

A causal convolutional neural network for multi-subject motion modeling and generation

Shuaiying Hou¹, Congyi Wang², Wenlin Zhuang³, Yu Chen¹, Yangang Wang³, Hujun Bao¹, Jinxiang Chai², and Weiwei Xu¹ (✉)

© The Author(s) 2023.

Abstract Inspired by the success of WaveNet in multi-subject speech synthesis, we propose a novel neural network based on causal convolutions for multi-subject motion modeling and generation. The network can capture the intrinsic characteristics of the motion of different subjects, such as the influence of skeleton scale variation on motion style. Moreover, after fine-tuning the network using a small motion dataset for a novel skeleton that is not included in the training dataset, it is able to synthesize high-quality motions with a personalized style for the novel skeleton. The experimental results demonstrate that our network can model the intrinsic characteristics of motions well and can be applied to various motion modeling and synthesis tasks.

Keywords deep learning; optimization; motion generation; motion denoising; motion control

1 Introduction

Human-motion generation is useful for many applications, such as human-action recognition [1, 2], motion prediction [3], and video synthesis [4]. Learning a powerful motion model from prerecorded human-motion data is challenging because highly nonlinear kinematic systems intrinsically govern

human motion. As a solution, motion models should scale up effectively to multi-subject motion datasets, synthesize various motions, and be amenable to multiple tasks, such as motion denoising, motion completion, and controllable motion synthesis.

Recent deep-learning-based motion-synthesis algorithms have shown great potential for human-motion generation. Autoregressive models, such as restricted Boltzmann machines and recurrent neural networks (RNNs) [5–7], have been applied to motion synthesis by predicting the possibility of motion in the future. Variational autoencoders (VAEs) [8–10] and generative adversarial networks (GANs) [11–15] have also been applied to motion modeling and synthesis. However, such models must employ careful training strategies to avoid error accumulation and mode collapse. Phase-functioned neural network (PFNN) and its successors [16–18] introduced a phase and local phase to reduce the difficulty of motion modeling.

Inspired by the success of causal-convolution-based WaveNet [19] in multisubject speech synthesis, we propose a novel neural network based on causal convolutions (CCNet) to address the aforementioned issues in motion modeling and synthesis. We added one-dimensional (1D) convolution layers to enable CCNet to accept skeleton configurations as an input, which is necessary for the network to handle the scale and style variations among the skeletons of different subjects. The output of CCNet is the probabilistic density function (PDF) of the motion at the next time step, which is conditioned using the motions at previous time steps, control signals, and skeleton configurations. Using a meticulously designed training strategy, CCNet can effectively capture the intrinsic characteristics of motions of

1 State Key Lab of CAD&CG, Zhejiang University, Hangzhou 310058, China. E-mail: S. Hou, 11721044@zju.edu.cn; Y. Chen, jianfenghou4@163.com; H. Bao, bao@cad.zju.edu.cn; W. Xu, xww@cad.zju.edu.cn (✉).

2 Xmov, Shanghai 200030, China. E-mail: C. Wang, artwang007@gmail.com; J. Chai, chajinxiang@xmov.ai.

3 School of Automation, Southeast University, Nanjing 210096, China. E-mail: W. Zhuang, wlzhuang@seu.edu.cn; Y. Wang, yangangwang@seu.edu.cn.

Manuscript received: 2022-05-25; accepted: 2022-08-02

different subjects and generate more than 20,000 frames of motion.

The freezing issues frequently encountered in prior studies can be effectively mitigated with a Gaussian loss that simultaneously penalizes the deviation of joint angles, positions, and velocities in training. After being trained on motion-capture (mocap) data across multiple subjects, CCNet can generate high-quality motions for different subjects. Furthermore, CCNet can synthesize motions for novel skeletons that are not in the training dataset. If the network is fine-tuned with a small motion dataset of the novel skeleton, it can generate high-quality motions that are similar to the skeleton's ground-truth mocap data. Although the topology of the novel skeletons is the same as that of the skeletons in the training dataset, their scale variations significantly influence the quality and style of the generated motions. CCNet can accommodate these unobserved variations, as shown in our experiments.

We built a new large-scale motion dataset based on 12 subjects with various motion types and transitional clips between different types of motions to model the intrinsic characteristics of multisubject motions. The dataset has 486,282 frames and will be made public together with our code.

To summarize, our main contributions are as follows. (1) We propose a novel neural network based on causal convolutions to model the intrinsic characteristics embodied in the motions of multiple subjects, which has rarely been explored in deep-learning-based motion synthesis methods. (2) A new high-quality motion dataset is constructed across multiple subjects for motion synthesis. (3) CCNet is trained on our new dataset and can efficiently generate high-quality motions (~ 65 fps) and achieve state-of-the-art results for synthesizing characteristic motions for different subjects.

2 Related work

Data-driven motion-modeling methods have become mainstream in recent decades, owing to the development of motion capture techniques and increased computing power of GPUs. Please refer to Refs. [20, 21] for a more comprehensive survey. In the following section, we review motion-synthesis methods that are related to our work.

2.1 Motion style

Our method exhibits superior performance in generating the motions of different intrinsic characteristics associated with different people, even when they possess the same type of human motions. For example, fat and thin men typically have different walking styles for the same motion type. A critical challenge is to model motion styles. Refs. [22, 23] implicitly parameterized the motion styles to synthesize diverse motions. Ref. [24] proposed a motion-style transfer method for a single person to address the problem of unlabeled heterogeneous motions. Wen et al. [25] applied normalizing flows to the task of unsupervised motion-style transfer and achieved impressive results that outperformed state-of-the-art methods. However, none of these methods have been extended to model the variation in human motion styles across different subjects. Aberman et al. [26] proposed a motion-retargeting method to address skeleton variations. However, the motions generated for a new subject were strictly limited to the source motion in terms of types and trajectories and had nearly the same style as the source motion. Ref. [27] proposed a method to address skeleton variations; however, it only models personalized style variations for particular human motions, such as walking or running. In contrast, our method can scale up to generate characteristic motions for multiple subjects of different types, and source motions do not need to be provided as in motion retargeting.

2.2 Deep-learning-based motion synthesis

Deep learning is a remarkable tool for learning a compact, low-dimensional motion space from a dataset. Previous studies have explored many neural-network structures for motion modeling and have made significant progress in this area, such as mixture-of-experts [9, 18], RNNs [28–30], fully connected networks [31, 32], graph networks [33, 34], VAEs [35, 36], and GANs [14, 15]. These methods achieved good performance for short-term motion prediction or periodic locomotion-synthesis tasks, such as walking or running generation. Various types of long-term motion generation often exhibit over-smooth motions or freezing poses [18]. Introducing a phase [16] or local phase [17] as an additional motion feature is helpful for avoiding such problems; however, obtaining these extra temporally related features requires effort.

Exact probability models [25, 37, 38] based on normalizing flows can also solve this problem and add diversity to the synthesized motions. However, these normalizing-flow-based motion-synthesis models are difficult to train and may cause noise and jerkiness.

Despite significant progress in deep-learning-based motion modeling and synthesis, constructing a model that is capable of accurately modeling the characteristics of motions across different subjects remains challenging. This is primarily owing to the lack of datasets containing features of multiple subjects. To address the challenges mentioned above, we propose CCNet, which models the intrinsic characteristics of the motions of different subjects by taking skeleton configurations as an input.

3 Our approach

3.1 Overview

The overall framework of the proposed system is illustrated in Fig. 1. The designed CCNet has three types of functional blocks: a motion-feature embedding encoder, a series of separate residual blocks (SRBs) (light-green blocks) used to capture the temporal correlations, and a decoder that maps the latent features to the probability distribution of the predicted motions. These blocks are discussed in Section 3.2.

