

# Learning Human Motion Models for Long-term Predictions

Partha Ghosh    Jie Song    Emre Aksan    Otmar Hilliges  
Advanced Interactive Technologies, ETH Zurich

pghosh@student.ethz.ch    {jsong, eaksan, otmar.hilliges}@inf.ethz.ch

## Abstract

*We propose a new architecture for the learning of predictive spatio-temporal motion models from data alone. Our approach, dubbed the Dropout Autoencoder LSTM (DAE-LSTM), is capable of synthesizing natural looking motion sequences over long-time horizons<sup>1</sup> without catastrophic drift or motion degradation. The model consists of two components, a 3-layer recurrent neural network to model temporal aspects and a novel autoencoder that is trained to implicitly recover the spatial structure of the human skeleton via randomly removing information about joints during training. This Dropout Autoencoder (DAE) is then used to filter each predicted pose by a 3-layer LSTM network, reducing accumulation of correlated error and hence drift over time. Furthermore to alleviate insufficiency of commonly used quality metric, we propose a new evaluation protocol using action classifiers to assess the quality of synthetic motion sequences. The proposed protocol can be used to assess quality of generated sequences of arbitrary length. Finally, we evaluate our proposed method on two of the largest motion-capture datasets available and show that our model outperforms the state-of-the-art techniques on a variety of actions, including cyclic and acyclic motion, and that it can produce natural looking sequences over longer time horizons than previous methods.*

## 1. Introduction

Predicting human motion over a significant time horizon is a challenging problem with applications in a variety of domains. For example in human computer interaction, human detection and tracking, activity recognition, robotics and image based pose estimation it is important to model and predict the most probable sequence of human motions in order to react accordingly and in a timely manner. Despite the inherent stochasticity and context dependency of natural motion, human observers are remarkably good at predicting what is going to happen next, exploit-

<sup>1</sup>> 10s for periodic motions, e.g. walking, > 2s for aperiodic motion, e.g. eating

ing assumptions about continuity and regularity in natural motion. However, formulating this domain knowledge into strong predictive models has been proven to be difficult. Integrating spatio-temporal information into algorithmic frameworks for motion prediction is hence either done via simple approximations such as optical flow [10, 27] or via manually designed and activity specific spatio-temporal graphs [9, 19]. Given the learning capability of deep neural networks and recurrent architectures in particular, there lies enormous potential but also many challenges in learning statistical motion models directly from data that can generalize over a range of activities and over long time horizons.

Embracing this challenge we propose a new augmented recurrent neural network (RNN) architecture, dubbed Dropout Autoencoder LSTM (DAE-LSTM). Our model is capable of extracting both structural and temporal dependencies directly from the training data and does not require expert designed and task dependent spatio-temporal graphs for input as is the case in prior work [19]. Our work treats the two aspects of the task, namely the inherent constraints imposed by the skeletal configuration and the constraints imposed by temporal coherence explicitly. Using a feed forward network for pose filtering and an RNN for temporal filtering, reduces drift due accumulation of error over time. We demonstrate this in a number of side-by-side comparisons to the state-of-the-art.

Specifically, we leverage de-noising autoencoders to learn the spatial structure and dependencies between different joints of the human skeleton while an LSTM network models temporal aspects of the motion. Contrary to related work that uses autoencoders to project the input data into a lower-dimensional manifold [11, 19], our model directly operates in the joint angle domain of the human skeleton. Although we use an autoencoder-like architecture it does not bear real resemblance to encoding-decoding in the usual sense of latent representation learning. We simply use the autoencoder to de-noise skeletal poses at every time step, i.e. our auto encoder takes a pose as input and produces the filtered version of it in the same domain. During training we perturb the inputs with random noise, as is common practice in de-noising tasks, but additionally use dropout layers on

the inputs to randomly remove entire joint positions from the training samples. Therefore, to be able to accurately reconstruct entire poses the network has to leverage information about the spatial dependencies between adjacent joints to correctly infer positions of the missing joints. Hence this training regime forces the network to implicitly recover the spatial configuration of the skeleton.

The proposed model learns to predict the most likely pose at time  $t + 1$  given the history of poses up to time  $t$ . Putting this model into recurrence allows for synthesis of novel and realistic motion sequences. We experimentally demonstrate that separating pose reconstruction and temporal modeling improves performance over settings where the autoencoder is primarily used for representation learning. While the architecture is simple, it captures both the spatial and temporal components of the problem well and improves prediction accuracy compared to the state-of-the-art on two publicly available datasets.

In the domain of generative motion models, the lack of appropriate evaluation protocols to assess the quality and naturalness of the generated sequences is a commonly faced issue. The generated sequences need to be perceptually similar to the training data but clearly one does not simply want to memorize and replicate the training data. To better assess this generative nature of the task we furthermore contribute an evaluation protocol that quantifies how natural a generated sequence is over arbitrarily long time horizons. To assess naturalness we propose to train a separate classifier to predict action class labels. Intuitively the longer a sequence can be classified to belong to the same action category as the seed sequence the higher the quality of the prediction.

