



UNIVERSITY OF LEEDS

This is a repository copy of *Dynamic Future Net: Diversified Human Motion Generation*.

White Rose Research Online URL for this paper:

<http://eprints.whiterose.ac.uk/163776/>

Version: Accepted Version

---

**Proceedings Paper:**

Chen, W, Wang, H [orcid.org/0000-0002-2281-5679](https://orcid.org/0000-0002-2281-5679), Yuan, Y et al. (2 more authors) (2020) *Dynamic Future Net: Diversified Human Motion Generation*. In: *MM '20: Proceedings of the 28th ACM International Conference on Multimedia*. ACM Multimedia 2020, 12-16 Oct 2020, Online. Association for Computing Machinery , pp. 2131-2139. ISBN 978-1-4503-7988-5

<https://doi.org/10.1145/3394171.3413669>

---

© 2020 ACM. This is an author produced version of a conference paper published in *MM '20: Proceedings of the 28th ACM International Conference on Multimedia*. Uploaded in accordance with the publisher's self-archiving policy.

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

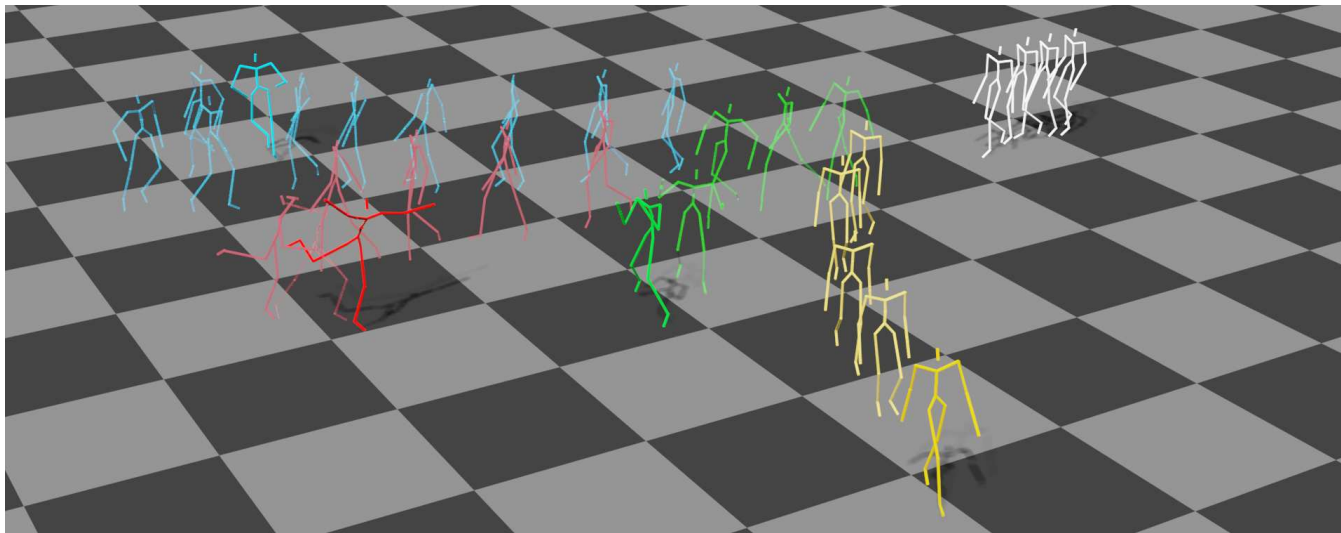
If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# Dynamic Future Net: Diversified Human Motion Generation

Anonymous Author(s)  
Submission Id: 1482



**Figure 1:** Given a 20-frame walking motion pre-fix (white), our model can generate diversified motion: walking (yellow), walking-to-running (blue), walking-to-boxing (green), and walking to dancing (red), with arbitrary duration. The corresponding animation can be found in `teaser.mp4` in supplementary video.

## ABSTRACT

Human motion modelling is crucial in many areas such as computer graphics, vision and virtual reality. Acquiring high-quality skeletal motions is difficult due to the need for specialized equipment and laborious manual post-posing, which necessitates maximizing the use of existing data to synthesize new data. However, it is a challenge due to the intrinsic motion stochasticity of human motion dynamics, manifested in the short and long terms. In the short term, there is strong randomness within a couple frames, e.g. one frame followed by multiple possible frames leading to different motion styles; while in the long term, there are non-deterministic action transitions. In this paper, we present *Dynamic Future Net*, a new deep learning model where we explicitly focus on the aforementioned motion stochasticity by constructing a generative model with non-trivial modelling capacity in temporal stochasticity. Given limited amounts of data, our model can generate a large number of high-quality motions with arbitrary duration, and visually-convincing variations in both space and time. We evaluate our model on a wide range of motions and compare it with the

state-of-the-art methods. Both qualitative and quantitative results show the superiority of our method, for its robustness, versatility and high-quality.

## CCS CONCEPTS

• Computing methodologies → Motion processing; Neural networks; Learning latent representations.

## KEYWORDS

human motion, neural networks, generative models

## ACM Reference Format:

Anonymous Author(s). 2020. Dynamic Future Net: Diversified Human Motion Generation. In *Multimedia '20: ACM International Conference on Multimedia, October 12–16, 2020, 2018, Seattle, WA*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/1122445.1122456>

## 1 INTRODUCTION

Modeling natural human motions is a central topic in several fields such as computer animation, bio-mechanics, virtual reality, etc, where high-quality motion data is a necessity. Despite the improved accuracy and lowered costs of motion capture systems, it is still highly desirable to make full use of existing data to generate diversified new data. One key challenge in motion generation is the dynamics modelling, where it has been shown that a latent space can be found due to the high coordination of body motions

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*ACM Multimedia '20, October 12–16, 2020, Seattle, WA*

© 2020 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/10.1145/1122445.1122456>

[21, 34, 39]. However, as much as the spatial aspect is studied, dynamics modelling, especially with the aim of diversified motion generation, still remains to be an open problem.

Human motion dynamics manifest several levels of short-term and long-term stochasticity. Given a homogeneous discretization of motions in time, the short-term stochasticity refers to the randomness in next one or few frames (pose-transition); while the long-term one refers to the random in the next or few actions (action-transition). Tradition methods model them by Finite State Machines with carefully organized data [31], which have limited model capacities for large amounts of data and require extensive pre-processing work. New deep learning methods either ignore them [21] or do not explicitly model them [39]. Very recently, dynamics modelling for diversified generation has just been investigated [40], but only from the perspective of the overall dynamics, rather than the detailed short/long term stochasticity.

In this paper, we propose a new deep learning model, Dynamic Future Net, or DFN, for automatic and diversified high-quality motion generation based on limited amounts of data. Given a motion, we assume that it can be discretized homogeneously in time and represented by a series of posture and instantaneous velocities. Following the observation that it is easier to learn the dynamics in a natural motion latent space [21], we first embed features in the data space into a latent space. Next, DFN learns explicitly the *history*, *current* and *future* state given any time, where we also model several conditional distributions for the influences of history and future state on the current state. The state-wise bidirectional modelling (extending into both the past and future) separates DFN from existing methods and endows us with the ability of modelling the short-term (next-frame) randomness and long-term (next-action) randomness. Last, for inference purposes, we propose new loss functions based on distributional similarities as opposed to point-wise estimation [39, 42], which captures the dynamics accurately but also keep the randomness that is crucial for diversified motion generation.

