

Learning Multimodal Latent Dynamics for Human-Robot Interaction

Vignesh Prasad¹, Lea Heitlinger², Dorothea Koert^{3,4}
 Ruth Stock-Homburg², Jan Peters^{4,5,6,7}, Georgia Chalvatzaki^{1,7}

Abstract—This article presents a method for learning well-coordinated Human-Robot Interaction (HRI) from Human-Human Interactions (HHI). We devise a hybrid approach using Hidden Markov Models (HMMs) as the latent space priors for a Variational Autoencoder to model a joint distribution over the interacting agents. We leverage the interaction dynamics learned from HHI to learn HRI and incorporate the conditional generation of robot motions from human observations into the training, thereby predicting more accurate robot trajectories. The generated robot motions are further adapted with Inverse Kinematics to ensure the desired physical proximity with a human, combining the ease of joint space learning and accurate task space reachability. For contact-rich interactions, we modulate the robot's stiffness using HMM segmentation for a compliant interaction. We verify the effectiveness of our approach deployed on a Humanoid robot via a user study. Our method generalizes well to various humans despite being trained on data from just two humans. We find that users perceive our method as more human-like, timely, and accurate and rank our method with a higher degree of preference over other baselines. We additionally show the ability of our approach to generate successful interactions in a more complex scenario of Bimanual Robot-to-Human Handovers.

Index Terms—Physical Human-Robot Interaction, Learning from Demonstration, Humanoid Robots, Probability and Statistical Methods

I. INTRODUCTION

ENSURING a synchronized and accurate reaction is an important aspect in Human-Robot Interaction (HRI) [2]. To do so, a human and a robot need to spatiotemporally coordinate their movements to reach a common target or perform a common task, that can be denoted as a *joint action* [79]. For realizing such coordinated jointly performed actions, spatial and temporal adaptation of motions in the shared physical space to accurately perform the action are key factors [80]. Such interpersonal coordination can enable a connection to one's interaction partner [57]. Therefore,

¹ Interactive Robot Perception and Learning group (PEARL), Department of Computer Science, TU Darmstadt, Germany.

² Chair for Marketing and Human Resource Management, Department of Law and Economics, TU Darmstadt, Germany.

³ Interactive AI Algorithms & Cognitive Models for Human-AI Interaction (IKIDA), Department of Computer Science, TU Darmstadt, Germany.

⁴ Centre for Cognitive Science, TU Darmstadt, Germany.

⁵ Institute for Intelligent Autonomous Systems (IAS), Department of Computer Science, TU Darmstadt, Germany.

⁶ Systems AI for Robot Learning, German Research Center for AI (DFKI).

⁷ Hessian Center for Artificial Intelligence (Hessian.AI), Darmstadt, Germany. Contact: vignesh.prasad@tu-darmstadt.de

Website: <https://sites.google.com/view/mild-hri>

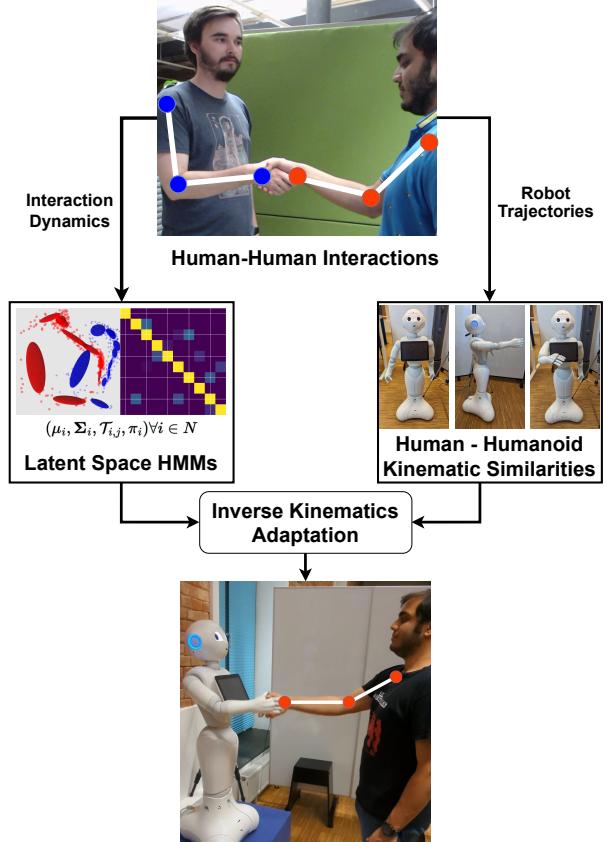


Fig. 1: We explore learning coordinated HRI behaviors using Hidden Markov Models (HMMs) to learn the interaction dynamics over a representation space spanned by a Variational Autoencoder (VAE). The VAEs are trained with the HMMs as a prior to better incorporate the understanding of the dynamics. Moreover, we learn directly from Human-Human interactions by leveraging the kinematic similarities of Humans and Humanoid robots. During testing, the predicted trajectories are adapted with Inverse Kinematics to mitigate errors in the spatial accuracy arising either from prediction errors or from shortcomings in transferring the human's actions to the robot.

well-executed and well-coordinated interactive behaviors can improve the perception of a social robot.

The paradigm of Learning from Demonstrations (LfD) is promising for HRI by learning joint distributions over human and robot trajectories in a modular, multimodal manner [15], [28], [29], [45], [70]. However, LfD approaches scale poorly with higher dimensions, which can be circumvented by incorporating Deep Learning for learning latent-space dynamics. Such Deep State-Space Models have shown good performance in capturing temporal dependencies of latent space trajectories, either using some kind of a forward propagation model [8],

[13], [23], [25], learning the parameters for a dynamics model with LfD [22], [43] or fitting parameterized LfD models to the underlying latent trajectories [27], [60]. In this paper, we further explore this direction of latent space dynamics in the context of Human-Robot Interaction.

Furthermore, when learning trajectories in the joint configuration space of a robot, minor deviations in the joint space can cause a perceivable deviation in the robot's task space. While learning task space trajectories can mitigate problems from errors in the joint space, to do so, the demonstrated trajectories need to fall within the reachability of the robot. When this reachability assumption is violated, the task space proximity in HRI scenarios may not be effectively learned. Moreover, learning purely task space trajectories doesn't necessarily ensure human-like joint configurations, which is especially relevant for Humanoid robots. To this end, the combination of joint space learning and task space reachability can be achieved using Inverse Kinematics to supplement the learned behaviors [18], [35], [72].

In this paper, as shown in Figure 1, we are interested in learning and adapting Deep LfD policies for spatiotemporally coherent interactive behaviors. We learn the interaction dynamics from HHI demonstrations using latent space Hidden Markov Models (HMMs). We show the efficacy of the learnt dynamics in real-world HRI scenarios. We do so in a manner that ensures spatially accurate physical proximity to the human partner and additionally enable compliant robot motions in contact-rich interactions like a handshake, thereby improving the perceived quality of the HRI behaviors.

A. Related Work

1) *Learning HRI from Demonstrations*: Early approaches for learning modular HRI policies modeled the interaction as a joint distribution with a Gaussian Mixture Model (GMM) learned over demonstrated trajectories of a human and a robot in a collaborative task [15]. The correlations between the human and the robot degrees of freedom (DoFs) can then be leveraged to generate the robot's trajectory given observations of the human for learning both proactive and reactive controllers [14], [70], [76]. Segmenting HRI demonstrations has also been shown using Graphical Models with Markov chain Monte Carlo [84], [85].

Along the lines of leveraging Gaussian approximations for LfD, Movement Primitives [68], which learn a distribution over underlying linear regression weight vectors, were extended for HRI by similarly learning a joint distribution over the weights of interacting agents [4], [16], [54]. Movement Primitive approaches can further be combined with GMMs for learning multiple task sequences seamlessly [29], [45], [50], [55], [65]. One drawback of the aforementioned primitive-based approaches is that they do not perform well on out-of-distribution data [96]. Additionally, for tasks that are loosely coupled in time, Movement Primitives can underperform [67]. These drawbacks therefore make them an unviable option for being able to generalize spatiotemporally in interactive behav-

iors. For a more extensive overview of Movement Primitive approaches for robot learning, we refer the reader to [90].

Vogt et al. [93] are similar in theory to our approach where they learn the interaction dynamics of Human-Human Interactions using an HMM over a low dimensional representation space of the human skeletons. In our approach, we show further performance improvements by incorporating human-conditioned reactive motion generation into the training pipeline. Additionally, Vogt et al. [93] train Interaction Meshes [41] for transferring the learned trajectories to a robot whose prediction is optimized during runtime to adapt to the human user's movements. As we work with Humanoid Robots, we follow a simpler approach by leveraging the structural similarity between a human and a humanoid. This allows us to map the joint motions directly [31] and adapt the motions to the interaction partner using Inverse Kinematics [72]. Through our user study, we find that our approach, while being simplistic, provides acceptable interactions with various human partners.

To make the learned distribution more robust, LfD methods are often learned via kinesthetic teaching which can be tedious for HRI tasks where one would need a variety of human partners. The need for extensive training data can potentially be circumvented by learning how humans adapt to a robot's trajectory [17]. A more general way, especially in the context of Humanoid robots, is by learning from human-human demonstrations by leveraging the kinematic similarities between humans and robots [31]. While such approaches of motion retargeting leave some room for error due to minor differences between human and robot geometries, adapting the trajectories in the robot's task space via Inverse Kinematics can improve the accuracy of such interactive behaviors [72], [92].

2) *Integrating LfD with Deep Learning*: Techniques at the intersection of LfD coupled with the use of Neural Network representations for higher dimensional data have grown in popularity for learning latent trajectory dynamics from demonstrations. Typically, an autoencoding approach, like VAEs, is used to encode latent trajectories over which a latent dynamics model is trained. In their simplest form, the latent dynamics can be modeled either with linear Gaussian models [43] or Kalman filters [8]. Other approaches learn stable dynamical systems, like Dynamic Movement Primitives [78] over VAE latent spaces [10], [20]–[22], [24]. Instead of learning a feed-forward dynamics model, Dermy et al. [27] model the entire trajectory's dynamics at once using Probabilistic Movement Primitives [68] achieving better results than [20].

Nagano et al. [60] demonstrated the use of Hidden Semi-Markov Models (HSMMs) as latent priors in a VAE for temporal action segmentation of motions involving a single human. They model each latent dimension independently, which is not favorable when learning interaction dynamics. To extend such an approach for learning HRI, one needs to consider the interdependence between dimensions. In our previous work, “MILD” [71], we extend this idea of using HSMMs as VAE priors for interactive tasks by exploiting the full rank of HSMM covariance matrices thereby capturing the dependencies between both agents by learning a joint distribution over the latent trajectories of interacting partners.

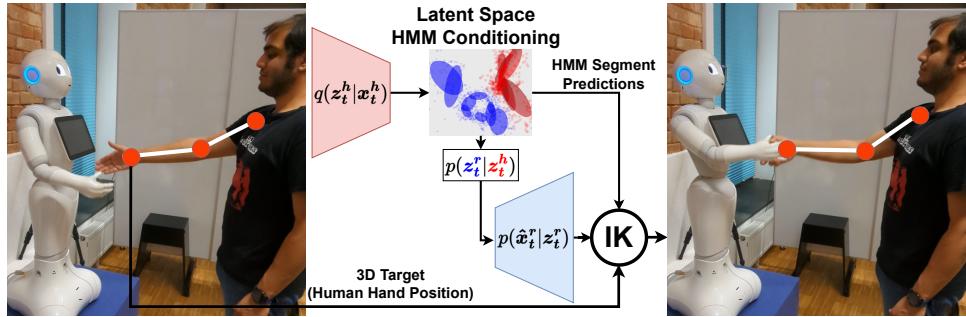


Fig. 2: Overview of our approach. We train VAEs to reconstruct the observations of the interacting agents ($\mathbf{x}_{1:t}^h, \mathbf{x}_{1:t}^r$) with an HMM prior to learn a joint distribution over the latent space trajectories $p(\mathbf{z}_{1:t}^h, \mathbf{z}_{1:t}^r)$ of the interacting agents. During test time, the observed agent's latent trajectory conditions the HMM to infer the robot agent's latent trajectory $p(\mathbf{z}_t^r | \mathbf{z}_{1:t}^h)$ which is decoded to generate the robot agent's joint trajectory $\hat{\mathbf{x}}_t^r$. To ensure the proximity of a robot's hand to the human's hand during test time, we additionally adapt the predicted trajectory using Inverse Kinematics to fulfill the contact-based nature of the interaction.

