

# On human motion prediction using recurrent neural networks

Julieta Martinez<sup>\*1</sup>, Michael J. Black<sup>2</sup>, and Javier Romero<sup>3</sup>

<sup>1</sup>University of British Columbia, Vancouver, Canada

<sup>2</sup>MPI for Intelligent Systems, Tübingen, Germany

<sup>3</sup>Body Labs Inc., New York, NY

julm@cs.ubc.ca, black@tuebingen.mpg.de, javier.romero@bodylabs.com

## Abstract

Human motion modelling is a classical problem at the intersection of graphics and computer vision, with applications spanning human-computer interaction, motion synthesis, and motion prediction for virtual and augmented reality. Following the success of deep learning methods in several computer vision tasks, recent work has focused on using deep recurrent neural networks (RNNs) to model human motion, with the goal of learning time-dependent representations that perform tasks such as short-term motion prediction and long-term human motion synthesis. We examine recent work, with a focus on the evaluation methodologies commonly used in the literature, and show that, surprisingly, state-of-the-art performance can be achieved by a simple baseline that does not attempt to model motion at all. We investigate this result, and analyze recent RNN methods by looking at the architectures, loss functions, and training procedures used in state-of-the-art approaches. We propose three changes to the standard RNN models typically used for human motion, which result in a simple and scalable RNN architecture that obtains state-of-the-art performance on human motion prediction.

## 1. Introduction

An important component of our capacity to interact with the world resides in the ability to predict its evolution over time. Handing an object to another person, playing sports, or simply walking in a crowded street would be extremely challenging without our understanding of how people move, and our ability to predict what they are likely to do in the following instants. Similarly, machines that are able to perceive and interact with moving people, either in physical or virtual environments, must have a notion of how people move. Since human motion is the result of both physical limitations (*e.g.* torque exerted by muscles, gravity, moment preservation) and the intentions of subjects (*how* to perform

<sup>\*</sup>Research carried out while Julieta was an intern at MPI.

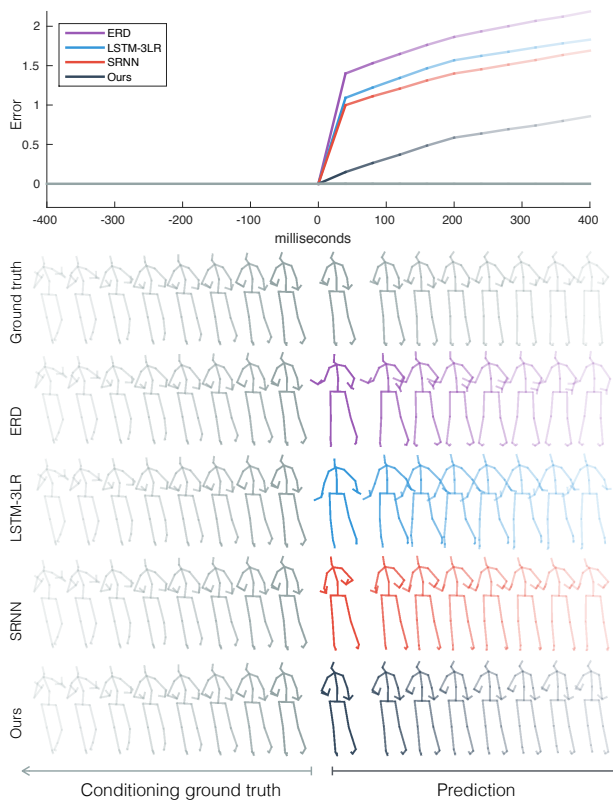


Figure 1. Top: Mean average prediction error for different motion prediction methods. Bottom: Ground truth passed to the network is shown in grey, and short-term motion predictions are shown in colour. Previous work, based on deep RNNs, produces strong discontinuities at the start of the prediction (middle column). Our method produce smooth, low-error predictions.

an intentional motion), motion modeling is a complex task that should be ideally learned from observations.

Our focus in this paper is to learn models of human motion from motion capture (mocap) data. More specifically, we are interested in human motion prediction, where we forecast the most likely future 3D poses of a person given their past motion. This problem has received interest in a

wide variety of fields, such as action prediction for socially-aware robotics [21], 3D people tracking within computer vision [13, 43], motion generation for computer graphics [22] or modeling biological motion in psychology [42].

Traditional approaches have typically imposed expert knowledge about motion in their systems in the form of Markovian assumptions [25, 32], smoothness, or low dimensional embeddings [48]. Recently, a family of methods based on deep recurrent neural networks (RNNs) have shown good performance on this task while trying to be more agnostic in their assumptions. For example, [10] uses curriculum learning and incorporates representation learning in the architecture, and [18] manually encodes the semantic similarity between different body parts. These approaches benefit from large, publicly available collections of motion capture data [16], as well as recent advances in the optimization of time-series modelling [9].

Recent work has validated its performance via two complementary methods: (1) *quantitative* prediction error in the short-term, typically measured as a mean-squared loss in angle-space, and (2) *qualitative* motion synthesis for longer time horizons, where the goal is to generate feasible motion. The first evaluation metric is particularly interesting for computer vision applications such as people tracking, where predictions are continually matched and corrected with new visual evidence. The second criterion, most relevant for open-loop motion generation in graphics, is hard to evaluate quantitatively, because human motion is a highly non-deterministic process over long time horizons. This problem is similar to the one found in recent research on deep generative networks [41], where the numerical evaluation based on the negative log-likelihood and Parzen window estimates are known to be far from perfect.

We have empirically observed that current deep RNN-based methods have difficulty obtaining good performance on both tasks. Current algorithms are often trained to minimize a quantitative loss for short-term prediction, while striving to achieve long-term plausible motion by tweaking the architectures or learning procedures. As a result, their long-term results suffer from occasional unrealistic artifacts such as foot sliding, while their short-term results are not practical for tracking due to clear discontinuities in the first prediction. In fact, the discontinuity problem is so severe, that we have found that state-of-the-art methods are quantitatively outperformed by a range of simple baselines, including a constant pose predictor. While this baseline does not produce interesting motion in the long-run, it highlights both a poor short-term performance, as well as a severe discontinuity problem in current deep RNN approaches. In this work, we argue that (a) the results achieved by recent work are not fully satisfactory for either of these problems, and (b) trying to address both problems at once is very challenging, especially in the absence of a proper quantitative

evaluation for long-term plausibility.

