

Anticipating many futures: Online human motion prediction and generation for human-robot interaction

Judith Bütepage, Hedvig Kjellström and Danica Kragic

Abstract—Fluent and safe interactions of humans and robots require both partners to anticipate the others’ actions. The bottleneck of most methods is the lack of an accurate model of natural human motion. In this work, we present a conditional variational autoencoder that is trained to predict a window of future human motion given a window of past frames. Using skeletal data obtained from RGB depth images, we show how this unsupervised approach can be used for online motion prediction for up to 1660 ms. Additionally, we demonstrate online target prediction within the first 300-500 ms after motion onset without the use of target specific training data. The advantage of our probabilistic approach is the possibility to draw samples of possible future motion patterns. Finally, we investigate how movements and kinematic cues are represented on the learned low dimensional manifold.

I. INTRODUCTION

Close interaction and collaboration with humans in shared workspaces requires robots to be able to coordinate their actions in space and time with their human partners. The robot’s actions need to be interpretable for the human partner to make the interaction intuitive. Similarly, human actions need to be identified on the macro level, e.g. grasping an object, and on the micro level, e.g. the exact trajectory and velocity of the movement. Online analysis of the current human motion is not sufficient to guarantee a fluent and safe interaction. Instead, the robot needs to be able to anticipate human motion and actions to be able to initiate actions in time and to render long-term planning possible.

A number of studies have addressed target prediction of human reaching movements which is often tackled with supervised trajectory classification [1], [2], [3], [4]. However, in most cases, these approaches only incorporate knowledge about the current arm configurations and do not anticipate human motion dynamics. Additionally, they depend on the training data and cannot generalize to novel targets. The bottleneck of an anticipatory approach is an accurate model of natural human motion. In order to make inferences about future movements, such a model is required to be able to distinguish between subtle kinematic cues about intentions which humans send during interaction [5], [6], [7].

Following concepts from neuroscience and psychology we address human motion modeling with a sensorimotor, bottom-up approach of internal forward models for action simulation and understanding in a social context [8]. Forward models in the common sense describe the sensory change induced by one’s own actions, see e.g. [9]. In a social context,

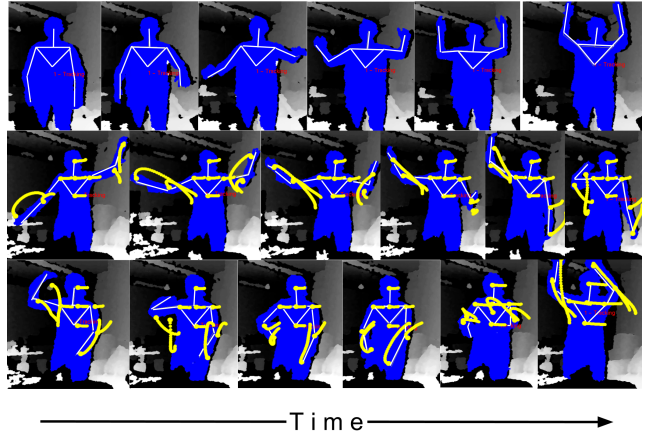


Fig. 1: Top row: As an example of training data we depict the depth image with the tracked upper body skeleton. Bottom rows: Given a window of past frames, we show the samples of future joint trajectories as yellow dots. These samples cover approximately 1660 ms of future motion. Each row depicts approximately 2 seconds of movements.

these models can be employed to predict the outcome of others’ actions, while treating the motor parameters of these actions as hidden variables that need to be approximated. As depicted in Figure 1, our data-driven approach allows to generate possible future trajectories of human motion given past observations. By focusing on generative models, we believe that a robotic system which integrates action generation and action understanding on a representational and modeling level will facilitate human-robot interaction.

A. Contributions and paper outline

This paper builds on our previous work [10] in which we demonstrated that an encoder-decoder model outperforms recurrent neural networks in human motion prediction tasks. In this work, we propose a model that is based on recent advances in deep generative models, namely variational autoencoders (VAE) [11]. By transforming the VAE into a conditional variational autoencoder (CVAE) [12], we are able to predict human skeletal motion extracted from RGB depth images such as depicted in the top row of Figure 1 online for up to 1660 ms. Without the need of motion capture systems and with an average processing time of 5 msec/frame, the predictions can be incorporated into online systems. CVAEs are well suited for predicting human motion as they can capture dynamic, non-linear dependencies under uncertainty. Furthermore, the Bayesian approach allows us to approximate the posterior distribution over hidden variables and a

predictive likelihood of trajectories under noisy observations. Thus, the CVAE provides us both with a measure of uncertainty over predictions and a possibility to draw samples of future motions given the current observed motion trajectory. The two bottom rows of Figure 1 depict such samples of possible upper body movements during online recordings. Each image shows a sampled trajectory of the joints for the next 1660 ms. We demonstrate how these inferences can be used to predict and sample general future human motion. Additionally, we show that the predictions made by our model can be used to determine the target position of reaching actions without the use of target specific training data. By investigating the learned representation, we find that legible and predictable motions [13] are automatically clustered by the model. The contribution of this work is a biologically inspired probabilistic forward model of human motion that is learned in an unsupervised manner from raw input data and can be employed in online human-robot interaction settings.