We show extensive experimental results to show DFN's robustness, versatility and high-quality. Unlike existing methods which have to be trained on one type of motions a time [40, 42], DFN can be trained using both single type of motions or mixed motions, which shows DFN's ability to capture multi-modal dynamics and therefore its versatility in diversified motion generation. Visual evaluation shows that DFN can generate high-quality motions with different dynamics.

In summary, our formal contributions include:

- (1) a new deep learning model, Dynamic Future Net, for automatic, diversified and high-quality human motion generation.
- (2) a new dynamic model that captures the transition stochasticity of the past, current and future states in motions.
- (3) insights of the importance of both short-term and long-term dynamics in human motion modelling.

## 2 BACKGROUND AND RELATED WORKS

### 2.1 Human pose and motion embedding

Given a human motion sequence, it is useful to find the low dimension representation of the whole sequence. Holden et al. [22] for

the first time use a convolution neural network to project the entire sequence to a low dimensional embedding space. Using this more abstract representation, one can blend motions or remove noises from corrupted motions. In [20, 21], the authors further make use of the power of the learned motion manifold and decoder to synthesize motions with constraints. Another important application of motion embedding is motion style generation [8], in which the embedding code can be tuned to matched the desired style. Although modeling a motion sequence with auto-encoders is straightforward, how it can model the dynamics of human motion is not clear. In [28], the authors model a motion sequence as a trajectory in pose embedding manifold, then use a bidirectional-RNN to encode the trajectory to model the dynamics of the trajectory, which can improve motion classification results. Moreover, they designed a graph-like network to better represent the human body components.

The existing methods focus on the embedding of the poses and dynamics. However, they do not explicitly model the distributions of these latent variables which governs the stochasticity of the dynamics. In this paper, we go a lever deeper and learn the latent variable distributions for the embedded poses and dynamics.

### 2.2 Deterministic human motion prediction and synthesis

In the effort of modelling motion dynamics, many methods employ deterministic transitions [2, 4, 9, 13, 17, 19, 24, 29, 36], especially in human motion prediction or generation. They either focus on short term dynamics modeling or spatial-temporal information of the overall dynamics. In [42], the authors propose a training technique for RNN to generate very long human motions. Although this technique solves the problem of the freezing phenomena of RNN, their model is deterministic, which makes the training difficult: given a past state, if multiple possible future motions are present in the data, the network will average them, which is a common problem in many human motion prediction methods.

One solution to this problem is to introduce control signals [14, 37]. They design several networks and make the character to follow a given trajectory in real-time. In [33], the control signal becomes the 3d human pose predicted by neural nets as a reference for an agent to imitate. In [1], the authors co-embed the language and corresponding motions to a share manifold, ignoring the fact that language-to-motion is a one-to-many mapping. Even with a specific control signal, one can still expect that there are different motions corresponding to the same control signal, essentially indicating the multi-modality nature of human motion dynamics.

Different from the existing methods, our paper focuses on the explicit modelling of the multi-modality nature of motion transitions in human motions. Further, we also aim to learn the stochasticity in those transitions.

### 2.3 Stochastic human motion synthesis

In [40] the authors combine RNNs and Generative Adversarial Networks (GANs) to generate stochastic human motions. They use a mixture density layer to model the stochastic property, and use an adversarial discriminator to judge whether the generated motion is natural or not. In MoGLow [18], the authors for the first time use normalizing flow to directly model the next frame distribution.

One advantage of this method is that it can capture complex distributions without learning an apparent latent space. Given the same initial poses and the same control signals or constraints, the model still generate different motion sequences. Chen et al. [5] combine Dynamic motion primitive and variational Bayesian filter to model the human motion dynamics. They show that the latent representations are self-clustering after training. However, in the transitions, it needs the information of the whole sequence, which separates it from being a pure generative model.

Our method differs from existing approaches in its treatment in the relations between the past, current and future states of human motions. Unlike the aforementioned methods, we explicitly model the current state based on both the past and the future. Also, we further model their randomness in the latent space that captures the transition multi-modality.

## 2.4 Stochastic RNN model

Modelling the stochasticity in time-series data has been a long-standing problem, such as music, hand writing and human voice [11, 25, 38]. The VRNN [7] for the first time combines Variational AutoEncoder (VAE) [27] and recurrent neural networks for this purpose. Later in [23], the authors disentangle the latent variables of the observation and the dynamics, with the observation latent being used to recover the full observation information, and the dynamic latent capturing the dynamics.

A key modelling choice in stochastic RNN models is the relations between the past, current and future. In early work, the posterior of current state is inferred from the past information, which makes it lack the ability to foresee the future. In [3, 10, 35], the authors show that the performance can be improved by incorporating the future state with a backward RNN in the inference stage. In [12], the authors design a model that can go beyond step-by-step modelling, and predict multiple steps up to a given horizon. Similar effort is also made in reinforcement learning, where the reward function takes the discounted future reward into consideration [12, 15]. In [16], the authors went further and designed a model that can predict multiple future scenarios, then choose the one with highest predicted reward from all the possibilities.

We observed that human motions follow a similar philosophy: the current state is a result of the past motion but also a particular choice for a certain planned future. Our research is inspired by Stochastic RNN models but focuses on human motion transition stochasticity.

## 3 METHOD OVERVIEW

Our method takes a homogeneous series of human pose representations as input. This representation contains the 3D joint coordinates relative to the root, and the root translation velocity over ground plane and the rotation velocity around the y-axis. We propose the Dynamic Future Net to model the motion dynamics as a future-guided transition and generate random natural human motions that transit between different actions.

As illustrated in Fig. 2, DFN is composed of three modules, a pose encoder, a pose trajectory encoder and a stochastic latent RNN. The pose encoder (Section 5.1) maps the high dimensional human pose to a latent space while the pose trajectory encoder (Section

5.2) embeds the trajectory in the latent pose space into a code. Such compact representations of pose sequences can facilitate the learning process [28]. As a key module, the stochastic latent RNN (Section 5.3) deploys a stochastic latent state and a deterministic latent state to learn two latent distributions for the pose-embedding and the future trajectory embedding. Such explicit learning of two different latent distributions on the one hand forces the model to learn strong temporal correlation and on the other hand generates motions with varied and natural transition. During inference we combine the past, current and future state to infer the current latent state distribution, and we combine the past and future to infer the future latent state distribution. In the generation stage, unlike existing methods [10] where the current state is generated from the past state only, we first generate the future state and combine it with the past state to generate the current latent state prior, from which we sample the current latent state then decode it to the pose-embedding and recover the current pose and velocity. We regard this process as a self-driving motion generation process guided by the envisioned dynamic future. In this way, the model can learn and generate rich and varied natural motions.