We focus on short-term prediction, which is the most relevant task for a visual tracking scenario. We investigate the reasons for the poor performance of recent methods on this task by analyzing several factors such as the network architectures and the training procedures used in state-of-the-art RNN methods. First, we consider the training schedule used in [10, 18]. It is a known problem in RNNs [5] and reinforcement learning [36] that networks cannot learn to recover from their own mistakes if they are fed only ground-truth during training. The authors of [10, 18] introduced increasing amounts of random noise during training to compensate for this effect. However, this noise is difficult to tune, makes it harder to choose the best model based on validation error, and has the effect of degrading the quality of the prediction in the first frame. Instead, we propose a simple approach that introduces realistic error in training time without any scheduling; we simply feed the predictions of the net, as it is done in test time. This increases the robustness of the predictions compared to a network trained only on ground truth, while avoiding the need of a difficult-to-tune schedule.

Unfortunately, this new architecture is still unable to accurately represent the conditioning poses in its hidden representation, which still results in a discontinuity in the first frame of the prediction. We borrow ideas from research on the statistics of hand motion [15], and model velocities instead of absolute joint angles, while keeping the loss in the original angle representation to avoid drift. Therefore, we propose a *residual* architecture that models first-order motion derivatives, which results in smooth and much more accurate short-term predictions.

Both of our contributions can be implemented using an architecture that is significantly simpler than those in previous work. In particular, we move from the usual multi-layer LSTM architectures (long short-term memory) to a single GRU (Gated Recurrent Unit), and do not require a spatial encoding layer. This allows us to train a single model on the entire Human 3.6M dataset [16] in a few hours. This differs from previous approaches [10, 18], which trained only action-specific models from that dataset. Our approach sets the new state of the art on short-term motion prediction, and overall gives insights into the challenges of motion modelling using RNNs. Our code is publicly available at <https://github.com/una-dinosauria/human-motion-prediction>.

## 2. Related work

Our main task of interest is human motion prediction, with a focus on recent deep RNN architectures [10, 18]. One of our findings is that, similar to [3, 17], a family of simple baselines outperform recent deep learning approaches. We briefly review the literature on these topics below.

**Modelling of human motion.** Learning statistical models of human motion is a difficult task due to the high-dimensionality, non-linear dynamics and stochastic nature of human movement. Over the last decade, and exploiting the latent low-dimensionality of action-specific human motion, most work has focused on extensions to latent-variable models that follow state-space equations such as hidden Markov models (HMMs) [26], exploring the trade-offs between model capacity and inference complexity. For example, Wang *et al.* [48] use Gaussian-Processes to perform non-linear motion prediction, and learn temporal dynamics using expectation maximization and Markov-chain Monte Carlo. Taylor *et al.* [40] assume a binary latent space and model motion using a conditional restricted Boltzmann machine (CRBM), which requires sampling for inference. Finally, Lehrmann *et al.* [26] use a random forest to non-linearly choose a linear system that predicts the next frame based on the last few observations.

**Applications of human motion models.** Motion is a key part of actions; therefore, the field of action recognition has paid special attention to models and representations of human motion. In their seminal work, Yacoob and Black [51] model motion with time-scale and time-shifted activity bases from a linear manifold of visual features computed with Principal Component Analysis (PCA). More complex models like mixtures of HMMs [33, 34], latent topic models of visual words [49] or LSTMs [29] are used in recent methods. Although their purpose (action classification from a sequence of poses) is different from ours, this field contains interesting insights for motion prediction, such as the importance of a mathematically sound orientation representation [46, 47] or how learned, compact motion representations improve action recognition accuracy [30].

Another popular use of motion models, specially short-term ones, is pose tracking. The use of simple linear Markovian models [37] or PCA models [44] has evolved to locally linear ones like factor analyzers [28], non-linear embeddings like Laplacian Eigenmaps [38], Isomap [19], dynamic variants of Gaussian Process Latent Variable Models (GPLVM) [48, 52], or physics-based models [8].

In animation, similar methods have been used for the generation of human pose sequences. Spaces of HMMs parameterised by style were used by Brand *et al.* [7] to generate complex motions. Arikian and Forsyth [2] collapse full sequences into nodes in a directed graph, connected with possible transitions between them, and in [24] cluster trees improve the path availability. More recently, motion models based on GPLVM have been used for controlling virtual characters in a physical simulator [27]. An overview of motion generation for virtual characters can be found in [45].

**Deep RNNs for human motion.** Our work focuses on recent approaches to motion modelling that are based on deep

RNNs. Fragkiadaki *et al.* [10] propose two architectures: LSTM-3LR (3 layers of Long Short-Term Memory cells) and ERD (Encoder-Recurrent-Decoder). Both are based on concatenated LSTM units, but the latter adds non-linear space encoders for data pre-processing. The authors also note that, during inference, the network is prone to accumulate errors, and quickly produces unrealistic human motion. Therefore, they propose to gradually add noise to the input during training (as is common in curriculum learning [6]), which forces the network to be more robust to prediction errors. This noise scheduling makes the network able to generate plausible motion for longer time horizons, specially on cyclic walking sequences. However, tuning the noise schedule is hard in practice.

More recently, Jain *et al.* [18] introduced structural RNNs (SRNNs), an approach that takes a manually designed graph that encodes semantic knowledge about the RNN as input, and creates a bi-layer architecture that assigns individual RNN units to semantically similar parts of the data. The authors also employ the noise scheduling technique introduced by Fragkiadaki *et al.*, and demonstrate that their network outperforms previous work both quantitatively in short-term prediction, as well as qualitatively. Interestingly, SRNNs produce plausible long-term motion for more challenging, locally-periodic actions such as eating and smoking, and does not collapse to unrealistic poses in aperiodic “discussion” sequences.

**Revisiting baselines amid deep learning.** The rise and impressive performance of deep learning methods in classical problems such as object recognition [23] has encouraged researchers to attack both new and historically challenging problems using variations of deep neural networks. For example, there is now a large body of work on visual question answering (VQA), *i.e.* the task of answering natural-language questions by looking at images, based almost exclusively on end-to-end trainable systems with deep CNNs for visual processing and deep RNNs for language modelling [31, 35, 50]. Recently, however, Zhou *et al.* [53] have shown that a simple baseline that concatenates features from questions’ words and CNN image features performs comparably to approaches based on deep RNNs. Moreover, Jabri *et al.* [17] have shown competitive performance on VQA with a simple baseline that does not take images into account, and state-of-the-art performance with a baseline that is trained to exploit the correlations between questions, images and answers.