3) Interaction Modeling with Recurrent Neural Networks:

When large datasets are available, Recurrent latent space models are powerful tools in approximating latent dynamics with some form of a forward propagating distribution [23], [30], [37], [46], [51]. Given their power of modeling temporal sequences, they yield themselves naturally for learning interactive/collaborative tasks in HRI [64], [81], [95]. To simplify the various scenarios of interaction tasks Oguz et al. [64] develop an ontology to categorize interaction scenarios and train an LSTM network for each case in a simple Imitation Learning paradigm. Similarly, Zhao et al. [95] explore using such LSTM-based policies for learning HRI in a simple collaboration scenario. Rather than simply regressing robot actions to human inputs, Bütepage et al. [13] explore the idea of learning shared latent dynamics in HRI in a systematic way. Bütepage et al. [13] initially learn latent embeddings of the human and robot motions individually using a VAE. Once the embeddings are learned, the latent interaction dynamics model of Human-Human Interactions (HHI) is learned using LSTMs. The learned HHI dynamics are used in conjunction with the learned robot embeddings to subsequently learn the robot dynamics from HRI demonstrations. However, we find in our experiments that the autoregressive nature of their approach, unfortunately, leads to a divergence in the performance since they only train it using ground truth demonstrations and not in an autoregressive manner. In contrast, rather than having an implicit shared representation, such as with an LSTM, the LfD approaches can be explicitly conditioned on the interaction partner's observations [4], [28] which we find leads to improved predictive performance.

4) Unsupervised Skill Discovery:

While skill discovery is not the main focus of our approach, there are some parallels between our approach and works on unsupervised skill discovery which we discuss below. In general, such approaches aim to partition a latent space into different skills which can then be sequenced for learning a variety of different tasks either from demonstrations or using Reinforcement Learning (an overview can be found in Sec. III-B in [59]). The works closest to our approach [61], [74], [89] explore this key idea of learning how to segment demonstrations into underlying recurring segments or “skills” by decomposing a latent representation of the trajectories and subsequently learning the temporal relations

between these skills to compose the observed trajectories. Similar approaches are further explored in literature using different skill representations such as the Options framework [83] or Autoregressive HMMs [48], [63]. While the main focus of such approaches is the decomposition of demonstrations to learn re-usable atomic “skills”, our main aim is to explore how to do so in HRI settings via a joint distribution over the actions of interacting partners and subsequently how the learned skills/segments can be inferred during test time based on the dynamic observations of a human, rather than reproducing a given task or based on a static goal.

B. Objectives and Contributions

In our previous work [71] on learning Multimodal Interactive Latent Dynamics, “MILD”, we explored how Deep LfD can be used for learning spatiotemporally coherent interaction dynamics. Rather than using an uninformed, stationary prior for learning HRI [13], we showed the use of Hidden Semi-Markov Models to model the latent space prior of a VAE as a joint distribution over the trajectories of both interacting agents. During test time, the robot trajectories can be conditionally generated based on the observed human trajectories in the latent space, showing that such an approach can capture the latent interaction dynamics suitably well. In this article, we extend MILD [71] by systematically improving the predictive abilities of our model by incorporating reactive motion generation into the training process. We do so by exploring two different sampling approaches to incorporate the conditional distribution of the HMM into the training, rather than using the HMM just as a VAE prior as done in [71]. Furthermore, to deploy the proposed approach on the Humanoid Social Robot Pepper [66], we leverage the similarity between the kinematic structures of Pepper and a Human to effectively learn HRI behaviors from Human-Human Interactions. We further show how the predictions from the latent space HMMs can help supplement the robot controller by modulating the target motions via Inverse Kinematics for a more accurate interaction and modulating the joint stiffnesses for ensuring compliant motions during a contact-rich interaction like a handshake (Figure 2).

Specifically, our contributions beyond MILD [71] are summarized below:

- (i) We systematically explore how to incorporate human-conditioned reactive motion generation of the robot trajectories in the training to improve the prediction accuracy of our model.
- (ii) We integrate Inverse Kinematics in our pipeline to spatially adapt the predicted robot motions to the interaction partner, thereby ensuring suitable physical proximity and improved spatial accuracy of the learned behaviors.
- (iii) For contact-rich interactions like handshaking, we perform stiffness modulation using the HMM's segment predictions to enable realistic and compliant robot motions.
- (iv) We validate our approach via a user study where participants interact with a Pepper robot controlled by our approach. We find that our method ranks highly compared to other baselines and provides a perceivably more human-like, natural, timely, and accurate interaction, thereby showcasing its effectiveness. We further show the applicability of our approach on a more complex task of Bimanual Robot-to-Human Handovers.

The rest of the paper is organized as follows. In Section II, we explain the foundations needed to understand our work. We then explain our approach in Section III. We highlight our experiments and results in Section IV and present some concluding remarks and directions for future work in Section V.

II. FOUNDATIONS

We first explain Variational Autoencoders (VAEs) (Section II-A) which is the main backbone of our approach, followed by Hidden Markov Models (HMMs) (Section II-B), which are key for learning the interaction dynamics, and finally, we give an introduction to Inverse Kinematics (IK) (Sec II-C) which is essential for improving the overall acceptance of our approach.

A. Variational Autoencoders

Variational Autoencoders (VAEs) [44], [75] are a type of neural network architecture that learns the identity function in an unsupervised, probabilistic way. The inputs “ x ” are encoded into lower dimensional latent space embeddings “ z ” that a decoder uses to reconstruct the original input. A prior distribution is enforced over the latent space during the learning, which is typically a standard normal distribution $p(z) = \mathcal{N}(z; \mathbf{0}, \mathbf{I})$. The goal is to estimate the true posterior $p(z|x)$, using a neural network $q(z|x)$ and is trained by minimizing the Kullback-Leibler (KL) divergence between them

$$KL(q(z|x)||p(z|x)) = \mathbb{E}_q[\log \frac{q(z|x)}{p(x, z)}] + \log p(x) \quad (1)$$

which can be re-written as

$$\log p(x) = KL(q(z|x)||p(z|x)) + \mathbb{E}_q[\log \frac{p(x, z)}{q(z|x)}]. \quad (2)$$

The KL divergence is always non-negative, therefore the second term in Eq. 2 acts as a lower bound. Maximizing it would effectively maximize the log-likelihood of the data

distribution or evidence, and is hence called the Evidence Lower Bound (ELBO), which can be written as

$$\mathbb{E}_q[\log \frac{p(x, z)}{q(z|x)}] = \mathbb{E}_q[\log p(x|z)] - \beta KL(q(z|x)||p(z)). \quad (3)$$

The first term aims to reconstruct the input via samples decoded from the posterior. The second term is the KL divergence between the prior and the posterior, which regularizes the learning. To prevent over-regularization, the KL divergence term is also weighted down with a factor β . Further information can be found in [38], [44], [75].

B. Hidden Markov Models

A Hidden Markov Model (HMM) is used to model a sequence of observations $z_{1:T}$ (robot joint angles, human skeletons, etc.) as a sequence of underlying hidden states such that they can “emit” the observations with a given probability. In mathematical terms, an HMM is characterized by a set of hidden states $i \in \{1, 2 \dots N\}$, each of which denotes a probability distribution, in our case a Gaussian with mean μ_i and covariance Σ_i , which characterize the emission probabilities of observations $\mathcal{N}(z_t; \mu_i, \Sigma_i)$. An initial state distribution π_i denotes the initial probabilities of being in each state, and the state transition probabilities $\mathcal{T}_{i,j}$ describe the probability of the model going from the i^{th} state to the j^{th} state. The sequential progression via the probability of each hidden state given an observed sequence $z_{1:t}$ is denoted by the forward variable of the HMM $\alpha_i(z_t)$

$$\begin{aligned} \alpha_i(z_t) &= \frac{\hat{\alpha}_i(z_t)}{\sum_{j=1}^N \hat{\alpha}_j(z_t)} \\ \hat{\alpha}_i(z_t) &= \mathcal{N}(z_t; \mu_i, \Sigma_i) \sum_{j=1}^N \alpha_j(z_{t-1}) \mathcal{T}_{j,i} \\ \hat{\alpha}_i(z_0) &= \pi_i \mathcal{N}(z_0; \mu_i, \Sigma_i) \end{aligned} \quad (4)$$

where $\hat{\alpha}_i(z_t)$ represents the non-normalized forward variable and π_i is the initial state distribution.

The HMM is trained using Expectation-Maximization over the parameters $(\pi_i, \mu_i, \Sigma_i, \mathcal{T}_{j,i})$ with the trajectory data. We refer the reader to [14], [70] for further details on HMMs in the context of robot learning.

In HRI scenarios, to encode the joint distribution between the human and the robot, we concatenate the Degrees of Freedom (DoFs) of both the human and the robot [15], [28], allowing the distribution to be decomposed as

$$\mu_i = \begin{bmatrix} \mu_i^h \\ \mu_i^r \end{bmatrix}; \Sigma_i = \begin{bmatrix} \Sigma_i^{hh} & \Sigma_i^{hr} \\ \Sigma_i^{rh} & \Sigma_i^{rr} \end{bmatrix} \quad (5)$$

where the superscript indicates the different agents (h -human, r -robot). Once the distributions are learned, given some obser-

vations of the human agent \mathbf{z}_t^h , the robot's trajectory can be conditionally generated using Gaussian Mixture Regression as

$$\mathbf{K}_i = \Sigma_i^{rh} (\Sigma_i^{hh})^{-1} \quad (6)$$

$$\hat{\boldsymbol{\mu}}_i^r = \boldsymbol{\mu}_i^r + \mathbf{K}_i (\mathbf{z}_t^h - \boldsymbol{\mu}_i^h) \quad (7)$$

$$\hat{\Sigma}_i^r = \Sigma_i^{rr} - \mathbf{K}_i \Sigma_i^{hr} + \hat{\boldsymbol{\mu}}_i^r (\hat{\boldsymbol{\mu}}_i^r)^T \quad (8)$$

$$\hat{\boldsymbol{\mu}}_t^r = \sum_{i=1}^N \alpha_i (\mathbf{z}_t^h) \hat{\boldsymbol{\mu}}_i^r \quad (9)$$

$$\hat{\Sigma}_t^r = \sum_{i=1}^N \alpha_i (\mathbf{z}_t^h) \hat{\Sigma}_i^r - \hat{\boldsymbol{\mu}}_t^r (\hat{\boldsymbol{\mu}}_t^r)^T \quad (10)$$

$$p(\mathbf{z}_t^r | \mathbf{z}_t^h) = \mathcal{N}(\mathbf{z}_t^r | \hat{\boldsymbol{\mu}}_t^r, \hat{\Sigma}_t^r) \quad (11)$$

where $\alpha_i(\mathbf{z}_t^h)$ is the forward variable calculated using the marginal distribution of the observed human agent.

C. Inverse Kinematics

Given a joint angle configuration \mathbf{y} , the end effector position can be calculated as $\mathbf{x}_{ee} = f(\mathbf{y})$ where $f(\mathbf{y})$ denotes the forward kinematics of the robot, which is based on the robot's geometry and calculates the end effector position through the hierarchy of intermediate transformations of each of the individual joints. For example, given the geometry of one's arm sizes, and the angles of the shoulder, elbow, and wrist joints, the position of the hand can be calculated through the position of the shoulder to the elbow to the hand.

Given a list of n joints $\mathbf{y} = \{y_1, y_2 \dots y_n\}$ and the corresponding relative transformations that the joints represent in the robot geometry $\mathbf{T}_1^0(y_1), \mathbf{T}_2^1(y_2) \dots \mathbf{T}_n^{n-1}(y_n)$ where $\mathbf{T}_j^i \in SE(3)$ is the relative transformation between links i and j , the end effector pose \mathbf{T}_{ee} can be calculated as

$$\mathbf{T}_{ee} = \begin{bmatrix} \mathbf{R}_{ee} & \mathbf{x}_{ee} \\ \mathbf{0} & 1 \end{bmatrix} = \mathbf{T}_n^0 = \prod_{i=1}^n \mathbf{T}_i^{i-1}(y_i) \quad (12)$$

where $\mathbf{x}_{ee}, \mathbf{R}_{ee}$ are the end effector position and rotation.