This paper is organized as follows. We begin the method section by presenting the data representation in Section III-A. As our model is based on variational autoencoders, we summarize the concepts behind VAEs in Section III-B before describing our temporal CVAE in Section III-C. Finally, we exemplify how the inferences made with help of the CVAE can be projected back into the context of the environment on a target prediction task in Section III-D. In the experiments, we shortly describe the data collection and model specifications in Section IV-A. Subsequently, we present the predictive performance of our approach in Section IV-B and the performance in end-point target classification in Section IV-C. Finally, we demonstrate the advantage of our probabilistic representation learning method by sampling future human motions in Section IV-D and by investigating how legible and predictable movements are represented on the learned low dimensional manifold in Section IV-E. We conclude the paper in Section V and outline future work.

II. RELATED WORK

Prediction of human activities in a HRI setting has been addressed by several groups. While the computer vision community has addressed the topic of skeletal human motion prediction with neural networks in a variety of works [14], [15], [16] in this section we will focus on related work with a HRI context. As our approach is trajectory based, we will also discard high-level approaches towards collaborative HRI such as affordance based coordination [17]. Among systems concerned with low-level dynamics, we can distinguish between supervised and unsupervised approaches. Supervised approaches commonly use information about the trajectory type, such as end-goal position, in order to estimate task specific models for prediction and classification. To account for differences in trajectory execution, in [1] the trajectories are aligned with help of dynamic time warping (DTW) before fitting multivariate Gaussian distributions to each time step. Task dependency is introduced by modeling the conditional likelihood of a trajectory given a motion class

and task. With help of this method, the goal of a reaching motion can be classified with 70 % accuracy within 500 ms after motion onset.

In addition to end-point prediction, the position of the human with respect to the robot needs to be taken into account. In order to estimate the area of the task space that will be occupied by the human during a goal-directed motion, Mainprice et al. [2] rely on Gaussian Mixture Models (GMM) and Gaussian Mixture Regression (GMR) to estimate trajectory and occupancy models of the human motion. Subsequently, these estimates are taken into account by the robot during action planning.

In a collaborative setting, the robot needs to consider the state of the human partner with respect to itself within the task space. In [18], interaction primitives are employed to anticipate human actions within task context. Interaction primitives [3] are dynamic motor primitives that model both the robot's and human's dynamics as well as the correlation between their actions in a specific task. By utilizing look-up tables of possible interaction sequences and by conditioning on the recent trajectories, [18] shows how the next human action can be anticipated and integrated into the robot's action selection process.

Unsupervised methods can be used to learn the dynamics of motion without explicitly defining high-level parameters that restrict the learning of structure. To accomplish this, in [4] GMMs and GMR are combined with a clustering algorithm to both learn the dynamics of the motions and to identify trajectories.

Finally, we include a work that is not concerned with 3D trajectory prediction but with general motion prediction in images as it is related by methodology and argumentation. Walker et al. [19] present a conditional variational autoencoder to predict the trajectory of all pixels in an image over a given time span. Given a certain image, the model can be used to generate future trajectories of all pixels. Being unsupervised, the information about e.g. human poses is not explicitly represented by the model. A similar idea relying on recurrent networks has been presented in [20]. Since these approaches are computationally exhaustive and do not operate in 3D, they are limited in their applicability to online HRI.

In comparison to the majority of the presented approaches, we do not rely on a model of the current state of the human body but anticipate future trajectories based on the past state. This allows us to infer an intended target without specific training data. As the evaluation of a neural network requires only feedforward matrix multiplications, it can be used in an online setting. The Bayesian approach is able to incorporate the noisy data received from the RGB depth images and to operate on a low frame rate. Except for [2], the approaches presented above [1], [3], [4] all rely on motion capture systems with higher frame rates.

III. METHOD

In this section we introduce a deep generative model for human motion prediction and generation that can be

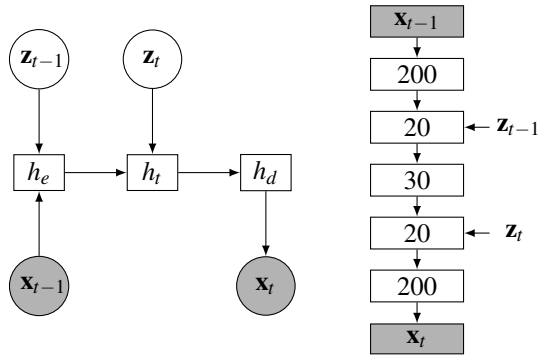


Fig. 2: Stochastic computational graph of the generative model CVAE. Circles represent random variables and squares represent parameterized mappings. Gray and white represent observed and unobserved entities respectively.

employed in online HRI tasks. The stochastic computational graph of this model is presented on the left in Figure 2. The main idea of our approach is to predict a window of frames at future time steps, given the last time steps. For this, we model a joint distribution over these two variables and a number of hidden variables which govern the unobserved dynamics. In Figure 2 the observed variables at time $t - 1$ and time t are given by \mathbf{x}_{t-1} and \mathbf{x}_t , while the hidden variables are represented by \mathbf{z}_{t-1} and \mathbf{z}_t . In order to allow for expressive models, we employ conditional variational autoencoders in which the encoder h_e , the inner layers h_t and the decoder h_d are represented by fully-connected neural networks with feedforward layers as shown on the right of Figure 2.

