance among various performance measures, and our method should not be evaluated only by the reconstruction accuracy. This trade-off will be discussed in Section 5.1.1.

The sequence-to-sequence model (Seq2seq) compared with ours is based on [MBR17]. The only difference is that a fully connected layer of 64 dimension is implemented between the encoder and the decoder to construct a motion manifold.

For ablation study, we prepare a set of variations of our model. The most basic model, denoted **S**, has only joint rotation decoder with reconstruction and Wasserstein regularizer losses, without the FK layer in the network. Next model **D** is the dual decoder model by adding the velocity decoder. From the dual model, we make variations by incrementally accumulating FK layer (**DK**), adversarial loss (**DKG**), manifold reconstruction loss (**DKGM**, our method). The last variation **DKGMZ** is made by concatenating the manifold vector to the decoder input, that is, $[\hat{\mathbf{q}}_{i+1}, \mathbf{Z}]$ is used instead of $\hat{\mathbf{q}}_{i+1}$ in Equations (3) and (5). The idea of this last variation is to prevent the decoder from forgetting the motion manifold vector. All variations have the same network weight dimensions and hyper-parameters as our model. Supplemental material includes details of implementing the compared models. All models are trained with datasets that include all action categories.

The accuracy of the motion reconstruction is evaluated for both the joint rotation decoder (Dec^r) and the joint velocity decoder (Dec^v). Both the Euclidean distances of joint angle errors (L_{ang} , also denoted as E_r) and joint position errors (L_{pos} or E_p) are used for each decoder for the reconstruction loss. As for the reconstruction quality of the motion manifold vector, we measure the L_1 -norm (E_z) of the difference between the motion manifold vector \mathbf{z} obtained by




Figure 4: Ground truth motions (green) and reconstruction results (coral) of our method from H3.6M dataset.

Table 1: Reconstruction errors of joint angles (E_r) and joint positions (E_p) at sample time frames, and the reconstruction error of the manifold vector (E_z). The error is measured with respect to the general actions (all the actions in the DB) in H3.6M dataset.

Model		1.2s		2.4s		3.6s		4.8s		6.0s		E_z
		E_r	E_p	E_r	E_p	E_r	E_p	E_r	E_p	E_r	E_p	
S	rot	0.889	0.957	0.971	0.978	0.990	1.040	1.097	1.078	1.195	1.181	0.317
	vel	–	–	–	–	–	–	–	–	–	–	–
D	rot	0.823	0.855	<u>0.868</u>	0.923	0.925	0.999	1.039	1.032	1.164	1.167	0.264
	vel	<u>0.856</u>	0.889	0.843	0.889	<u>0.877</u>	0.961	1.008	1.081	1.127	1.212	0.259
DK	rot	1.020	0.561	1.099	0.682	1.110	0.706	1.195	0.761	1.261	0.822	0.196
	vel	1.347	0.600	1.353	0.698	1.323	0.723	1.382	0.756	1.391	0.809	0.288
DKG	rot	0.986	<u>0.549</u>	1.077	0.657	1.094	<u>0.679</u>	1.180	0.726	1.251	0.810	0.188
	vel	1.343	0.589	1.345	0.682	1.332	<u>0.702</u>	1.405	0.765	1.415	0.834	0.307
DKGM (ours)	rot	0.997	0.541	1.066	<u>0.659</u>	1.084	0.668	1.162	0.696	1.258	0.780	0.182
	vel	1.356	0.590	1.381	<u>0.673</u>	1.338	0.694	1.400	0.735	1.406	0.792	0.293
DKGMZ	rot	0.906	0.629	0.909	0.730	0.886	0.724	<u>0.954</u>	0.754	<u>1.053</u>	<u>0.788</u>	<u>0.164</u>
	vel	0.877	0.635	0.883	0.703	0.848	0.689	0.916	<u>0.706</u>	1.030	0.815	0.157
Seq2seq	rot	–	–	–	–	–	–	–	–	–	–	–
	vel	0.875	0.863	0.870	0.954	0.891	1.059	1.039	1.177	1.154	1.258	0.216

encoding a motion sequence and the reconstructed vector $\hat{\mathbf{z}}^r$ obtained by sequentially decoding \mathbf{z} and encoding it.

Table 1 shows the reconstruction errors of our method and others for the datasets containing all action categories (15 actions in H3.6M dataset). The reported errors are the average of 30 motions randomly selected from a test dataset. A total of 150 frames are divided into five intervals, and errors (E_r , E_p) are measured for each interval to investigate the temporal characteristic. The lowest and the next lowest errors are marked in bold and with underline, respectively.