We represent the n th frame in our training data as $\mathbf{x}_n = \{\mathbf{x}_n^e, \mathbf{x}_n^o, \mathbf{x}_n^p, \mathbf{x}_n^v, \mathbf{x}_n^f\}$, where \mathbf{x}_n^e denotes the vector of the relative joint rotations, which is represented using exponential coordinates [39], \mathbf{x}_n^o is the vector of the relative angular velocities of the joints, \mathbf{x}_n^p are the 3D joint positions relative to the previous frame, and \mathbf{x}_n^v is the vector of the joint linear velocities. The foot-contact information in the n th frame is represented as a 2D binary

vector \mathbf{x}_n^f . The skeleton configuration combined with the direction, velocity, and motion type form our control signals $\mathbf{c}_n = \{\mathbf{c}_n^s, \mathbf{c}_n^d, \mathbf{c}_n^t\}$, where \mathbf{c}_n^s is a 12D vector formed by sparsely sampling the points on the motion trajectory, starting from the n th frame in a one-second motion clip, and \mathbf{c}_n^t is a 10D vector using one-hot encoding, which represents all ten types of motions in our dataset. Specifically, the skeleton configuration \mathbf{c}_n^s can be represented as $\mathbf{c}_n^s = \{h_r, t_1^x, t_1^y, t_1^z, \dots, t_m^x, t_m^y, t_m^z\}$, where h_r is the height of the root joint, and the 3D positions of non-root joints, that is, $\{t_1^x, t_1^y, t_1^z, \dots, t_m^x, t_m^y, t_m^z\}$, are set to be relative to the root. Given the training data and corresponding control signals, our goal is to train CCNet F_θ parameterized by θ to model the PDF of the predicted motion for the n th frame:

$$p(\mathbf{x}_n | \mathbf{X}, \mathbf{c}_n) = F_\theta(\mathbf{X}, \mathbf{c}_n) \quad (1)$$

where $\mathbf{X} = \{\mathbf{x}_{n-l-1}, \dots, \mathbf{x}_{n-1}\}$ and \mathbf{c}_n are the motion data of the previous l frames and the control signals of the n th frame, respectively.

We used Gaussian loss to encourage CCNet to output ground-truth motion with high probability, foot-contact loss to facilitate the removal of foot sliding in the generated motions, and smoothness loss to reduce jerkiness (see Section 3.3). Noise is added to the sampled training motion data; thus, the network is robust to the accumulated error in the motion synthesis and can produce high-quality, non-freezing motions. The slight foot sliding in the generated motions was removed using an inverse kinematic (IK) algorithm based on the predicted foot-contact labels.

3.2 CCNet architecture

3.2.1 Encoder

Encoder ψ_E has a simple ‘‘Conv1D-ReLU’’ structure, where the kernel size of the 1D convolution is 1. Conv1D layers with a kernel size of 1 ensure

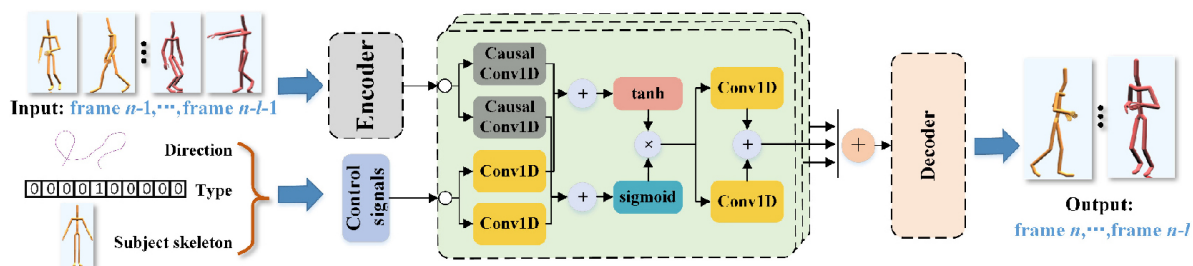


Fig. 1 The proposed CCNet can scale up to a large-scale motion dataset across multiple subjects. Left: Examples of motion-capture data and control signals. Middle: A schematic of our network. It consists of a motion-feature embedding encoder, a series of separate residual blocks (light-green blocks) used to capture temporal correlations, and a decoder that maps the latent features to the probability distribution of predicted motions. We omit the skip connections here for simplicity. Right: Examples of motions sampled from our network outputs.

that the motion feature at each input frame is independent. Formally, the encoder ψ_E takes the motion representation \mathbf{X} of the previous l frames as the input and maps the observed sequence to a latent vector \mathbf{z} :

$$\mathbf{z} = \psi_E(D(\mathbf{X})) \quad (2)$$

The dropout layer before the encoder, denoted by D , is used to resolve the possible overfitting problem, and its drop probability is set to 0.5.

3.2.2 Separate residual blocks

The core component of CCNet is the set of SRBs ψ_R^i , which is illustrated in Fig. 2. The SRBs are similar to the residual blocks used in WaveNet [19], which uses dilated causal convolution to guarantee the temporal ordering of the input motion data. The difference is that we add the Conv1D layers to every SRB to extract the features of the control signals to enhance the model's ability to capture the motion characteristics of different subjects, and we fuse these features with the motions through summation. Our network includes 20 SRBs, which are executed recursively. Each i th block takes the output of the $(i-1)$ th block and the control signals as its input, whereas the inputs of the first block ψ_R^0 are the outputs from the encoder and control signals. The kernel size of the Conv1D layers is 1. Zeros are padded before the feature of $\psi_E(\mathbf{x}_{n-l-1})$; therefore, the output of the causal convolution of a frame n depends only on the frames before it. The padding size can be computed as $(k-1)d$, where k is the kernel size, and d is the dilation size. The causal receptive

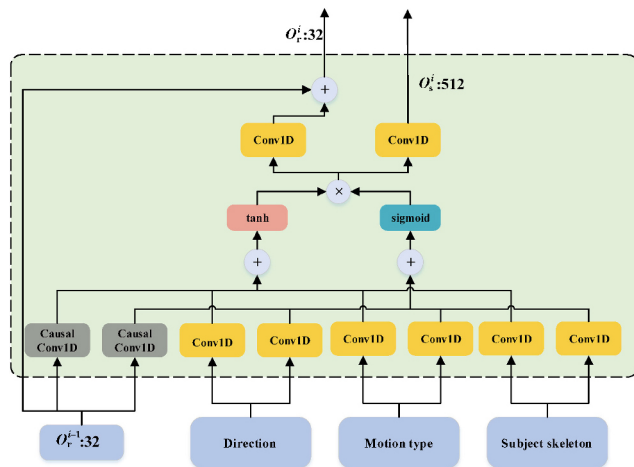


Fig. 2 Detailed architecture of the separate residual block. Each type of control signal is input to its own Conv1D layer, and the kernel size of Conv1D is 1. The numbers beside O_r^{i-1} , O_r^i , and O_s^i indicate their number of channels.

field length CRL of CCNet can be computed using k and d as Eq. (3):

$$\text{CRL} = (k-1) + \sum_{i=0}^{19} d_i(k-1) \quad (3)$$

We can set different dilation sizes d_i for different SRBs to adjust CRL. The kernel and dilation sizes are set to 2 in all SRBs; accordingly, the CRL of CCNet is 41. We also experimented with other CRLs by setting different d_i values; however, we observed that a CRL of 41 is optimal (see Section 4.8). All SRBs generate features $\tilde{\mathbf{z}}$ for the decoder as Eq. (4):

$$\tilde{\mathbf{z}} = \psi_R^0(\mathbf{z}, \mathbf{c}_n) + \sum_{i=1}^{19} \psi_R^i(\dots \psi_R^1(\psi_R^0(\mathbf{z}, \mathbf{c}_n), \mathbf{c}_n)) \quad (4)$$

3.2.3 Decoder

Decoder ψ_D is a simple “ReLU-Conv1D” structure, where the convolution kernel size is set to 1. It maps the summed features from the SRBs to the PDF of the predicted motion, as Eq. (5):

$$\hat{\mu}_n, \sigma_n = \psi_D(\tilde{\mathbf{z}}) \quad (5)$$

where $\hat{\mu}_n$ is a vector of channel-wise mean values. Vector σ_n is used to compute the final standard deviation values $\hat{\sigma}_n = e^{-\sigma_n}$. This element-wise operation ensures that we always obtain positive standard deviation values for $\hat{\sigma}_n$. Subsequently, the poses at frame n can be obtained by directly using the mean $\hat{\mu}_n$ (default setting in our implementation), or they can be sampled from the predicted PDF. Note that the decoder can output n frames each time during the training iterations, owing to the fully convolutional operations.