Our work is somewhat similar to that of Jabri *et al.* [17], in that we have found a very simple baseline that outperforms sophisticated state-of-the-art methods based on deep RNNs for short-term motion prediction. In particular, our baseline outperforms the ERD and LSTM-3LR models by Fragkiadaki *et al.* [10], as well as the structural RNN (SRNN) method of Jain *et al.* [18]. Another example of baselines

outperforming recent work in the field of pose models can be found in [25], where a Gaussian pose prior outperforms the more complicated GPLVM.

### 3. Method

Recent deep learning methods for human pose prediction [10,18] offer an agnostic learning framework that could potentially be integrated with video data [10] or used for other forecasting applications [18]. However, for the specific task of motion forecasting, we note that they have a few common pitfalls that we would like to improve.

#### 3.1. Problems

**First frame discontinuity.** While both methods generate continuous motion, a noticeable jump between the conditioning ground truth and the first predicted frame is present in their results (see Figure 1). This jump is particularly harmful for tracking applications, where short-term predictions are continuously updated with new visual evidence.

**Hyper-parameter tuning.** These methods add to the typical set of network hyper-parameters an additional one, particularly hard to tune: the noise schedule.

In time series modelling, it is often necessary to model noise as part of the input, in order to improve robustness against noisy observations. For example, in Kalman filtering, a small amount of Gaussian noise is modelled explicitly as part of the standard state-space equations. In applications such as motion synthesis, exposing the method to the errors that the network will make at test time is crucial to prevent the predicted poses from leaving the manifold of plausible human motion. Algorithms like DAGGER [36], used in reinforcement learning, use queries to an “expert” during training so that the predictor learns how to correct its own errors. It is, however, not straightforward how one would use this approach for pose prediction.

The basic architectures that we use, RNNs, typically do not consider this mismatch between train and test input, which makes them prone to accumulate errors at inference time. To alleviate this problem, Fragkiadaki *et al.* propose to use noise scheduling; that is, to inject noise of gradually increasing magnitude to the input in training time (see Fig. 2, left), which corresponds to a type of curriculum learning. Jain *et al.* [18] similarly adopt this idea, and have found that it helps stabilizing long-term motion synthesis. The downsides are, (1) that both noise distribution and magnitude scheduling are hard to tune, (2) that while this noise improves long-term predictions, it tends to hurt performance in short-term predictions, as they become discontinuous from previous observations, and (3) that the common rule for choosing the best model, based on lowest validation error, is not valid anymore, since lowest validation error typically corresponds to the validation epoch

without injected noise.

**Depth and complexity of the models.** LSTM-3LR, ERD and SRNN use more than one RNN in their architectures, stacking two or three layers for increased model capacity. While deeper models have empirically shown the best performance on a series of tasks such as machine translation [39], deep networks are known to be hard to train when data is scarce (which is the data regime for action-specific motion models). Moreover, recent work has shown that shallow RNNs with minimal representation processing can achieve very competitive results in tasks such as the learning of sentence-level embeddings [20], as long as a large corpus of data is available. Finally, deeper models are computationally expensive, which is an important factor to consider in the context of large-scale training datasets.

**Action-specific networks.** Although the vision community has recently benefited from large-scale, publicly available datasets of motion capture data [16], motion modelling systems have been typically trained on small action-specific subsets. While restricting the training data to coherent subsets makes modelling easier, it is also well-known that deep networks work best when exposed to large and diverse training datasets [20,23]. This should specially apply to datasets like Human3.6M, where different actions contain large portions of very similar data (*e.g.* sitting or walking).

#### 3.2. Solutions

**Sequence-to-sequence [39] architecture.** We address short-term motion prediction as the search for a function that maps an input sequence (conditioning ground truth) to an output (predicted) sequence. In this sense, the problem is analogous to machine translation, where sequence-to-sequence (seq2seq) architectures are popular. In seq2seq, two networks are trained; (a) an encoder that receives the inputs and generates an internal representation, and (b), a decoder network, that takes the internal state and produces a maximum likelihood estimate for prediction. Unlike the common practice in machine translation, we enforce the encoder and the decoder to share weights, which we found to accelerate convergence. A benefit of this architecture is that the encoding-decoding procedure during training is more similar to the protocol used at test-time. Moreover, there are multiple variations of seq2seq architectures (*e.g.*, with attention mechanisms [4], or bi-directional encoders [35]), that could potentially improve motion prediction.

**Sampling-based loss.** While it is often common in RNNs to feed the ground truth at each training time-step to the network, this approach has the downside of the network not being able to recover from its own mistakes. Previous work has addressed this problem by scheduling the rate at which the network sees either the ground truth or its own predictions [5], or by co-training and adversarial network to force