As the name suggests, Inverse Kinematics (IK) is the reverse process of estimating the robot's joint angles from a given end effector pose, which can be solved via optimization. We aim to estimate $\mathbf{y} = f^{-1}(\mathbf{x}_{ee})$, which can be solved by finding the optimal joint configuration that minimizes the distance between the expected and predicted positions

$$\mathbf{y}^* = \arg \min_{\mathbf{y}} \|f(\mathbf{y}) - \mathbf{x}_{ee}\|^2. \quad (13)$$

III. MULTIMODAL INTERACTIVE LATENT DYNAMICS

In this section, we introduce our overall approach, which can be seen in Figure 3. We start by explaining the key idea of our previous work MILD [71], which enhances the representation learning abilities of VAEs with HMMs as the prior distribution for modeling the interaction dynamics (Section III-A). After introducing our previous work, we go through the key improvements on top of MILD, namely incorporating the human-conditioned HMM distribution for reactive motion generation into the training process in Section III-B followed by the adaptation of the learned trajectories via Inverse Kinematics for improving the generated robot

motions as well as the incorporation of stiffness control in Section III-C. For HRI scenarios, we use the superscripts h and r to denote the human and the robot variables respectively. For consistency, in the HHI scenarios, we use h and r to denote the variables of the first and second human partners respectively.

A. Learning Interaction Dynamics using HMMs as VAE priors

Typically in VAEs, the prior $p(\mathbf{z})$ is modeled as a stationary distribution. When it comes to learning trajectory dynamics, having meaningful priors can help learn temporally coherent latent spaces [22]. To this end, we explore using HMMs to learn latent-space dynamics in a modular manner by breaking down the trajectories into multiple phases and learning the sequencing between them. Since the HMMs learn a joint distribution over the latent trajectories of both interacting agents, we can conditionally generate the motion of the robot after observing the human agent. We do so by using the HMM distribution as the prior for the VAE. We then update the HMM at the end of each epoch by running expectation-maximization on the VAE embeddings, alternating between training the VAE and then subsequently updating the HMM.

The VAE prior at each time-step is calculated as the respective marginal distribution of the most likely HMM component at that timestep, given by the HMM forward variable (Eq. 4). However, the recurrent nature of estimating the forward variable, and consequently its gradients, leads to numerical instabilities from backpropagation through time. Hence, we use an approximation in the form of an unobserved forward variable by setting the likelihood term in Eq. 4 to unity. This unobserved forward variable provides a good approximation of the sequential progression of the hidden states based on the learned transitions, which can be written as

$$\begin{aligned} \bar{\alpha}_i^t &= \frac{\hat{\alpha}_i^t}{\sum_{j=1}^N \hat{\alpha}_j^t} & \hat{\alpha}_i^t &= \sum_{j=1}^N \bar{\alpha}_j^{t-1} \mathcal{T}_{j,i} & \hat{\alpha}_i^0 &= \pi_i \\ i_t^* &= \arg \max_i \bar{\alpha}_i^t \end{aligned} \quad (14)$$

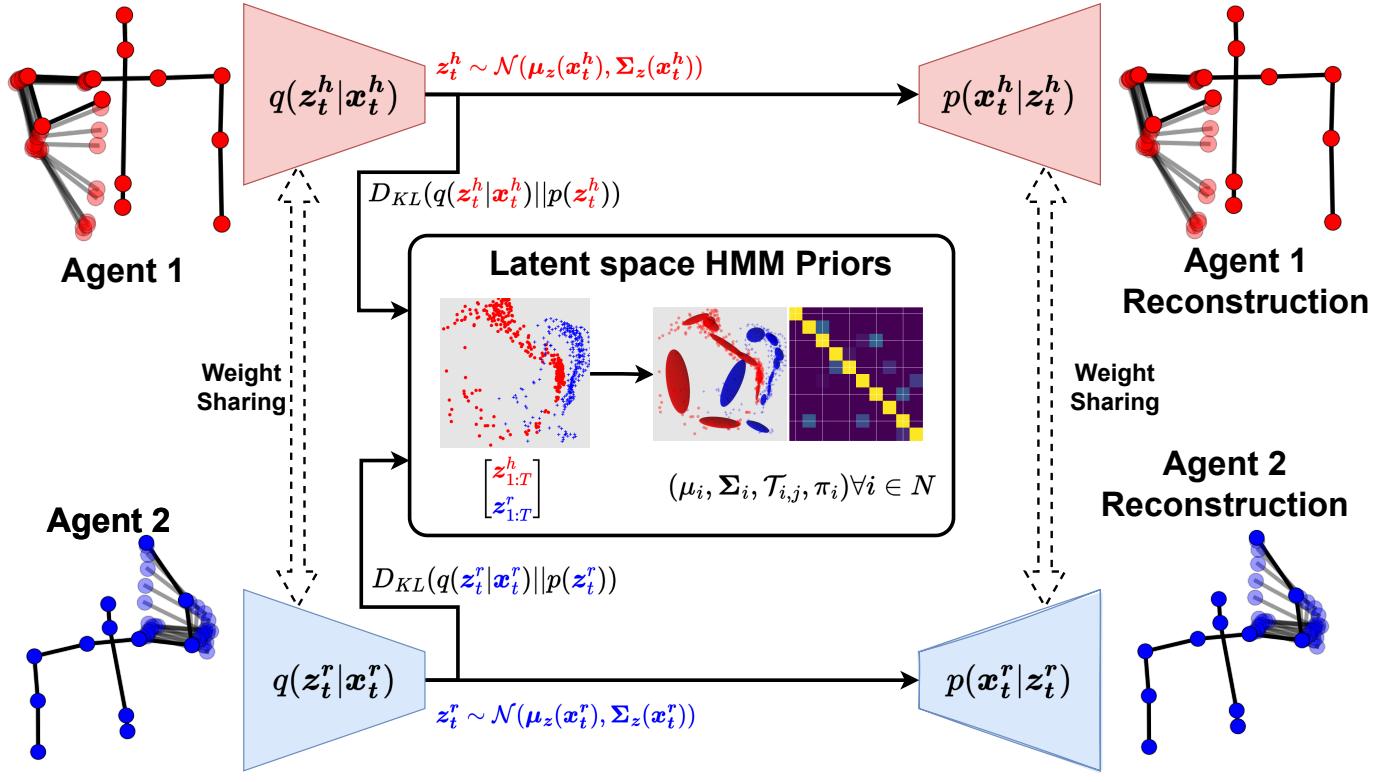
$$KL_t^h = KL(q(\mathbf{z}_t^h | \mathbf{x}_t^h) || \mathcal{N}(\mathbf{z}_t^h; \boldsymbol{\mu}_{i_t^*}^h, \Sigma_{i_t^*}^{hh}))$$

$$KL_t^r = KL(q(\mathbf{z}_t^r | \mathbf{x}_t^r) || \mathcal{N}(\mathbf{z}_t^r; \boldsymbol{\mu}_{i_t^*}^r, \Sigma_{i_t^*}^{rr}))$$

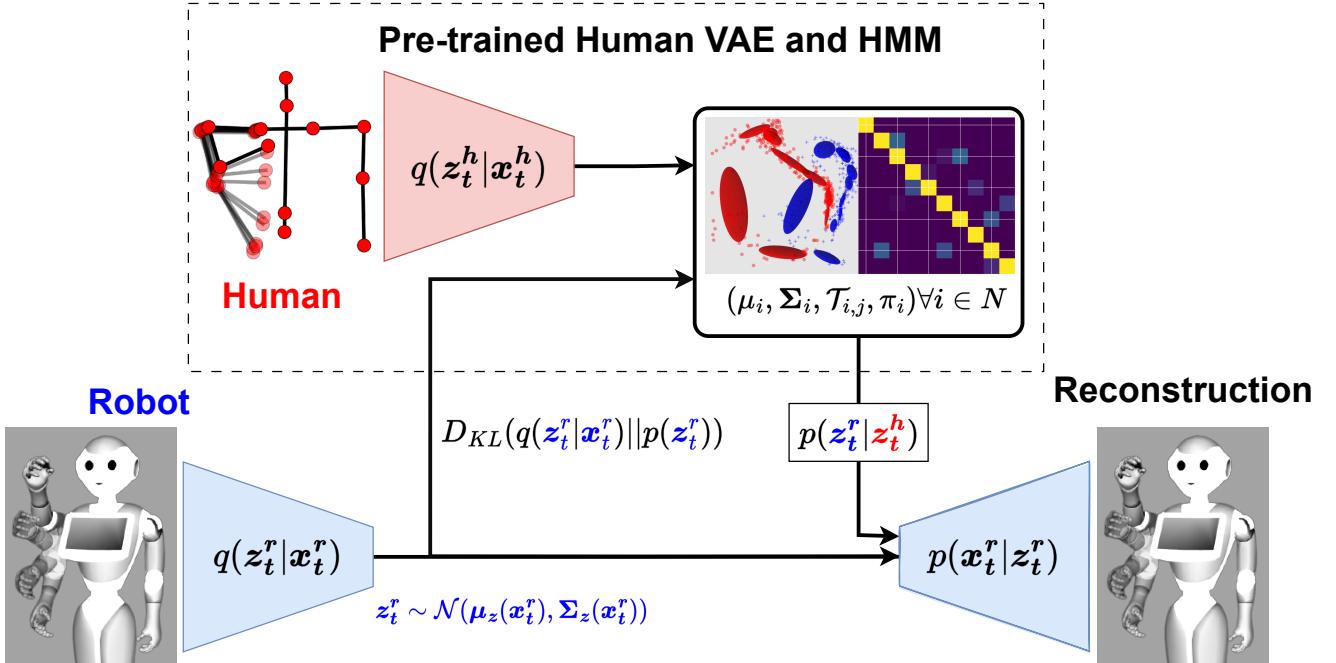
where $\mathcal{N}(\mathbf{z}_t^h; \boldsymbol{\mu}_{i_t^*}^h, \Sigma_{i_t^*}^{hh})$ and $\mathcal{N}(\mathbf{z}_t^r; \boldsymbol{\mu}_{i_t^*}^r, \Sigma_{i_t^*}^{rr})$ are the marginal distributions of the most likely HMM component i_t^* at the given timestep t for each agent. The ELBO can then be reformulated as

$$\begin{aligned} \mathbb{E}_q [\log \frac{p(\mathbf{x}_t^h, \mathbf{x}_t^r, \mathbf{z}_t^h, \mathbf{z}_t^r)}{q(\mathbf{z}_t^h, \mathbf{z}_t^r | \mathbf{x}_t^h, \mathbf{x}_t^r)}] &= \mathbb{E}_q [\log p(\mathbf{x}_t^h | \mathbf{z}_t^h)] \\ &+ \mathbb{E}_q [\log p(\mathbf{x}_t^r | \mathbf{z}_t^r)] \\ &- \beta (KL_t^h + KL_t^r). \end{aligned} \quad (15)$$

The overall training procedure, shown in Algorithm 1 and Figure 3a, is as follows. We train the VAE to reconstruct the input trajectory and use the marginal distributions of each agent from the HMM (Eq. 14) to regularize the latent posterior distribution. During a given epoch in the VAE training, the



(a) Learning Interaction Models from Human-Human Interactions.



(b) Using Human-Human Interaction models for Human-Robot Interaction models.

Fig. 3: Overview of our training approach. We model the interaction dynamics in the latent space of a VAE using HMMs to model a joint distribution over the latent trajectories of both agents. The use of HMMs for the prior, as opposed to a unimodal Gaussian, enforces a multimodal modularized latent space with the HMM forward variable incorporating the learned transitions between the multiple modes. We first learn the interaction dynamics from Human-Human demonstrations by training both the VAEs and HMMs (Figure 3a). We subsequently use the HMM prior learned from Human-Human demonstrations for learning the latent dynamics for HRI. We do so by regularizing the robot VAEs and, additionally, by training the robot decoder to reconstruct samples from the HMM's conditional distribution after observing the human partner (Figure 3b). This further incorporates the interdependence between the human and the robot during training, thereby improving the predicted robot trajectories.