3.3 Training loss

The training loss consists of four terms: Gaussian loss L_G , motion smoothness loss L_s , foot-contact label loss L_f , and direction-control loss L_d . The training loss can be formulated as Eq. (6):

$$L = L_G + \lambda_1 L_s + \lambda_2 L_f + \lambda_3 L_d \quad (6)$$

where the weights λ_1 , λ_2 , and λ_3 are empirically set to 10.0, 2.0, and 1.0 in all our experiments, respectively.

3.3.1 Gaussian loss

This term follows the Gaussian mixture loss in the work of Fragkiadaki et al., whereas we used only one mode and set the covariance matrix as a diagonal to reduce the number of parameters. It can be written as Eq. (7):

$$L_G = -\ln(p(\mathbf{x}_n | \hat{\mu}_n, \hat{\sigma}_n)) \quad (7)$$

where \mathbf{x}_n is the motion representation extracted from the n th frame, and the binary foot-contact label in \mathbf{x}_n is addressed in L_f ; thus, it is not included in this term. The Gaussian loss learns to maximize the probability of the motion representation vector of the ground-truth mocap data during training; therefore, the captured motion data are of high probability. We add a constraint to ensure that the standard deviation $\hat{\sigma}_n$ is greater than the threshold (1×10^{-4}) using a clipping operation. After training, we observed that the standard deviations that were output by the trained CCNet were typically between 1×10^{-4} and 1×10^{-3} , and their mean value was approximately 2.449 times the threshold of 1×10^{-4} . Consequently, we can sample a motion according to a Gaussian distribution to enrich the variations in the synthesized motion. The joint positions and linear velocities included in this term can help model the correlations between the rotational degrees of freedom of different joints because such quantities are affected by all the parent joints on the kinematic chain connected to the joints.

3.3.2 Smoothness loss

This term is a soft constraint that prevents a sudden change in velocities at joints and smoothens the synthesized motion, which can be formulated as

$$L_s = \sum_{n=2}^N (\hat{\mu}_{n-2} + \hat{\mu}_n - 2\hat{\mu}_{n-1}) \quad (8)$$

The smoothness loss is only optimized for the mean of the predicted Gaussian distributions, because the motion generated by the network is typically close to the mean at each frame.

3.3.3 Foot-contacts loss

We adopt the binary cross-entropy (BCE) loss function to train the network to predict whether the foot is in contact with the supporting plane in the n th frame:

$$L_f = \text{BCE}(\mathbf{x}_n^f, \hat{\mathbf{x}}_n^f) \quad (9)$$

where \mathbf{x}_n^f is the ground-truth foot-contact label for the data, and $\hat{\mathbf{x}}_n^f$ is the network prediction. Foot-contact labels can be used to trigger IK algorithms to remove foot sliding in the synthesized motions.

3.3.4 Direction-control loss

For simplicity, we represent this term as a Gaussian loss and integrate it into L_G . It enables CCNet to predict the direction and velocity control signals, and we only use the mean of the predicted PDF, $\hat{\mathbf{c}}_n^d$ and $\hat{\mathbf{c}}_n^v$,

when generating motions. Thus, the final Gaussian loss becomes

$$L_G = -\ln(p(\mathbf{x}_n, \mathbf{c}_n^d, \mathbf{c}_n^v | \hat{\mu}_n, \hat{\mathbf{c}}_n^d, \hat{\mathbf{c}}_n^v, \hat{\sigma}_n, \hat{\sigma}_n^d, \hat{\sigma}_n^v)) \quad (10)$$

where $\hat{\sigma}_n^d$ and $\hat{\sigma}_n^v$ are the predicted standard deviations of the direction and velocity control signals, respectively, which are computed in the same manner as $\hat{\sigma}_n$ in Eq. (5) and are discarded after training. This term is helpful in interactive motion control when control signals are occasionally inputted by the user. In this case, the predicted control-signal values are fed into the network to continue motion synthesis.

4 Results and discussion

4.1 Dataset and baselines

4.1.1 Datasets

Using the mocap technique, we built a human-motion dataset for 12 different subjects. Three subjects were women, and the rest were men, as shown in Fig. 3. The dataset includes 10 types of motion: walking, running, jumping with the left foot, jumping with the right foot, jumping with both feet, walking backward, zombie walking, kicking, punching, and kicking while punching. All subjects were asked to perform the first seven types of motion, and five were asked to perform the last three types of motion. We recorded approximately 20 min of motion for each subject and asked them to perform two types of motion in one sequence to facilitate the learning of transitions between different motion types. Finally, we obtained 486,282 frames of poses (comprising 27 distinct nonfinger joints) from our datasets. The test dataset was formed by all the motion sequences of subject 7, who was randomly selected from the subjects. Additionally, one motion sequence was randomly




Fig. 3 Skeletons and meshes in our dataset. The skeletons of subjects 0–11 are from left to right, and their heights are 1.6 m, 1.7 m, 1.68 m, 1.83 m, 1.72 m, 1.65 m, 1.95 m, 1.7 m, 1.78 m, 1.66 m, 1.55 m, 1.54 m. Subjects 2, 9, and 10 are women.

selected from the sequences of the remaining subjects. The test dataset was used to test how our network handles skeleton variations after being trained on multisubject motion data. It contained 41 motion sequences and 88,649 frames.

4.1.2 Baselines

We primarily focused on human-motion modeling that captures the intrinsic characteristics embodied in the motions of multiple subjects. This is a unique and rarely explored task compared to most existing methods. There are few methods that are similar to our method proposed in this study; hence, we compare our model to three classic models that are most similar: ERD in the work of Fragkiadaki et al. implemented using four LSTM layers as in Ref. [40], called ERD-4LR; DAE-LSTM [41]; and PFNN, which is an MLP-based network [16]. To test the performances of these three network structures in modeling the motion characteristics of multiple subjects, we added parameters to their first layers to accept skeleton configurations as inputs. Please refer to the Electronic Supplementary Material (ESM) for the detailed network parameters of the models.

4.2 Implementation details

We implemented our algorithm using PyTorch version 1.6. The RMSProp optimizer [42] was employed with an initial learning rate of 1×10^{-4} , which decayed to 1×10^{-6} until 2000 epochs. The batch size was set to 256 with each sample containing a motion sequence of 240 consecutive frames. There are two steps to generating the training batches: (1) randomly selecting a motion clip from the dataset and then the starting frame index in the clip, (2) repeatedly using a one-frame interval for a sample of 240 frames in the clip, that is, the starting frame index, f_{s+1} , of the next 240 frame sequence is $f_s + 1$. For an input sequence, $\mathbf{X} = \{\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_{n-1}\}$, we add independent identically distributed Gaussian noise (with 0 mean and 0.03 standard deviation) to train the network to address accumulated errors in motion synthesis. During training, CCNet can produce the output $\mathbf{Y} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n\}$, owing to the guaranteed ordering in all dilated causal convolutions, which is helpful for speeding up the training procedure.

Although the network can generate high-quality motion, slight foot sliding may still occur. If not mentioned, the IK algorithm is adopted to remove foot sliding in generated motions according to the

predicted foot-contact labels. We refer to the initial frames that are input to CCNet to begin motion generation as seed frames hereafter.

4.3 Quantitative and qualitative evaluation on test dataset

We benchmark our CCNet with baseline models for the motion-denoising error (except for PFNN because it is primarily designed for controllable motion synthesis) and trajectory-following accuracy. We also present examples to demonstrate the quality of the generated motion. In motion denoising, we use the mean of the predicted PDF as the frame poses. In all other experiments, we sample poses from the predicted PDF.