A. Data representation

We represent the human upper body as a vector of Cartesian coordinates x, y, z for each joint. Given a number of N_j joints and a time step t , this results in a matrix $\mathbf{f}_t = [f_{t,i,x}, f_{t,i,y}, f_{t,i,z}]_{i=1:N_j}$ of dimension $(N_j, 3)$. A time window covering the interval $[t : t + \Delta t]$ is given by the concatenation of frames $\mathbf{F}_{t:(t+\Delta t)} = [\mathbf{f}_t, \mathbf{f}_{t+1}, \dots, \mathbf{f}_{t+\Delta t}]$ of dimension $(\Delta t, N_j, 3)$. For simplicity, the past Δt frames are denoted by $\mathbf{F}_{t-\Delta t} := \mathbf{F}_{(t-\Delta t):t}$ and the future Δt frames are denoted by $\mathbf{F}_{t+\Delta t} := \mathbf{F}_{(t+1):(t+1+\Delta t)}$. Thus, the past observation \mathbf{x}_{t-1} as shown in Figure 2 is given by a vectorization of the matrix $\mathbf{F}_{t-\Delta t}$ and the next observation \mathbf{x}_t is given by a vectorization of the matrix $\mathbf{F}_{t+\Delta t}$. We will discuss the CVAE in terms of the more general \mathbf{x}_{t-1} and \mathbf{x}_t while using $\mathbf{F}_{t-\Delta t}$ and $\mathbf{F}_{t+\Delta t}$ in application specific sections.

B. Background of variational autoencoders

Generative models are characterized by a joint distribution $p_\theta(\mathbf{x}, \mathbf{z})$ over observed random variables \mathbf{x} and hidden random variables \mathbf{z} which can be characterized by parameters θ . We denote a distribution p that is parameterized by θ by p_θ and the i th observation of a random variable \mathbf{x} by \mathbf{x}_i . In order to make inferences over the hidden variables and the parameters, we require a tractable posterior $p_\theta(\mathbf{z}|\mathbf{x})$ for any instantiation of \mathbf{z} and \mathbf{x} which can be acquired through Bayes' theorem. However, the model evidence




Fig. 3: The joints of the skeleton used in this work.

$p_\theta(\mathbf{x}) = \int p_\theta(\mathbf{z})p_\theta(\mathbf{x}|\mathbf{z})d\mathbf{z}$ is often intractable. Variational autoencoders [11] have been proposed to circumvent this intractability by approximating the true posterior $p_\theta(\mathbf{z}|\mathbf{x})$ with help of a probabilistic recognition model, the encoder, $q_\phi(\mathbf{z}|\mathbf{x})$ with parameters ϕ . Likewise, the data likelihood $p_\theta(\mathbf{x}|\mathbf{z})$ can be modeled by a probabilistic decoder model. It can be shown that the marginal likelihood of sample \mathbf{x}_i can be approximated by maximizing the lower bound

$$\mathcal{L}(\phi, \theta, \mathbf{x}_i) = -D_{KL}(q_i||p_\theta(\mathbf{z})) + \mathbb{E}_{q_i}[\log(p_\theta(\mathbf{x}_i|\mathbf{z}))], \quad (1)$$

where q_i denotes $q_\phi(\mathbf{z}|\mathbf{x}_i)$, $D_{KL}(q||p)$ is the Kullback-Leibler divergence between the distributions q and p and $\mathbb{E}_{p_\theta}[f(\mathbf{x})]$ is the expectation of a function $f(\mathbf{x})$ under the distribution $p_\theta(\mathbf{x})$. Variational autoencoders model $q_\phi(\mathbf{z}|\mathbf{x}_i)$ with help of a differentiable transformation g_ϕ such that $\mathbf{z} \sim g_\phi(\boldsymbol{\varepsilon}, \mathbf{x}_i)$ where $\boldsymbol{\varepsilon}$ is a noise variable $\boldsymbol{\varepsilon} \sim p(\boldsymbol{\varepsilon})$. Using ideas from Stochastic Gradient Variational Bayes, the lower bound can then be approximated by averaging over S samples as follows

$$\hat{\mathcal{L}}(\phi, \theta, \mathbf{x}_i) = \frac{1}{S} \sum_{s=1}^S \log(p_\theta(\mathbf{x}_i, \mathbf{z}_i^s)) - \log(q_\phi(\mathbf{z}_i^s|\mathbf{x}_i)), \quad (2)$$

where the hidden variables are given by $\mathbf{z}_i^s = g_\phi(\boldsymbol{\varepsilon}_s, \mathbf{x}_i)$ and $\boldsymbol{\varepsilon}_s \sim p(\boldsymbol{\varepsilon})$. Modeling the encoder and decoder with neural networks allows for an expressive model that can describe non-linear mappings between the hidden and the observed variables while leveraging the advantages of the Bayesian approach. With help of the so called reparameterization trick the parameters θ and ϕ can be efficiently optimized by backpropagation.