HMM parameters are fixed. After an epoch, the new HMM parameters are estimated using the learned encodings $\mathbf{z}_{1:T}^h, \mathbf{z}_{1:T}^r$ of the training trajectories. The HMMs are then fixed as the VAE prior for the next epoch. Currently, we learn a separate HMM per interaction. We defer learning the HMM selection via activity recognition to future work. For further details on training HMMs with expectation-maximization, or the use of HMMs in robot learning, we refer the reader to [14], [26].

Algorithm 1: Learning Latent Dynamics from Human-Human Interactions

Data: A set of trajectories with action labels
 $\mathbf{X} = \{\mathbf{X}_{1:T}^h, \mathbf{X}_{1:T}^r, c\}$ for $|\mathcal{C}|$ interactions

Result: VAE weights and $|\mathcal{C}|$ HMM parameters

Initialize VAE weights randomly

```

for  $c \in [1, |\mathcal{C}|]$  do
  for  $i \in [1, N]$  do
     $\mu_i^c \leftarrow \mathbf{0}$ 
     $\Sigma_i^c \leftarrow \mathbf{I}$ 
  end
end
while not converged do
  for  $\mathbf{x}_{1:T}^h, \mathbf{x}_{1:T}^r, c \in \mathbf{X}$  do
    Compute VAE Posterior  $q(\mathbf{z}_t^h | \mathbf{x}_t^h)$  and  $q(\mathbf{z}_t^r | \mathbf{x}_t^r)$ 
    Reconstruct posterior samples
    Maximize ELBO (Eq. 15) to update VAE weights
  end
  for  $c \in [1, |\mathcal{C}|]$  do
     $\mathbf{X}^c \leftarrow$  set of demonstrations of Interaction  $c$ 
     $\mathbf{Z}^c \leftarrow \emptyset$ 
    for  $\mathbf{x}_{1:T}^h, \mathbf{x}_{1:T}^r, c \in \mathbf{X}^c$  do
       $\mathbf{z}_{1:T}^h \sim q(\cdot | \mathbf{x}_{1:T}^h); \mathbf{z}_{1:T}^r \sim q(\cdot | \mathbf{x}_{1:T}^r)$ 
       $\mathbf{Z}^c \leftarrow \mathbf{Z}^c \cup \begin{bmatrix} \mathbf{z}_{1:T}^h \\ \mathbf{z}_{1:T}^r \end{bmatrix}$ 
    end
    Train the  $c^{th}$  HMM with  $\mathbf{Z}^c$ 
  end
end

```

B. Conditional Training of HRI Dynamics from HHI

Based on the initial idea of [71] presented above in Section III-A, in this section, we first explore how the HMMs learned from the Human-Human demonstrations can be used to regularize learning the robot motions, after which we highlight the incorporation of the conditional generation of the robot motions. The HMMs learned from Human-Human demonstrations capture the overall latent interaction dynamics between two agents. These HMMs can, therefore, be used as an informative prior to learn robot motions to perform the given interaction. Hence, we use the marginal distribution of the second agent from the HMMs trained on the HHI demonstrations as the latent space prior for the robot VAE. Although we regularize the VAEs with the HMM marginals (Eq. 14), during test time the decoder would see samples from the conditional distribution of the HMM after observing the human agent (Eq. 6-11) which would be finitely divergent from what the decoder would be trained to reconstruct in a normal autoencoding approach.

When the output spaces of both interaction partners are similar, such as in the HHI scenarios, sharing the weights of

the VAEs additionally enables the decoder to learn to reconstruct the target distribution. However, given the difference in output spaces of the human and the robot in HRI scenarios, we instead train the robot decoder to additionally reconstruct samples from conditional distribution (Figure 3b).

Moreover, the VAE provides a confidence estimate of the posterior probability, which can additionally be incorporated into the conditional distribution $p(\mathbf{z}_t^r | \mathbf{z}_t^h)$ as

$$\mathbf{K}_i = \Sigma_i^{rh} (\Sigma_i^{hh} + \Sigma_{\mathbf{z}}(\mathbf{x}_t^h))^{-1} \quad (16)$$

$$\hat{\mu}_i^r = \mu_i^r + \mathbf{K}_i (\mu_{\mathbf{z}}(\mathbf{x}_t^h) - \mu_i^h) \quad (17)$$

$$\hat{\Sigma}_i^r = \Sigma_i^{rr} - \mathbf{K}_i \Sigma_i^{hr} + \hat{\mu}_i^r (\hat{\mu}_i^r)^T \quad (18)$$

$$\hat{\mu}_t^r = \sum_{i=1}^N \bar{\alpha}_i^t \hat{\mu}_i^r \quad (19)$$

$$\hat{\Sigma}_t^r = \left[\sum_{i=1}^N \bar{\alpha}_i^t \hat{\Sigma}_i^r \right] - \hat{\mu}_t^r (\hat{\mu}_t^r)^T \quad (20)$$

$$p(\mathbf{z}_t^r | \mathbf{q}_t^h) = \mathcal{N}(\mathbf{z}_t^r; \hat{\mu}_t^r, \hat{\Sigma}_t^r) \quad (21)$$

where the terms in magenta are from the human VAE posterior $q_t^h = q(\mathbf{z}_t^h | \mathbf{x}_t^h)$ and the terms in orange are from the HMM. We then reconstruct samples from $p(\mathbf{z}_t^r | \mathbf{q}_t^h)$ thereby allowing the robot VAE's decoder to be trained with samples drawn from a distribution similar to what it would encounter during test time, instead of just using samples from the posterior. Our modified ELBO at each timestep for training the robot VAE in the HRI scenario (Algorithm 2) can be written as

$$\begin{aligned} \mathbb{E}_q [\log \frac{p(\mathbf{x}_t^h, \mathbf{x}_t^r, \mathbf{z}_t^r)}{q(\mathbf{z}_t^r | \mathbf{x}_t^h, \mathbf{x}_t^r)}] &= \mathbb{E}_{q_r} \log p(\mathbf{x}_t^r | \mathbf{z}_t^r) \\ &\quad - \beta KL_t^r \\ &\quad + \mathbb{E}_{\mathbf{z}_t^r \sim p(\mathbf{z}_t^r | \mathbf{q}_t^h)} \log p(\mathbf{x}_t^r | \mathbf{z}_t^r) \end{aligned} \quad (22)$$

where the first term is the reconstruction term, KL_t^r is the regularization term calculated according to Eq. 14, and the third term denotes the reconstruction of the conditional samples drawn from $p(\mathbf{z}_t^r | \mathbf{q}_t^h)$ (Eq. 21). Both the reconstruction term and the conditional reconstruction are estimated in a Monte Carlo fashion by averaging the loss over multiple samples drawn from the corresponding latent distributions.

Algorithm 2: Learning HRI Dynamics

Data: A set of trajectories with action labels
 $\mathbf{X} = \{\mathbf{X}_{1:T}^h, \mathbf{X}_{1:T}^r, c\}$ for $|\mathcal{C}|$ actions, Human VAE and HMMs

Result: Robot VAE weights

Initialize VAE weights randomly

```

while not converged do
  for  $\mathbf{x}_{1:T}^h, \mathbf{x}_{1:T}^r, c \in \mathbf{X}$  do
    Compute VAE Posterior  $q(\mathbf{z}_t^r | \mathbf{x}_t^r)$ 
    Compute Latent Conditional  $p(\mathbf{z}_t^r | \mathbf{q}_t^h)$  (Eq. 16-21)
    Reconstruct posterior and conditional samples
    Maximize ELBO (Eq. 22) to update VAE weights
  end
end

```

During testing, we first encode the observations of the human agent \mathbf{x}_t^h and then condition the HMM $p(\mathbf{z}_t^r | \mathbf{q}_t^h)$ to

generate the latent trajectory of the second agent \hat{z}_t^r using Eq. 16 - 21, which is then decoded to obtain the actions of the second agent $p(\hat{x}_t^r | z_t^r)$.

C. Inverse Kinematics Adaptation and Stiffness Modulation

In typical LfD approaches, robot motions are learned via kinesthetic teaching [68], [78]. In HRI scenarios, kinesthetic teaching gives good results [4], [13], [16], [28] but can become tedious when trying to generalize to multiple human interaction partners. One way to circumvent this could be to execute randomized open-loop trajectories for the robot and learn how the human adapts to the robot motions [17].

Since we are dealing with a Humanoid robot, we learn the robot motions using the kinematic similarities between the human and the robot [31]. From the 3D positions of an arm, we extract the shoulder angles (yaw, pitch, and roll) and the bending of the elbow by using the geometry of the tracked skeleton, as seen in Figure 4. We defer the calculation of the wrist angle to our future work and instead use a fixed value for the wrist for each interaction.

Since we extract joint angles from the human demonstrations, some inaccuracies are present due to slight differences between the kinematic structure of the human skeleton and the robot. Given the difference in the arm dimensions of the Pepper robot and a human, the spatial accuracy of the learned behaviors is quite limited. Additionally, given the smaller size of the Pepper robot, its reachability is also limited, thereby limiting the ability to learn task space trajectories.

To bridge the mismatch in motion retargeting, we adapt the predicted motions during test time with Inverse Kinematics to reach the human partner's hand while using the predicted motions as a prior. The use of Inverse Kinematics enables the physical proximity needed by the interaction while keeping the robot configuration close to the demonstrated behaviors [35]. Moreover, we do not need to use Inverse Kinematics all the time, but only in the segments involved in physical contact. We therefore inspect the underlying segments of the HMM manually to see which segments comprise of the contact-based portion of the trajectory and which ones do not, and subsequently perform the IK adaptation only in the contact-based segments.

Given an initial prediction of the robot's joint angle distribution (μ_q, Σ_q) and a task space goal in 3D space (μ_x, Σ_x) (which in this case is the human partner's hand position), we aim to find a joint angle configuration q^* that reaches μ_x , but does not stray too far from μ_q [35] which boils down to the following optimization problem

$$\begin{aligned} q^* &= \arg \max_q \mathcal{N}(f(q) | \mu_x, \Sigma_x) \mathcal{N}(q | \mu_q, \Sigma_q) \\ &= \arg \min_q \lambda_x (\mu_x - f(q))^T \Sigma_x^{-1} (\mu_x - f(q)) + \\ &\quad \lambda_q (\mu_q - q)^T \Sigma_q^{-1} (\mu_q - q) \end{aligned} \quad (23)$$

where $f(q)$ is the forward kinematics model of the robot to estimate the end-effector location given the joint angles q , and λ_x, λ_q are relative weights balancing whether it is more important to stay close to the initial distribution or to reach the task space goal. In practice, the joint configuration prior

μ_q comes from the VAE decoder $p(x_t^r | z_t^r)$, and μ_x comes from the human observation x_t^h , and we set Σ_q and Σ_x to identity. We can then simplify Eq. 23 to

$$q^* = \arg \min_q \lambda_x \| \mu_x - f(q) \|^2 + \lambda_q \| \mu_q - q \|^2 \quad (24)$$

where $\|\cdot\|$ denotes the Euclidean norm.

Moreover, we do not need to employ IK during every segment of the interaction. Once the training converges, we identify the segments from the underlying HMM that correspond to the contact-based phases of the interaction. Based on the HMM's forward variable prediction from the human observation, we adapt the human-conditioned robot motions using IK if the current segment is contact-based. The incorporation of IK is further detailed in Algorithm 3.

Algorithm 3: Conditioning on Human Observations and IK Adaptation

Data: An observation of the human agent $x_{1:t}^h$, trained VAE and HMM Models, a set of contact-based HMM states \mathcal{I} for the given interaction

Result: Conditioned Trajectory for the second agent $\hat{x}_{1:t}^r$

```

for  $t \in [1, T]$  do
  Encode the human observation  $q(z_t^h | x_t^h)$ 
  Compute Latent Conditional  $p(z_t^r | q_t^h)$  (Eq. 16-21)
  Decode the conditioned prediction  $\mu_q = p(\hat{x}_t^r | z_t^r)$ 
  if  $\arg \max_i \alpha_i(z_t^h)$  is a contact-based segment then
     $\hat{q}_t = \arg \min_q \lambda_x \| \mu_x - f(q) \|^2 + \lambda_q \| \mu_q - q \|^2$ 
  else
     $\hat{q}_t = \mu_q$ 
  end
  Send  $\hat{q}_t$  to the robot controller.
end

```

In a contact-rich interaction like handshaking, along with moving in a coordinated manner, the robot must be compliant with the human's motion when in contact. When executing handshaking trajectories on the robot, we use the forward variable to detect when to lower the target joint stiffness once the robot enters the contact-based segments of the interaction. Without such a mechanism, the Pepper robot's arm would not comply with the human's hand, and the human would be unable to displace Pepper's hand for executing a proper handshake.