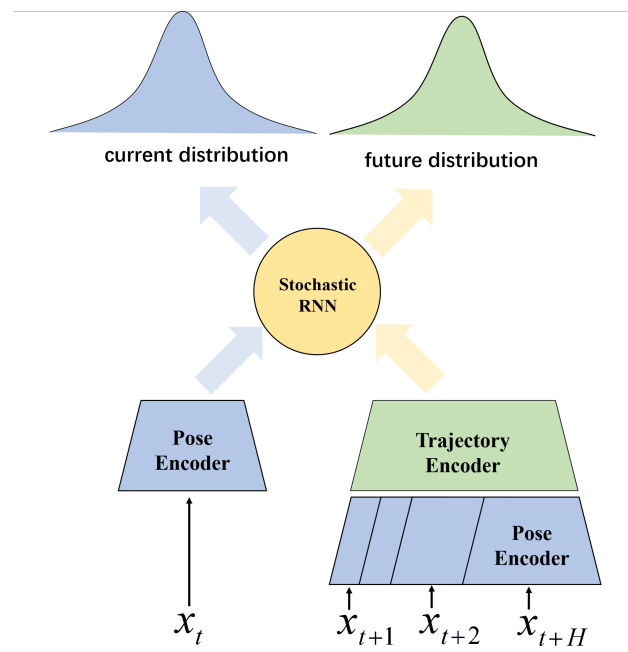
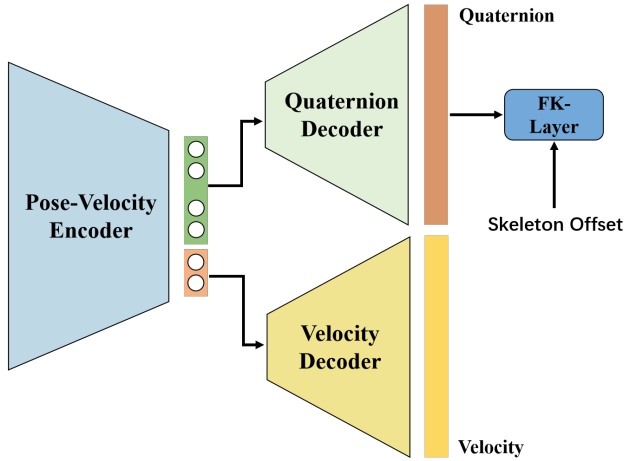


Figure 2: Overview of the proposed Dynamic Future Network. In the learning process, the network take human motion sequence as input and predict the long term distribution and the short term(next frame) distribution.

## 4 DATA PREPARATION

We train our model on the CMU Human motion capture dataset. As the skeletons in the origin dataset are different, we first retarget all motion to a chosen skeleton as in [21]. The skeleton contains 24 joints, we first extract the X and Z global coordinate of the root, and rotate the human pose to the Y-axis direction as in [22], the

global position and angle of human pose can be recovered from the X-Z velocity and the rotation velocity around the Y axis. Finally the original human pose vector contains 76 degrees of freedom, 72 for 3D joint positions, 4 for the global translation and rotation velocity.



**Figure 3: The pose-velocity auto-encoder network. The input of the encoder is the 76 dimensional pose-velocity vector. The encoder outputs two code, one for pose, the other for velocity. The pose code and velocity code is fed into a quaternion decoder and velocity decoder separately. The 3D joint positions are recovered from quaternions by the Forward Kinematics (FK) layer.**

## 5 METHODOLOGY

Formally, we start by describing a motion as a homogeneous time series:  $\{X_0, \dots, X_T\}$ , where  $X_t$  is the motion frame at time  $t$  and contains the joint positions and global velocities. Starting with a joint distribution  $P(X_{<t}, X_t, X_{t+1:t+H})$ , we model the influence of the past frames  $X_{<t}$  and future frames  $X_{t+1:t+H}$  on the current frame  $X_t$  by transition probabilistic distributions  $P(X_t|X_{t+1:t+H})$  and  $P(X_t|X_{<t})$ , where  $H$  is the duration of a short-horizon future. The key reason of such a modelling choice is based two observations: the current frame is a result of the past motion and therefore conditioned on it, captured by  $P(X_t|X_{<t})$ . Meanwhile, the current frame is also a choice made for certain planned future, e.g. needing to stop swing the legs aiming for a transition from walk to standing, captured by  $P(X_t|X_{t+1:t+H})$ . In addition, since the past motion will also limit the possibilities of the future motion, there is also a impact of the past on the future,  $P(X_{t+1:t+H}|X_{<t})$ . Overall, the joint probability:

$$P(X_{<t}, X_t, X_{t+1:t+H}) \propto P(X_t|X_{t+1:t+H})P(X_{t+1:t+H}, X_t|X_{<t}) \quad (1)$$

Not that the two probabilities on the right side play different roles.  $P(X_{t+1:t+H}, X_t|X_{<t})$  is the probability of unrolling from the past to the future. Given a known past, this is a joint probability of both the current and the future, containing all the possible transitions. On top of it,  $P(X_t|X_{t+1:t+H})$  dictates that if the future is also known, then the current can be inferred. This *explicit* modelling of the transition probabilistic distributions between the past, current and

future helps capturing the transition stochasticity, which facilitates diversified motion generation as shown in the experiments.

Learning the transitional probabilities in the data space, however, is difficult due to the curse of dimensionality. We therefore project motions to a latent space, which involves embedding the frames as well as the dynamics. We then learn the transition distributions in the latent space. During inference, we then recover motions from sampled states in the latent space to the original data space. DFN is naturally divided into three components: Spatial (frame) embedding, dynamics embedding and dynamics modelling.

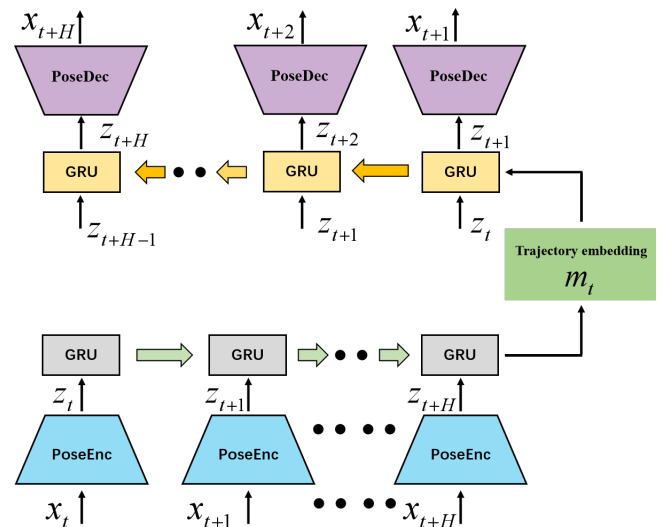
### 5.1 Spatial Embedding

We use an auto-encoder for frame embedding,  $z_t = \text{PoseEnc}(X_t)$  and  $\hat{X}_t = \text{PoseDec}(z_t)$ , shown in Figure 3. *PoseEnc* is multi-layer perceptron network to project the data into the latent space. Then we separate the latent feature into two components to represent the pose code and the global velocity code. *PoseDec* contains two components, the quaternion decoder and the velocity decoder. The quaternion decoder takes the pose latent feature as input and outputs joint angles (represented by quaternions), and the velocity decoder takes the latent velocity feature as input and outputs the velocity. The quaternion decoder essentially is a differential Inverse Kinematics module. As stated in [32], using joint rotations instead of joint positions maintains the bone lengths. After the reconstruction, we use a Forward Kinematics layer to compute the 3D joint positions. To train the auto-encoder, we use a Mean Squared Error loss function:

$$L_{sl} = \frac{1}{T} \sum_{t=0}^T \|X_t - \hat{X}_t\|_2^2 \quad (2)$$

where  $T$  is the number of frames in a motion.