C. Conditional variational autoencoders

While VAEs are pure autoencoders, i.e. the likelihood of the data depends only on a single observed variable, we are interested in predicting future time steps given past observations. Thus, we require a sequential generative model that encodes the dynamics of the time series. In order to allow for online evaluations of the model, we aim at a feedforward model under the markov assumption. We propose an approach similar to the ideas presented in [12] and [21] but do not consider recurrent connections and extensions. Instead, we assume that the generative process consists of three components: an encoder h_e , a transitioner h_t and a decoder h_d , depending on parameters θ^e , θ^t and θ^d respectively. In general terms, the encoding and decoding of an observation \mathbf{x}_{t-1} to predict the next observation \mathbf{x}_t is

described by

$$\text{latent variables } \mathbf{z}_t \sim \mathcal{N}(\mathbf{z}_t | \mathbf{0}, \mathbf{I}) \quad (3)$$

$$\text{encoded } \mathbf{e}_t = h_e(\mathbf{z}_{t-1}, \mathbf{x}_{t-1}, \theta^e) \quad (4)$$

$$\text{transitioned } \mathbf{t}_t = h_t(\mathbf{z}_t, \mathbf{e}_t, \theta^t) \quad (5)$$

$$\text{observation } \hat{\mathbf{x}}_t \sim p_{\theta^d}(\mathbf{x}_t | h_d(\mathbf{t}_t, \theta^d)). \quad (6)$$

At each time step t , the encoder maps the last observation \mathbf{x}_{t-1} and a sample from the prior to an encoded state \mathbf{e}_t , see Equation 4. This encoded state is passed through the transitioner to encode the dynamics in the hidden variables. The resulting state \mathbf{t}_t is decoded into the probability distribution over observations \mathbf{x}_t . A sample from this output distribution is denoted by $\hat{\mathbf{x}}_t$. Both the output of the encoder and the transitioner map the sample of the prior onto a non-linear feature space with help of a mean and covariance estimate. Note that the hidden sequence does only depend on the last observation and not the previous hidden states for $t' < t - 1$. Thus, this model is not a hidden markov or recurrent model but assumes conditional independence $\mathbf{x}_t \perp \mathbf{x}_{t'} | \mathbf{x}_{t-1} \quad \forall t' < t - 1$. The encoder-decoder structure facilitates the learning of a low dimensional manifold of the data dynamics. When fixing all parameters, the distribution of \mathbf{x}_t can be viewed as a conditional distribution $p_\lambda(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{z}_{t-1:t})$, where $\lambda = (\theta^e, \theta^t, \theta^d)$. In order to model continuous outputs, we define the output of the decoder $p_d(\mathbf{x}_t | h_d(\mathbf{t}_t, \theta^d))$ to be normal distributed with a mean μ^d and diagonal covariance matrix Σ^d which are functions of \mathbf{t}_t parameterized by θ^d . As $p_{\theta^d}(\mathbf{x}_t | h_d(\mathbf{t}_t, \theta^d))$ treats the output dimensions independently, we can make predictions for a single joint j , by considering only the corresponding dimensions.

D. Target inference of reaching motions

While our model is unsupervised, it is applicable to target prediction. For this, we do not gather target specific training data but use the predictive power of our model to infer future positions. Given a number of N_g , equally likely targets, let μ^g be the position of target $g \in [1, N_g]$ in Cartesian coordinates. Furthermore, let $\hat{\mathbf{f}}_{\tau,j} \sim \mathcal{N}(\mathbf{f}_{\tau,j} | \mu_{I_j}^d, \Sigma_{I_j}^d)$ be the predictive distribution of the j th joint at time $\tau \in [t + 1 : t + 1 + \Delta t]$, determined by $p_\lambda(\mathbf{f}_{\tau,j} | \mathbf{F}_{t-\Delta t}, \mathbf{z}_{t-1:t})$. Here I_j is the index set of all dimensions that correspond to joint j and $\hat{\mathbf{f}}_{\tau,j}$ indicates a sample of the distribution.

We model the location of a target g by a multivariate Gaussian distribution with mean μ^g and isotropic covariance function Σ^g . Thus, the likelihood of a single prediction $\hat{\mathbf{f}}_{\tau,j}$ at time τ of joint j being aimed at goal g is given by $p(\hat{\mathbf{f}}_{\tau,j} | g) \sim \mathcal{N}(\hat{\mathbf{f}}_{\tau,j} | \mu^g, \Sigma^g)$. The probability of a goal given an observation is therefore given by

$$p(g | \hat{\mathbf{f}}_{\tau,j}) p_\lambda(\hat{\mathbf{f}}_{\tau,j} | \mathbf{F}_{t-\Delta t}, \mathbf{z}_{t-1:t}) \propto p(\hat{\mathbf{f}}_{\tau,j} | g) p_\lambda(\hat{\mathbf{f}}_{\tau,j} | \mathbf{F}_{t-\Delta t}, \mathbf{z}_{t-1:t}) \\ \propto \mathcal{N}(\hat{\mathbf{f}}_{\tau,j} | \mu^g, \Sigma^g) \mathcal{N}(\hat{\mathbf{f}}_{\tau,j} | \mu_{I_j}^d, \Sigma_{I_j}^d). \quad (7)$$

For an entire sequence, we have

$$p(g | \hat{\mathbf{F}}_{t+\Delta t, j}) p_{\theta^d}(\hat{\mathbf{F}}_{t+\Delta t, j} | \mathbf{F}_{t-\Delta t}, \mathbf{z}_{t-1:t}) \\ \propto \prod_{\tau=t+1}^{t+1+\Delta t} p(g | \hat{\mathbf{f}}_{\tau, j}) p_{\theta^d}(\hat{\mathbf{f}}_{\tau, j} | \mathbf{F}_{t-\Delta t}, \mathbf{z}_{t-1:t}). \quad (8)$$

We compute the probability of target g by normalizing over all targets.