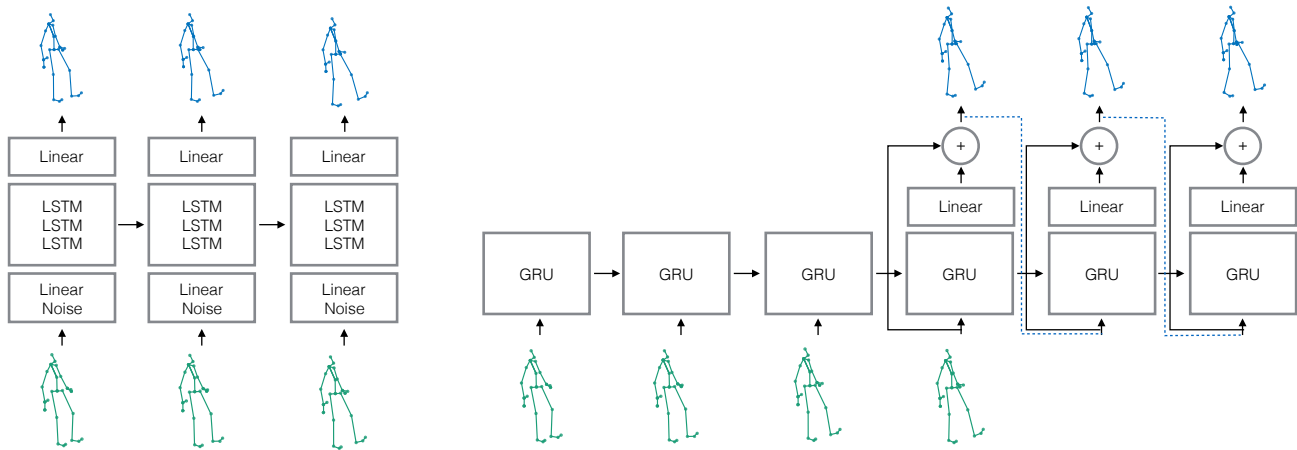


Figure 2. Training procedure as done in previous work, and our proposed sequence-to-sequence residual architecture. Green stick figures represent ground truth, and blue stick figures represent predictions. Left: LSTM-3LR architecture, introduced by Fragkiadaki *et al.* [10]. During training, ground truth is fed to the network at each time-step, and noise is added to the input. Right: Our sequence-to-sequence architecture; During training, the ground truth is fed to an *encoder* network, and the error is computed on a *decoder* network that feeds its own predictions. The decoder also has a residual connection, which effectively forces the RNN to internally model angle velocities.

the internal states of the RNN to be similar during train and test time [11]. These approaches, however, rely heavily on hyper-parameter tuning, which we want to avoid. Striving for simplicity, during training we let the decoder produce a sequence by always taking as input its own samples. This approach requires absolutely no parameter tuning. Another benefit of this approach is that we can directly control the length of the sequences that we train on. As we will see, training to minimize the error on long-term motions results in networks that produce plausible motion in the long run, while training to minimize error the short-term reduces the error rate in the first few predicted frames.

**Residual architecture.** While using a seq2seq architecture trained with a sampling-based loss can produce plausible long-term motion, we have observed that there is still a strong discontinuity between the conditioning sequence and prediction. Our main insight is that motion continuity, a known property of human motion, is easier to express in terms of velocities than in poses. While it takes considerable modelling effort to represent all possible conditioning poses so that the first frame prediction is continuous, it only requires modeling one particular velocity (zero, or close to zero velocity) to achieve the same effect. This idea is simple to implement in current deep learning architectures since it translates into adding a residual connection between the input and the output of each RNN cell (see Fig. 2, right). We note that, although residual connections have been shown to improve performance on very deep convolutional networks [14], in our case they help us model prior knowledge about the statistics of human motion.

**Multi-action models.** We also explore training a single model to predict motion for multiple actions, in contrast

to previous work [10, 18], which has focused on building action-specific models. While modelling multiple actions is a more difficult task than modelling single-action sets, it is now a common practice to train a single, conditional model, on multiple data modalities, as this allows the network to exploit regularities in large datasets [12]. Semantic knowledge about each activity can be easily incorporated using one-hot vectors; *i.e.*, concatenating, in the input, a 15-dimensional vector that has zeros everywhere, but a value of one in the index of the indicated action.

## 4. Experimental setup

We consider three main sets of experiments to quantify the impact of our contributions:

### 1. Seq2seq architecture and sampling-based loss.

First, we train action-specific models using our proposed sequence-to-sequence architecture with sampling-based loss, and compare it to previous work, which uses noise scheduling, and to a baseline that feeds the ground truth at each time-step. The goal of these experiments is to verify that using a sampling-based loss, which does not require parameter tuning, performs on par with previous work on short-term motion prediction, while still producing plausible long-term motion. In these experiments, the network is trained to minimize the loss over 1 second of motion.

### 2. Residual architecture.

The second set of experiments explore the effects of using a residual architecture that models first-order motion derivatives, while keeping the loss in the original angle space. Here, we are interested in learning whether a residual architecture improves short term prediction; therefore, in these experiments, the network is trained

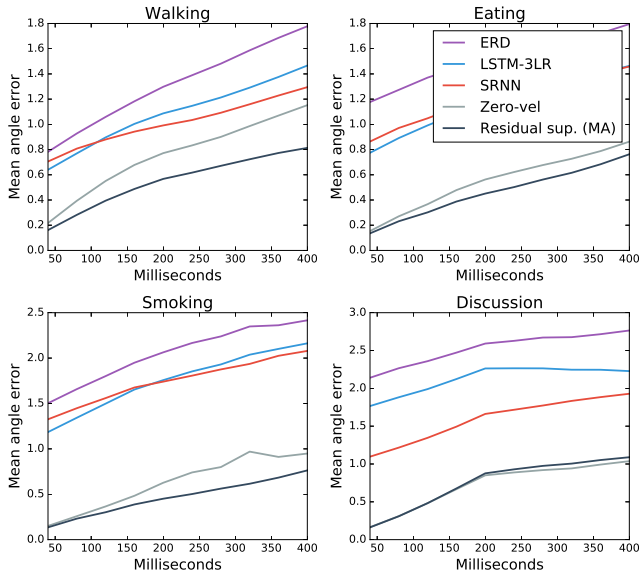


Figure 3. Error curves comparing ERD [10], LSTM-3LR [10], SRNN [18] and our method (**Residual sup. (MA)** in Table 1) with residual connections, sampling-based loss and trained on multiple actions, as well as a zero-velocity baseline.

to minimize the prediction error over 400 milliseconds.

**3. Multi-action models.** Our last round of experiments quantifies the benefits of training our architecture on the entire Human 3.6M dataset, as opposed to building action-specific models. We consider both a supervised and an unsupervised variant. The *supervised* variant enhances the input to the model by concatenating one-hot vectors with the 15 action classes. In contrast, the *unsupervised* variant does not use one-hot input during training nor prediction. In these experiments we also train the network to minimize the prediction error over the next 400 milliseconds.

**Dataset and data pre-processing.** Following previous work, we use the Human 3.6M (H3.6M) dataset by Ionescu *et al.* [16], which is currently the largest publicly available dataset of motion capture data. H3.6M includes seven actors performing 15 varied activities such as walking, smoking, engaging in a discussion, taking pictures, and talking on the phone, each in two different trials. For a fair comparison, we adopt the pose representation and evaluation loss from [10, 18]. Pose is represented as an exponential map representation of each joint, with a special pre-processing of global translation and rotation (see [40] for more details). For evaluation, similar to [10, 18], we measure the Euclidean distance between our prediction and the ground truth in angle-space for increasing time horizons. We report the average error on eight randomly sampled test sequences, and use the sequences of subject five for testing, while the rest of the sequences are used for training.