We evaluate the proposed model on the H3.6m dataset of Ionescu et al. [18] and the more recent dataset of Holden et al. [17] in a pose forecasting task. Our model outperforms the 3-layer LSTM baseline and two state-of-the-art models [11, 19] both in terms of short and long horizon predictions. Furthermore, we detail results from the proposed evaluation protocol and demonstrate that this can be used to analyze the performance of such generative tasks.

## 2. Related Work

Here we provide an overview of recent literature that deals with human motion modeling. This is one of the core problems in computer vision and machine intelligence and has hence received much attention in the literature (for surveys see [1, 22, 29]). Recently deep learning based approaches have outperformed traditional methods on many body skeleton based tasks [14] and hence we focus our discussion on motion prediction via deep learning methods.

Spatio-temporal modeling of human activity is a crucial aspect in many problem domains including activity recognition from videos [21], human-object interaction [20] and robotics [5]. Manually designed spatio-temporal graphs (st-

graphs) are typically applied to represent such problems, where nodes of the graph represent the interaction components, and the edges capture their spatio-temporal relationship. However, creating these models requires expertise and domain knowledge. Holden *et al.* [17] and Judith *et al.* [4] propose a generative model for the automation of character animation in graphics. However, this approach is not predictive in the sense of prior poses and hence is not suitable for many vision tasks.

In particular the activity and action recognition communities have explored the use of spatio-temporal models for image based action recognition [8, 21, 33] and human object interaction [20, 13]. Often several different networks are trained separately and connected manually whereas we learn spatial structure and the spatio-temporal aspects in an end-to-end trainable model and directly from data. The task of motion prediction or motion synthesis is a relatively recent development and has seen comparatively little attention in the literature [11, 19].

Generally speaking there are two main directions in modeling temporal dependencies and state transitions. Namely, explicit parametric filtering such as Gaussian processes or other variants of Bayesian filtering such as HMMs or the Kalman Filter [36, 37]. Alternatively, various flavors of deep learning methods and in particular recurrent neural networks (RNNs) have been proposed for a variety of tasks [12, 15, 16, 35]. These methods currently outperform traditional methods in many domains including that of motion prediction, with the two methods proposed in [11, 19] being the most closely related to ours.

Fragkiadaki *et al.* [11] propose to jointly learn a representation of pose data and its time variance from images. An autoencoder is used for representation learning while the time variance is learned through an RNN which is sandwiched between the encoder and the decoder. The main focus of the work is to extract motion from video frames where representation learning step is crucial. However, for body pose based motion prediction the joint angle space of the human skeleton is already relatively low dimensional and the sequences are smooth. Hence, in cases where the input is already available in joint angle form, we argue that an additional representation learning step is not necessary. In consequence, our method employs a spatio-temporal component that directly operates in joint-angle space, whereas the work in [11] operates on the transformed latent space. Specifically, we separate concerns where the autoencoder is used as a spatial filter and the RNN as temporal predictor. Furthermore, we propose a different learning strategy and architecture to minimize correlation between the predictor and the filter.

Using over complete autoencoders to model kinematic dependencies imposed by human skeleton has been proposed for image based cases [32]. In contrast, our approach does

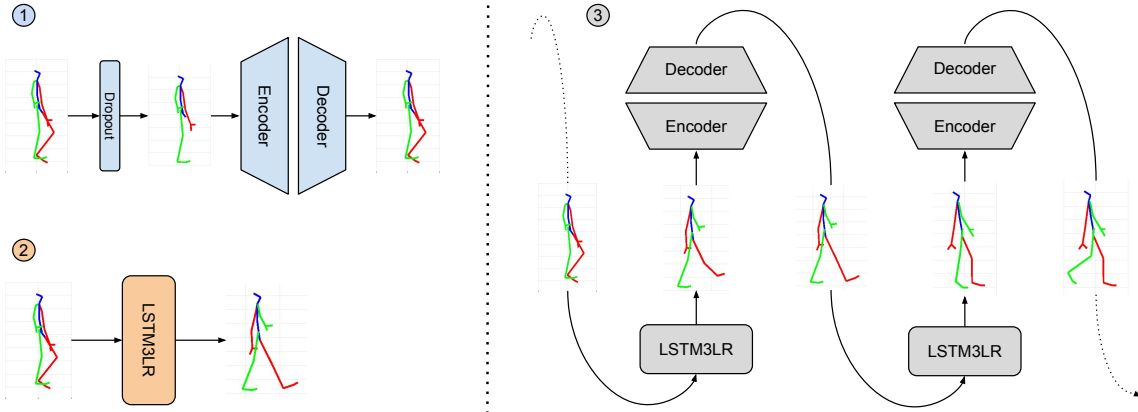


Figure 1. Schematic overview over the proposed method. (1) A variant of de-noising autoencoders learns the spatial configuration of the human skeleton via training with dropouts, removing entire joints at random which have to be reconstructed by the network. (2) We train a 3-layer LSTM recurrent neural network to predict skeletal configurations over time. (3) At inference time both components are stacked and the dropout autoencoder filters the noisy predictions of the LSTM layers, preventing accumulation of error and hence pose drift over time.

not model the temporal dependency in the latent space of the autoencoder. This is motivated by the observation that unlike image data, mocap data in its original representation is smooth and continuous while there exists no guarantees of these properties in the learnt latent space.