### 5.2 Dynamics Embedding



**Figure 4: A seq-to-seq network for trajectory embedding.**

After learning the pose latent space, we project the dynamics as trajectories in this space using a Recurrent Neural Network shown

in Figure 4. We employ a sequence-to-sequence architecture as it forces the model to learn long-term dynamics. The RNN consists of Gated Recurrent Unit (GRU) [6], encodes a sequence of encoded frames  $\{z_t, \dots, z_{t+H}\}$  into a latent representation  $m_t$ , and then unroll to reconstruct the same sequence  $\{z'_{t+1}, \dots, z'_{t+H}\}$  from  $m_t$  given  $z_t$ . To  $z_t = PoseEnc(X_t)$ ,  $m_t$  is a future summary over multiple frames. We use the following loss function:

$$L_{tl} = L_{rec} + L_{smooth} \quad (3)$$

$$\text{where } L_{rec} = \frac{1}{T} \sum_{t=0}^T \|z_t - z'_t\|_2^2$$

$$L_{smooth} = \frac{1}{T} \sum_{t=0}^T \|V_t - \hat{V}_t\|_2^2$$

where  $T$  is the frame number of a motion,  $V_t$  and  $\hat{V}_t$  are the original and reconstructed joint velocities. To facilitate training, we use Eq. 2 to pre-train the posture auto-encoder and fix its weights when training the RNN module.

### 5.3 Dynamics modelling

**Generative Model.** After embedding the poses and dynamics into a latent space, we now explain the dynamics modelling, which is the key technical contribution of this paper. We propose a new dynamics model that captures the joint distribution  $P(X_{<t}, X_t, X_{t+1:t+H})$ . Rather than directly learning the distribution in the data space, we aim to learn the latent joint distribution  $P(z_{<t}, z_t, z_{t+1:t+H})$ , where we abstract the past, current and future features separately. First, given the Markov property, we assume that all past information  $z_{<t}$  is encoded into  $h_t$  which is a deterministic (known) past state. Next, we assume that the of future information  $z_{t+1:t+H}$  can be summarized into a future state  $f_t$ , and  $f_t$  is drawn from a distribution over all possible future states conditioned on  $h_t$ . Last, we also assume that there is a current state  $s_t$  which captures the current information  $z_t$  and  $s_t$  is drawn from a distribution of all possible current states. Then we can therefore assume:

$$\begin{aligned} P(z_{<t}, z_t, z_{t+1:t+H}) &= P(z_{t+1:t+H}, z_t | z_{<t}) P(z_{<t}) \\ &\propto P(z_t | z_{t+1:t+H}) P(z_t | z_{<t}) P(z_{t+1:t+H} | z_{<t}) \\ &= P(s_t | f_t) P(s_t | h_t) P(f_t | h_t) \\ &= P(s_t | f_t, h_t) P(f_t | h_t) \end{aligned} \quad (4)$$

where we directly use  $s_t$ ,  $h_t$  and  $f_t$  to replace the corresponding  $z$  variables by assuming mappings between them which will be explained later. Different from existing methods [10, 35], our direct conditioning of the current state on the future and past state  $P(s_t | f_t, h_t)$ , and the future state on the past state  $P(f_t | h_t)$  allows us great flexibility in modelling the stochasticity in transitions. The generation model is shown in Fig.5 a.

**Future Feature Prior.** Given  $h_t$ , we first predict the future state, via  $P(f_t | h_t)$ . Here, we assume a diagonal multivariate Gaussian prior over  $f_t$  [5, 14, 26]

$$\begin{aligned} p(f_t | h_t) &= \text{Gaussian}(f_t; \mu_t^f, \sigma_t^f), \\ \text{where } [\mu_t^f, \sigma_t^f] &= g_f^p(h_t) \end{aligned} \quad (5)$$

where  $\mu_t^f$  and  $\sigma_t^f$  are the mean and covariance.  $g_f^p$  is a three-layer MLP with hidden dimension 256 and LeakyReLU activation. It contains all the possible future state given the past. It can represent a goal or a driving signal for the generative process. Also it forces the model to learn rich motion transitions and long term correlations, overcoming the freezing problem of traditional RNN [42].

**Current Feature Prior.** Next, we explain  $P(s_t | f_t, h_t)$ . Although  $h_t$  is a known (deterministic) past,  $f_t$  is random. We therefore first sample a specific future state  $f_t$ , then decode it to an unrolled future summary  $m_t$ , and finally condition the current state  $s_t$  on  $h_t$  and  $m_t$ . We therefore have:

$$P(s_t | f_t, h_t) \propto P(s_t | m_t, h_t) P(m_t | f_t, h_t) \quad (6)$$

$$P(s_t | m_t, h_t) = \text{Gaussian}(s_t; \mu_t^s, \sigma_t^s) \quad (7)$$

$$[\mu_t^s, \sigma_t^s] = g_s^p(h_t, m_t^p)$$

where  $P(m_t | f_t, h_t)$  is parameterized by a four layer MLP with hidden dimension 128 with LeakyReLU activation.  $\mu_t^s$  and  $\sigma_t^s$  are the mean and covariance.  $g_s^p$  is a two layer MLP with hidden dimension 128. After being able to sample the current state  $s_t$ , we can compute the current feature  $z_t$  via  $z_t = MLP(s_t, h_t)$ , where the MLP has three hidden layer with 128 dimensions and LeakyReLU activation.

Finally, given the current and future state, the past state is updated as follows (Fig.5 c):

$$h_{t+1} = \text{GRU}(h_t, s_t, f_t) \quad (8)$$

where the GRU has two stacked layer and hidden state of dimension 128. Now the generation model in Fig.5 a is complete.

Different from existing methods [7, 10] where the prior of current state is a function of past state  $h_t$  only, and where the future state is shared with the current state, we let the model learn two different distributions for current and future states. The prior of current state is also a function of the future state, which will force the model to make use of the future information.