**A scalable seq2seq architecture.** In all our experiments, we use a single gated recurrent unit [9] (GRU) with 1024 units, as a computationally less-expensive alternative to LSTMs, and we do not use any time-independent layers for representation learning. Experimentally, we found that stacking recurrent layers makes the architecture harder to train, while it also makes it slower; we also found that the best performance is obtained without a spatial encoder. We do, however, use a spatial decoder to back-project the 1024-dimensional output of the GRU to 54 dimensions, the number of independent joint angles provided in H3.6M.

We use a learning rate of 0.005 in our multi-action experiments, and a rate of 0.05 in our action-specific experiments; in both cases, the batch size is 16, and we clip the gradients to a maximum L2-norm of 5. During training as well as testing, we feed 2 seconds of motion to the encoder, and predict either 1 second (for long-term experiments) or 400 milliseconds (for short-term prediction) of motion from the decoder. We implemented our architecture using TensorFlow [1], which takes 75ms for forward processing and back-propagation per iteration on an NVIDIA Titan GPU.

**Baselines.** We compare against two recent approaches to human motion prediction based on deep RNNs: LSTM-3LR and ERD by Fragkiadaki *et al.* [10], and SRNN by Jain *et al.* [18]. To reproduce previous work, we rely on the pre-trained models and implementations of ERD, LSTM-3LR and SRNN publicly available<sup>1</sup>. These implementations represent the best efforts of the SRNN authors to reproduce the results of the ERD and LSTM-3LR models reported by Fragkiadaki *et al.* [10], as there is no official public implementation for that work. We found that, out of the box, these baselines produce results slightly different (most often better) from those reported by Jain *et al.* [18].

We also consider an agnostic zero-velocity baseline which constantly predicts the last observed frame. For completeness, we also consider running averages of the last two and four observed frames. While these baselines are very simple to implement, they have not been considered in recent work that uses RNNs to model human motion.

## 5. Results

Figure 3 shows a summary of the results obtained by ERD, LSTM-3LR and SRNN, as well as a zero-velocity baseline and our method, on four actions of the Human 3.6 dataset. Tables 1 and 2 describe these results in more detail, and include results on the rest of the actions. In the remainder of the section we analyze these results.

**Zero-velocity baseline.** The first striking result is the comparatively good performance of the baselines, specially

<sup>1</sup><https://github.com/asheshjain399/RNNexp>

milliseconds	Walking				Eating				Smoking				Discussion			
	80	160	320	400	80	160	320	400	80	160	320	400	80	160	320	400
ERD [10]	0.93	1.18	1.59	1.78	1.27	1.45	1.66	1.80	1.66	1.95	2.35	2.42	2.27	2.47	2.68	2.76
LSTM-3LR [10]	0.77	1.00	1.29	1.47	0.89	1.09	1.35	1.46	1.34	1.65	2.04	2.16	1.88	2.12	2.25	2.23
SRNN [18]	0.81	0.94	1.16	1.30	0.97	1.14	1.35	1.46	1.45	1.68	1.94	2.08	1.22	1.49	1.83	1.93
Running avg. 4	0.64	0.87	1.07	1.20	0.40	0.59	0.77	0.88	0.37	0.58	1.03	1.02	0.60	0.90	1.11	1.15
Running avg. 2	0.48	0.74	1.02	1.17	0.32	0.52	0.74	0.87	0.30	0.52	0.99	0.97	0.41	0.74	0.99	1.09
Zero-velocity	0.39	0.68	0.99	1.15	0.27	0.48	0.73	<u>0.86</u>	<u>0.26</u>	<u>0.48</u>	<u>0.97</u>	<u>0.95</u>	<b>0.31</b>	<b>0.67</b>	<b>0.94</b>	<b>1.04</b>
Zero noise (SA)	0.44	0.71	1.16	1.34	0.39	0.65	1.13	1.36	0.51	0.83	1.48	1.62	0.57	1.47	2.08	2.30
Sampling-based loss (SA)	0.92	0.98	1.02	1.20	0.98	0.99	1.18	1.31	1.38	1.39	1.56	1.65	1.78	1.80	1.83	1.90
Residual (SA)	0.34	0.60	0.95	1.09	0.30	0.53	0.92	1.13	0.36	0.66	1.17	1.27	0.44	0.93	1.45	1.60
Residual unsup. (MA)	<b>0.27</b>	<b>0.47</b>	<b>0.70</b>	<b>0.78</b>	<u>0.25</u>	<u>0.43</u>	<u>0.71</u>	0.87	0.33	0.61	1.04	1.19	<b>0.31</b>	0.69	1.03	1.12
Residual sup. (MA)	<u>0.28</u>	<u>0.49</u>	<u>0.72</u>	<u>0.81</u>	<b>0.23</b>	<b>0.39</b>	<b>0.62</b>	<b>0.76</b>	<b>0.23</b>	<b>0.39</b>	<b>0.62</b>	<b>0.76</b>	<b>0.31</b>	<u>0.68</u>	<u>1.01</u>	<u>1.09</u>
Untied (MA)	0.33	0.54	0.78	0.91	0.28	0.45	0.65	0.83	0.35	0.62	1.03	1.14	0.35	0.71	1.01	1.09

Table 1. Detailed results for motion prediction, measured in mean angle error for walking, eating, smoking and discussion activities of the Human 3.6M dataset. The top section corresponds to previous work based on deep recurrent neural networks. “Zero noise” is a model trained by feeding ground truth at each time step. “Sampling-based loss” is trained by letting the decoder feed its own output. **SA** stands for “Single action”, and **MA** stands for “Multi-action”. Finally “Untied” is the same model as Residual sup (MA), but with untied weights between encoder and decoder.