4.3.1 Motion denoising and completion

The trained CCNet can be directly applied to motion denoising and motion completion. For motion denoising, we randomly select the motion sequence \mathbf{X} of a subject and add independent identically distributed Gaussian noise (mean 0, standard deviation 0.01–0.1) to obtain noisy motion data $\hat{\mathbf{X}}$. We use the mean of the predicted PDF as the denoised motion \mathbf{Y} by feeding $\hat{\mathbf{X}}$ to CCNet. Each frame of the denoised poses is not fed back into the network. Frames with indices less than CRL were denoised based on all the frames before them. Figures 4(a) and 4(b) show the denoising results. The standard deviation of the noise in this experiment was set to 0.08. Before denoising, the trajectories of the right hand and right toes fluctuated, and the foot was underneath the ground in some frames. It can be observed that these artifacts are

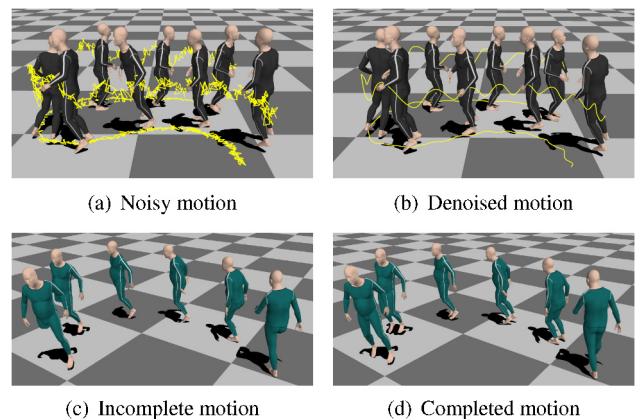


Fig. 4 (a, b) A motion denoising result for subject 7 (from 0 min 9 s to 0 min 20 s in the video in the ESM). (c, d) A motion completion result for subject 5 (from 0 min 24 s to 0 min 35 s in the video in the ESM).

significantly reduced in the denoised motion. We compared CCNet to baseline networks in terms of the quality of the denoised motions. We added Gaussian noise to the test data with standard deviations of 0.03, 0.05, and 0.1 and then used CCNet, DAE-LSTM, and ERD-4LR, respectively, to denoise the noisy motion data. The error between the ground-truth motion and denoising result was computed as the Euclidean distance between their motion-representation vectors. We also trained these three models on a selected CMU mocap dataset to further compare their performance on motion denoising (please refer to the ESM for details on the selection of CMU mocap data). As shown in Table 1, the error of the denoised motion generated by CCNet was less than those of the motions denoised by DAE-LSTM and ERD-4LR.

The motion-completion procedure was similar to that of motion denoising. The experimental results are shown in Figs. 4(c) and 4(d). We first select a 700-frame motion sequence containing walking and jumping with both feet and then set the rotations of the joints of the right legs of 30% of the frames to 0. CCNet accepts incomplete motion as an input and outputs a complete, natural-looking motion. Moreover, it can be observed from Figs. 4(b) and 4(d) that the poses of jumping with both feet vary for different subjects, which means that CCNet can capture the intrinsic styles of different subjects.

4.3.2 Following user-specified trajectories

Synthesizing different types of motions along a specified trajectory is a desirable function in motion planning. We allow users to specify a motion trajectory \mathbf{J} on the XOZ plane with additional velocity and motion-type information. We then map

the trajectory into the control signals \mathbf{c}_n^d and \mathbf{c}_n^t that are supported in our system. Please refer to the ESM for further details.

As shown in Fig. 5, CCNet can synthesize motion using two user-specified trajectories. Figure 7(a) shows that the synthesized motions can follow a trajectory with large curvatures and frequently changing motion types. In Fig. 7(b), we show that CCNet can generate various motions, such as the kicking and punching present in our training dataset, when the user specifies these two types along a trajectory.

We leveraged the average distance between the user-specified and root trajectories on the XOZ plane as the criterion for comparison for the trajectory-following accuracy. In this experiment, we used six different trajectories that were manually specified by users, extracted the direction control signals, and randomly assigned motion types to the trajectory segments. Subsequently, we synthesized motions using the first 120 frames of the 33 locomotion sequences in the test dataset as the seed frames for each specified trajectory and obtained 198 motion-synthesis results. The trajectory distance is computed by summing the closest distance between the projected root position and target trajectory in each frame. The means and standard deviations of the averaged trajectory distances are as follows: 27.878 ± 8.516 cm for CCNet, 158.67 ± 30.94 cm for PFNN, and 171.973 ± 31.862 cm for ERD-4LR; an example is shown in Fig. 6. The results of the CCNet model are more accurate than those of the baseline models. We present the six trajectories in the ESM.

4.4 Interactive control

CCNet can easily be integrated into interactive applications. We demonstrate this capability by

Table 1 Motion-denoising comparisons. STD: standard deviation; D: our dataset; S-CMU: selected CMU mocap dataset. IK is disabled in this experiment. The errors of the motions that are denoised by CCNet are less than those of the motions denoised by ERD-4LR and DAE-LSTM

Noise STD	Denoising error (mean \pm std)			
	ERD-4LR	DAE-LSTM	CCNet	
D	0.03	0.768 \pm 0.357	0.687 \pm 0.384	0.528\pm0.126
	0.05	0.768 \pm 0.357	0.687 \pm 0.384	0.539\pm0.125
	0.1	0.770 \pm 0.361	0.687 \pm 0.384	0.584\pm0.117
S-CMU	0.03	2.248 \pm 1.132	2.214 \pm 1.364	1.789\pm1.250
	0.05	2.251 \pm 1.135	2.214 \pm 1.365	1.789\pm1.250
	0.1	2.246 \pm 1.132	2.217 \pm 1.364	1.788\pm1.251




Fig. 5 Trajectory-following results of two subjects (from 0 min 49 s to 1 min 7 s in the video in the ESM). Left: A synthesized motion transitioning from walking to running and then to zombie walking for subject 10. Right: A synthesized motion transitioning from jumping with the right foot to jumping with the left foot and then to jumping with both feet for subject 11. We use different colors to represent different motion types (refer to the video in the ESM for details).




Fig. 6 Comparisons against ERD-4LR and PFNN (from 4 min 24 s to 4 min 46 s in the video in the ESM). The character starts by jumping with the left foot and then changes to jumping with the right foot till the end. The total errors (3000 frames) between the synthesized trajectories (yellow lines) and input trajectories (green lines) of ERD-4LR, PFNN, and CCNet are 177.143 cm, 156.604 cm, and 27.043 cm, respectively. IK is disabled in this experiment.




Fig. 7 CCNet can synthesize (a) motion heading along a complex trajectory for subject 7 (from 3 min 1 s to 3 min 15 s in the video in the ESM) and (b) the motion of kicking while punching for subject 5 (from 1 min 8 s to 1 min 15 s in the video in the ESM).

developing a demo that allows the user to control direction, velocity, and motion type through a keyboard. Direction and velocity signals are used to generate future motion trajectories \mathbf{c}_n^d online, similar to PFNN. We used the LibTorch API to ease the implementation of CCNet in C++.

Specifically, the user can control the motion type with the number keys, from 1 to 5, to select from five motion types: walking, running, jumping with the left foot, jumping with the right foot, and jumping with both feet. Once a key (for example, 2) is pressed, we update the motion-type label by interpolating the new type label with the previous one in 20 frames, which means that the character can smoothly transition from the previous motion type to the new one. The user can also control the velocity by pressing the up and down keys and the heading direction of the character by pressing the left and right keys. Once the left key is pressed, the trajectory turns to the left. This is achieved by first computing a small offset vector $\mathbf{o}_n = [1, 0] * h * 0.015$, where h is the height of the root. This offset is added to \mathbf{c}_n^d by $\mathbf{o}_n * w_i$, where $w_i = i/5, i = 0.5$. Thus, the offset is added to the six points in the predicted control signal $\hat{\mathbf{c}}_n^d$ through the corresponding w_i . The distance between the 2D points in the updated $\hat{\mathbf{c}}_n^d$ is then adjusted according to the user-specified velocity v_u .