### 5.4 Inference

In the generation model in Fig.5 a, the key variables to be inferred are  $s_t$  and  $f_t$ , shown in Fig.5 b. The posterior of the future state  $f_t$  is dependent on past state  $h_t$  and its unrolled future summary  $m_t$ . The posterior of the current state  $s_t$  is dependant on the feature  $z_t$ , the past state  $h_t$  and the future summary  $m_t$ . We first factorize the dynamics as follow:

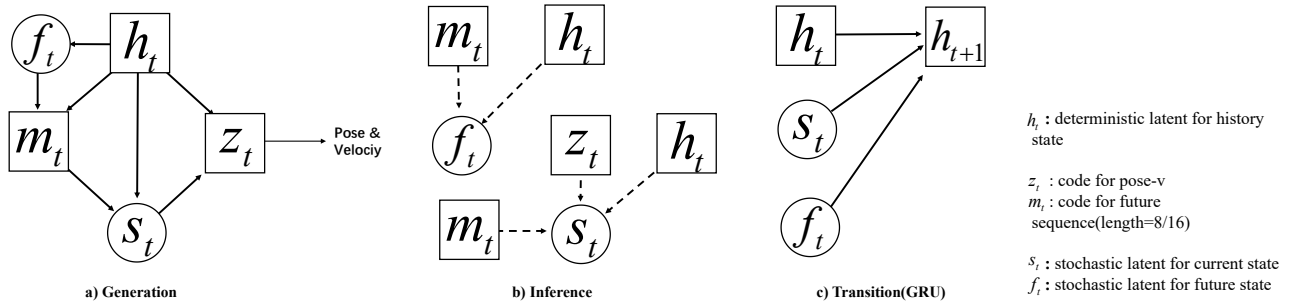
$$q(s_{\leq T} | z_{\leq T}) \approx \prod_{t=1}^T q(s_t | z_{\leq t-1}, z_t, z_{\leq t+H}) = \prod_{t=1}^T q(s_t | h_t, z_t, m_t) \quad (9)$$

where  $T$  is the total length of motion. Here we approximate the  $q(s_t | z_{\leq T})$  with  $q(s_t | z_{\leq t+H})$ , as the correlations between  $s_t$  and the far future is likely to be small, so we only consider up to  $t+H$ . Then for each time step, we use a MLP to parameterize the posterior:

$$q(s_t | h_t, z_t, m_t) = \text{Gaussian}(\mu_t^s, \sigma_t^s), [\mu_t^s, \sigma_t^s] = MLP(h_t, z_t, m_t) \quad (10)$$

where the MLP has two hidden layers with 32 dimensions and LeakyReLU activation. For the future state we approximate its posterior as follow:

$$q(f_t | h_t, m_t) = \text{Gaussian}(\mu_t^f, \sigma_t^f), [\mu_t^f, \sigma_t^f] = MLP(h_t, m_t) \quad (11)$$



**Figure 5: The stochastic latent RNN. a) Generation Model. The current pose embedding feature  $z_t$  depends on the current and past latent state  $s_t$  and  $h_t$ .  $s_t$  depends on the past state  $h_t$  and the future summary  $m_t$  which depends on the future latent state  $f_t$  and past state  $h_t$  b) Inference on  $s_t$  and  $f_t$ . c) Transition of  $h_t$**

where the MLP has two hidden layers with 512 dimensions and LeakyReLU activation.

### 5.5 Temporal difference loss

Besides inferring  $s_t$  and  $f_t$ , we also constrain the dynamics of  $s$ . We assume a relation between two states  $s_{t_1}$  and  $s_{t_2}$  at  $t_1$  and  $t_2$  where  $t_1 < t_2$ , similar to [12]:

$$p(s_{t_1}, s_{t_2}, t_1, t_2) \propto p(s_{t_2} | s_{t_1}, h_{t_1}, h_{t_2}) \quad (12)$$

where we parameterize the posterior:

$$q(s_{t_1} | s_{t_2}, h_{t_1}, h_{t_2}) = \text{Gaussian}(s_{t_1}; \mu_{s_{t_1}}, \sigma_{s_{t_1}}) \quad (13)$$

where  $\mu_{s_{t_1}}, \sigma_{s_{t_1}} = \text{MLP}(s_{t_2}, h_{t_1}, h_{t_2})$

where  $\mu_{s_{t_1}}$  and  $\sigma_{s_{t_1}}$  are the mean and covariance. Here the MLP has two hidden layers with 32 dimensions and LeakyReLU activation. This way we can sample  $s_{t_1}^{\text{posterior}}$  during inference. Meanwhile, we hope to reconstruct  $s_{t_2}$  with the time difference  $\delta t = |t_1 - t_2|$

$$s_{t_2}^{\text{rec}} = \text{MLP}(s_{t_1}^{\text{posterior}}, \delta t) \quad (14)$$

where the MLP for skip prediction has three hidden layers each with 32 dimensions and LeakyReLU activation.

### 5.6 Learning

Finally, we compose all terms for the loss function:

$$\max_{\Phi} L_d + L_T \quad (15)$$

where  $\Phi$  is the set of all learnable parameters in our networks and

$$L_d = \sum_{t=1}^T [-KL(q(s_t | h_t, z_t, m_t) || p(s_t | h_t, m_t)) - KL(q(f_t | h_t, m_t) || p(f_t | h_t)) + \log(p(z_t | s_t, h_t)) + \log(p(m_t | f_t, h_t))]$$

$$L_T = \sum_{t_1, t_2} [-KL(q(s_{t_1} | s_{t_2}, h_{t_1}, h_{t_2}) || p(s_{t_1})) + p(s_{t_2} | s_{t_1}, t_1, t_2)]$$

$KL$  is the Kullback–Leibler divergence.

After training the pose auto-encoder and the sequence auto-encoder (Section 5.1-5.2), we freeze their parameters and train the dynamics model (Section 5.3).

## 6 EXPERIMENT AND RESULTS

For all our experiments, we use CMU MoCap database<sup>1</sup>. CMU dataset is a high-quality dataset acquired using optical motion capture systems, containing 2605 trials in 6 categories and 23 sub-categories. Its high-quality serves our purposes well as it provides good data ‘seeds’ for motion generation. Also, the tremendous effort that went into capturing the data shows the need for tools such as DFN for data augmentation. To carefully evaluate DFN, we select with different motion classes with different features and dynamics, shown in Tab.1, to show that DFN can generate new high-quality motions with arbitrary lengths using different motion prefixes. Next, we evaluate DFN on data with single and mixed motion classes to see its ability to learn the different transition stochasticity on data with a single type of dynamics and mixed types of dynamics. Last, we push the limit of DFN by reducing the training data, to show that DFN can make use of a small amount of data to generate high-quality and diversified data, which is crucial for data augmentation. More example can be found in the supplementary video.

### 6.1 Open-loop Motion Generation

We first show open-loop motion generation, where we do not moderate accumulative errors. We use a 8 to 20-frame motion prefix to start motion generation to get 900 frames (dfn\_run2box\_2char and dfn\_boxing\_3char in the video). The motion stability indicates that DFN does not suffer from the problem of cumulative error that is common in time-series generation [39]. Given the same prefix, the diversity is shown in their transitions between different postures (short-term) and different actions (long-term).