We first compare with respect to E_r and E_p errors. Comparing **S** and **D**, the latter has lower E_r and E_p errors, which suggests that the joint rotation and velocity decoders complement with each other to reduce the errors. Comparing **D** and **DK**, the latter reduces E_p error significantly while only mildly sacrificing E_r error. **DKG** has lower E_r and E_p errors than **DK**, but higher errors than **D** and **S**. This shows that adversarial loss slightly reduces reconstruction error. However, it turns out that the adversarial loss helps reconstruct the original behaviours, as will be discussed in Section 5.1.2. Examining the error of **DKGM** and **DKGMZ**, we can see that adding manifold reconstruction loss does not significantly affect the reconstruction errors while explicitly feeding the manifold vector to the decoder helps reduce the errors.

Next, we examine manifold reconstruction error, E_z ($= L_M$). Comparing **D** and **S**, it is remarkable that **D** reduces E_z error even without any manifold-related loss term. However, adding FK layer to reduce joint position error slightly increases E_z for the velocity decoder while it is decreased for the rotation decoder. Comparing **DK** and **DKG**, we can see that adversarial loss has negligible effect to the manifold reconstruction error. Subsequently, **DKGM** reduces E_z slightly by adding the manifold reconstruction error, and **DKGMZ** achieves the lowest E_z error by explicitly feeding the manifold vector to the decoder.

Seq2seq [MBR17] shows less E_r than our model, but E_p is higher. In addition, our model shows better E_z errors with respect to rotation

decoder. Figure 4 visualizes the reconstruction results with our model over time in comparison with the ground truth input motion.

5.1.1. Tradeoff between joint angle, joint position and motion manifold

This experiment examines the effect of different settings of the weight λ_w for the regularization on the reconstruction errors (L_{ang} and L_{pos}) and on motion manifold (L_M) on the test set. We employed **D** model for this experiment to exclude the effect of other loss terms. Figures 5(a) and (b) show that the joint reconstruction errors decrease as λ_w becomes smaller, which makes **D** model closer to a pure autoencoder, sacrificing the ability to enforce a prior over the motion manifold space while obtaining better reconstruction loss. For the same reason, Figure 5(c) shows that the motion manifold reconstruction error L_M decreases as λ_w becomes larger.

As our goal is to obtain an effective motion manifold that is able to generate realistic motions, it is important to find a suitable set of weight parameters that compromise among different qualities.

5.1.2. Adversarial loss and explicit feeding manifold vector

Here, we discuss the effects of adversarial loss (Section 3.3) and explicitly feeding motion manifold vector to the decoders on motion quality. First, Table 1 shows that **DKG** decreases E_r from **DK** only slightly. However, Figure 6 shows that **DK** cannot properly reconstruct the original motion, reconstructing only posing motion from the original motion of posing with walking. In contrast, **DKG** improves the overall motion quality by better reconstructing the behaviours in the original motion. Comparing our method (**DKGM**) and **DKGMZ**, the latter results in lower E_r and E_p than our method as shown in Table 1. However, Figure 6 reveals that **DKGMZ** fails to capture walking motion. We conjecture that directly feeding manifold vector to decoder reduces reconstruction loss by explicitly retaining the motion manifold vector, but tends to converge to mean pose. In contrast, our method successfully reconstructs the




Figure 5: Reconstruction errors of joint angle, joint position and manifold according to training step while adjusting λ_W for H3.6M dataset.




Figure 6: Reconstruction results of different loss combinations for a posing while walking motion. Supplementary video includes full motions.




Figure 7: Results of randomly sampling motions from the motion manifold P_Z .

original posing with walking behaviour. This observation suggests that, while the joint reconstruction error is an important indicator of motion quality, it may not appropriately assess the motion quality in terms of reconstructing the original behaviours.

5.2. Random motion samples

To verify whether the latent motion manifold can create meaningful motions, we randomly sampled P_Z and decoded to obtain motions. We extracted 30 random samples from the motion manifold learned with H3.6M dataset. Figure 7 is the results of random sampling from

P_Z , and one can see that our method can create various actions including sitting, crossing the legs and resting on the wall. This result suggests that our motion manifold and decoder can create a wide range of plausible behaviours.