Because $\hat{\mathbf{c}}_n^d$ represents the future motion trajectory within one second, we can adjust the velocity by multiplying the distance between the 2D points by the ratio v_u/v_{cur} . The current scalar velocity of the character, v_{cur} , is computed using the length of the 2D points in \mathbf{c}_n^d . The velocity is changed from the current velocity to the user-specified velocity within a 20-frame interval. Figure 8 illustrates the user interface used in interactive control, and further results can be observed between 1 min 32 s and 1 min 51 s in the accompanying video in the ESM.

4.5 User study

To measure the visual quality of the generated motions, we followed the advice of a researcher in the field of human interaction to conduct a two-alternative forced-choice user study. We selected 16 participants (six women and ten men) with experience in 3D animation or games because they are able to judge motion quality. Then, we gave the participants five clear and detailed criteria, which are described in the ESM, and showed them some examples for each criterion before the user study. The procedure of the user study is as follows. First, we presented the 16 participants with all the groups of motion



Fig. 8 User interface for interactive control. The green dots on the ground represent the direction-control signal. IK is disabled in this experiment.

sequences: four groups for CCNet and the baseline models. Each group contained 16 pairs of motion sequences. In each pair, one is the mocap sequence, and the other is generated by CCNet or one of the baseline models. Second, we asked the participants to answer the question “which motion sequence in the pair is of better motion quality?” according to the five criteria.

After obtaining the user study results, we checked them. First, we checked the time that each participant spent on completing the questionnaire. If the time was less than 10 min (the shortest time needed to judge all motion sequences), we discarded the questionnaire. Second, if a questionnaire had blank responses, it was discarded. Third, if a questionnaire had conflicting choices, it was also discarded. For example, if a participant chose A , B , and C as the better sequence from three pairs, (A, B) , (B, C) , and (A, C) , we treated these as conflicting choices. From the first two choices, we can infer that A is better than C ; however, the participant chose C as the better sequence from the third pair. Finally, we obtained 15 valid questionnaires for DAE-LSTM and 16 for the other models.

We performed a t-test on the user study results to verify the hypothesis that CCNet can generate motions of better quality than the baseline models, and the results are shown in Table 2. The P values of CCNet versus other baseline models were all less than the selected threshold (0.05). Therefore, the motions generated by CCNet were significantly different from those generated by the baselines. Based on the average number of motion sequences selected by the participants (mean in Table 2), the number of choices for CCNet is larger than those of the other baselines, which verifies that CCNet

Table 2 T-test of user-study results (confidence interval = 0.95). VS: performing t-test between the results of CCNet and all the results of the baseline models in the second row. Mean: the average number of generated sequences selected by all the participants compared to mocap sequences in the same group. Std: the standard deviation of the number that is selected

	DAE-LSTM	ERD-4LR	PFNN
VS	(mean: 1.75 std: 0.968)	(mean: 2.75 std: 1.199)	(mean: 2.375 std: 0.992)
CCNet (mean: 6.625 std: 1.165)	P-value: 6.1188E-13 t-value: -11.9588	P-value: 2.7365E-8 t-value: -7.4383	P-value: 1.0127E-9 t-value: -8.7165

can better capture the intrinsic characteristics of the motions of different subjects. Furthermore, we prepared another three-group dataset that contained pairs of motion sequences generated by CCNet and each baseline model. As listed in Table 3, the number of CCNet-generated motion sequences selected by the participants was still higher than that of the sequences generated by the baseline models. We also performed an ANOVA test on the user study results, as illustrated in Table 4, which verifies the user study’s statistical significance.

4.6 Generalization to unseen skeletons

After training CCNet with multisubject motion data, the model can generate motions for skeletons that are not in the training dataset. As illustrated in Figs. 3, 4(a), 4(b), and 7(a), we applied the trained CCNet to automatically generate motions for the skeleton of subject 7, which was unseen during training. It can be observed in Fig. 9 that the baseline models cannot differentiate the variations exhibited by different skeletons as effectively as CCNet can. We further tested the generalization ability of CCNet by applying it to a specially designed skeleton that was generated by scaling the skeleton of subject 7. The topology of the skeleton remained the same as that of other skeletons but varied substantially in the lower-body scale. Because there is no mocap data for the skeleton, we utilized the motion-retargeting [43]

Table 3 Average selected numbers for CCNet-generated motion sequences. Baseline vs. CCNet: a group of 16 pairs of motion sequences generated by a baseline model and CCNet. Mean±std: mean and variance of the numbers of CCNet-generated motion sequences selected by all the participants

Groups	Numbers for CCNet (mean±std)
DAE-LSTM vs. CCNet	12.31±2.34
ERD-4LR vs. CCNet	11.63±2.87
PFNN vs. CCNet	11.45±1.87

Table 4 ANOVA-test of user study for confidence interval = 0.95. SS: sum-of-squares for variability between groups. df: degrees of freedom. MS: mean squares. F: F ratio. —: not applicable

Source of variation	Between groups	Within groups	Total
SS	346.6875	76.7500	423.4375
df	3	60	63
MS	115.5625	1.2792	—
F	90.3420	—	—
P-value	3.1855E-22	—	—
F-critical	2.7581	—	—




Fig. 9 Foot-ground penetrations in the motions generated by DAE-LSTM, ERD-4LR, and PFNN. Left: 611-th frame generated by DAE-LSTM for subject 9. Middle: 450-th frame generated by ERD-4LR for subject 7. Right: 666-th frame generated by PFNN for subject 8. DAE-LSTM, ERD-4LR, and PFNN cannot effectively differentiate the variations in different skeletons and lead to foot-ground penetrations, as indicated by the red rectangles.

algorithm to generate 120 seed frames for it. Figure 10 illustrates that CCNet can effectively generalize the new skeleton. In addition, we used ERD-4LR and PFNN to generate motions for the skeleton. The results show that both motions contain large, sharp changes between the seed frames and generated frames, which is inferior to the motions generated by CCNet. Please refer to the accompanying video in the ESM from 4 min 16 s to 4 min 21 s for the relevant results.

Given a part of the motion data of a novel skeleton, CCNet can learn to generate motions for the skeleton that are similar to its ground-truth mocap data. Table 5 shows that, after fine-tuning the network using the walking and running motion of subject 7, the relative pose difference rel_p for all mocap data of this subject in the test dataset can be significantly reduced (refer to the accompanying video in the ESM from 2 min 31 s to 2 min 47 s for the comparison

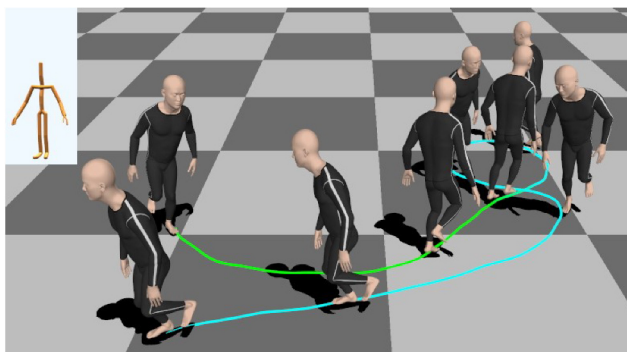


Fig. 10 Trajectory-following results generated by CCNet for an unseen skeleton subject 7b (from 2 min 3 s to 2 min 15 s in the video in the ESM). The skeleton of subject 7b is generated by scaling the lower body of subject 7's skeleton by 0.8.