### 6.2 Dynamics Multi-modality

We investigate how well DFN can capture different transition stochasticity in different motions, using several types of motions with different properties (shown in Tab.1). We first train DFN on them separately then jointly. The results can be found in dfn\_walk\_top, dfn\_walk\_close, walking1-walking2, running1-running3, dancing1

<sup>1</sup><http://mocap.cs.cmu.edu/>

Motion	Cyclic	Main Body Part	Rhythmic	Dynamics
Walking	Yes	Lower	No	Low
Boxing	No	Upper	No	High
Dancing	No	Full	Yes	High
Running	Yes	Lower	No	High

**Table 1: Motion types and their features.**

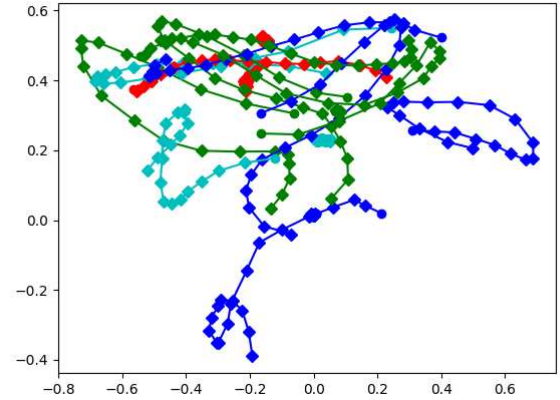
and boxing1 in the video. We observe that DFN can learn the transition stochasticity well when trained on single type of motions. The diversity can be found in short-term and long-term transitions, which are two-levels of multi-modality captured well by DFN. In walking (dfn\_walk\_top and dfn\_walk\_close in video), the short-term stochasticity is shown in within-cycle motion randomness, which enriches the walking style. The long-term stochasticity is shown when a turning is generated. The action-level transition has also been captured and generated. Similar observations are also found in other motions.

When trained on mixed data (combining all motion in Tab.1), DFN learns higher-level action transitions between different motion classes. We can see examples (action\_transition in video and frame-level image in supplementary file) that transit from dancing to running, from boxing to dancing or from slow walk to running, showing the modelling capacity at two levels. Within a single action, diversified styles are learned well. Between different actions, transitions are learned well too. This demonstrates the benefits of modelling randomness explicitly between the past, current and future state, which would otherwise make it hard to capture the multi-modality and lead it to average over all types of dynamics, resulting in meaningless mean poses and motions.

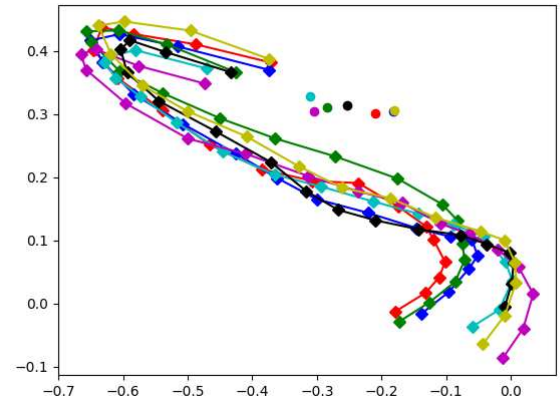
### 6.3 Diversified Generation

Although visually it is clear that DFN can generate diversified motions, we also numerically show the diversity especially when the duration of generation becomes long. First, we randomly select training data of different classes, and show their latent feature trajectory in Fig. 6, with embedding dimension of 16. We then use a PCA model to project the embedding features to 2D. The trajectories are continuous and smooth without extra constraints on the auto-encoder. It shows that the motion dynamics are captured well in the latent space, which is critical in motion generation. Next, we show a group of generated motions in Fig. 7. Even with the same motion prefix, motions start to diversify from the beginning, which is a distinct property lacking in deterministic generators in most action-prediction models such as [30], and our sequence is in 3d which is more difficult than 2d [41].

**Distributional Shift in Time.** The motion diversity increases in time. To show that there is a distributional shift of poses, we randomly pick an initial sequence from training data, then randomly generate 4096 sequences each with 128 frames. We visualize the current latent state  $s$  and latent feature  $z$  at  $t = 32$  and  $t = 128$  in Fig. 8. Note that the distributions of  $s$  and  $z$  capture the stochasticity at two different levels, one at the stochastic state level and the other at the latent feature level. The red dots represent them at  $t = 32$  and the yellow dots at  $t = 128$ .



**Figure 6: Randomly selected training motions in the latent space. Color indicate different motion class. Smooth trajectories are universally obtained by embedding.**

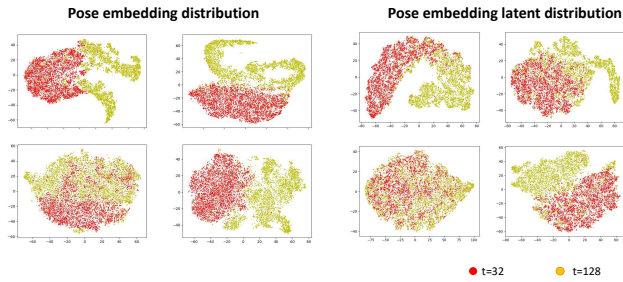


**Figure 7: Pose embedding trajectory of random generated sequences given same initialization with 20 frame. The circle mark the first frame. We see that as time goes, these trajectory depart from each other.**

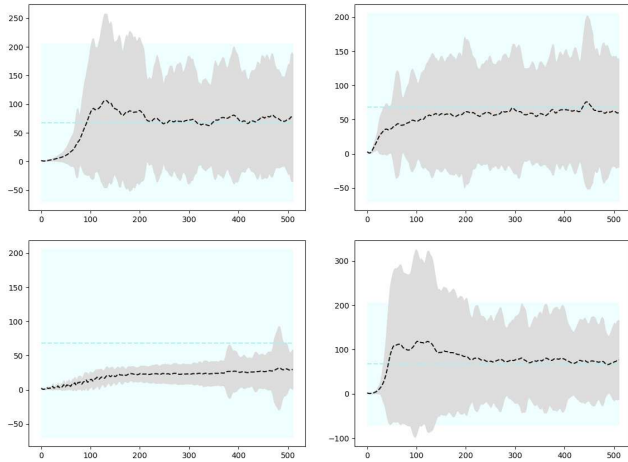
For both  $z$  and  $s$ , the red ( $t = 32$ ) are concentrated more, showing that the difference between the 4096 generated motions are still somewhat similar in the early stage. However, the yellow ( $t = 128$ ) show that the generated motions start to diversify later. Not only they shift out of the original red region, indicating that they are now in different pose regions, they also start to diverge more, shown by different modes in yellow areas, meaning they have diverged into several different pose regions.

**Distribution Matching in Time.** Another way to test the diversity of generated motions is to see their statistical similarity to the training motions. Since the motion prefix is from one particular motion, the more similar the generated motions are to the whole





**Figure 8: Four groups of motions generated from four different motion prefixes, each group with 4096 motions, and their  $z$  (Left) and  $s$  (Right) at  $t = 32$  and  $t = 128$ . We can see that the earlier distributions are more concentrated and diverge fast as time passes.**



**Figure 9: Four groups of motions generated from four different motion prefixes, each group with 4096 motions. The  $x$  axis represents the time dimension, and the  $y$  axis represent the mean distance to average pose at each time step. The band represents the variations.**

training dataset, the more diverse they are, because the generated motions have leave the original motion region where the motion prefix is.