To examine the importance of Wasserstein Auto-Encoder (WAE), we experimented random sampling by replacing the WAE regularizer with a simple L_2 -norm $\|z\|^2$ loss. Sampled motions from this method, as shown in Figure 7 (right), often show unnatural poses and extreme joint rotations. This experiment shows that the WAE regularizer not only helps achieve the desired motion manifold distribution but also improves quality of motion sampling.

5.3. Motion interpolation with latent motion manifold

We can interpolate two different motions by encoding them into the latent motion manifold and then performing linear interpolation between the encoded motion manifold vectors. The resulting interpolated motion created by our method is not just frame-by-frame interpolation, but may contain meaningful transition between the input motions. For example, interpolating sitting down motion and photo taking motion creates hand raising motion to prepare to take a picture from sitting posture. When waiting and smoking motions are interpolated, an interesting motion that a character seems tired of waiting and starts to smoke is created. The capability of creating such meaningful motions is due to the Wasserstein regularizer that shortens the distance between the encoded vectors by matching the motion manifold to the multivariate normal prior. Figure 1 and the supplemental video show the interpolated motions.

Figure 8 compares our model with [HSKJ15] with respect to interpolation. See supplementary material for the implementation of [HSKJ15]. For the interpolation from sitting to walking (top) and from sitting down to taking photo (bottom), our model shows a natural transition between two motions while [HSKJ15] creates somewhat averaged motion between the two motions.

5.4. Denoising motion data

Our motion model can denoise motion data by projecting it to the latent motion manifold and decoding the motion manifold vector to obtain a reconstructed motion. Since the motion manifold is constructed only from human motion capture data, any element in the




Figure 8: Interpolation from sitting to walking (top) and from sitting down to taking photo (bottom) made by our model (left) and [HSKJ15] (right).


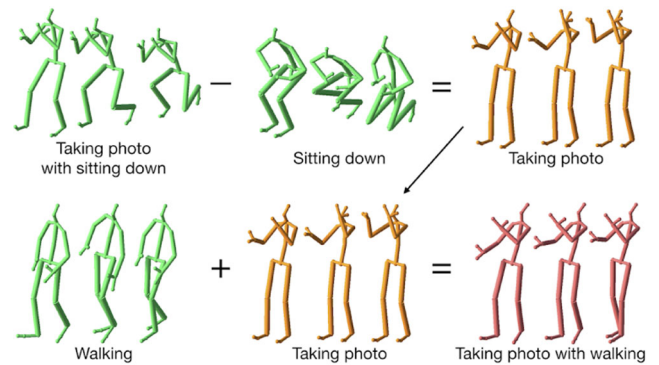


Figure 9: Denoising experiment. Three poses are shown from the noise corrupted motion (orange), denoised motion by our method (coral) and the ground truth motion (green). Two motions (top and bottom) are shown.

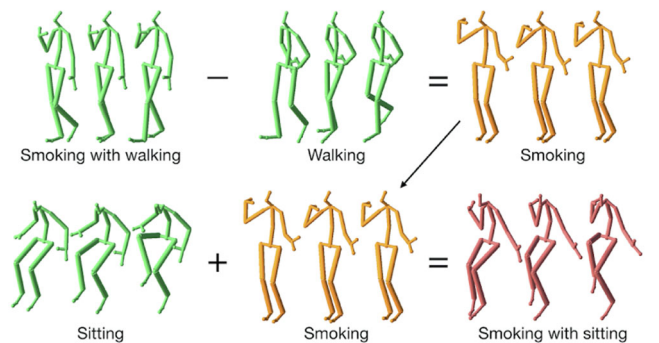
manifold is likely to be decoded to natural motion. Therefore, denoising effect occurs when noisy motion data is projected to the motion manifold. We experiment on the denoising capability of our method in the similar manner as in [HSKJ15]. We generate noise corrupted motion by randomly setting joint angles to zero with a probability of 0.5, which makes half of the joint angle information meaningless. Figure 9 shows the denoised results which are quite similar to the ground truth motions.

5.5. Motion analogy

Through motion analogy, we can understand how our model organizes motion manifold to represent the feature of actions. Details about analogy can be found in [Whi16]. We perform vector algebraic operations with the latent vectors encoded from different motions and explore how the model organizes the latent space to represent motions. Figure 10(a) shows that subtracting a motion manifold vector for ‘sitting down’ motion from ‘taking photo with sitting down’ motion creates a vector representing ‘taking photo’ motion. The character is standing because a zero vector in our motion manifold corresponds to an idle standing motion. Subsequently,



(a) Motion analogy among “Walking with posing”, “Walking” and “Sitting” actions.