milliseconds	Directions				Greeting				Phoning				Posing				Purchases				Sitting			
	80	160	320	400	80	160	320	400	80	160	320	400	80	160	320	400	80	160	320	400	80	160	320	400
Zero-velocity	<b>0.25</b>	<b>0.44</b>	<b>0.61</b>	<b>0.68</b>	0.80	1.23	1.81	1.87	0.80	1.23	1.81	1.87	<b>0.32</b>	<b>0.63</b>	<b>1.16</b>	<b>1.45</b>	0.72	1.03	1.46	1.49	<u>0.43</u>	1.12	<b>1.41</b>	<b>1.58</b>
Res. (SA)	0.44	0.95	1.27	1.55	0.87	1.40	2.19	2.26	0.31	0.57	0.88	1.04	0.50	0.96	1.64	1.96	0.74	1.60	1.57	1.72	0.44	<u>1.05</u>	1.51	1.69
Res. unsup. (MA)	0.27	<u>0.47</u>	0.73	0.87	<u>0.77</u>	<u>1.18</u>	<b>1.74</b>	<u>1.84</u>	<u>0.24</u>	<b>0.43</b>	<b>0.68</b>	0.83	0.40	0.77	1.32	1.62	0.62	1.10	<b>1.07</b>	<b>1.14</b>	0.68	<b>1.04</b>	<u>1.43</u>	1.65
Res. sup. (MA)	<u>0.26</u>	<u>0.47</u>	<u>0.72</u>	<u>0.84</u>	<b>0.75</b>	<b>1.17</b>	<b>1.74</b>	<b>1.83</b>	<b>0.23</b>	<b>0.43</b>	<u>0.69</u>	<b>0.82</b>	<u>0.36</u>	<u>0.71</u>	<u>1.22</u>	<u>1.48</u>	<b>0.51</b>	<b>0.97</b>	<b>1.07</b>	<u>1.16</u>	<b>0.41</b>	<u>1.05</u>	1.49	<u>1.63</u>
Untied (MA)	0.31	0.52	0.77	0.89	0.79	1.19	1.72	<b>1.83</b>	0.27	0.46	0.68	0.85	0.42	0.77	1.29	1.58	<u>0.52</u>	<u>1.01</u>	<b>1.07</b>	<u>1.16</u>	0.51	1.13	1.56	1.74
milliseconds	Sitting down				Taking photo				Waiting				Walking Dog				Walking together				Average			
	80	160	320	400	80	160	320	400	80	160	320	400	80	160	320	400	80	160	320	400	80	160	320	400
Zero-velocity	<b>0.27</b>	<b>0.54</b>	<b>0.93</b>	<b>1.05</b>	<b>0.22</b>	<b>0.47</b>	<b>0.78</b>	<b>0.89</b>	<b>0.27</b>	<b>0.49</b>	<b>0.96</b>	<b>1.12</b>	0.60	0.96	1.27	<u>1.33</u>	<u>0.33</u>	<u>0.60</u>	0.96	1.03	0.42	0.74	1.12	<u>1.20</u>
Res. (SA)	<u>0.38</u>	<u>0.77</u>	<u>1.36</u>	<u>1.59</u>	0.37	0.66	1.30	1.70	0.36	0.73	1.31	1.51	0.62	1.02	1.55	1.65	0.44	0.81	1.25	1.36	0.46	0.88	1.35	1.54
Res. unsup. (MA)	0.41	0.80	1.43	1.63	0.27	0.56	0.98	1.16	0.32	0.62	1.13	1.30	<u>0.58</u>	<u>0.95</u>	<u>1.37</u>	1.45	0.35	0.62	<u>0.87</u>	<b>0.87</b>	<u>0.39</u>	<u>0.72</u>	<u>1.08</u>	1.22
Res. sup. (MA)	0.39	0.81	1.40	1.62	<u>0.24</u>	<u>0.51</u>	<u>0.90</u>	<u>1.05</u>	<u>0.28</u>	<u>0.53</u>	<u>1.02</u>	<u>1.14</u>	<b>0.56</b>	<b>0.91</b>	<b>1.26</b>	<b>1.40</b>	<b>0.31</b>	<b>0.58</b>	<b>0.87</b>	<u>0.91</u>	<b>0.36</b>	<b>0.67</b>	<b>1.02</b>	<b>1.15</b>
Untied (MA)	0.47	0.89	1.57	1.72	0.30	0.56	0.95	1.12	0.38	0.64	1.18	1.41	0.61	0.98	1.42	1.54	0.40	0.69	0.98	1.03	0.42	0.74	1.11	1.26

Table 2. Prediction results for our zero-velocity baseline and our main prediction methods on the remainder 11 actions of the H3.6m dataset.

the zero-velocity one. They clearly outperform state-of-the-art results, highlighting the severity of the discontinuities between conditioning and prediction in previous work. The good performance of the baseline also means that deterministic losses are not suitable to evaluate motion forecasting with a long time horizon.

**Sampling-based loss.** In Table 1, using our sampling-based loss consistently achieves motion prediction error competitive with or better than the state of the art. Moreover, since we have trained our model to minimize the error over a 1-second time horizon, the network retains the ability to generate plausible motion in the long run. Figure 4 shows a few qualitative examples of long-term motion using this approach. Given that our proposed sampling-based loss

does not require any hyper-parameter tuning, we would argue that this is a fast-to-train, interesting alternative to previous work for long-term motion generation using RNNs.

**Residual architecture and multi-action models.** Finally, we report the performance obtained by our architecture with sampling-based loss, residual connections and trained on single (SA) or multiple actions (MA) in the bottom subgroup of Table 1. We can see that using a residual connection greatly improves performance and pushes our method beyond the state of the art, which highlights the fact that velocity representations are easier to model by our network. Importantly, our method obtains its best performance when trained on multiple actions; this result, together with the simplicity of our approach, uncovers the