IV. EXPERIMENTS

In this section, we investigate the performance of the proposed temporal CVAE model. Our aim is to anticipate human motion in an online setting and to use these predictions to infer the human's intention. For HRI, the understanding of the future human actions is of advantage for safe and task-oriented action planning by the robot.

A. Data collection and model specifications

We collected skeletal data of the upper body with a Kinect sensor at 30 fps [22]. The upper body consists of nine joints as indicated and enumerated by Figure 3. In total, we collected 100 minutes of arbitrary, general movements. Among others, these movements consisted of walking, grasping, waving and dancing. Note that the types of movement in the training data set are irrelevant as long as they represent the dynamics of human motion sufficiently. We used 80 min of the recordings for training and separate 10 min for validating and testing respectively. The joints $I_j = [1 : 8]$ in each frame \mathbf{f}_{t, I_j} were translated into a local reference frame which was centered at the root joint $\mathbf{f}_{t, [0]}$. The index [...] indicates the index locations of the joints as enumerated in Figure 3. Furthermore, we normalized the limb length to be of length one. This data set was further processed into input and output data by concatenating frames of each recording r into sample pairs $\{\mathbf{F}_{t-\Delta t}, \mathbf{F}_{t+\Delta t}\}_{t \in [\Delta t: T_r - \Delta t]}$ where T_r is the length of recording r . We set $\Delta t = 50$, i.e. that each time window covers approximately 1660 ms.

As the structure of the human body is hierarchical, we use this prior information and train four separate models. First of all, in order to convert a prediction back into the original context, we require a model that predicts the translation of the root joint for every time step $I_{root} = [0]$. Thus, the first model is trained on the input-output pairs $\{\mathbf{F}_{t-\Delta t, I_{root}}, \mathbf{F}_{t+\Delta t, I_{root}}\}$. Furthermore, we train separate models for the main body, left arm and right arm with the index sets $I_{torso} = [1, 2, 3, 4]$, $I_{right} = [3, 5, 7]$ and $I_{left} = [4, 6, 8]$. For simplicity, let $limb \in [root, torso, right, left]$ denote any of these parts. Each of these models is equipped with a parameter set $\theta_{limb} = [\theta_{limb}^e, \theta_{limb}^t, \theta_{limb}^d]$ and we denote the set of all parameters by $\lambda = [\theta_{root}, \theta_{torso}, \theta_{right}, \theta_{left}]$.

The right in Figure 2 depicts the structure of our models. The hidden variables and the output of the decoder require estimates of both a mean and a covariance for each dimension. We train these models with the Adam optimizer [23] using a learning rate of 0.001 and mini-batches of size 1000. The models are trained until the log-likelihood converges.

After training the models, the prediction and generation of future motion is performed on an Intel(R) Core(TM) i7-4930K CPU 3.40GHz 64 bits machine, taking on average 5 msec per frame.

Linear baseline: In the following sections we are making

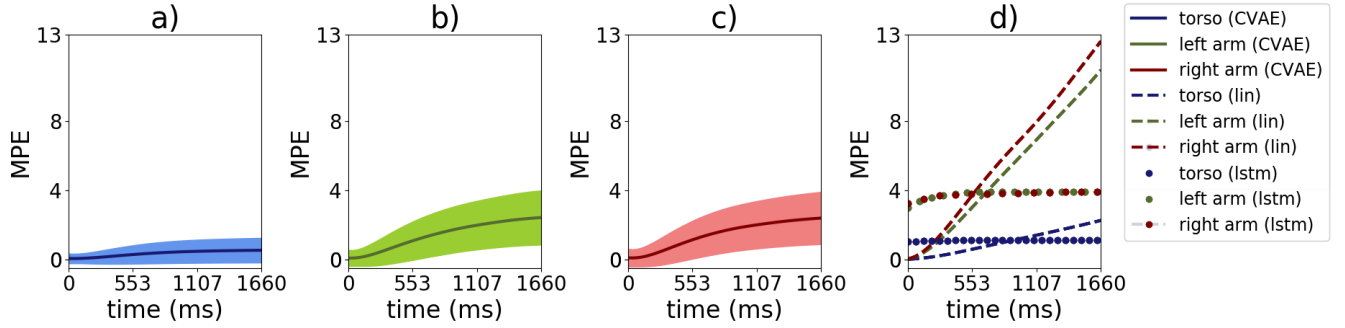


Fig. 4: The motion prediction error computed on the test set. The mean and variance of the prediction error are presented for the CVAE’s trained for the a) torso, b) left arm and c) right arm. The prediction error made by the baseline linear model and the baseline LSTM model are depicted for all three limbs separately in d).

use of a linear model for comparison. Under the assumption of constant velocity this model predicts future time steps by estimating the current velocity \mathbf{v} as the average velocity during the last 20 frames. The velocity \mathbf{v} is of dimension $(|I_{limb}|, 3)$. A prediction of τ steps into the future at time t is given by $\hat{\mathbf{f}}_{t+\tau, I_{limb}} = \mathbf{f}_{t, I_{limb}} + \tau \mathbf{v}$.

LSTM baseline: We also compare the prediction error to a recurrent neural network with the following structure: Input–FC(500)–FC(300)–LSTM(200)–FC(300)–FC(500)–Output. FC stands for fully connected layer and LSTM stands for Long Short Term Memory layer. As for the CVAE, we train a separate model for each limb on the same training data and the same time windows $\{\mathbf{F}_{t-\Delta t}, \mathbf{F}_{t+\Delta t}\}$. We apply a dropout rate of 0.2 to each layer.