We employ the mean-distance distribution as a measure, as in [41]. For each time step, we calculate the mean pose of all generated motions, then calculate the Euclidean distances between the mean pose and all other poses at that time step. We then plot the mean distance and variance in Figure 9. The blue background indicate the mean and variance of mean-distance distribution of the training dataset. It shows that as time goes, the mean-distance distribution of generated poses gradually matches that of the training data. This further shows the generation diversity.

## 6.4 Generation on Limited Training Data

DFN aims to solve the problem of data scarcity, so it should only require as little data as possible for generation. We therefore push DFN to its limit by reducing the training data, to see the minimal amounts of data needed. To investigate each individual type of motions, we train DFN on walking, running, and boxing data separately. We start from full training data where the longest sequence lasts for around 10 minutes, and gradually reduce the duration by sampling until the quality of the generated motions start to deteriorate. Although DFN responds to reduced training data slightly differently on different motions, we finally able to reduce the training data to a tiny amount, with the longest sequence being only 15 seconds (12 second for walking, 15 second for boxing and 7 second for running). DFN can still generate stable motions even when trained on merely a 7-second long motion. (The result can be seen in `reduced_data` in video) The impact of reducing the training data is mainly on the diversity of the motion. (However we can see in supplementary video that the generated boxing motion still has a certain of diversity). Less training data contains fewer transition diversities (both short-term and long-term). The generated motions therefore are less diverse. This is understandable as DFN cannot deviate too much from the original data distribution to ensure the motion quality.

## 6.5 Comparison

To our best knowledge, the only similar paper to ours is [40] which also focuses on diversified motion generation. However, the biggest difference is that DFN explicitly models the influence of the future on the current. This enables DFN to explicitly model the transition randomness at different stages and levels. This is the key reason why DFN can be trained well on multiple types of motions, separately and jointly, which has not been shown in [40]. However, a direct numerical comparison is difficult due to the lack of widely accepted metrics for diversified motion generation. In addition, the method in [40] uses heavy post-processing while DFN does not.

## 7 CONCLUSION AND DISCUSSIONS

In this paper, we propose a new generative model, DFN, for diversified human motion generation. DFN can generate motions with arbitrary lengths. It successfully captures the transition stochasticity in short and long term, and capable of learning the multi-modal randomness in different motions. The training data needed is small. We have conducted extensive evaluation to show DFN's robustness, versatility and diversity in motion generation.

There are two main limitations in our method. There is no control signal, and sometimes it can overly smooth high-frequency motions. We will address them in the future. Our explicit modelling of the future makes it convenient to introduce desired future as control signals; while replacing some of the Gaussian components with multi-modal priors might mitigate the over-smoothing issue.

## REFERENCES

- [1] Chaitanya Ahuja and Louis-Philippe Morency. 2019. Language2Pose: Natural Language Grounded Pose Forecasting. In *2019 International Conference on 3D Vision (3DV)*. IEEE, 719–728.
- [2] Federico Bartoli, Giuseppe Lisanti, Lamberto Ballan, and Alberto Del Bimbo. 2018. Context-aware trajectory prediction. In *2018 24th International Conference on*