During a handshaking interaction with the robot, to prevent sudden changes in stiffness arising from the misclassification of the segment, we disable back-transitions into the initial reaching segment. Additionally, since the forward variable is calculated using only the human partner's latent state during test time, we found some mismatches in the segment prediction compared to using the full joint human-robot states for timesteps near the transition boundary between the initial reaching segment and the subsequent segments. Therefore, taking a leaf out of Transition State Clustering [36], [47], we learn an additional distribution over the states that get misclassified at the transition boundary between the reaching and contact phase. We use this additional transition state distribution to detect when the interaction proceeds into the contact-based segments when the probability of either the

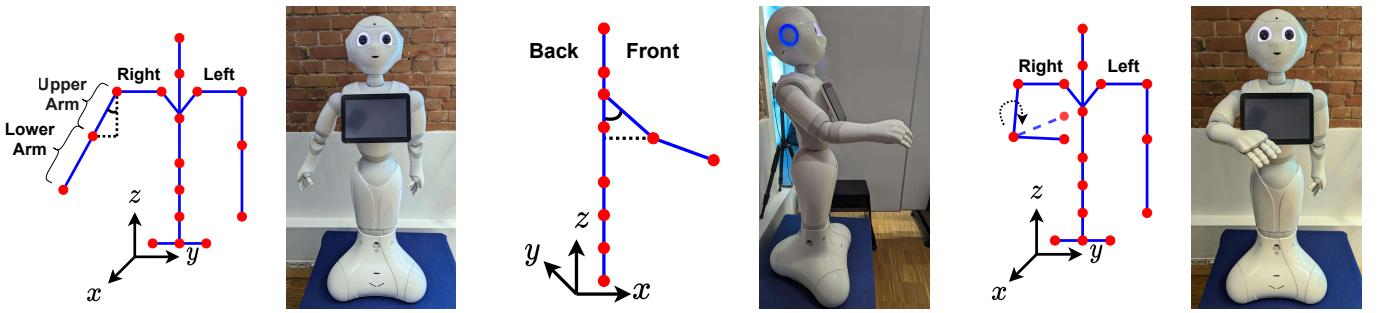


Fig. 4: Geometric similarities between the degrees of freedom of a human's upper body and the humanoid robot Pepper. The shoulder roll, pitch, yaw, and elbow angle of a human can be directly mapped to Pepper's joint angles.

contact segment or the transition state exceeds that of the reaching segment. Doing so gives a better indication of when to lower the robot's stiffness and provides a more suitable interaction. Without such a scheme, due to the misclassification of the active segment, the joint stiffness would not always be lowered correctly, thereby resulting in a rigid and non-compliant handshake.

IV. EXPERIMENTS AND RESULTS

In this section, we first provide our implementation details (Section IV-A) and the datasets used (Section IV-B). We then present the results of predicting the controlled agent's trajectories after conditioning on the observed agent in Section IV-C and finally, we discuss our user-study in Section IV-D and show some additional results for a Bimanual Robot-to-Human Handover scenario in Section IV-E.

A. Experimental Setup

The VAEs are implemented using PyTorch [69] with 2 hidden layers in the encoder and decoder each with sizes (40, 20) and (20, 40) respectively and a 5-dimensional latent space with Leaky ReLU activations [53] at all layers except the output layer. The weights are initialized using Xavier's initialization [33]. The networks are trained with $\beta = 5 \times 10^{-3}$, 10 Monte Carlo samples per input data and using the Adam optimizer with weight decay [52] with a learning rate of 5×10^{-4} . In the HHI scenarios, we share the parameters of the VAEs for both human agents as the inputs are structurally similar. The networks were trained for 400 epochs. The best out of 4 seeds were used to report the results.

A separate HMM is trained for each interaction. The HMMs are implemented using a custom PyTorch version of Pb-DLlib¹ [70]. Each HMM has 6 hidden states, which we found as the most numerically stable. The HMMs are initialized by splitting each of the latent trajectories in the training set into equally sized segments over time. To prevent numerical instabilities arising from vanishing values in the HMM covariance matrices, we add a small positive regularization constant of 10^{-4} to the diagonal elements. When sampling from the conditional distribution, due to numerical errors in computing the Cholesky decomposition of the covariance matrices, rather than adding a constant value to all elements, we add linearly increasing regularization based on the dimensions, going from

9.1×10^{-5} for the first dimension to 10^{-4} for the last dimension. Additionally, we adapt the covariance matrices of the conditional distribution by iteratively adding an even smaller constant that is proportional to the absolute value of the smallest eigenvalue of the covariance matrix to the diagonal elements [39].

We implement our approach on the Pepper robot [66], which is a 1.2m tall humanoid robot that has the same degrees of freedom in its arm as a human. For Inverse Kinematics, we use a modified version of the IKPy library [56]². We set $\lambda_x = 1$ and $\lambda_q = 0.01$ for the objective function in Eq. 24.

B. Dataset

1) Bütepage et al. [13]: ³ record HHI and HRI demonstrations of 4 interactions: Waving, Handshaking, and two kinds of fist bumps. The first fistbump called "Rocket Fistbump", involves bumping fists at a low level and then raising them upwards while maintaining contact with each other. The second is called "Parachute Fistbump" in which partners bump their fists at a high level and bring them down while simultaneously oscillating the hands sideways, while in contact with each other. Since our testing scenario involves the humanoid robot Pepper [66], we additionally extract the joint angles from one of the human partner's skeletons from the above-mentioned HHI data for the Pepper robot using the similarities in DoFs between a human and Pepper [31], [72] (Figure 4), which we denote this as "HRI-Pepper". Bütepage et al. [13] additionally record demonstrations of these actions with a human partner interacting with an ABB YuMi-IRB 14000 robot controlled via kinesthetic teaching. We call this scenario as "HRI-Yumi".

In the HHI scenario (and the HRI-Pepper scenario), there are 181 trajectories (32 - Waving, 38 - Handshake, 70 - Rocket Fistbump, 49 - Parachute Fistbump) of which 80% of the trajectories (149 trajectories) are used for training and the rest (32 trajectories) for testing. In the HRI-Yumi scenario, there are 41 trajectories (10 each for Waving, Handshaking, and Rocket Fistbump and 11 for Parachute Fistbump) of which we use a similar split with 32 trajectories for training and 9 for testing. We use a time window of 5 observations as the input for a given timestep as done in [13]. We downsample the data to 20Hz to match our testing scenario.

We use the 3D positions and the velocities (represented as position deltas) of the right arm joints (shoulder, elbow,

¹<https://git.ias.informatik.tu-darmstadt.de/prasad/pbdlib-torch>

²<https://github.com/souljaboy764/ikpy>

³https://github.com/souljaboy764/human_robot_interaction_data

and wrist), with the origin at the shoulder, leading to an input size of 90 dimensions (5x3x6: 5 timesteps, 3 joints, 6 dimensions) for a human partner. For the HRI-Pepper and HRI-Yumi scenarios, we use a similar window of joint angles, leading to an input size of 20 dimensions (5x4) for the 4 joint angles of Pepper’s right arm and an input size of 35 dimensions (5x7) for the 7 joint angles of Yumi’s right arm.

2) *Nuitrack Skeleton Interaction Dataset*: (NuiSI)⁴ is a dataset which we collected ourselves of the same 4 interactions as in [13], namely Waving, Handshaking, Rocket Fistbump and Parachute Fistbump. While the dataset in Bütepage et al. [13] has clean motions of the interaction partners, they require high-accuracy inertial sensors in whole-body suits. Therefore, for a more realistic setting, we capture the interaction partners using low-cost hardware.

While other datasets exist that perform skeleton tracking with a Kinect Camera [82], [86], [91], the participants in these datasets are recorded from the side. Therefore, due to partial occlusions, there is a high degree of noise in skeleton tracking, which made it difficult to train a suitable model for the interactions. Therefore, we record two human partners interacting with one another using one Intel Realsense D435 camera per partner (two in total), such that each camera captures a full frontal view of the interaction. We use Nuitrack⁵ for tracking the upper body skeleton joints in each frame at 30Hz. As done above with [13], we additionally extract the joint trajectories for the Pepper robot, which leaves us finally with both the HHI and HRI-Pepper scenarios. We have 44 trajectories (12 - Waving, 11 - Handshaking, 12 - Rocket Fistbump, 9 - Parachute Fistbump), of which we similarly use 80% of the trajectories for training (33 trajectories) and the rest (11 trajectories) for testing.

The skeletons are then rotated such that they follow the orientation shown in Figure 4 with the x -axis in the forward direction, the y -axis going from right to left, and the z -axis in the upward direction. For training, the data is processed similarly as mentioned above with a window size of 5 time steps. Given the relatively low amount of samples, we finetune the network trained on the aforementioned dataset of [13]. As done with [13], we extract the joint angles for the Pepper robot from the skeleton trajectories for the HRI-Pepper scenario [31], [72].

C. Conditioned Prediction Results

We test the conditioning ability of our approach compared to [13]⁶ to evaluate the accuracy of the generated motions of the controlled agent after observing the human interaction-partner. We evaluate the approaches over the four interactions of the dataset in [13] and our collected data (NuiSI). We calculate the Mean Squared Error (MSE) averaged over each joint and the 5-time step window.

We do not train with the conditional loss in the Human-Human scenarios since we use shared weights, the decoder already learns to reconstruct the ground truth samples for the

TABLE I: Prediction MSE (in cm) for the second human partner’s trajectories after observing the first human partner averaged over all joints and timesteps. (* – $p < 0.05$, ** – $p < 0.01$, lower is better)

Dataset	Action	MILD v1	[13]
[13]	Hand Wave	0.788 ± 1.226**	4.121 ± 2.252
	Handshake	1.654 ± 1.549*	1.181 ± 0.859
	Rocket Fistbump	0.370 ± 0.682	0.544 ± 1.249
	Parachute Fistbump	0.537 ± 0.579**	0.977 ± 1.141
NuiSI	Hand Wave	0.408 ± 0.538**	3.168 ± 3.392
	Handshake	0.311 ± 0.259**	1.489 ± 3.327
	Rocket Fistbump	1.142 ± 1.375**	3.576 ± 3.082
	Parachute Fistbump	0.453 ± 0.578**	2.008 ± 2.024

conditional distribution. Additionally, we found that the interplay between the shared weights and the conditional training discussed in Section III-B would cause the HMM posterior to collapse into a single unimodal distribution. Therefore, we only show comparisons of our vanilla approach presented in Section III-A, which we denote as “MILD v1”.

The results of MILD v1 in predicting the interaction partner’s trajectories in the HHI scenarios can be seen in Table I. It can be seen that MILD v1 with its simplistic nature of using an HMM for the latent dynamics performs significantly better (as verified with a Mann-Whitney U Test) than [13] where the VAEs are trained with an uninformative standard normal distribution as a prior. Although additional LSTMs are employed in [13] to learn the latent dynamics, since the VAEs are trained with an uninformative prior, their approach fails to accurately reconstruct motions the learnt latent dynamics. In contrast, the latent dynamics of each segment of the interaction is captured well by the HMM which therefore acts as an informative prior for the VAEs which is reflected in the improved prediction accuracy of MILD v1.