Martinez *et al.* [23] treat the problem of human motion modeling, focusing on short term action prediction and conclude that achieving both long and short term accuracy remains challenging. This is accredited to side-effects of curriculum learning, degrading short term prediction results. [23] avoids long term prediction and only reports results for a maximum of 400 ms into the future which is arguably sufficient for articulated object tracking but may not be sufficient for other tasks. In our work, decoupling spatial and temporal filtering during training improves robustness of the network over long time horizons, while maintaining short term prediction accuracy.

Integration of structural information in the form of spatio-temporal structural elements into deep learning models has been attempted before in [19, 21, 30]. This often requires manual design of structural elements. The main focus of Jain *et al.* [19] is to automate the transformation of manually created st-graphs into an LSTM architecture. Although this process removes much manual labor it introduces a multitude of hyper parameters, such as network architecture and design for every independent node and edge. Further more due to inherent constraints in such networks they are usually less powerful than an unstructured network of similar size. This necessitates [19] to train different models for different activities even within the H3.6M dataset. While our work also leverages the spatial structure of the data, we propose a method that does not require expert designed nor action specific st-graphs as input but instead

learns the spatial structure of the skeleton directly from the data. The key idea is to train a deep autoencoder network to implicitly capture the inter-joint dependencies by randomly removing individual joints from the inputs at training time. The temporal evolution of the motion sequences is captured by an LSTM network operating directly on reconstructed and de-noised poses. Contrary to previous work [17, 19] we train a single, unified model to perform all actions and do not require task specific networks.

### 3. Method

Figure 1 illustrates our proposed architecture. The method comprises of two main components, namely, a Dropout Autoencoder (DAE) and 3-layer LSTM (LSTM3LR). These components serve distinct purposes but accomplish a common task, that of predicting human motion into the future. More precisely the model predicts a pose configuration for the time step  $X_{t+1}$  given all  $X_{1:t}$  prior poses up to time step  $t$ . Each pose at time  $t$  consists of joint angles  $X_t = [x_1, x_2, \dots, x_n]$  of the human skeleton. Hereby the LSTM3LR outputs the most probable pose  $X_{t+1}$  given  $X_{1:t}$ . The result is then fed to an autoencoder acting as a filter to refine the prediction based on the implicit constraints imposed by the human skeleton.

The main motivation and novelty in our approach is threefold. First, the data underlying our task has a well defined spatial structure in the human skeleton and integrating this domain knowledge into the model is important. We focus on data driven recovery of the spatial configuration unlike previous attempts [19] which model it manually. Second, we observe that human motion is typically smooth, low dimensional and displays consistent spatio-temporal patterns. Hence we make no effort to perform representation

learning which can potentially introduce detrimental artifacts. Third, at inference time the predicted poses recursively serve as input for the next time step and hence even small errors in the prediction will quickly accumulate and degrade the prediction quality over long time horizons. To avoid this we de-correlate errors in each time step with output from two networks with widely different characteristics.

With these observations in place we propose a simple yet effective network architecture comprising of two main components dedicated to learning and modeling the structural aspects of the task and the spatio-temporal patterns respectively. An autoencoder learns to model the configuration of the human skeleton and is used to filter noisy predictions of the RNN but only operates in the spatial domain.

During training, both the autoencoder and 3-layer LSTM networks are pre-trained independently. In a subsequent fine-tuning step both models are trained further in an end-to-end fashion.

### 3.1. Learning spatial joint angle configurations

The Dropout Autoencoder (DAE) component is based on de-noising autoencoders, used for the learning of representations that are robust to noisy data [34]. More formally, a de-noising autoencoder learns the conditional distribution  $P_{\theta_D}(X|\tilde{X})$  where  $P_{\theta_D}$  is represented by a neural network with parameters  $\theta_D$ , to recover the data sample  $X$  given a corrupted sample  $\tilde{X}$ . During training,  $X$  is perturbed by a stochastic corruption process  $C$  where  $\tilde{X} \sim C(\tilde{X}|X)$  [2].

Similarly to prior work we perturb our input data with random noise but importantly extend the architecture to more explicitly reason about the spatial configuration of the human skeleton. We introduce dropout layers directly after the input layer with the effect of randomly removing joints entirely from the skeleton rather than simply perturbing their position and angles. The only way to recover the full pose  $X$  is then to reconstruct the missing joint angle information via inference from the adjacent joints. Importantly, during pre-training of the DAE we do not use any temporal information but for consistencies sake keep the time subscript  $t$  in this section. For a pair of clean and corrupted pose samples  $(X_t, \tilde{X}_t)$  we minimize the squared Euclidean loss:

$$\mathcal{L}(\cdot) = \left\| X_t - \tilde{X}_t \right\|^2 = \sum_n^N (x_n - \tilde{x}_n)^2 \quad (1)$$

During training of DAE the corruption process  $C$  is implicitly modeled in the network by the dropout layer just after the input layer. Introducing the dropout layer directly after the input layer forces the network to implicitly learn the spatial correlation of joints and our experiments suggest that this scheme produces better results than using the more standard multivariate Gaussian de-noising scheme only.