- 929 *Pattern Recognition (ICPR)*. IEEE, 1941–1946.
- 930 [3] Justin Bayer and Christian Osendorfer. 2014. Learning Stochastic Recurrent  
931 Networks. *stat 1050* (2014), 27.
- 932 [4] Judith Butepage, Michael J Black, Danica Kragic, and Hedvig Kjellstrom. 2017.  
933 Deep representation learning for human motion prediction and classification.  
934 In *Proceedings of the IEEE conference on computer vision and pattern recognition*.  
935 6158–6166.
- 936 [5] Nutan Chen, Maximilian Karl, and Patrick Van Der Smagt. 2016. Dynamic  
937 movement primitives in latent space of time-dependent variational autoencoders.  
938 In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*.  
939 IEEE, 629–636.
- 940 [6] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio.  
941 2014. On the Properties of Neural Machine Translation: Encoder–Decoder  
942 Approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and  
943 Structure in Statistical Translation*. Association for Computational Linguistics,  
944 Doha, Qatar, 103–111. <https://doi.org/10.3115/v1/W14-4012>
- 945 [7] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville,  
946 and Yoshua Bengio. 2015. A recurrent latent variable model for sequential data.  
947 In *Advances in neural information processing systems*. 2980–2988.
- 948 [8] Han Du, Erik Herrmann, Janis Sprenger, Noshaba Cheema, Somayeh Hosseini,  
949 Klaus Fischer, and Philipp Slusallek. 2019. Stylistic Locomotion Modeling with  
950 Conditional Variational Autoencoder. In *Eurographics (Short Papers)*. 9–12.
- 951 [9] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. 2015.  
952 Recurrent network models for human dynamics. In *Proceedings of the IEEE Inter-  
953 national Conference on Computer Vision*. 4346–4354.
- 954 [10] Anirudh Goyal Alias Parth Goyal, Alessandro Sordoni, Marc-Alexandre Côté,  
955 Nan Rosemary Ke, and Yoshua Bengio. 2017. Z-forcing: Training stochastic  
956 recurrent networks. In *Advances in neural information processing systems*. 6713–  
957 6723.
- 958 [11] Alex Graves. 2013. Generating Sequences With Recurrent Neural Networks.  
959 *CoRR abs/1308.0850* (2013). [http://dblp.uni-trier.de/db/journals/corr/corr1308.  
960 html#Graves13](http://dblp.uni-trier.de/db/journals/corr/corr1308.html#Graves13)
- 961 [12] Karol Gregor, George Papamakarios, Frederic Besse, Lars Buesing, and Theo-  
962 phane Weber. 2019. Temporal Difference Variational Auto-Encoder. In *Inter-  
963 national Conference on Learning Representations*. [https://openreview.net/forum?id=  
964 S1x4ghC9tQ](https://openreview.net/forum?id=S1x4ghC9tQ)
- 965 [13] Xiao Guo and Jongmoo Choi. 2019. Human Motion Prediction via Learning  
966 Local Structure Representations and Temporal Dependencies. In *Proceedings of  
967 the AAAI Conference on Artificial Intelligence*, Vol. 33. 2580–2587.
- 968 [14] Ikhsanul Habibie, Daniel Holden, Jonathan Schwarz, Joe Yearsley, Taku Komura,  
969 Jun Saito, Ikuo Kusajima, Xi Zhao, Myung-Geol Choi, Ruizhen Hu, et al. 2017. A  
970 Recurrent Variational Autoencoder for Human Motion Synthesis. In *BMVC*.
- 971 [15] Danijar Hafner, Timothy Lillicrap, Jimmy Ba, and Mohammad Norouzi. 2020.  
972 Dream to Control: Learning Behaviors by Latent Imagination. In *International  
973 Conference on Learning Representations*. [https://openreview.net/forum?id=  
974 S1IOTC4tDS](https://openreview.net/forum?id=S1IOTC4tDS)
- 975 [16] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak  
976 Lee, and James Davidson. 2019. Learning Latent Dynamics for Planning from  
977 Pixels. In *International Conference on Machine Learning*. 2555–2565.
- 978 [17] Félix G Harvey and Christopher Pal. 2018. Recurrent transition networks for  
979 character locomotion. In *SIGGRAPH Asia 2018 Technical Briefs*. 1–4.
- 980 [18] Gustav Eje Henter, Simon Alexanderson, and Jonas Beskow. 2019. Moglow:  
981 Probabilistic and controllable motion synthesis using normalising flows. *arXiv  
982 preprint arXiv:1905.06598* (2019).
- 983 [19] Alejandro Hernandez, Jurgen Gall, and Francesc Moreno-Noguer. 2019. Human  
984 motion prediction via spatio-temporal inpainting. In *Proceedings of the IEEE  
985 International Conference on Computer Vision*. 7134–7143.
- 986 [20] Daniel Holden, Taku Komura, and Jun Saito. 2017. Phase-functioned neural  
987 networks for character control. *ACM Transactions on Graphics (TOG)* 36, 4 (2017),  
988 1–13.
- 989 [21] Daniel Holden, Jun Saito, and Taku Komura. 2016. A deep learning framework  
990 for character motion synthesis and editing. *ACM Transactions on Graphics (TOG)*  
991 35, 4 (2016), 1–11.
- 992 [22] Daniel Holden, Jun Saito, Taku Komura, and Thomas Joyce. 2015. Learning  
993 motion manifolds with convolutional autoencoders. In *SIGGRAPH Asia 2015  
994 Technical Briefs*. 1–4.
- 995 [23] Wei-Ning Hsu, Yu Zhang, and James Glass. 2017. Unsupervised learning of  
996 disentangled and interpretable representations from sequential data. In *Advances  
997 in neural information processing systems*. 1878–1889.
- 998 [24] Junfeng Hu, Zhencheng Fan, Jun Liao, and Li Liu. 2019. Predicting Long-Term  
999 Skeletal Motions by a Spatio-Temporal Hierarchical Recurrent Network. *arXiv  
1000 preprint arXiv:1911.02404* (2019).
- 1001 [25] Cheng-Zhi Anna Huang, Ashish Vaswani, Jakob Uszkoreit, Ian Simon, Curtis  
1002 Hawthorne, Noam Shazeer, Andrew M. Dai, Matthew D. Hoffman, Monica Din-  
1003 culescu, and Douglas Eck. 2019. Music Transformer. In *International Conference  
1004 on Learning Representations*. <https://openreview.net/forum?id=rJe4ShAcF7>
- 1005 [26] Maximilian Karl, Maximilian Soelch, Justin Bayer, and Patrick van der Smagt.  
1006 2017. DEEP VARIATIONAL BAYES FILTERS: UNSUPERVISED LEARNING OF  
1007 STATE SPACE MODELS FROM RAW DATA. *stat 1050* (2017), 3.
- 1008 [27] Diederik P Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes.  
1009 *stat 1050* (2014), 1.
- 1010 [28] Jogendra Nath Kundu, Maharshi Gor, Phani Krishna Uppala, and Venkatesh Babu  
1011 Radhakrishnan. 2019. Unsupervised feature learning of human actions as trajec-  
1012 tories in pose embedding manifold. In *2019 IEEE Winter Conference on Applications  
1013 of Computer Vision (WACV)*. IEEE, 1459–1467.
- 1014 [29] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. 2019. Learning  
1015 trajectory dependencies for human motion prediction. In *Proceedings of the IEEE  
1016 International Conference on Computer Vision*. 9489–9497.
- 1017 [30] Julieta Martinez, Michael J Black, and Javier Romero. 2017. On human motion  
1018 prediction using recurrent neural networks. In *Proceedings of the IEEE Conference  
1019 on Computer Vision and Pattern Recognition*. 2891–2900.
- 1020 [31] Jianyuan Min and Jinxiang Chai. 2012. Motion Graphs++: A Compact Generative  
1021 Model for Semantic Motion Analysis and Synthesis. *ACM Trans. Graph.* 31, 6,  
1022 Article 153 (Nov. 2012), 12 pages. <https://doi.org/10.1145/2366145.2366172>
- 1023 [32] Dario Pavllo, David Grangier, and Michael Auli. 2018. QuaterNet: A Quaternion-  
1024 based Recurrent Model for Human Motion.
- 1025 [33] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. 2018.  
1026 Deepmimic: Example-guided deep reinforcement learning of physics-based char-  
1027 acter skills. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–14.
- 1028 [34] A. Safonova, Jessica Hodgins, and Nancy Pollard. 2004. Synthesizing physically  
1029 realistic human motion in low dimensional. *ACM Transactions on Graphics - TOG*  
1030 (01 2004).
- 1031 [35] Dmitriy Serdyuk, Nan Rosemary Ke, Alessandro Sordoni, Adam Trischler, Chris  
1032 Pal, and Yoshua Bengio. 2018. Twin Networks: Matching the Future for Sequence  
1033 Generation. In *International Conference on Learning Representations*. [https://  
1034 openreview.net/forum?id=BydLzGb0Z](https://openreview.net/forum?id=BydLzGb0Z)
- 1035 [36] Xiangbo Shu, Liyan Zhang, Guo-Jun Qi, Wei Liu, and Jinhui Tang. 2019. Spa-  
1036 tiotemporal Co-attention Recurrent Neural Networks for Human-Skeleton Motion  
1037 Prediction.
- 1038 [37] Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. 2019. Neural state  
1039 machine for character-scene interactions. *ACM Transactions on Graphics (TOG)*  
1040 38, 6 (2019), 1–14.
- 1041 [38] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol  
1042 Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu.  
1043 [n.d.]. WaveNet: A Generative Model for Raw Audio. In *9th ISCA Speech Synthesis  
1044 Workshop*. 125–125.
- 1045 [39] He Wang, Edmond SL Ho, Hubert PH Shum, and Zhanxing Zhu. 2019. Spatio-  
1046 temporal Manifold Learning for Human Motions via Long-horizon Modeling.  
1047 *IEEE transactions on visualization and computer graphics* (2019).
- 1048 [40] Zhiyong Wang, Jinxiang Chai, and Shihong Xia. 2019. Combining recurrent  
1049 neural networks and adversarial training for human motion synthesis and control.  
1050 *IEEE transactions on visualization and computer graphics* (2019).
- 1051 [41] Zhenyi Wang, Ping Yu, Yang Zhao, Ruiyi Zhang, Yufan Zhou, Junsong Yuan,  
1052 and Changyou Chen. 2019. Learning Diverse Stochastic Human-Action Gener-  
1053 ators by Learning Smooth Latent Transitions. *CoRR abs/1912.10150* (2019).  
1054 [arXiv:1912.10150 http://arxiv.org/abs/1912.10150](http://arxiv.org/abs/1912.10150)
- 1055 [42] Yi Zhou, Zimo Li, Shuangjiu Xiao, Chong He, Zeng Huang, and Hao Li. 2018.  
1056 Auto-Conditioned Recurrent Networks for Extended Complex Human Motion  
1057 Synthesis. In *International Conference on Learning Representations*. [https://  
1058 openreview.net/forum?id=r11Q2SIRW](https://openreview.net/forum?id=r11Q2SIRW)