B. Online human motion prediction

As our models are trained to predict future human motion, we present the predictive performance for a held out test data set containing 10 minutes of arbitrary upper body movements. All errors are computed based on normalized data samples. Let T_{test} be the number of frames of the test data. For each time $t \in [\Delta t : T_{test} - \Delta t]$ and each limb index set I_{limb} , we predict the next frames $\mathbf{F}_{t+\Delta t, I_{limb}}$ based on the

past frames $\mathbf{F}_{t-\Delta t, I_{limb}}$. For every time step t the decoder returns mean estimates $\mu_{t, I_{limb}}^d$ and variance estimates $\Sigma_{t, I_{limb}}^d$. We approximate the average motion prediction error (MPE) for each time step in the predicted time window Δt by taking the mean squared error over all test samples and summing over the joint and spatial dimensions as follows

$$MPE_{limb} = \sum_{j \in I_{limb}} \sum_{k \in [x, y, z]} \frac{1}{N} \sum_{t=\Delta t}^{T_{test}-\Delta t} (\mu_{t, I_{limb}}^d - \mathbf{F}_{t+\Delta t, I_{limb}})_{j, k}^2, \quad (9)$$

where $N = T_{test} - 2\Delta t$. Additionally, we report a variance estimate for each time step in the predicted time window Δt as the average sum of variances of the limb and spatial dimensions. In Figure 4 a)–c) we visualize the motion prediction errors of the torso, right arm and left arm model for the duration of 1660 ms. Since the skeleton is represented in a local reference frame, any natural movement of the torso is restricted to rotations. Therefore, the prediction error is comparatively low. The MPE for both arms is similar and grows more strongly than for the torso. The model expresses less uncertainty for over the initial predictions. Note that due to the stochasticity in human motion, an accurate longterm prediction (> 560 ms) is often not possible [24]. For HRI it is important to represent these uncertainties about motion predictions such that the robot can take these into account during motion planning. In comparison to our CVAE models, the LSTM models exhibit a discontinuity at the first prediction step, as also discussed in [16], which results in larger errors. The simple linear extrapolation in Figure 4 d) showcases the importance of modeling dynamics. While the initial error is comparable to the CVAE, the error increases rapidly after the initial frames.

C. Target prediction

As described in the related work, accurate online target prediction of human motion has been mostly addressed with GMMs in supervised settings. Although our approach is unsupervised, we investigate whether we can infer target locations based on predictions of our models. For this, we collected ten reaching trajectories of the right arm towards each of four different target locations. As depicted in Figure 7, these locations were approximately 25–30 cm apart as

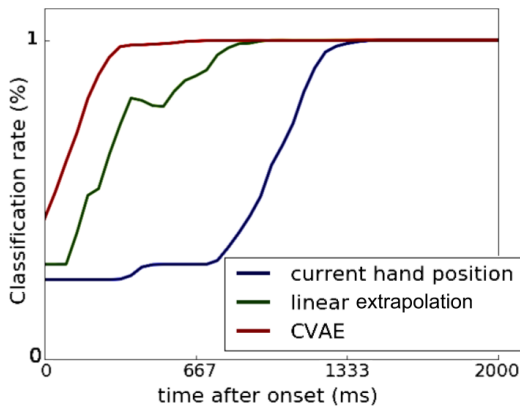


Fig. 5: End-point classification based on four different target locations. We compare the performance of the current hand position, a linear extrapolation and the CVAE model over time after movement onset.




Fig. 6: Five different samples of possible futures. The samples were generated by propagating the past motion window through the network, sampling from the encoder and transitioner distributions and visualizing the mean output of the decoder. We depict the past 800 ms and samples of the next 800 ms.

in [1] while the initial position was held constant. We classify the target location by using data and predictions of the right hand joint, i.e. $I_j = 7$. First, we benchmark our predictive approach on two different ideas, using the current hand position and linear extrapolation. The classification is performed as described in Section III-D by replacing the sample frame by the corresponding signal and setting the covariance to be the identity matrix. The initial predictions rely on the initial pose before motion onset.

To verify that anticipation does improve classification in the early stages of a reaching movement, we classify the target by using the current last twenty frames at each time step. As shown in Figure 5, this approach requires the hand to be at the target position, at around 1300 ms, for a 99.9 % correct classification rate. A simple linear extrapolation for twenty frames as described in Section IV-A, reaches 99.9 % within the first 600 ms after motion onset. Nevertheless, it appears that the dynamics learned by the CVAE are necessary for accurate prediction after around 300 ms.

Furthermore, we compare our method to the trajectory alignment method with probabilistic flow tubes (PFT) and GMMs presented in [1]. The emphasis in [1] lies on fast target prediction based on the current hand trajectory and relies on training data from four targets and three initial positions to classify the resulting twelve classes. Given a general model of human motion, the task of end-point classification is independent of the initial position, which renders the functionality of our experiment and the experiments presented in [1] to be the same. The results are

presented in Table I as the classification rate based on different percentages of the trajectory from the initial position until the goal is reached (approx 1300 ms). Our results indicate that anticipating future trajectories is of advantage as a simple linear extrapolation produces only slightly lower rates compared to GMMs [1]. The CVAE outperforms all methods and gives accurate classification rates within the first 200–300 ms after motion onset. Early inferences about the human’s intention are required for the robot to plan ahead. As our model performs in real time, these predictions can be integrated into an online system. Note that our models were trained on a different set of data, see Section IV-A, and have not been fine-tuned with trajectory specific examples.