### 3.2. Learning temporal structure

Our goal is to recursively predict natural human poses into the future given a seed-sequence of motion capture data. This task shares similarities with other time-sequence data such as handwriting synthesis for which RNNs augmented with LSTM memory cells [16] have been shown to work well [12]. Similar to prior work [11, 19] we leverage a 3-layer LSTM network to model the temporal aspects of the task and to predict the poses forward over the time horizon. Each predicted pose  $X_{t+1}$  is filtered by the DAE component before feeding it back into the LSTM3LR network, improving the prediction quality and reducing drift over time.

The LSTM3LR network can either be utilized as a Mixture of Density Network (MDN) for probabilistic or as usual for deterministic predictions [3]. In the probabilistic case the output is modeled by a distribution family  $P_{\theta_L}(X_{t+1}|X_{1:t})$  such as a Gaussian Mixture Model (GMM). The network is then used to parametrize the predictive distribution and trained by minimizing the negative log-likelihood. In the deterministic case the predictive distribution  $P_{\theta_L}(X_{t+1}|X_{1:t})$  is implicitly modeled by the LSTM3LR network with parameters  $\theta_L$ . The network is trained by minimizing the Euclidean loss between target and predicted pose configuration.

$$\mathcal{L}(\cdot) = \left\| X_{t+1} - \hat{X}_{t+1} \right\|^2 = \sum_n^N (x_n - \hat{x}_n)^2, \quad (2)$$

where  $X_{t+1}$  and  $\hat{X}_{t+1}$  are the ground truth and predicted pose for time step  $t + 1$  respectively.

In the case of handwriting synthesis [12] the inputs are low-dimensional and sampling from a GMM distribution has been shown to prevent collapse to the mean sample. For higher dimensional data such as human poses used in this work it is only practical to use very few mixture models which furthermore have to be restricted to diagonal covariances for each component. The deterministic and probabilistic prediction configurations did not show any significant differences in our qualitative and quantitative experiments. Prior work reports similar relative performance of deterministic and probabilistic prediction [11]. Concurring with [11] we conclude that the expressive power of a mixture model with few components for high dimensional tasks such as the human motion prediction is actually inadequate and hence all but one mixture model component collapses essentially making just unimodal prediction and we hence chose the deterministic parametrization. Our experiments show that our model can produce more realistic locomotion sequence over longer time horizons than the state-of-the-art (cf. 4.5).

### 3.3. Training and inference

As outlined above it is fair to expect that the LSTM3LR component will start to predict at least somewhat noisy poses after a sufficiently large number of time steps. We therefore assume that the corruption process  $C$  is implicitly attached to the LSTM network. Consequently we leverage the DAE component to filter and improve the prediction by counteracting the corruption process. Our final architecture is then formalized as:

$$\text{LSTM3LR: } \tilde{X}_t \sim P_{\theta_L}(X_t|X_{1:t-1}) \quad (3)$$

$$\text{DAE: } X_t \sim P_{\theta_D}(X_t|\tilde{X}_t) \quad (4)$$

Because  $\tilde{X}_t \sim C(\tilde{X}_t|X_t)$  and the LSTM are assumed to be coupled, the predictions drawn from the LSTM3LR network (Eq. 3) are also assumed to be corrupted. This assumption can be verified experimentally.

After the separate pre-training phase we stack the LSTM3LR and DAE components together and continue training with a brief fine-tuning phase (i.e., training for  $\sim 2$  epochs) using both losses from Eq. 1 & 2. We experimentally found that removing the dropouts during this fine-tuning process improves the performance. Inline with the literature [28] we experimentally confirmed that annealing the dropout rate for both the input and intermediate dropout layers to zero yields the best performance. Finally, in a departure from prior work [11] the input and output representations of both the DAE and the LSTM3LR are in the original joint angle space rather than the latent space of the autoencoder.

At inference time (Figure 1, (3)) the DAE component refines each of the LSTM3LR’s pose predictions, leveraging the implicitly learned spatial structure of the skeleton. Our experiments show that this architecture leads to better sequence predictions across a variety of actions.

## 4. Experiments

We evaluate our proposed model extensively on two large publicly available datasets by Ionescu *et al.* [18] and Holden *et al.* [17], These datasets contain a large number of subjects, activities and serve as good testbed for natural human motion in varied conditions.

### 4.1. Datasets

As we train one network that generalizes to all the action categories as opposed to our most closely related work [11, 19] where a new model is trained for every activity, it is slightly unfair to compare the test errors directly. Yet to facilitate ease of comparison with the state-of-the art we evaluate our method on the H3.6M dataset following [11, 19] and conduct additional experiments on the dataset accumulated by Holden *et al.* Further more since with our implementation of SRNN following the protocol outlined in

[19] we did not manage to obtain competitive results in the Holden dataset [17], we exclude this model from our experiments in the following sections. This could partially be because of lack of action labels in this dataset and hence we tried to train one SRNN model for all of the activities as opposed to action specific models.