Table 5 Influence of the fine-tuning of CCNet with partial motion data of an unseen skeleton for subject 7. FT: fine-tuning, W: walking; R: running; CCNet-D1: CCNet trained on dataset 1, which contains motions of subjects 1, 3, 4, and 8; CCNet-D2: CCNet trained on dataset 2, which contains motions of subjects 0, 5, 6, and 11. Fine-tuning CCNet trained on our entire training dataset with walking and running mocap data of subject 7 (first row) achieves the lowest relative pose difference. IK is disabled in this experiment

Model	Relative pose difference (mean \pm std)		
	No FT	FT with W	FT with W and R
CCNet	0.0830 \pm 0.0925	0.0536 \pm 0.0365	0.0483\pm0.0544
CCNet-D1	0.0861 \pm 0.1550	0.0543 \pm 0.0646	0.0525 \pm 0.0594
CCNet-D2	0.0914 \pm 0.1080	0.0598 \pm 0.0770	0.0590 \pm 0.0713

of generated jumping motions of subject 7 before and after fine-tuning). This implies that CCNet can capture the intrinsic characteristics embodied in the motions of the new subject better than other models. We compute the relative pose difference as

$$\text{rel}_p = \frac{1}{N} \sum_{n=0}^N (\|\hat{\mathbf{x}}_n - \mathbf{x}_n\|_2 / \|\mathbf{x}_n\|_2)$$

where N is the number of frames, and $\hat{\mathbf{x}}_n$ and \mathbf{x}_n are the motion-representation vectors of motions generated by CCNet and the corresponding mocap data. The ability to generalize to new skeletons is crucial because it can reduce efforts to capture a large amount of mocap data for a new skeleton in motion-synthesis applications.

To evaluate how the number of subjects in the dataset influences the generalization ability of CCNet, we intentionally put the motions of subjects 1, 3, 4, and 8 into dataset 1 and the motions of subjects 0, 5, 6, and 11 into dataset 2. Table 5 lists the rel_p values of the motion generated by CCNet trained on dataset 1 (CCNet-D1) and dataset 2 (CCNet-D2). Because the heights of subjects 1, 3, 4, and 8 are closer to subject 7's height, the rel_p value of CCNet-D1 is less than that of CCNet-D2; however, this value is still larger than that of CCNet trained on the entire training dataset. Thus, to improve the generalization ability of CCNet for new skeletons, it is better to construct a dataset with more subjects so that the network learns how to process their skeleton and style variations.

4.7 Motion prediction on H3.6M dataset

We directly trained CCNet on the H3.6M [44] dataset without any additional modifications to test its motion-prediction ability. We followed the same data representation and reported the mean angle error

(MAE) in the same test dataset, as Fragkiadaki et al. and Liu et al. do. We compared CCNet with ERD-4LR, HP-GAN [12], QuaterNet [45], and AM-GAN [15], and the results are reported in Table 6. Because Liu et al. have not yet released their code, the MAEs of AM-GAN in Table 6 are sourced from their paper. CCNet is primarily designed for long-term multi-subject motion generation (typically at least 10 s) instead of motion prediction; however, it also achieves average performance among these methods. We believe that with meticulous adjustments, the performance of CCNet can be improved for motion prediction.

4.8 Evaluation of network hyper-parameters and training settings

In this section, we present ablation-study experiments to determine the following hyper-parameters and training settings selected for CCNet: (1) the causal receptive field length (CRL), (2) number of consecutive frames of each sample in a batch (NCF), (3) joint rotations and angular velocities in the

data representation (ROT), (4) smoothness loss term (Smooth), (5) skeleton configurations (SK), (6) advantages of SRBs over LSTM (SRB2LSTM), and (7) seed frame length when synthesizing. We chose these values or settings to minimize loss L in Eq. (6) for both the training and test datasets. For a better visualization, we plotted the logarithmic loss curves using the formula $\log_{10}(L + 320)$ in Figs. 11 and 12. Because L is typically approximately -300 , a bias of 320 is necessary to obtain a positive result.

4.8.1 CRL and NCF

We conducted experiments to choose the CRL value among three settings: 31 (dilation sizes of SRBs are repeatedly 1 and 2), 41 (dilation sizes of SRBs are 2), and 46 (dilation sizes of SRBs are repeatedly 1, 2, and 4). The number of SRBs is fixed at 20. We also tested three different NCFs, 60, 120, and 240, which correspond to 1 s, 2 s, and 4 s of motion for our 60-fps motion dataset, respectively. Note that we keep the NCF fixed at 240 when computing the loss on the test dataset for fair comparisons. Based

Table 6 Motion prediction comparisons (in MAE) over four different action types in the H3.6M dataset

Method	80 ms	160 ms	320 ms	400 ms	1000 ms
	Walking				
ERD-4LR	0.93	1.18	1.59	1.78	2.24
HP-GAN	0.95	1.17	1.69	1.79	2.47
QuaterNet	0.35	0.64	1.19	1.38	1.58
AM-GAN	0.23	0.51	0.62	0.66	0.84
CCNet	1.34	1.40	1.45	1.46	1.55
Eating					
ERD-4LR	1.27	1.45	1.66	1.80	2.02
HP-GAN	1.28	1.47	1.70	1.82	2.51
QuaterNet	0.31	0.49	0.82	0.97	1.89
AM-GAN	0.20	0.31	0.49	0.66	1.15
CCNet	1.35	1.45	1.58	1.65	2.04
Smoking					
ERD-4LR	1.66	1.95	2.35	2.42	3.14
HP-GAN	1.71	1.89	2.33	2.42	3.20
QuaterNet	0.32	0.55	0.96	1.07	1.37
AM-GAN	0.25	0.46	0.88	0.88	1.10
CCNet	1.81	1.92	2.14	2.21	2.65
Discussion					
ERD-4LR	2.27	2.47	2.68	2.76	3.11
HP-GAN	2.29	2.61	2.79	2.88	3.67
QuaterNet	0.31	0.67	0.94	1.04	1.96
AM-GAN	0.28	0.55	0.81	0.92	1.58
CCNet	1.55	1.72	1.83	1.83	2.22




Fig. 11 Logarithmic loss curves obtained using different hyper-parameters of CCNet. Left: training. Right: test. We modify each hyper-parameter, including CRL, NCF, with/without skeleton configuration (SK or w/o.SK), and with/without the joint rotations and angular velocities in the data representation (ROT or w/o.ROT).

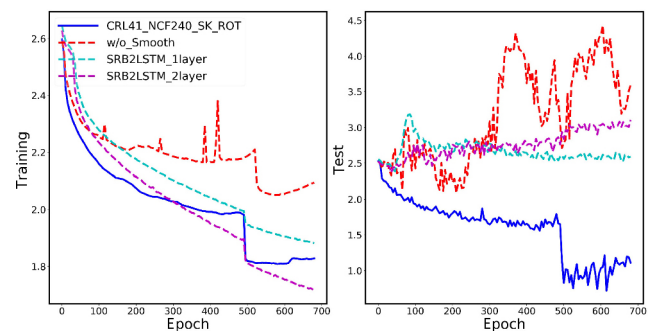


Fig. 12 Ablation study of smoothness loss term and SRBs. We remove the smoothness loss term, replace the SRBs with a 1-layer LSTM and 2-layer LSTM, and evaluate the logarithmic losses of the corresponding re-trained models. Left: training. Right: test.

on Fig. 11, we chose $CRL = 41$ and $NCF = 240$ for CCNet because these settings led to the lowest loss.

4.8.2 ROT, smooth, and SK

To verify their importance, we removed the joint rotations and angular velocities from the input and removed the smoothness loss term from Eq. (6). Additionally, the skeleton configurations were removed by disconnecting their corresponding 1D convolution modules from CCNet. We can observe from Figs. 11 and 12 that without the joint rotations and angular velocities, smoothness loss term, or skeleton configurations, the corresponding networks overfit the training set. Therefore, we can conclude that encoding joint rotations and angular velocities in the data representation and the smoothness loss term are essential for CCNet to converge to a better result. Additionally, skeleton configuration is important in the network to disambiguate the motions of different subjects.

4.8.3 LSTM vs. SRB

We determined the advantages of the SRBs over LSTM by replacing SRBs with LSTM layers with different numbers of layers and hidden state channels. The results show that only the 1-layer and 2-layer LSTM settings with 512 hidden channels using a smaller initial learning rate, 1×10^{-5} , can converge to but overfit the training set, as depicted in Fig. 12.