D. Sampling future motion trajectories

The advantage of CVAEs over common autoencoders is that they model a probability distribution over future poses instead of a point estimate. Therefore, we can sample natural human motion from our model. In HRI it is of importance to be able to anticipate more than a single future. Even if the target of a human motion is fixed, the trajectories towards this goal can vary over time. The structure of the temporal CVAE allows us to sample from three different distributions, the encoder, the transitioner and the decoder. When propagating a window of past motion frames through the network at each of these layers, we can either use the mean estimate or a sample of the resulting distribution. When applying the mean in all layers, the output of the CVAE is a prediction. When sampling of at least one layer in the hierarchy, the output is a sample of the predictive distribution. To be of




Fig. 7: Experimental setup for target prediction. The four targets (green) are 25-30 cm apart.

Method	Classification rate			
	% of traj.	20	43	60
PFT [1]	73.26	89.55	95.76	98.06
GMM [1]	57.08	85.83	96.94	99.24
lin. extr.	53.47	79.25	95.4	99.91
CVAE	89.17	98.93	99.89	99.97

TABLE I: End-point classification rate compared to results presented in [1].




Fig. 8: Examples of the legible and predictable motion patterns to target (1) and (2).

value for the robot, these samples should be in accordance with the current body configuration and represent natural motion patterns. In Figure 6 we showcase five samples of future motion given a past sequence. These samples were generated by propagating the input through the three models corresponding to the torso, the left arm and the right arm and sampling at both the encoder and transitioner stage. The resulting samples are in accordance with the input data, while varying slightly from sample to sample. Online demonstrations of this method can be found in the at <https://www.youtube.com/watch?v=L93X3zh1sQo>.

E. Disentangling legibility & predictability of human motion

Due to the encoder–decoder structure of temporal CVAEs, these models do not only provide a mechanism to predict and generate natural human motion but they also learn a low dimensional manifold of the data. Ideally, this manifold should disentangle factors of variation. For example, when moving an arm up along a line, the underlying dynamics could be explained by a single dimension. When the body rotates and the arm movement is kept along the same line, the first dimension must not change, but an additional dimension needs to account for the rotation. The same principle should hold for more subtle kinematic cues such as legible and predictable motion. A *legible* motion lets an observer infer the goal given an observed trajectory, while a *predictable*

motion is the trajectory that an observer would expect given a goal [13]. Being able to distinguish these kinematic cues could facilitate a natural human-robot interaction. As discussed above, predictable and legible motion of a reaching movement should live in the same low dimensional space but be distinguishable within this space. A well separated low dimensional space would facilitate the robot’s decision making compared to noisy signals of high dimensional arm trajectories.

In this experiment, we investigate whether these theoretical assumptions hold true for the manifold learned by our model by scrutinizing the representation of legible and predictable motion. We recorded legible and predictable reaching movements towards two targets, (1) and (2), with five recordings each, as exemplified in Fig. 8. The targets were positioned an arm length from the left (1) and the right (2) shoulder. In the lower right corner of Fig. 9 we depict the position of the right hand for all trajectories. Note, that we base all results on whole arm trajectories over time which corresponds to 450 dimensions for a single time window. To visualize the learned representation, we apply principal component analysis to the output of the right arm encoder layer of all testing samples. In Fig. 9 we show the projection of 2000 samples onto the two main components. Furthermore, we project the neural output corresponding to the legible and predictable motion trajectories in different colors.

It becomes apparent that the movement types towards goal (2) are well separated while the movements towards target (1) are overlapping. Intuitively, the legible movement towards (2) reaches rather naturally away from the body and towards the target, see legible (2) in Fig. 8. Thus, both the hand and the elbow follow a different trajectory than during a straight reach as shown by predictable (2) in Fig. 8. Since target (1) lies to the left, the legible movements require an unnatural turn of the arm along the body, see legible (1) in




Fig. 9: PCA of the activity of neurons in the inner layer of the model corresponding to the right arm. The majority of dots represents random test samples. The colored dots visualize the identity of testing samples corresponding to the reaching trajectories. The picture in the lower right depicts the x position of the right hand during legible and predictable movements towards two targets.

Traj. type	Classification rate				
% of traj.	20	43	60	80	100
legible (1)	73.89	99.99	100	100	100
predictable (1)	18.13	63.69	93.78	97.02	100
legible (2)	100	100	100	100	100
predictable (2)	95.79	89.96	94.7	97.9	100

TABLE II: End-point classification rate for legible and predictable motion trajectories to target (1) and (2).

Fig. 8. As the elbow does not significantly deviate from the predictable trajectory, shown by predictable (1) in Fig. 8, the two movements are less distinguishable in the latent space. An additional factor could be that the model has not learned to represent unnatural human motion as this is not contained in the training data.

In order to verify whether legible motion has a positive impact on target inference, we classify the end-location of all movement types following the procedure described in Sec. III-D. As Table II shows, legible motion indeed facilitates target inference for both targets. Especially during the initial phase of the movement, legibility seems to be of high importance.