(b) Motion analogy among “Smoking with sitting”, “Sitting” and “Walking” actions.

Figure 10: Motion analogy experiments performing arithmetic operations in the motion manifold.

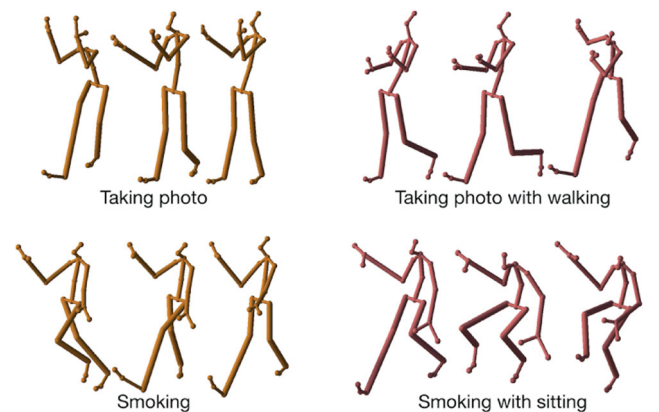


Figure 11: Motion analogy experiment with [HSKJ15].

when an encoded ‘walking’ motion manifold vector is added, the motion vector becomes a vector for ‘taking photo with walking’ motion. Figure 10(b) shows a similar analogy among ‘walking’, ‘smoking with walking’ and ‘sitting’ motions.

Figure 11 shows the experiments of performing analogy with [HSKJ15]. Figure 11 (top) is the result of taking photo (left) and

taking photo with walking (right) that correspond to Figure 10(a), and Figure 11 (bottom) shows smoking and smoking with sitting to compare with Figure 10(b). One can see that the motion manifold obtained with [HSKJ15] does not support analogy on the motion manifold.

6. Conclusion and Future Work

In this paper, we presented a novel sequential network for constructing a latent motion manifold for modelling human motion. The main contributions of our method are the combined decoder for the joint rotation and joint velocity, and considering both the joint rotations and positions by adding the FK layer in both decoders, which improve the reconstruction accuracy. In addition, we composed a set of loss functions, each of which contribute to enhancing the quality of motions generated from the motion manifold space from different aspects. The capabilities of our model have been examined through various experiments such as random sampling, motion interpolation, denoising and motion analogy.

Our method has several limitations. First, as a sequence-to-sequence framework, the performance of our model degrades if trained to produce motions longer than 10 s. The supplementary video shows randomly generated motions with our network being trained to learn 300 frames (approx. 13 s). Resulting motions tend to lose details. This limitation may be alleviated by employing an attention mechanism [LPM15, BCB14]. Second, the encoded motions tend to be smoothed in the process of matching the latent motion manifold to the prior distribution through the regularizer. For example, motions that contain frequent hand shaking, such as ‘walking with dog’ or ‘discussion’ motions in H3.6M dataset, lose fine details when reconstructed. Overcoming these limitations will be important future work.

We only considered joint rotations in the encoder, but incorporating additional information, such as joint positions and velocities, may be beneficial to achieve better motion qualities. In addition, in the process of learning a motion manifold, loss terms to check validity of motions, such as joint limit, velocity limit and foot sliding, are not needed as all input motion data are considered valid. However, when an actual motion is sampled from the manifold and applied to an environment, such criteria may need to be checked.

Most studies on motion space learning have focused on representing a wide range of motion categories with a compact representation. In fact, the range of motion categories is only one aspect of the variedness of human motions. Even a single motion category such as walking exhibits widely different styles depending on gender, body scale, emotion and personality. Developing a motion manifold that can generate stylistic variations of motion is another important future research direction.

Acknowledgements

This work was supported by Giga Korea Project (GK17P0200) and Basic Science Research Program (NRF-2020R1A2C2011541) funded by Ministry of Science and ICT, Korea.