**Human3.6M** [6, 18] is currently the largest single dataset of high quality 3D joint positions. It consists of 15 action categories, performed by seven different professional actors and contains cyclic motions such as walking and non-cyclic activities. The actors are recorded with a Vicon motion capture system, providing high quality 3D body joint locations in the global coordinate frame sampled at 50 frames per second (fps). We follow [19, 11] and treat subject 5 in a leave-one-subject-out evaluation setting. The dataset is down sampled by 2 in time domain in order to obtain an effective fps rate of 25

**Holden *et al.*** [17] accumulated a large motion dataset from many freely available databases [7, 24, 26] and augmented these with their own data. The dataset contains around six million frames of high quality motion capture data for a single character sampled at 120 fps. While the dataset does not contain action labels it covers an even wider range of poses and hence serves well as complementary test set. We follow the training preprocessing settings reported in [17] and reserve 20% of the dataset for testing. Similar to preprocessing of H3.6M dataset we down sample this dataset by 4 to get an effective fps of 30

### 4.2. Implementation Details

**Data preprocessing** The above datasets have been preprocessed [17, 19] to normalize skeleton size i.e. height difference across all actors. The H3.6M data is further preprocessed so that the relative joint angles are taken as input as detailed in [31]. This ensures direct comparability with [11, 19]. Finally, we normalize each feature into the range of  $[0, 1]$  separately and scale inputs during prediction time with the shift and scale values computed from the training data.

**Training details** The auto encoder uses 3 dense layers with 3000 units each. We do not enforce weight sharing between layers. We use Relu nonlinearity to encourage sparsity and use dropout and  $l_2$  regularization to prevent overfitting. The learning rates are initialized as 0.005 for the first stage of training and dropped by a factor of 2 every time when validation loss flattens out. In end-to-end training, a lower learning rate 0.0001 is used. The dropout rate is set to 0.5 for the first stage and slowly annealed when validation error stops decreasing. The DAE and LSTM3LR networks are initially trained for 20 epochs. The unified end-to-end model typically starts to converge after two epochs of fine-tuning. As is common we also make use of curriculum learning to train both the autoencoder and the




Figure 2. Pose reconstruction error (L2) over dropout rate (subplot above) and Effect of RNN output filtering (L2 Error vs timestep subplot below). Both standard curriculum learning and ours improve reconstruction but ours is more robust to large perturbations.

LSTM network. A Gaussian noise with variance schedule of  $[0.01, 0.05, 0.1]$  was used while a dropout schedule of  $[0.01, 0.02, 0.04, 0.08, 0.1]$  was used. During the fine turning phase a reverse schedule is used. The scheduling hyper parameters did not impact the final model quality significantly.

### 4.3. Impact of the Dropout Autoencoder

As proposed in the method section we provide direct evidence here that the dropout learning scheme makes predictions more resilient against noise introduced by the RNN over time.

In Figure 2-a we compare pose reconstruction performance under different amounts of input corruption for three different autoencoder settings: our proposed model DAE (DropoutNoise), a standard de-noising autoencoder GAE (GaussianNoise) and a vanilla autoencoder (Vanilla). Our Dropout Autoencoder configuration is more robust to increasing amount of corruption and recovers the noisy input with lower error rates.

Similarly, we compare the performance of these autoencoders by stacking them with a pre-trained LSTM3LR network. The autoencoders are expected to filter out noisy predictions of the LSTM component. The filtered predictions are then compared with the ground-truth data. Figure 2-b shows that our model DAE-LSTM yields better performance and DAE improves the prediction quality by effectively removing the noise introduced by the LSTM3LR network.

Further, we assess the impact of the DAE component on overall prediction accuracy in our second evaluation dataset. Table 1 compares Euclidean distance to ground-truth averaged across the Holden dataset. Note that both rows result from the same model, however the top-row is the error of the unfiltered LSTM output and the bottom row is the average

Methods	Short-term (ms)		Long-term (ms)		
	80	160	320	560	1000
Ours without Filtering	2.72	3.39	4.44	3.96	4.02
<b>Ours</b>	2.42	3.14	4.37	4.09	4.03

Table 1. Comparison of average error in joint position at different time horizons on Holden. Error in  $cm$  for unfiltered LSTM predictions (top) and that obtained via filtering with the DAE network (bottom). Filtering via the DAE network at every time-step improves accuracy and reduces long-term drift.

error after filtering these predictions with the DAE component. The LSTM3LR produces noisy predictions which are improved by the DAE (note that these accuracies are identical to bottom row in Table 3).

### 4.4. Short-term motion prediction

We first report quantitative results on H3.6M dataset using the same experiment configuration with [11, 19]. The metric is taken from [11], simply calculating the Euclidean distance between the predicted MOCAP vector and the ground truth. Please note that while this metric is useful to evaluate the short-term<sup>2</sup> predictions, it may yield deceiving assessments in longer time horizons since good models generate novel sequences, where deviations from ground truth are desired indeed. As reported previously [11, 19], it is worthwhile noting that the metric does not always correspond directly with the observed motion quality. The direct Euclidean error computation between the predicted and the ground truth MOCAP vector makes this metric less intuitive. In other words, a minor error in base frame (e.g., hip joint) angle can cause a large visual error, while the same error at a child node (e.g., wrist or ankle joints) would cause an insignificant effect. We therefore report results only on the action classes that have been reported on in the literature.

Furthermore since we were not able to replicate the performance of SRNN [19] with one model for all actions<sup>3</sup> we avoid reporting suboptimal results and only compare with results previously presented in [19].

Analogously to the literature we also include a comparison to 3-layer LSTM architecture (LSTM3LR) as a baseline. In all our motion prediction experiments we initialize each model with 50 seed frames and then start predicting 300 frames (12s) into the future.