V. CONCLUSIONS AND FUTURE WORK

In this work we presented a probabilistic deep learning approach to online human motion prediction for HRI. To overcome the drawbacks of using only the current body configurations for classification, we predict human motion for up to 1660 ms. Target inference of reaching movements based on these predictions are accurate within 300-500 ms after motion onset. Our experiments suggest that our CVAE models are able to capture the underlying dynamics of human motion. The Bayesian representation learning approach provides a method for sampling future trajectories and to investigate the underlying low dimensional structure of kinematic cues. Although this approach is data-intensive during training, the use of skeletal data extracted from RGB depth images allows for low-cost recordings of natural, everyday movements. Due to the short processing time, the anticipatory signal can be integrated into the robot's action planning.

Instead of top-down approaches such as explicit cost functions or target specific training data, our approach is a bottom-up, data-driven model. The sensorimotor theory of predictive motion understanding in humans suggests that anticipation of sensory changes is important for online social interaction. Importantly, these two approaches should not be viewed as opposing forces but as complementary methods. We see our approach as a first step into this direction. For future work, we are envisioning two paths. Firstly, we would like to include context information into the prediction in those cases when it is available. Secondly, we are aiming at integrating our system into online HRI tasks to investigate the importance of anticipating many futures in physical human-robot interaction.

ACKNOWLEDGMENT

This work was supported by the EU through the project socSMCs (H2020-FETPROACT-2014) and the Swedish Foundation for Strategic Research.

REFERENCES

- [1] C. Pérez-D'Arpino and J. A. Shah, "Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 6175–6182.
- [2] J. Mainprice and D. Berenson, "Human-robot collaborative manipulation planning using early prediction of human motion," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013, pp. 299–306.
- [3] G. Maeda, A. Maloo, M. Ewerton, R. Lioutikov, and J. Peters, "Anticipative interaction primitives for human-robot collaboration," in *2016 AAAI Fall Symposium Series*, 2016.
- [4] R. Luo and D. Berenson, "A framework for unsupervised online human reaching motion recognition and early prediction," in *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*. IEEE, 2015, pp. 2426–2433.
- [5] L. Sartori, C. Becchio, B. G. Bara, and U. Castiello, "Does the intention to communicate affect action kinematics?" *Consciousness and cognition*, vol. 18, no. 3, pp. 766–772, 2009.
- [6] L. M. Sacheli, E. Tidoni, E. F. Pavone, S. M. Aglioti, and M. Candidi, "Kinematics fingerprints of leader and follower role-taking during cooperative joint actions," *Experimental brain research*, vol. 226, no. 4, pp. 473–486, 2013.
- [7] C. Scorolli, M. Miatton, L. A. Wheaton, and A. M. Borghi, "I give you a cup, i get a cup: a kinematic study on social intention," *Neuropsychologia*, vol. 57, pp. 196–204, 2014.
- [8] G. Pezzulo and H. Dindo, "What should i do next? using shared representations to solve interaction problems," *Experimental Brain Research*, vol. 211, no. 3-4, pp. 613–630, 2011.
- [9] M. Haruno, D. M. Wolpert, and M. Kawato, "Mosaic model for sensorimotor learning and control," *Neural computation*, vol. 13, no. 10, pp. 2201–2220, 2001.
- [10] J. Bütepage, M. Black, D. Kragic, and H. Kjellström, "Deep representation learning for human motion prediction and classification," 2017.
- [11] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.
- [12] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," in *Advances in Neural Information Processing Systems*, 2015, pp. 3483–3491.
- [13] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2013, pp. 301–308.
- [14] G. W. Taylor, G. E. Hinton, and S. T. Roweis, "Modeling human motion using binary latent variables," in *Advances in Neural Information Processing Systems*, 2006.
- [15] A. Jain, A. R. Zamir, S. Savarese, and A. Saxena, "Structural-RNN: Deep learning on spatio-temporal graphs," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [16] J. Martinez, M. J. Black, and J. Romero, "On human motion prediction using recurrent neural networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017, pp. 4674–4683.
- [17] H. S. Koppula, A. Jain, and A. Saxena, "Anticipatory planning for human-robot teams," in *Experimental Robotics*. Springer, 2016, pp. 453–470.
- [18] H. B. Amor, G. Neumann, S. Kamthe, O. Kroemer, and J. Peters, "Interaction primitives for human-robot cooperation tasks," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 2831–2837.
- [19] J. Walker, C. Doersch, A. Gupta, and M. Hebert, "An uncertain future: Forecasting from static images using variational autoencoders," in *European Conference on Computer Vision*. Springer, 2016, pp. 835–851.
- [20] C. Finn, I. Goodfellow, and S. Levine, "Unsupervised learning for physical interaction through video prediction," in *Advances In Neural Information Processing Systems*, 2016, pp. 64–72.
- [21] D. Rezende, I. Danihelka, K. Gregor, D. Wierstra *et al.*, "One-shot generalization in deep generative models," in *Proceedings of The 33rd International Conference on Machine Learning*, 2016, pp. 1521–1529.
- [22] Z. Zhang, "Microsoft kinect sensor and its effect," *IEEE multimedia*, vol. 19, no. 2, pp. 4–10, 2012.
- [23] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [24] K. Fragkiadaki, S. Levine, P. Felsen, and J. Malik, "Recurrent network models for human dynamics," in *IEEE International Conference on Computer Vision*, 2015.