4.8.4 Seed frame length

The influence of seed-frame length on the quality of the generated motions was measured using rel_p . Specifically, we extracted seed frames from the mocap data in the test dataset and then used the networks to predict a frame for comparison with the corresponding mocap frame. Table 7 shows the rel_p values for different seed-frame lengths. The lower value indicates that the generated motion has more

Table 7 Ablation study on the length of seed frames. We synthesized motion sequences using different seed-frame lengths to check their influences on the generated motions. IK is disabled

Seed frame numbers	Relative pose difference (mean \pm std)		
	ERD-4LR	DAE-LSTM	CCNet
1	0.214 \pm 0.302	0.560 \pm 0.572	0.130\pm0.179
5	0.180 \pm 0.145	0.325 \pm 0.508	0.091\pm0.128
10	0.140 \pm 0.159	0.303 \pm 0.490	0.084\pm0.136
30	0.145 \pm 0.179	0.182 \pm 0.270	0.074\pm0.128
60	0.120 \pm 0.138	0.178 \pm 0.253	0.074\pm0.124
120	0.126 \pm 0.160	0.115 \pm 0.144	0.074\pm0.125

similarity with the mocap data; thus, it is of better quality. It can be observed that CCNet is robust to variations in the length of the seed frames compared to ERD-4LR and DAE-LSTM, and it does not require excessively long seed frames to synthesize high-quality motions. However, we observe that jitters between the seed frames and generated frames are slightly more obvious when the seed-frame numbers are one and five (please refer to the accompanying video in the ESM from 3 min 23 s to 3 min 46 s for details). We hypothesize that CCNet cannot obtain sufficient information to generate smooth motions for such short seed frames.

5 Conclusions

We designed a novel neural network for motion generation, CCNet, to synthesize high-quality motions for multiple subjects. The trained CCNet can capture the motion characteristics of different subjects well and synthesize various types of motion, such as punching. Moreover, CCNet can generate motion for novel skeletons. Given a few sample motions of a novel skeleton, the pretrained CCNet can be fine-tuned to synthesize motions that better reproduce the intrinsic characteristics of the motions of the skeleton. In the future, we plan to extend our method to include skeletons with different typologies.

Acknowledgements

We thank the anonymous reviewers for their constructive comments. Weiwei Xu is partially supported by the National Natural Science Foundation of China (No. 61732016).

Declaration of competing interest

The authors have no competing interests to declare that are relevant to the content of this article. The author Hujun Bao is the Associate Editor of this journal.

Electronic Supplementary Material

There are three types of electronic supplementary materials. The first is a video named “ccnet_cvm.mp4”, which contains the generated motions rendered into the video. The second is a file named “CCNet_supplementary_material.pdf”, which describes the details of the model parameters, user

study criteria, and trajectories used in the trajectory-following experiment. The last includes python codes and a training script named “run_wavenet.sh”. They are available in the online version of this article at <https://doi.org/10.1007/s41095-022-0307-3>.

References

- [1] Zhang, P. F.; Lan, C. L.; Zeng, W. J.; Xing, J. L.; Xue, J. R.; Zheng, N. N. Semantics-guided neural networks for efficient skeleton-based human action recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 1109–1118, 2020.
- [2] Chen, Z.; Li, S. C.; Yang, B.; Li, Q. H.; Liu, H. Multi-scale spatial temporal graph convolutional network for skeleton-based action recognition. *Proceedings of the AAAI Conference on Artificial Intelligence* Vol. 35, No. 2, 1113–1122, 2021.
- [3] Gui, L. Y.; Wang, Y. X.; Liang, X. D.; Moura, J. M. F. Adversarial geometry-aware human motion prediction. In: *Computer Vision – ECCV 2018. Lecture Notes in Computer Science, Vol. 11208*. Ferrari, V.; Hebert, M.; Sminchisescu, C.; Weiss, Y. Eds. Springer Cham, 823–842, 2018.
- [4] Wang, T. C.; Liu, M. Y.; Tao, A.; Liu, G. L.; Kautz, J.; Catanzaro, B. Few-shot video-to-video synthesis. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems, Article No. 451, 5013–5024, 2019.
- [5] Taylor, G. W.; Hinton, G. E. Factored conditional restricted Boltzmann Machines for modeling motion style. In: Proceedings of the 26th Annual International Conference on Machine Learning, 1025–1032, 2009.
- [6] Fragkiadaki, K.; Levine, S.; Felsen, P.; Malik, J. Recurrent network models for human dynamics. In: Proceedings of the IEEE International Conference on Computer Vision, 4346–4354, 2015.
- [7] Martinez, J.; Black, M. J.; Romero, J. On human motion prediction using recurrent neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 4674–4683, 2017.
- [8] Yan, X. C.; 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: *Computer Vision – ECCV 2018. Lecture Notes in Computer Science, Vol. 11209*. Ferrari, V.; Hebert, M.; Sminchisescu, C.; Weiss, Y. Eds. Springer Cham, 276–293, 2018.
- [9] Ling, H. Y.; Zinno, F.; Cheng, G.; Van De Panne, M. Character controllers using motion VAEs. *ACM Transactions on Graphics* Vol. 39, No. 4, Article No. 40, 2020.
- [10] Ghosh, A.; Cheema, N.; Oguz, C.; Theobalt, C.; Slusallek, P. Synthesis of compositional animations from textual descriptions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 1376–1386, 2021.
- [11] Wang, Z. Y.; Chai, J. X.; Xia, S. H. Combining recurrent neural networks and adversarial training for human motion synthesis and control. *IEEE Transactions on Visualization and Computer Graphics* Vol. 27, No. 1, 14–28, 2021.
- [12] Barsoum, E.; Kender, J.; Liu, Z. C. HP-GAN: Probabilistic 3D human motion prediction via GAN. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 1499–149909, 2018.
- [13] Kundu, J. N.; Gor, M.; Babu, R. V. BiHMP-GAN: Bidirectional 3D human motion prediction GAN. *Proceedings of the AAAI Conference on Artificial Intelligence* Vol. 33, No. 1, 8553–8560, 2019.
- [14] Wang, Z. Y.; Yu, P.; Zhao, Y.; Zhang, R. Y.; Zhou, Y. F.; Yuan, J. S.; Chen, C. Y. Learning diverse stochastic human-action generators by learning smooth latent transitions. *Proceedings of the AAAI Conference on Artificial Intelligence* Vol. 34, No. 7, 12281–12288, 2020.
- [15] Liu, Z. G.; Lyu, K. D.; Wu, S.; Chen, H. P.; Hao, Y. B.; Ji, S. L. Aggregated multi-GANs for controlled 3D human motion prediction. *Proceedings of the AAAI Conference on Artificial Intelligence* Vol. 35, No. 3, 2225–2232, 2021.
- [16] Holden, D.; Komura, T.; Saito, J. Phase-functioned neural networks for character control. *ACM Transactions on Graphics* Vol. 36, No. 4, Article No. 42, 2017.
- [17] Starke, S.; Zhao, Y. W.; Komura, T.; Zaman, K. Local motion phases for learning multi-contact character movements. *ACM Transactions on Graphics* Vol. 39, No. 4, Article No. 54, 2020.
- [18] Starke, S.; Zhao, Y. W.; Zinno, F.; Komura, T. Neural animation layering for synthesizing martial arts movements. *ACM Transactions on Graphics* Vol. 40, No. 4, Article No. 92, 2021.
- [19] Van den Oord, A.; Dieleman, S.; Zen, H.; Simonyan, K.; Vinyals, O.; Graves, A.; Kalchbrenner, N.; Senior, A.; Kavukcuoglu, K. WaveNet: A generative model for raw audio. In: Proceedings of the 9th ISCA Speech Synthesis Workshop, 125, 2016.
- [20] Wang, X.; Chen, Q. D.; Wang, W. L. 3D human motion editing and synthesis: A survey. *Computational and Mathematical Methods in Medicine* Vol. 2014, 104535, 2014.