Table 2 summarizes results from the walking, eating, smoking and discussion activities for short- and long-term periods. The simple baseline (LSTM3LR) is surprisingly

<sup>2</sup>We have indicated in the comparison tables what can be considered as short term.

<sup>3</sup>The original work on SRNN [19] implements different model for every action

Methods	Short-term (ms)		Long-term (ms)		
	80	160	320	560	1000
Walking activity					
LSTM3LR[19]	1.18	1.50	1.67	1.81	2.20
ERD [19]	1.30	1.56	1.84	2.00	2.38
S-RNN [19]	1.08	1.34	1.60	1.90	2.13
<b>Ours</b>	<b>1.00</b>	<b>1.11</b>	<b>1.39</b>	<b>1.55</b>	<b>1.39</b>
Eating activity					
LSTM3LR[19]	1.36	1.79	2.29	2.49	2.82
ERD [19]	1.66	1.93	2.28	2.36	2.41
S-RNN [19]	1.35	1.71	2.12	2.28	2.58
<b>Ours</b>	<b>1.31</b>	<b>1.49</b>	<b>1.86</b>	<b>1.76</b>	<b>2.01</b>
Smoking activity					
LSTM3LR[19]	2.05	2.34	3.10	3.24	3.42
ERD [19]	2.34	2.74	3.73	3.68	3.82
S-RNN [19]	1.90	2.30	2.90	3.21	3.23
<b>Ours</b>	<b>0.92</b>	<b>1.03</b>	<b>1.15</b>	<b>1.38</b>	<b>1.77</b>
Discussion activity					
LSTM3LR[19]	2.25	2.33	2.45	2.48	2.93
ERD [19]	2.67	2.97	3.23	3.47	2.92
S-RNN [19]	1.67	2.03	2.20	2.39	2.43
<b>Ours</b>	<b>1.11</b>	<b>1.20</b>	<b>1.38</b>	<b>1.53</b>	<b>1.73</b>

Table 2. Comparison of short-term predictions (<1s) of the different models over four different activities on the H3.6M dataset. We report error as the Euclidean norm (L2) of un-normalized ground truth and predicted MOCAP vectors.

Methods	Short-term (ms)		Long-term (ms)		
	80	160	320	560	1000
LSTM3LR	2.76	3.41	<b>4.23</b>	<b>3.89</b>	4.12
ERD	2.87	3.88	5.64	6.08	6.96
<b>Ours</b>	<b>2.42</b>	<b>3.14</b>	4.37	4.09	<b>4.03</b>

Table 3. Comparison of short-term predictions (<1s) of the different models on the Holden dataset. We report average Euclidean norm (L2) error of ground truth and predicted MOCAP vectors (expressed in cm/joint unit). The height of the skeleton (1.7m) was used to convert errors to metric scale.

competitive in making short-term predictions. However, a qualitative inspection on Figure 3 shows that the baseline quickly converges to the safe mean pose, whereas the other models generate diverse and natural poses. Since ERD does not explicitly model the skeletal structure it starts to generate unnatural poses quickly. Our model, on the other hand, continues to predict smooth and natural looking poses especially over the longest horizon (1000ms).

Table 3 shows the results obtained from the same experiment conducted on the Holden dataset. Because there are no action labels we average the error across all test sequences. Note that here the error has a unit of *cm-per-joint* as opposed to *unit distance* in the exponential domain in Table 2.

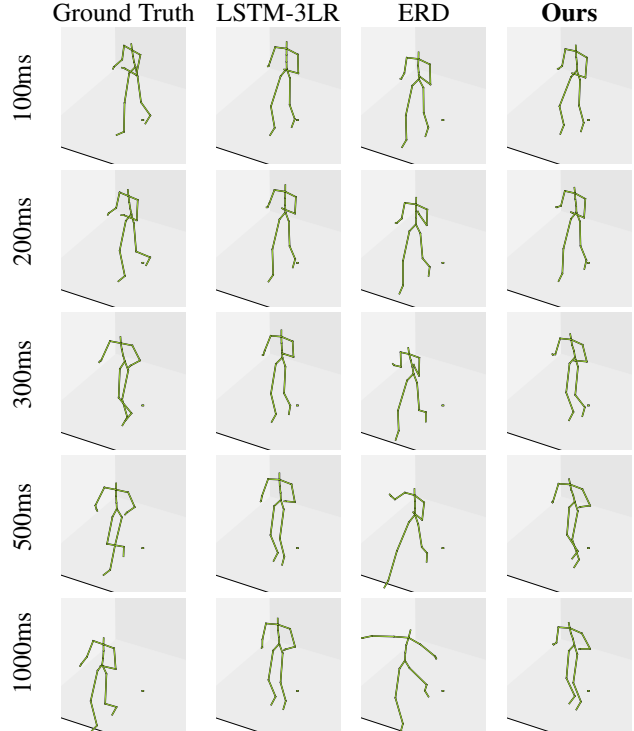


Figure 3. Qualitative comparison of our model with the state-of-the-art on the “walking” activity. The baseline LSTM quickly converges to a safe mean configuration. ERD produces unnatural poses for longer horizons. Ours produces natural looking configurations without collapsing to an average pose.