Coming to HRI scenarios, starting from the initial incorporation of the HMM prior as shown in Section III-A (denoted as “MILD v1”), we explore two variants of the last conditional training term in Eq. 22. The first is with reconstructing the conditional latent predictions (i.e. the mean in Eq. 19) of samples drawn from the posterior distribution both with and without the use of the VAE posterior covariance (in Eq. 16). This can be summarized mathematically as $\mathbb{E}_{q(z_t^h|x_t^h)} \log p(x_t^r|\mu_t^r)$ where μ_t^r is calculated using Eq. 9 with samples drawn from $q(z_t^h|x_t^h)$. We explore this variant both without and with the posterior covariance in Eq. 16, denoted as “MILD v2.1” and “MILD v2.2” respectively

The second variant uses the posterior mean and covariance to calculate the conditional distribution and subsequently reconstruct samples drawn from the conditional distribution. This can be summarized mathematically as $\mathbb{E}_{p(z_t^r|q_t^h)} \log p(x_t^r|z_t^r)$ where $p(z_t^r|q_t^h)$ is calculated using Eq. 16 - 21 and the samples drawn from this distribution are reconstructed. We explore this variant both without and with the posterior covariance in Eq. 16, which we denote as “MILD v3.1” and “MILD v3.2” respectively.

The key differences between the different variants are highlighted in Table III. To summarize, the prior means and covariances from the HMM are used by both variants in the conditioning, denoted by the orange terms in Eq. 16 - 21. The key difference between the variants is how the posterior

⁴<https://github.com/souljaboy764/nuisi-dataset>

⁵<https://nuitrack.com/>

⁶Results reported using our implementation of [13].

TABLE II: Prediction MSE (in radians) for robot trajectories after observing the human interaction partner averaged over all joints and timesteps. (Lower is better, lowest values highlighted in bold, significance shown in Table A.1)

Dataset	Action	MILD v1	MILD v2.1	MILD v2.2	MILD v3.1	MILD v3.2	Bütepage et al. [13]
HRI-Yumi [13]	Hand Wave	1.705 \pm 0.521	1.349 \pm 1.972	1.641 \pm 1.968	1.033 \pm 1.204	1.143 \pm 1.330	0.225 \pm 0.302
	Handshake	0.290 \pm 0.148	0.073 \pm 0.040	0.068 \pm 0.052	0.104 \pm 0.056	0.123 \pm 0.069	0.133 \pm 0.214
	Rocket Fistbump	0.428 \pm 0.175	0.236 \pm 0.167	0.183 \pm 0.122	0.130 \pm 0.074	0.128 \pm 0.071	0.147 \pm 0.119
	Parachute Fistbump	0.425 \pm 0.150	0.028 \pm 0.042	0.033 \pm 0.034	0.028 \pm 0.034	0.028 \pm 0.035	0.181 \pm 0.155
HRI-Pepper [13]	Hand Wave	0.267 \pm 0.152	0.161 \pm 0.228	0.165 \pm 0.189	0.103 \pm 0.103	0.106 \pm 0.105	0.664 \pm 0.277
	Handshake	0.327 \pm 0.253	0.111 \pm 0.092	0.153 \pm 0.154	0.061 \pm 0.048	0.056 \pm 0.041	0.184 \pm 0.141
	Rocket Fistbump	0.161 \pm 0.095	0.035 \pm 0.068	0.035 \pm 0.068	0.021 \pm 0.037	0.018 \pm 0.035	0.033 \pm 0.045
	Parachute Fistbump	0.265 \pm 0.178	0.116 \pm 0.176	0.112 \pm 0.181	0.095 \pm 0.151	0.088 \pm 0.148	0.189 \pm 0.196
HRI-Pepper (NuiSI)	Hand Wave	0.760 \pm 0.325	0.050 \pm 0.084	0.060 \pm 0.087	0.046 \pm 0.059	0.049 \pm 0.059	0.057 \pm 0.093
	Handshake	0.225 \pm 0.114	0.025 \pm 0.022	0.025 \pm 0.020	0.021 \pm 0.015	0.020 \pm 0.014	0.083 \pm 0.075
	Rocket Fistbump	0.354 \pm 0.238	0.077 \pm 0.095	0.080 \pm 0.088	0.077 \pm 0.067	0.079 \pm 0.072	0.101 \pm 0.086
	Parachute Fistbump	0.201 \pm 0.072	0.032 \pm 0.038	0.028 \pm 0.040	0.025 \pm 0.028	0.022 \pm 0.027	0.049 \pm 0.040

TABLE III: Differences in conditioning and sampling strategies of the variants of MILD.

Variant	Conditional Training Inputs	Samples given to Decoder
MILD v1	None	Only Posterior Samples
MILD v2.1	Posterior Samples	Posterior Samples and their corresponding conditioned outputs
MILD v2.2	Posterior Samples and Covariance	
MILD v3.1	Posterior Mean	Posterior Samples, Conditional Samples
MILD v3.2	Posterior Mean and Covariance	

distribution terms (denoted in magenta in Eq. 16 - 21) are used. In MILD v3.1 and MILD v3.2, we directly use the posterior mean and covariance as shown in Eq. 16 - 21. In the case of MILD v2.1 and MILD v2.2, we do not directly use the posterior mean and covariance, but first draw samples from the posterior distribution $z_t^h \sim \mathcal{N}(\mu_z(\mathbf{x}_t^h), \Sigma_z(\mathbf{x}_t^h))$ which are then used in the conditioning in Eq. 17 instead of the posterior mean $\mu_z(\mathbf{x}_t^h)$ which can be written as $\hat{\mu}_i^r = \mu_i^r + K_i(z_t^h - \mu_i^h)$ to directly get the conditional samples that then get reconstructed.

As seen in Table II, reconstructing samples from the conditional distribution (MILD v3.1 and v3.2) provides much better results as compared to conditioning samples drawn from the posterior (MILD v2.1 and v2.2), both of which perform better than MILD v1 that does not use conditional training. The samples drawn from the conditional distribution would be more representative of the type of samples that the decoder would see during run time. While enough samples from the posterior, when conditioned, can also estimate this distribution, empirically, this fails to match up to sampling from the conditional distribution. One argument for this is that reconstructing conditional samples enables learning a joint latent space more susceptible to the HMM conditioning.

Furthermore, the overall improvement in performance compared to MILD v1 and [13] highlights the advantage of incorporating conditional prediction into the training process for reactive motion generation. The importance of incorporating reactive motion generation into the training can also be seen in the improved performance of [13] over MILD v1 in the HRI scenarios as compared to the HHI scenarios. This improvement comes from the fact that MILD uses the HMMs just as a latent prior, whereas in [13], the authors explicitly train a separate HRI dynamics model for predicting the robot motions from the latent trajectories of both the human and the robot. During testing, however, the HRI dynamics network does not have access to the ground truth target of the robot which is used to train the network. Therefore, the HRI dynamics is predicted in an autoregressive manner, which deteriorates the performance due to out-of-distribution data. In this regard, it

can be seen that explicitly incorporating such conditional out-of-distribution samples can lead to better results, as seen in the improved performance of the variants of MILD (v2.1 - v3.2).

Some examples of the learned behaviors, along with the progression of the HMM in the latent space can be seen in Figure 6, where we show a sample interaction for handshake (Figure 6a) and rocket fistbump (Figure 6b) on the Pepper robot. As it can be seen, the HMM captures the sequencing between the multiple modes of the latent space to generate suitable motions for real-world HRI scenarios. This is additionally validated via a user study (Section IV-D) which shows the ability of our approach to generalize well to various users, despite being trained on demonstrations of just two partners.

D. HRI User Study

To see the effectiveness of our approach in producing acceptable physical behaviors, we conducted a user study where participants interacted with the robot. We evaluate our proposed approach both with and without IK adaptation, denoted as “MILD-IK” and “MILD” (while we use “MILD v3.2”, for ease of notation, we shorten it to “MILD”) against a baseline IK algorithm (Eq. 13) that uses the human’s hand pose as the target location (“Base-IK”). We run this study for two interactions, handshake and rocket fistbump. The study design was given a positive vote by the Ethics Commission at TU Darmstadt (Application EK 48/2023).

1) *Procedure*: The study took place in a laboratory setting (Figure 5) where participants were first guided to a desk to fill out a consent form and a pre-questionnaire which included demographic information, prior experiences with robots, their attitude towards robots [88], attitude towards physical interactions, and personality questions to gauge extroversion [73]. We assessed these measures to gain additional information about our sample. Participants were then shown an instruction video of two humans performing an interaction (either a handshake or a rocket fistbump). Participants were then given the experiment protocol to read wherein they were instructed that they had to lead the interaction with the robot. Additionally, some general instructions were provided to the participants about



Fig. 5: Setup of the User Study for interacting with the Pepper robot. As Pepper is quite short (1.2m), it is placed on a pedestal to match the height of a human partner. The camera behind Pepper tracks the human partner's motion, which is used to generate Pepper's motions. how they should position themselves to maintain their posture and the limit of the arm movements they can perform. This was done to prevent occlusions in the skeleton tracking and to maintain the robot's reachability limits.

Participants were then guided behind a barrier where they would see the robot for the first time and would then stand at an adequate distance from the robot. When they were ready, an initial interaction was performed with a hard-coded motion. This is to get the participant habituated to the way the robot moves and counter any novelty effects that may occur and to have the participant better understand how to perform the interaction with the robot. After the initial run, the participant was informed that the experimental trials would begin. The sequence of each trial was as follows:

- 1) Before each trial, the participant would see a video stream of the skeleton tracker to better position themselves.
- 2) Once the position was set and the tracking was stable, the experimenter would signal them to start the trial.
- 3) The participant would start the interaction and the robot would respond reactively.
- 4) Once the participant goes back to a neutral position with their hands by their side, the robot arm would go back to a neutral position, marking the end of the trial.
- 5) This process would then repeat once again after which the participant was asked to fill out a questionnaire.

This process constitutes a single session and was repeated 3 times (3 sessions in total) wherein the robot was controlled by one of the three aforementioned algorithms (Base-IK, MILD, MILD-IK). The algorithms were shown in randomized order to all participants to avoid sequence effects. The participant was neither informed of the randomization nor the algorithm they were interacting with. After all 3 sessions, the participant was asked to fill out a final questionnaire where they had to rank the sessions based on their preference and answer some open-ended questions about the sessions.

2) *Participant Sample:* A total of 20 users (8 female, 12 male) participated in our study and were recruited through the university environment. Of the 20 participants, 10 performed the Rocket fistbump (4 female, 6 male) and the rest performed a handshake. The mean age of the participants was 27.85 years (SD: 3.51). Participants had an average level of experience

with robots overall on a scale from 1 (no experience at all) to 5 (a lot of experience) with a mean of 2.90 (SD: 1.45). They had quite a positive attitude towards robots on a scale from 1 (very negative) to 7 (very positive) with a mean of 5.90 (SD: 1.16). This is also consistent with the ratings regarding the attitudes towards robots for the following items on a scale of 1 (Strongly Agree) to 7 (Strongly Disagree). Participants largely had a high agreement towards feeling relaxed when interacting with robots (Mean: 5.60, SD: 1.05) and high levels of disagreement towards being paranoid when interacting with a robot (Mean: 1.85, SD: 0.93) and towards feeling nervous standing in front of a robot (Mean: 1.80, SD: 1.06). Participants had a positive outlook towards physical interactions in general on a scale from 1 (distant) to 7 (open) with a mean of 6.10 (SD: 0.84). This was also confirmed with the Big 5 extroversion scale [73] with a mean extroversion of 4.75 (SD: 1.34) out of 7.

3) *Methodology:* The study followed a within-subject design where participants interacted with a Pepper robot controlled by each of the algorithms (Base-IK, MILD and MILD-IK) twice in a randomized order, leading to 6 HRIs per participant. Since humans perceive different types of physical touch differently [12], [32], [77], we aimed to remove any influence the type of interaction (handshake or fistbump) might have on how participants view the robot during different interactions. We wanted to keep the focus on understanding how our proposed algorithm is perceived by users. Therefore, each participant had to either perform a handshake or a fistbump, not both. Additionally, to prevent any sudden jumps in the motion of the robot, a weighted moving average filter was used to smoothen out the predicted motions of the robot.

We break the 6 trials into 3 sessions, where each session corresponds to two trials of a given algorithm. Each session was evaluated with 16 different items adapted from the Godspeed [7] and the SASSI [42] questionnaires, each rated on a 5-Point scale (1 - Strongly disagree, 5 - Strongly agree):

- The interaction with the robot was pleasant.
- The interaction with the robot was exciting.
- The interaction with the robot was human-like.
- The interaction with the robot was natural.
- The interaction with the robot was friendly.
- The interaction with the robot was comfortable.
- The interaction with the robot was well-timed.
- The interaction with the robot was accurate.
- The interaction with the robot was annoying.
- The interaction with the robot was awkward.
- The interaction with the robot was scary.
- The robot interacted in an aggressive way.
- I am satisfied with the way the robot interacted with me.
- The second trial in the session was more effortless than the first.