References

- [BBKK17] BÜTEPAGE J., BLACK M. J., KRAGIC D., KJELLSTRÖM H.: Deep representation learning for human motion prediction and classification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017), IEEE, p. 2017.
- [BCB14] BAHDANAU D., CHO K., BENGIO Y.: Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014).
- [BVJS15] BENGIO S., VINYALS O., JAITLEY N., SHAZEER N.: Scheduled sampling for sequence prediction with recurrent neural networks. *Advances in Neural Information Processing Systems* (2015), 1171–1179.
- [CH05] CHAI J., HODGINS J. K.: Performance animation from low-dimensional control signals. *ACM Transactions on Graphics (ToG)* 24 (2005), 686–696.
- [DWW15] DU Y., WANG W., WANG L.: Hierarchical recurrent neural network for skeleton based action recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Boston, USA, 2015), pp. 1110–1118.
- [FLFM15] FRAGKIADAKI K., LEVINE S., FELSEN P., MALIK J.: Recurrent network models for human dynamics. In *Proceedings of the IEEE International Conference on Computer Vision* (Las Condes Araucano Park, Chile, 2015), pp. 4346–4354.
- [HSK16] HOLDEN D., SAITO J., KOMURA T.: A deep learning framework for character motion synthesis and editing. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 138.
- [HSKJ15] HOLDEN D., SAITO J., KOMURA T., JOYCE T.: Learning motion manifolds with convolutional autoencoders. In *SIGGRAPH Asia 2015 Technical Briefs* (2015), ACM, p. 18.
- [IPOS14] IONESCU C., PAPAVALA D., OLARU V., SMINCHISESCU C.: Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36, 7 (jul 2014), 1325–1339.
- [JZSS16] JAIN A., ZAMIR A. R., SAVARESE S., SAXENA A.: Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Las Vegas, USA, 2016), pp. 5308–5317.
- [Law04] LAWRENCE N. D.: Gaussian process latent variable models for visualisation of high dimensional data. *Advances in Neural Information Processing Systems* (2004), 329–336.
- [LPM15] LUONG M.-T., PHAM H., MANNING C. D.: Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025* (2015).
- [LTH*18] LEE H.-Y., TSENG H.-Y., HUANG J.-B., SINGH M., YANG M.-H.: Diverse image-to-image translation via disentangled representations. In *Proceedings of the European Conference on Computer Vision (ECCV)* (Munich, Germany, 2018), pp. 35–51.

- [LWB*10] LEE Y., WAMPLER K., BERNSTEIN G., POPOVIĆ J., POPOVIĆ Z.: Motion fields for interactive character locomotion. *ACM Transactions on Graphics (TOG)* 29 (2010), 138.
- [LZX*17] LI Z., ZHOU Y., XIAO S., HE C., HUANG Z., LI H.: Auto-conditioned recurrent networks for extended complex human motion synthesis. *arXiv preprint arXiv:1707.05363* (2017).
- [MBR17] MARTINEZ J., BLACK M. J., ROMERO J.: On human motion prediction using recurrent neural networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (Honolulu, Hawaii, 2017), IEEE, pp. 4674–4683.
- [MKSL14] MITTELMAN R., KUIPERS B., SAVARESE S., LEE H.: Structured recurrent temporal restricted Boltzmann machines. In *International Conference on Machine Learning* (Beijing, China, 2014), pp. 1647–1655.
- [PGA18] PAVLLO D., GRANGIER D., AULI M.: Quaternet: A quaternion-based recurrent model for human motion. *arXiv preprint arXiv:1805.06485* (2018).
- [RMW14] REZENDE D. J., MOHAMED S., WIERSTRA D.: Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082* (2014).
- [SMS15] SRIVASTAVA N., MANSIMOV E., SALAKHUDINOV R.: Unsupervised learning of video representations using lstms. In *International Conference on Machine Learning* (Lille, France, 2015), pp. 843–852.
- [TBGS17] TOLSTIKHIN I., BOUSQUET O., GELLY S., SCHOELKOPF B.: Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558* (2017).
- [THR07] TAYLOR G. W., HINTON G. E., ROWEIS S. T.: Modeling human motion using binary latent variables. *Advances in Neural Information Processing Systems* (2007), 1345–1352.
- [VYCL18] VILLEGAS R., YANG J., CEYLAN D., LEE H.: Neural kinematic networks for unsupervised motion retargetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Utah, USA, 2018), pp. 8639–8648.
- [Whi16] WHITE T.: Sampling generative networks. *arXiv preprint arXiv:1609.04468* (2016).
- [WS14] WU D., SHAO L.: Leveraging hierarchical parametric networks for skeletal joints based action segmentation and recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Ohio, USA, 2014), pp. 724–731.
- [WW17] WANG H., WANG L.: Modeling temporal dynamics and spatial configurations of actions using two-stream recurrent neural networks. In *e Conference on Computer Vision and Pattern Recognition (CVPR)* (Honolulu, Hawaii, 2017).

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Video s1

Video s2

Supplementary Material