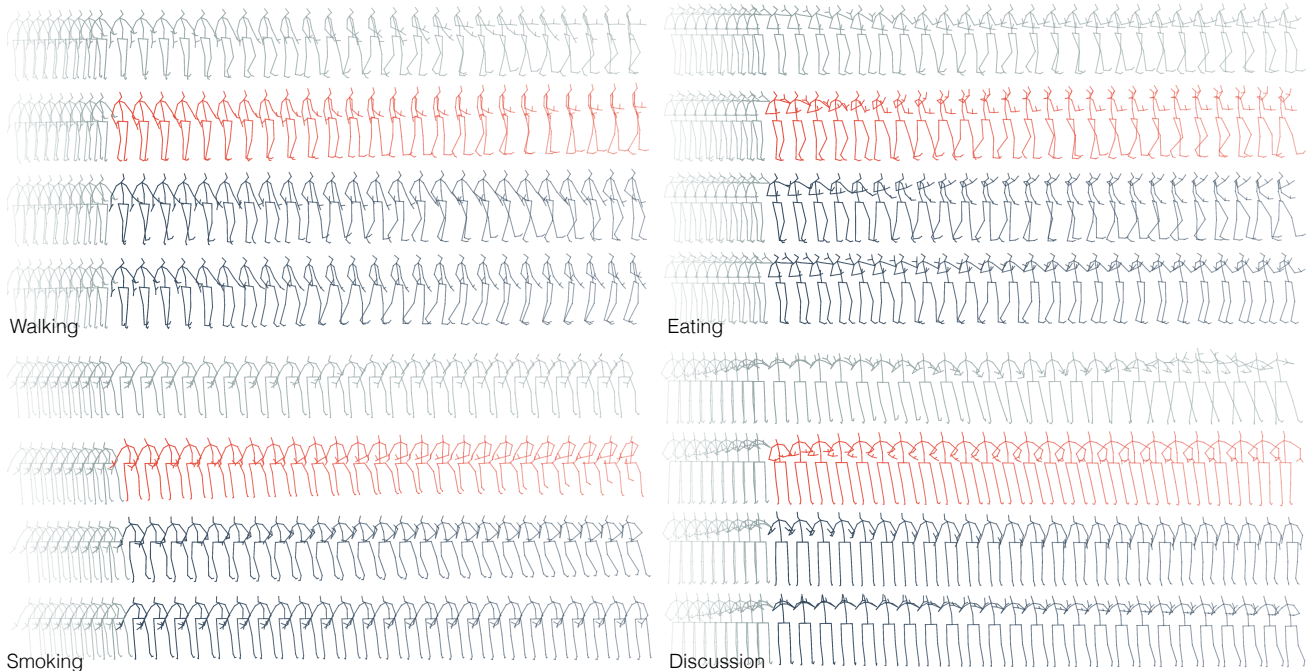


Figure 4. Qualitative long-term motion generation, showing two seconds of motion prediction on different activities. The gray top sequence corresponds to ground truth and the red one to SRNN. The first dark blue sequence corresponds to our method, trained on specific actions, and without residual connection, but using sampling-based loss (**Sampling-based loss (SA)** on Table 1). This model produces plausible motion in the long term, but does suffer from discontinuities in short-term predictions. The last blue sequence corresponds to our full model, including residual connections, and trained on multiple actions (**Residual sup. (MA)** on Table 1); this model produces smooth, continuous predictions in the short term, but converges to a mean pose.

importance of large amounts of training data when learning short-term motion dynamics. We also note that highly aperiodic classes such as discussion, directions and sitting down remain very hard to model.

Moreover, we observe that adding semantic information to the network in the form of action labels helps in most cases, albeit by a small margin. Likely, this is due to the fact that, for short-term motion prediction, modelling physical constraints (*e.g.* momentum preservation) is more important than modelling high-level semantic intentions.

When analysing Fig. 4, it becomes obvious that the best numerical results do not correspond to the best qualitative long-term motion – a result that persists even when trained to minimize loss over long horizons (*e.g.* 1 second). One can hardly blame the method though, since our network is achieving the lowest loss in an independent validation set. In other words, the network is excelling in the task that has been assigned to it. In order to produce better qualitative results, we argue that a different loss that encourages other similarity measures (*e.g.* adversarial, entropy-based *etc.*) should be used instead. Our results suggest that it is inherently hard to produce both accurate short-term predictions – which are relatively deterministic and seem to be properly optimized with the current loss – and long-term forecasting using RNNs.

## 6. Conclusions and future work

We have demonstrated that previous work on human motion modelling using deep RNNs has harshly neglected the important task of short-term motion prediction, as we have shown that a zero-velocity prediction is a simple but hard-to-beat baseline that largely outperforms the state of the art. Based on this observation, we have developed a sequence-to-sequence architecture with residual connections which, when trained on a sample-based loss, outperforms previous work. Our proposed architecture, being simple and scalable, can be trained on large-scale datasets of human motion, which we have found to be crucial to learn the short-term dynamics of human motion. Finally, we have shown that providing high-level supervision to the network in the form of action labels improves performance, but an unsupervised baseline is very competitive nonetheless. We find this last result particularly encouraging, as it departs from previous work in human motion modelling which has typically worked on small, action-specific datasets. Future work may focus on exploring ways to use even larger datasets of motion capture in an unsupervised manner.

**Acknowledgements.** The authors thank Laura Sevilla-Lara for proofreading our work. We also thank Nvidia for the donation of some of the GPUs used in this research. This research was supported in part by NSERC.