At the end of the experiment, we asked the participants to rank the 3 sessions (algorithms) in their order of preference.

We run a one-way Repeated Measures ANOVA to compare the responses of the different algorithms followed by paired sample t-tests for a post hoc analysis (Table IV).

4) *Study Results:* As seen in Figure 7, MILD-IK was by far ranked in the first place much more than MILD which was by far the least ranked, and Base-IK which was mostly

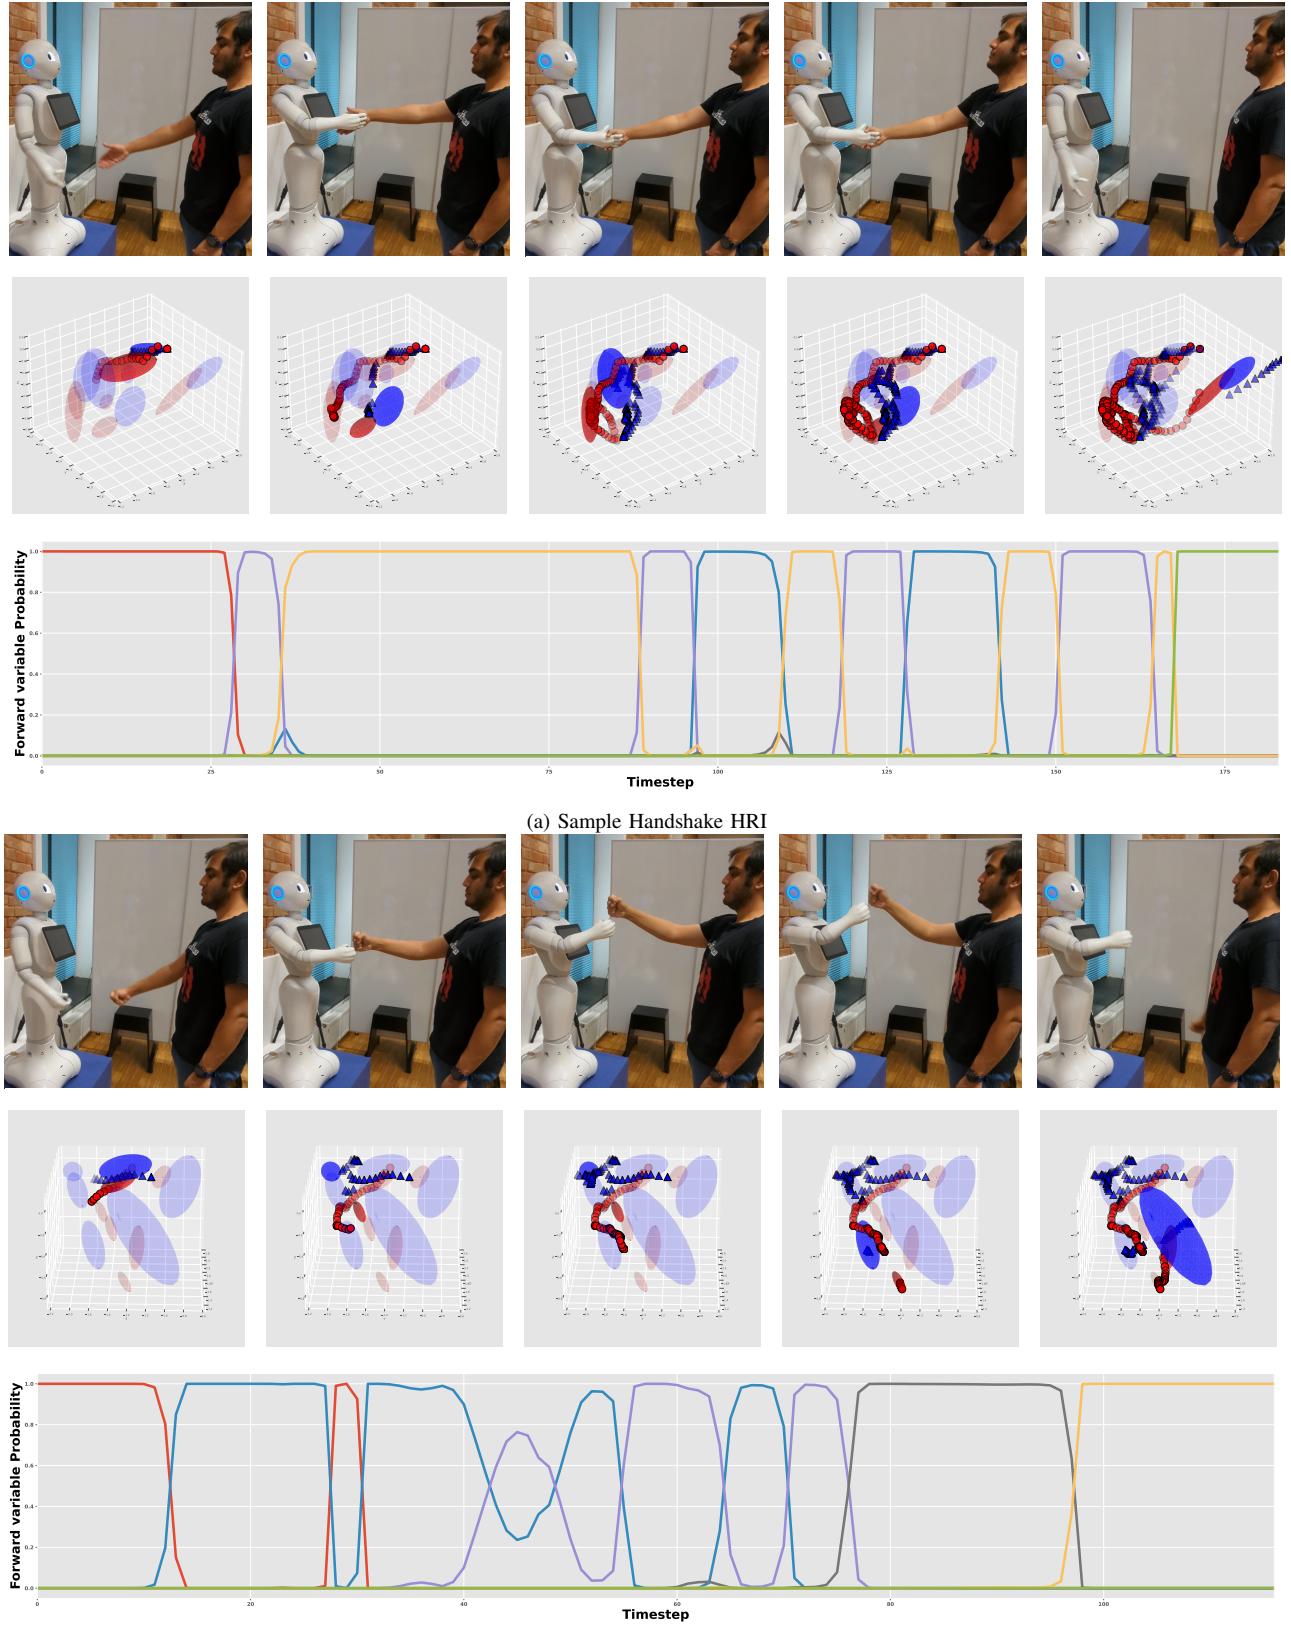


Fig. 6: Sample sequences for the Handshake and Rocket fistbump interactions on the Pepper robot. The top row of each interaction shows the result of the generated reactive motion after observing the human partner's skeleton. The middle row shows the latent trajectories and the HMM segments over the first 3 dimensions of the latent space (Red - Human, Blue - Robot). The opacity of each cluster is the corresponding cluster probability, given by the HMM forward variable. The progression of the HMM forward variable is shown in the bottom row, with different colors denoting the different segments.

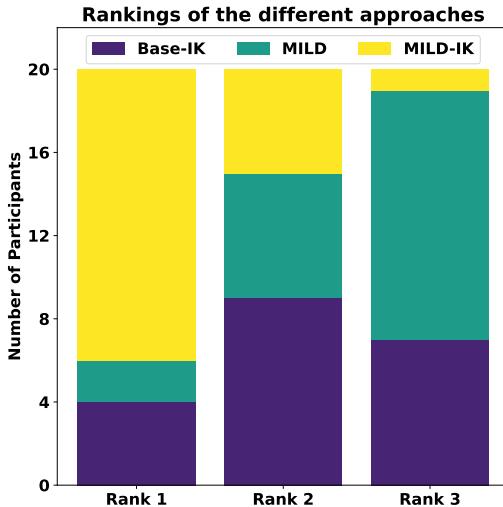


Fig. 7: Participant's ranking of the algorithms. (MILD-IK - yellow stars, MILD - green lines, Base-IK - purple dots). Most participants ranked MILD-IK first, better than both Base-IK and MILD.

ranked second. Just using MILD without IK was the least preferred among the algorithms, even though it was programmed keeping the interactiveness in mind, but due to the motion retargeting issues mentioned in Section III-C, it is unable to reach the participant's hand accurately, leading to a low acceptance by the participants which can further be seen in the results in Figure 10. All approaches have similarly high levels of pleasantness, excitement, and friendliness and similarly low levels of annoyance, aggression, and scariness. This goes to show that the overall interaction scenario was well perceived by the participants and can additionally be attributed to an overall positive attitude toward robots.

Since the Base-IK approach does not have any prior over the IK solutions, it would lead to poses where the elbow is pointed outwards and the hand is turned sideways, which participants remarked was unnatural and awkward. Even though it was able to reach the human's hand location, this unnatural pose of the robot hand prevented the Base-IK approach from being able to reach the human hand in a graspable manner, due to the orientation of the end-effector. This is also reflected via a significantly lower accuracy rating for Base-IK and can further be seen in relatively lower (although non-significant) trends for naturalness, comfort, and satisfaction.

Participants found MILD awkward at times since the predicted motions would sometimes fall short of the partner's hand and the robot would thereby not reach the participants' hand correctly. This can especially be seen in the low accuracy that is given to MILD. This reachability issue gets mitigated with MILD-IK as it generates more natural poses. However, with MILD-IK, since the prediction of when to start the IK adaptation comes from the HMM, there would be a distinct moment when the robot hand would reach the human's hand, which we attribute as a possible reason for a relatively higher awkwardness rating.

Overall, the high acceptability of MILD-IK is also reflected via significantly better ratings of the human-likeness, timing, and perceived accuracy. The effortlessness between two trials in an interaction was also rated significantly higher for MILD-

IK showing that it was easier for participants to get habituated to the movements of the robot. MILD-IK also achieves a higher rating for timing, which is attributed to the ability of the HMM to generate the receding motion in a reactive manner, unlike Base-IK which would still try to reach the hand until the participant would go back to a neutral pose. Some of these peculiarities can be seen in Figure 11.

5) *Objective Performance Metrics*: We evaluate some objective performance metrics relating to the human partner's motions. We analyze the ergonomics of the arm movements of the participants using the REBA [40] and RULA [58] scores (Figure 8). We additionally analyze the smoothness of the human movements using SPARC [6] (Figure 9).

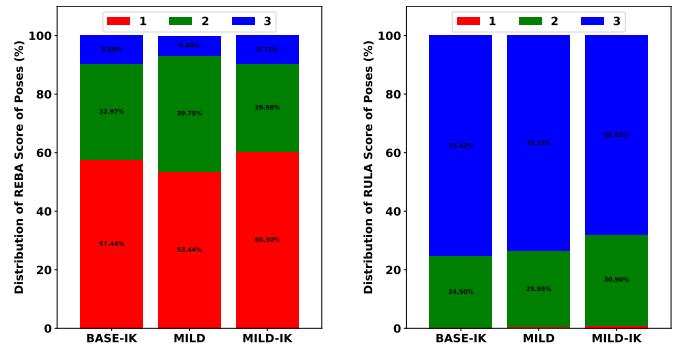


Fig. 8: Distribution of REBA and RULA scores across all the poses of the human partner for different algorithms.