Similarly to the H3.6M case our model either outperforms others or performs similarly with the baseline model.

#### 4.5. A metric for motion synthesis

In order to better differentiate model capabilities, especially for long-term prediction horizons, we leverage a pre-trained activity classifier for the evaluation of synthetic motion sequences. Intuitively, a high quality synthetic sequence should be assigned the same action label as the seed, whereas drift and motion degradation should impact the classification outcome negatively. This evaluation protocol is similar to the evaluation of generative adversarial networks [25]. The evaluation by a separate classifier network is highly correlated to human judgment of action quality. Please refer to the supplementary videos that visualize action class probabilities of the classifier alongside an animated skeleton.

Given the success of the classifier in evaluating the synthetic sequences accurately, reformulating the problem as a auto-regressive generative adversarial network (GAN) holds potential. However, it arguably requires significant modifications and we leave this as an interesting direction for future work. Here we provide a visual representation of




Figure 4. Comparison of class probabilities of long term sequence generation using pre-trained classifier. Our model generates sequences which belong to the same (correct) class for long  $> 10s$  horizons. The eating activity is initially confused with highly similar sitting activity but ours still yields best results.)

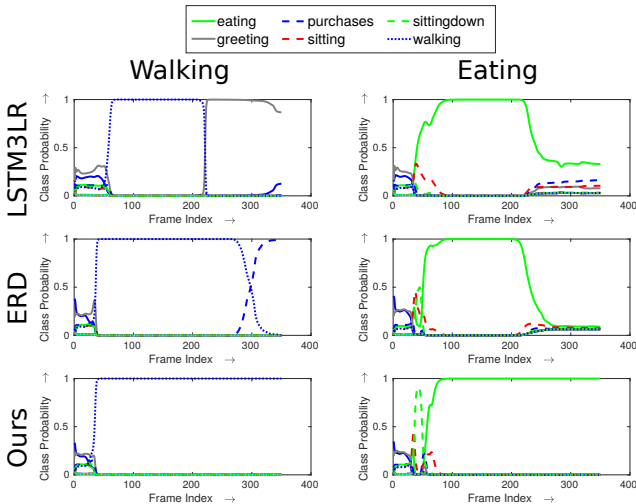


Figure 5. Multi-class probabilities for eating and walking. Initially in eating activity generated by our model is confused with sitting but eventually the arm motions lead to correct classification, whereas the baseline converges to a mean pose and hence remains ambiguous.

the action classification probabilities and not their numerical values since their precise values are dependent upon the training details of the classifier, consequently making the precise probability values unimportant or even misleading.

In our experiments we train a separate LSTM network performing on par with state-of-the-art action recognition methods [30]. It is used to assign class probabilities to the synthetic pose sequences generated by the baseline, the ERD network and our model.

Figure 4 plots class probabilities of “walking” and “eating” categories. Our model produces sequences that are

classified correctly for longer time horizons than the baseline and ERD networks especially for cyclic motions such as “walking”. Note that for the non-cyclic “eating” motion (Figure 4, bottom) the performance is degraded. Inspecting Figure 5 closely reveals that the output from our model is initially confused with a very similar “sitting” activity which only become distinguishable from “eating” when the hands start moving. This effect is best viewed in the video provided along with the supplementary material<sup>4</sup>.

## 5. Conclusion

In this paper we have proposed the Dropout Autoencoder LSTM (DAE-LSTM) model for prediction of natural and realistic human motion given a short seed sequence. Our proposed model consists of a 3-layer LSTM (LSTM3LR) and a dropout autoencoder (DAE). We train the autoencoder by randomly removing individual joints from the training poses in order to learn the spatial dependencies of the human skeleton implicitly. Furthermore, we have introduced an evaluation protocol that can be used to better analyze the quality of synthetic motion sequences in particular over long-time horizons where the simple Euclidean distance to the seed sequence does not provide a meaningful assessment anymore. Finally, we have experimentally demonstrated that our method outperforms the LSTM3LRbaseline as well as the most closely related work ERD in a variety of experiments performed on two datasets. However if scrutinized closely one can notice that the animated skeleton leans slightly backwards. We attribute this to the fact that no physics based feedback is given to the model. Hence the model has no concept of mass, balance or gravity, which prevents it from identifying small error in overall orientation which strikes to human evaluator as physically impossible or improbable pose. Incorporating a physical model is left for future work.

## References

- [1] J. K. Aggarwal and Q. Cai. Human motion analysis: A review. In *Nonrigid and Articulated Motion Workshop, 1997. Proceedings., IEEE*, pages 90–102. IEEE, 1997. 2
- [2] Y. Bengio, L. Yao, G. Alain, and P. Vincent. Generalized denoising auto-encoders as generative models. In *Advances in Neural Information Processing Systems*, pages 899–907, 2013. 4
- [3] C. M. Bishop. Mixture density networks. 1994. 4
- [4] J. Bütepage, M. J. Black, D. Kragic, and H. Kjellström. Deep representation learning for human motion prediction and classification. *CoRR*, abs/1702.07486, 2017. 2

<sup>4</sup>As opposed to previous attempts we do not include this analysis for discussion and smoking action as they are indistinguishable from other activities in the action set (e.g. directions, waiting, walking and sitting) which confuses the classifier.