## References

- [1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. 6
- [2] O. Arikan and D. A. Forsyth. Interactive motion generation from examples. *TOG*, 21(3):483–490, 2002. 3
- [3] S. Arora, Y. Liang, and T. Ma. A simple but tough-to-beat baseline for sentence embeddings. *Under review at ICLR*, 2017. 2
- [4] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *ICLR*, 2015. 4
- [5] S. Bengio, O. Vinyals, N. Jaitly, and N. Shazeer. Scheduled sampling for sequence prediction with recurrent neural networks. In *NIPS*, 2015. 2, 4
- [6] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *ICML*, pages 41–48, 2009. 3
- [7] M. Brand and A. Hertzmann. Style machines. In *SIGGRAPH*, 2000. 3
- [8] M. A. Brubaker, D. J. Fleet, and A. Hertzmann. Physics-based person tracking using the anthropomorphic walker. *IJCV*, 87(1-2):140–155, 2010. 3
- [9] K. Cho, B. Van Merriënboer, D. Bahdanau, and Y. Bengio. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014. 2, 6
- [10] K. Fragkiadaki, S. Levine, P. Felsen, and J. Malik. Recurrent network models for human dynamics. In *CVPR*, 2015. 2, 3, 4, 5, 6, 7
- [11] A. Goyal, A. Lamb, Y. Zhang, S. Zhang, A. Courville, and Y. Bengio. Professor forcing: A new algorithm for training recurrent networks. In *NIPS*, 2016. 5
- [12] A. Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013. 5
- [13] A. Gupta, J. Martinez, J. J. Little, and R. J. Woodham. 3D Pose from Motion for Cross-view Action Recognition via Non-linear Circulant Temporal Encoding. In *CVPR*, 2014. 2
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [15] J. N. Ingram, K. P. Körding, I. S. Howard, and D. M. Wolpert. The statistics of natural hand movements. *Experimental Brain Research*, 188(2):223–236, 2008. 2
- [16] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human 3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *TPAMI*, 36(7), 2014. 2, 4, 6
- [17] A. Jabri, A. Joulin, and L. van der Maaten. Revisiting visual question answering baselines. In *ECCV*, 2016. 2, 3
- [18] A. Jain, A. R. Zamir, S. Savarese, and A. Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *CVPR*, 2016. 2, 3, 4, 5, 6, 7
- [19] O. C. Jenkins and M. J. Matarić. A spatio-temporal extension to isomap nonlinear dimension reduction. In *ICML*, 2004. 3
- [20] R. Kiros, Y. Zhu, R. R. Salakhutdinov, R. Zemel, R. Urtasun, A. Torralba, and S. Fidler. Skip-thought vectors. In *NIPS*, 2015. 4
- [21] H. S. Koppula and A. Saxena. Learning spatio-temporal structure from rgb-d videos for human activity detection and anticipation. In *ICML*, 2013. 2
- [22] L. Kovar, M. Gleicher, and F. Pighin. Motion graphs. *TOG*, 21(3), 2002. 2
- [23] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012. 3, 4
- [24] J. Lee, J. Chai, P. S. A. Reitsma, J. K. Hodgins, and N. S. Pollard. Interactive control of avatars animated with human motion data. *TOG*, 21(3):491–500, 2002. 3
- [25] A. M. Lehrmann, P. V. Gehler, and S. Nowozin. A non-parametric bayesian network prior of human pose. In *ICCV*, 2013. 2, 4
- [26] A. M. Lehrmann, P. V. Gehler, and S. Nowozin. Efficient nonlinear markov models for human motion. In *CVPR*, 2014. 3
- [27] S. Levine, J. M. Wang, A. Haraux, Z. Popović, and V. Koltun. Continuous character control with low-dimensional embeddings. *ACM Transactions on Graphics*, 31(4):28, 2012. 3
- [28] R. Li, M. H. Yang, S. Sclaroff, and T. P. Tian. Monocular tracking of 3D human motion with a coordinated mixture of factor analyzers. In *ECCV*, 2006. 3
- [29] J. Liu, A. Shahroudy, D. Xu, and G. Wang. Spatio-temporal lstm with trust gates for 3d human action recognition. In *ECCV*, 2016. 3
- [30] B. Mahasseni and S. Todorovic. Regularizing long short term memory with 3D human-skeleton sequences for action recognition. In *CVPR*, 2016. 3
- [31] M. Malinowski, M. Rohrbach, and M. Fritz. Ask your neurons: A neural-based approach to answering questions about images. In *ICCV*, 2015. 3
- [32] P. Pavlovic, J. M. Rehg, and J. MacCormick. Learning switching linear models of human motion. In *NIPS*, 2000. 2
- [33] L. L. Presti, M. L. Cascia, S. Sclaroff, and O. Camps. Hannelet-based dynamical systems modeling for 3D action recognition. *Image and Vision Computing*, 44:29–43, 2015. 3
- [34] N. Raman and S. Maybank. Action classification using a discriminative multilevel HDP-HMM. *Neurocomputing*, 154:149–161, 2015. 3
- [35] M. Ren, R. Kiros, and R. Zemel. Exploring models and data for image question answering. In *NIPS*, 2015. 3, 4
- [36] S. Ross, G. J. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *AISTATS*, 2011. 2, 4
- [37] H. Sidenbladh, M. J. Black, and D. J. Fleet. Stochastic tracking of 3D human figures using 2D image motion. In *ECCV*, 2000. 3

- [38] C. Sminchisescu and A. Jepson. Generative modeling for continuous non-linearly embedded visual inference. In *ICML*, 2004. 3
- [39] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *NIPS*, 2014. 4
- [40] G. W. Taylor, G. E. Hinton, and S. T. Roweis. Modeling human motion using binary latent variables. In *NIPS*, 2006. 3, 6
- [41] L. Theis, A. v. d. Oord, and M. Bethge. A note on the evaluation of generative models. In *ICLR*, 2016. 2
- [42] N. F. Troje. Decomposing biological motion: A framework for analysis and synthesis of human gait patterns. *Journal of Vision*, 2(5):2, 2002. 2
- [43] R. Urtasun, D. J. Fleet, and P. Fua. 3d people tracking with gaussian process dynamical models. In *CVPR*, volume 1, 2006. 2
- [44] R. Urtasun, D. J. Fleet, and P. Fua. Temporal motion models for monocular and multiview 3D human body tracking. *CVIU*, 103(2-3), 2006. 3
- [45] H. van Welbergen, B. J. H. van Basten, A. Egges, Z. Rutkay, and M. H. Overmars. Real time animation of virtual humans: A trade-off between naturalness and control. *Comput. Graph. Forum*, 29(8):2530–2554, 2010. 3
- [46] R. Vemulapalli, F. Arrate, and R. Chellappa. Human action recognition by representing 3D skeletons as points in a Lie group. In *CVPR*, 2014. 3
- [47] R. Vemulapalli and R. Chellappa. Rolling rotations for recognizing human actions from 3D skeletal data. In *CVPR*, 2016. 3
- [48] J. Wang, A. Hertzmann, and D. M. Blei. Gaussian process dynamical models. In *NIPS*, 2005. 2, 3
- [49] Y. Wang, P. Sabzmeydani, and G. Mori. Semi-latent dirichlet allocation: A hierarchical model for human action recognition. In *Workshop on Human Motion*, pages 240–254, 2007. 3
- [50] H. Xu and K. Saenko. Ask, attend and answer: Exploring question-guided spatial attention for visual question answering. In *ECCV*, 2016. 3
- [51] Y. Yacoob and M. J. Black. Parameterized modeling and recognition of activities. In *ICCV*, 1998. 3
- [52] A. Yao, J. Gall, L. J. V. Gool, and R. Urtasun. Learning probabilistic non-linear latent variable models for tracking complex activities. In *NIPS*, 2011. 3
- [53] B. Zhou, Y. Tian, S. Sukhbaatar, A. Szlam, and R. Fergus. Simple baseline for visual question answering. *arXiv preprint arXiv:1512.02167*, 2015. 3