- [21] Xia, S. H.; Gao, L.; Lai, Y. K.; Yuan, M. Z.; Chai, J. X. A survey on human performance capture and animation. *Journal of Computer Science and Technology* Vol. 32, No. 3, 536–554, 2017.
- [22] Brand, M.; Hertzmann, A. Style machines. In: Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, 183–192, 2000.
- [23] Xu, J. W.; Xu, H. Z.; Ni, B. B.; Yang, X. K.; Wang, X. L.; Darrell, T. Hierarchical style-based networks for motion synthesis. In: *Computer Vision – ECCV 2020. Lecture Notes in Computer Science, Vol. 12356*. Vedaldi, A.; Bischof, H.; Brox, T.; Frahm, J. M. Eds. Springer Cham, 178–194, 2020.
- [24] Xia, S. H.; Wang, C. Y.; Chai, J. X.; Hodgins, J. Realtime style transfer for unlabeled heterogeneous human motion. *ACM Transactions on Graphics* Vol. 34, No. 4, Article No. 119, 2015.
- [25] Wen, Y. H.; Yang, Z. P.; Fu, H. B.; Gao, L.; Sun, Y. N.; Liu, Y. J. Autoregressive stylized motion synthesis with generative flow. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 13607–13607, 2021.
- [26] Aberman, K.; Li, P. Z.; Lischinski, D.; Sorkine-Hornung, O.; Cohen-Or, D.; Chen, B. Q. Skeleton-aware networks for deep motion retargeting. *ACM Transactions on Graphics* Vol. 39, No. 4, Article No. 62, 2020.
- [27] Min, J. Y.; Liu, H. J.; Chai, J. X. Synthesis and editing of personalized stylistic human motion. In: Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, 39–46, 2010.
- [28] Nie, Q.; Liu, Z. W.; Liu, Y. H. Unsupervised 3D human pose representation with viewpoint and pose disentanglement. In: *Computer Vision – ECCV 2020. Lecture Notes in Computer Science, Vol. 12364*. Vedaldi, A.; Bischof, H.; Brox, T.; Frahm, J. M. Eds. Springer Cham, 102–118, 2020.
- [29] Corona, E.; Pumarola, A.; Alenyà, G.; Moreno-Noguer, F. Context-aware human motion prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 6990–6999, 2020.
- [30] Ghorbani, S.; Wloka, C.; Etemad, A.; Brubaker, M. A.; Troje, N. F. Probabilistic character motion synthesis using a hierarchical deep latent variable model. *Computer Graphics Forum* Vol. 39, No. 8, 225–239, 2020.
- [31] Holden, D.; Saito, J.; Komura, T. A deep learning framework for character motion synthesis and editing. *ACM Transactions on Graphics* Vol. 35, No. 4, Article No. 138, 2016.
- [32] Mao, W.; Liu, M. M.; Salzmann, M. History repeats itself: Human motion prediction via motion attention. In: *Computer Vision – ECCV 2020. Lecture Notes in Computer Science, Vol. 12359*. Vedaldi, A.; Bischof, H.; Brox, T.; Frahm, J. M. Eds. Springer Cham, 474–489, 2020.
- [33] Li, M. S.; Chen, S. H.; Zhao, Y. H.; Zhang, Y.; Wang, Y. F.; Tian, Q. Dynamic multiscale graph neural networks for 3D skeleton based human motion prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 211–220, 2020.
- [34] Cui, Q. J.; Sun, H. J.; Yang, F. Learning dynamic relationships for 3D human motion prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 6518–6526, 2020.
- [35] Yuan, Y.; Kitani, K. M. Diverse trajectory forecasting with determinantal point processes. In: Proceedings of the 8th International Conference on Learning Representations, 2020.
- [36] Petrovich, M.; Black, M. J.; Varol, G. Action-conditioned 3D human motion synthesis with transformer VAE. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, 10965–10975, 2021.
- [37] Henter, G. E.; Alexanderson, S.; Beskow, J. MoGlow: Probabilistic and controllable motion synthesis using normalising flows. *ACM Transactions on Graphics* Vol. 39, No. 6, Article No. 236, 2020.
- [38] Valle-Pérez, G.; Henter, G. E.; Beskow, J.; Holzapfel, A.; Oudeyer, P. Y.; Alexanderson, S. Transflower: Probabilistic autoregressive dance generation with multimodal attention. *ACM Transactions on Graphics* Vol. 40, No. 6, Article No. 195, 2021.
- [39] Grassia, F. S. Practical parameterization of rotations using the exponential map. *Journal of Graphics Tools* Vol. 3, No. 3, 29–48, 1998.
- [40] Lee, K.; Lee, S.; Lee, J. Interactive character animation by learning multi-objective control. *ACM Transactions on Graphics* Vol. 37, No. 6, Article No. 180, 2018.
- [41] Ghosh, P.; Song, J.; Aksan, E.; Hilliges, O. Learning human motion models for long-term predictions. In: Proceedings of the International Conference on 3D Vision, 458–466, 2017.
- [42] Graves, A. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.
- [43] Buss, S. R. Introduction to inverse kinematics with Jacobian transpose, pseudoinverse and damped least squares methods. *IEEE Journal of Robotics and Automation* Vol. 17, 16, 2004.
- [44] Ionescu, C.; Papava, 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* Vol. 36, No. 7, 1325–1339, 2014.


- [45] Pavllo, D.; Feichtenhofer, C.; Auli, M.; Grangier, D. Modeling human motion with quaternion-based neural networks. *International Journal of Computer Vision* Vol. 128, No. 4, 855–872, 2020.



Shuaiying Hou is currently a Ph.D. student at State Key Lab of CAD&CG at Zhejiang University. He received his B.S. degree in software engineering from Northwestern Polytechnical University, in 2017. His research interests include computer animation and computer graphics.




Congyi Wang received his Ph.D. degree in computer science from Institute of Computing Technology, Chinese Academy of Sciences in 2017. Since 2018, he has been a research scientist at Xmov, a startup company aiming at AI-powered virtual production. His research interests include computer animation, computer graphics, computer vision, and speech signal processing.



Wenlin Zhuang received his M.S. degree from the School of Automation, Southeast University, in 2021. His research interests include human pose estimation and 3D human animation.




Yu Chen received his master degree in computer science from Zhejiang University, in 2016. His research interests include computer animation, computer graphics, computer vision, and speech signal processing.



Yangang Wang is currently an associate professor in the School of Automation at Southeast University. He received his Ph.D. degree from Department of Automation at Tsinghua University, in 2014. His research interests include motion capture and animation, 3D reconstruction, and image processing.



Hujun Bao is a Cheung Kong professor in the School of Computer Science and Technology in Zhejiang University, and the director of State Key Laboratory of CAD&CG. He received his B.S. and Ph.D. degrees in applied mathematics from Zhejiang University in 1987 and 1993, respectively. His research interests include geometry computing, vision computing, real-time rendering and virtual reality.



Jinxiang Chai received his Ph.D. degree in computer science from Carnegie Mellon University. He is currently an associate professor in the Department of Computer Science and Engineering at Texas A&M University. His primary research is in the area of computer graphics and vision with broad applications in other disciplines such as virtual and augmented reality, robotics, human computer interaction, and biomechanics. He received an NSF CAREER award for his work on theory and practice of Bayesian motion synthesis.



Weiwei Xu is currently a professor at the State Key Lab of CAD&CG in Zhejiang University. He was a Qianjiang Professor at Hangzhou Normal University and a researcher in the Internet Graphics Group at Microsoft Research Asia from 2005 to 2012. He was a post-doc researcher at Ritsmeikan University in Japan for over one year. He received his Ph.D. degree in computer graphics from Zhejiang University, and B.S. and master degrees in computer science from Hohai University in 1996 and 1999, respectively.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Other papers from this open access journal are available free of charge from <http://www.springer.com/journal/41095>. To submit a manuscript, please go to <https://www.editorialmanager.com/cvmj>.