- [5] J. Bütepage, H. Kjellström, and D. Kragic. Anticipating many futures: Online human motion prediction and synthesis for human-robot collaboration. *arXiv preprint arXiv:1702.08212*, 2017. 2
- [6] C. S. Catalin Ionescu, Fuxin Li. Latent structured models for human pose estimation. In *International Conference on Computer Vision*, 2011. 5
- [7] CMU. Carnegie-mellon mocap database. 5
- [8] Y. Du, W. Wang, and L. Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1110–1118, 2015. 2
- [9] V. Ferrari, M. Marin-Jimenez, and A. Zisserman. Progressive search space reduction for human pose estimation. In *CVPR*, pages 1–8, 2008. 1
- [10] K. Fragkiadaki, H. Hu, and J. Shi. Pose from flow and flow from pose. In *CVPR*, pages 2059–2066, 2013. 1
- [11] K. Fragkiadaki, S. Levine, and J. Malik. Recurrent network models for kinematic tracking. *CoRR*, abs/1508.00271, 2015. 1, 2, 4, 5, 6
- [12] A. Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013. 2, 4
- [13] A. Gupta, A. Kembhavi, and L. S. Davis. Observing human-object interactions: Using spatial and functional compatibility for recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(10):1775–1789, Oct 2009. 2
- [14] F. Han, B. Reily, W. Hoff, and H. Zhang. Space-time representation of people based on 3d skeletal data: A review. *arXiv preprint arXiv:1601.01006*, 2016. 2
- [15] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6):82–97, 2012. 2
- [16] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, Nov. 1997. 2, 4
- [17] D. Holden, J. Saito, and T. Komura. A deep learning framework for character motion synthesis and editing. In *SIGGRAPH 2016*, 2016. 2, 3, 5
- [18] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, jul 2014. 2, 5
- [19] A. Jain, A. R. Zamir, S. Savarese, and A. Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. *CoRR*, abs/1511.05298, 2015. 1, 2, 3, 4, 5, 6, 7
- [20] H. S. Koppula and A. Saxena. Anticipating human activities using object affordances for reactive robotic response. *IEEE Trans. Pattern Anal. Mach. Intell.*, 38(1):14–29, Jan. 2016. 2
- [21] J. Liu, A. Shahroudy, D. Xu, and G. Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In *European Conference on Computer Vision*, pages 816–833. Springer, 2016. 2, 3
- [22] C. Mandery, Ö. Terlemez, M. Do, N. Vahrenkamp, and T. Asfour. Unifying representations and large-scale whole-body motion databases for studying human motion. *IEEE Transactions on Robotics*, 32(4):796–809, 2016. 2
- [23] J. Martinez, M. J. Black, and J. Romero. On human motion prediction using recurrent neural networks. *CoRR*, abs/1705.02445, 2017. 3
- [24] M. Müller, T. Röder, M. Clausen, B. Eberhardt, B. Krüger, and A. Weber. Documentation mocap database hdm05. Technical Report CG-2007-2, Universität Bonn, June 2007. 5
- [25] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier gans. *arXiv preprint arXiv:1610.09585*, 2016. 7
- [26] F. Ofii, R. Chaudhry, G. Kurillo, R. Vidal, and R. Bajcsy. Berkeley mhad: A comprehensive multimodal human action database. In *2013 IEEE Workshop on Applications of Computer Vision (WACV)*, pages 53–60, Jan 2013. 5
- [27] D. Ramanan, D. A. Forsyth, and A. Zisserman. Strike a pose: Tracking people by finding stylized poses. In *CVPR*, volume 1, pages 271–278, 2005. 1
- [28] S. J. Rennie, V. Goel, and S. Thomas. Annealed dropout training of deep networks. In *Spoken Language Technology Workshop (SLT), 2014 IEEE*, pages 159–164. IEEE, 2014. 5
- [29] K. M. Sagayam and D. J. Hemanth. Hand posture and gesture recognition techniques for virtual reality applications: a survey. *Virtual Reality*, 21(2):91–107, 2017. 2
- [30] A. Shahroudy, J. Liu, T. Ng, and G. Wang. NTU RGB+D: A large scale dataset for 3d human activity analysis. *CoRR*, abs/1604.02808, 2016. 3, 8
- [31] G. W. Taylor, G. E. Hinton, and S. T. Roweis. Modeling human motion using binary latent variables. In P. B. Schölkopf, J. C. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 1345–1352. MIT Press, 2007. 5
- [32] B. Tekin, I. Katircioglu, M. Salzmann, V. Lepetit, and P. Fua. Structured prediction of 3d human pose with deep neural networks. *CoRR*, abs/1605.05180, 2016. 2
- [33] V. Veeriah, N. Zhuang, and G.-J. Qi. Differential recurrent neural networks for action recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4041–4049, 2015. 2
- [34] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103. ACM, 2008. 4
- [35] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3156–3164, 2015. 2
- [36] J. M. Wang, D. J. Fleet, and A. Hertzmann. Gaussian process dynamical models for human motion. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(2):283–298, Feb. 2008. 2
- [37] D. Wu and L. Shao. Leveraging hierarchical parametric networks for skeletal joints based action segmentation and recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 724–731, 2014. 2