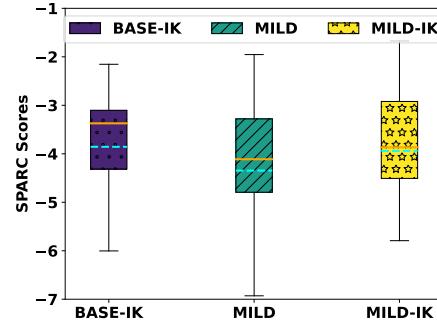


Fig. 9: Distribution of SPARC scores across the wrist trajectories of the human partner for different algorithms. The orange line depicts the median and the dotted cyan line depicts the mean.

We found the distribution of the REBA and RULA scores to be very similar across the different algorithms, and within the range of Negligible- to Low-Risk posture scores. Similarly, there were no significant differences in the SPARC scores as well. This lack of significance comes mainly from the instructions that participants received for maintaining their posture and arm location, which was needed to minimize occlusions in the skeleton tracking and to comply with the robot's joint limits. Consequently, the deviation in the postures of the participants was negligible across the different trials. Regarding the smoothness of the human partner's motion, given the limited range of motion, in addition to the noise in skeleton tracking, we found no significant differences between the different algorithms.

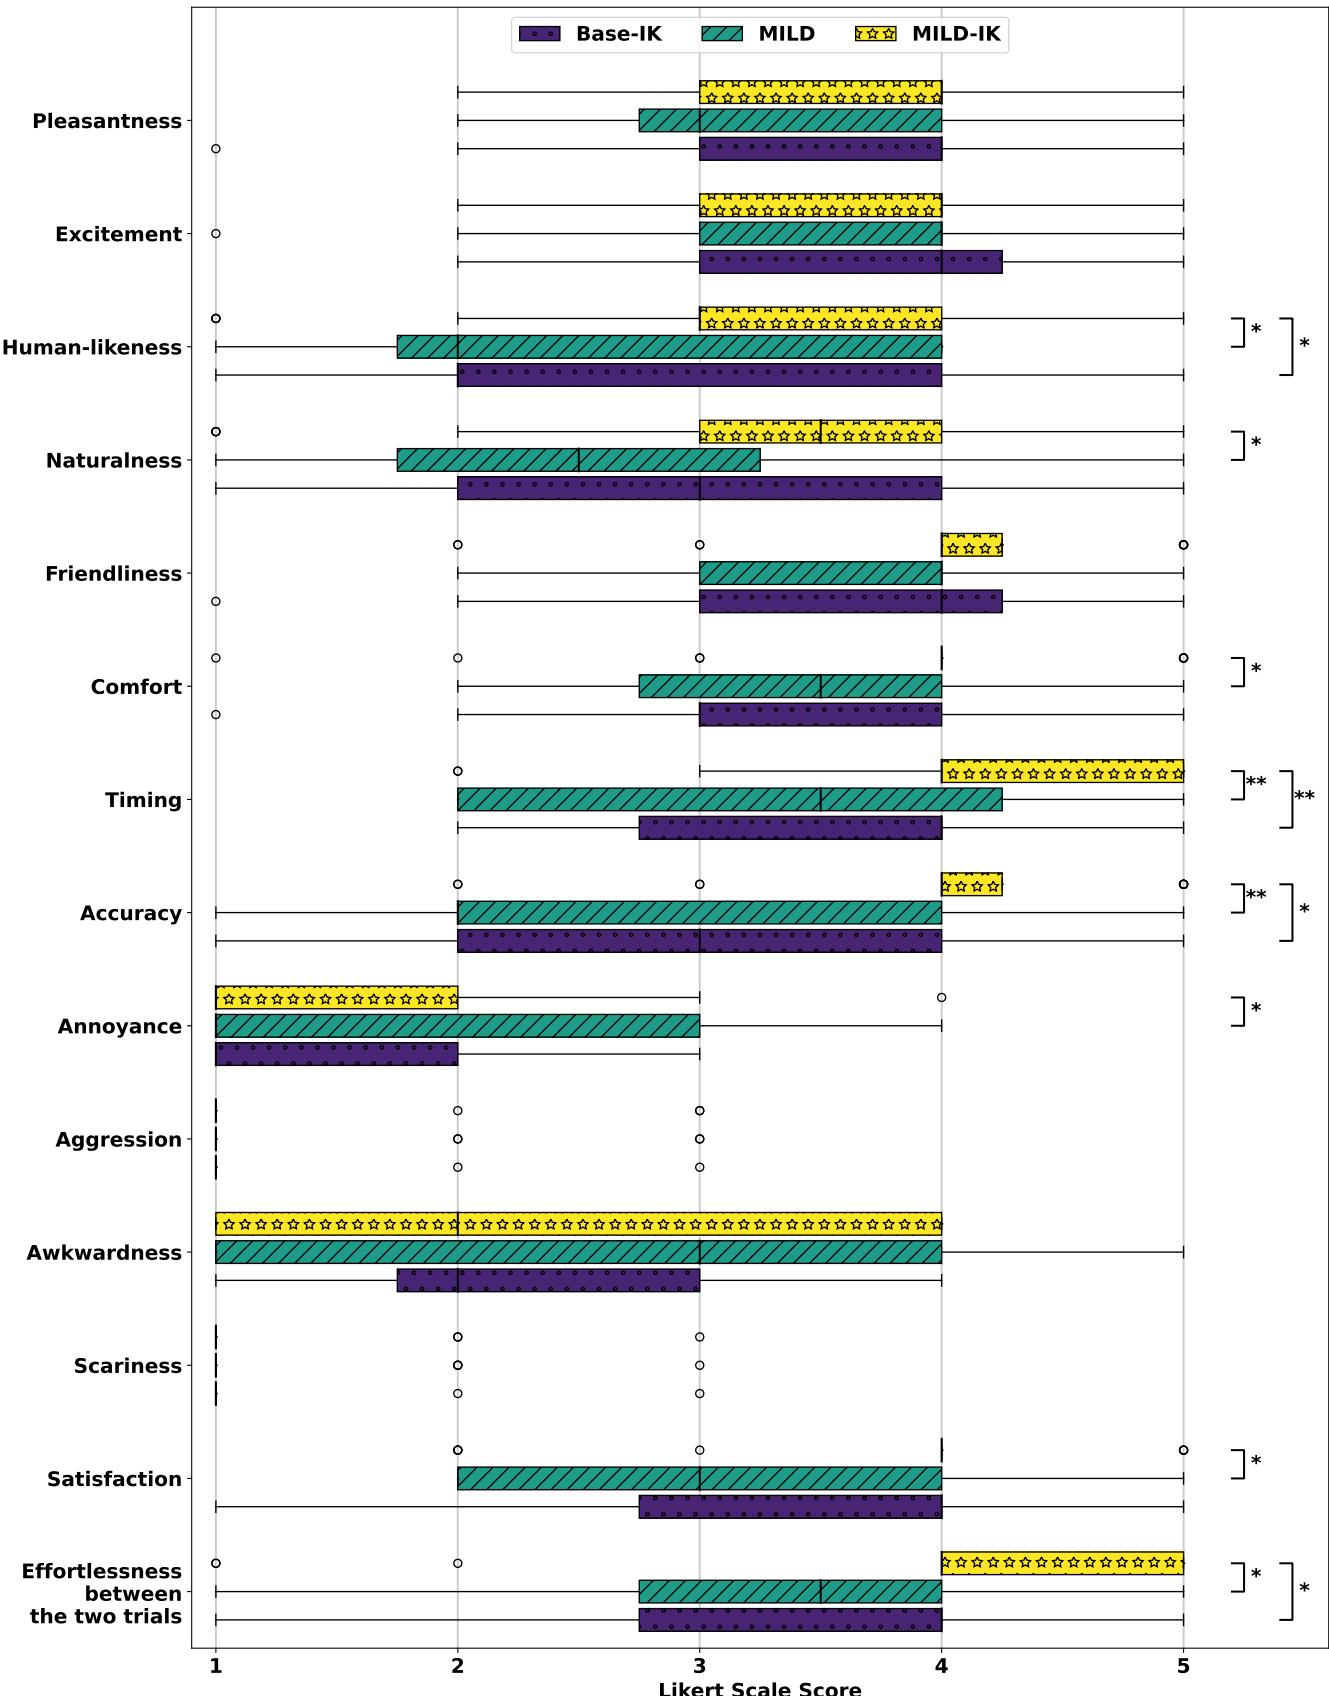
Fig. 10: Boxplot of the comparison between the three different approaches. (* - $p < 0.05$, ** - $p < 0.01$)

TABLE IV: Results of a one-way Repeated Measures ANOVA for each of the survey items. Values less than 0.01 are reported as 0. The last three columns show the mean differences between the different algorithms (1 - Base-IK, 2 - MILD, 3 - MILD-IK) which were analyzed using paired sample t-tests (* - $p < 0.05$, ** - $p < 0.01$).

Survey Item	Anova Results		Base-IK		MILD		MILD-IK		$\Delta\mu_{1,2}$	$\Delta\mu_{3,2}$	$\Delta\mu_{3,1}$
	$F_{2,38}$	p -value	Mean	SD	Mean	SD	Mean	SD			
Pleasantness	0.86	0.43	3.50	1.07	3.40	1.07	3.75	0.89	0.10	0.35	0.25
Excitement	1.32	0.28	3.85	0.85	3.45	1.02	3.75	0.89	0.40	0.30	-0.10
Human-likeness	4.89	0.01	2.60	1.20	2.45	1.16	3.25	1.22	0.15	0.80*	0.65*
Naturalness	3.46	0.04	2.95	1.24	2.55	1.20	3.25	1.18	0.40	0.70*	0.30
Friendliness	0.56	0.58	3.75	1.04	3.75	0.83	3.95	0.86	0.00	0.20	0.20
Comfort	2.79	0.07	3.35	1.06	3.35	0.96	3.85	0.96	0.00	0.50*	0.50
Timing	5.77	0.01	3.60	1.07	3.45	1.16	4.15	0.96	0.15	0.70**	0.55*
Accuracy	10.15	0.00	3.15	1.28	2.55	1.24	3.95	0.86	0.60	1.40	0.80
Annoyance	2.08	0.14	1.60	0.80	1.95	1.20	1.50	0.81	-0.35	-0.45*	-0.10
Aggression	0.77	0.47	1.15	0.48	1.30	0.64	1.25	0.62	-0.15	-0.05	0.10
Awkwardness	2.29	0.12	2.30	1.05	2.75	1.34	2.35	1.24	-0.45	-0.40	0.05
Scareness	0.59	0.56	1.15	0.48	1.25	0.54	1.20	0.51	-0.10	-0.05	0.05
Satisfaction	2.54	0.09	3.30	1.14	3.15	0.96	3.80	0.87	0.15	0.65*	0.50
Effortlessness between the two trials	4.15	0.02	3.35	1.19	3.40	1.16	4.05	1.24	-0.05	0.65*	0.70*

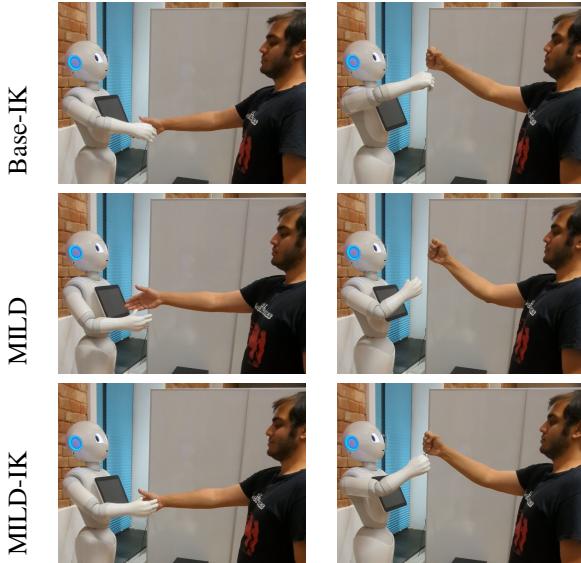


Fig. 11: Peculiarities of the different algorithms. The “Base-IK” approach reaches the human’s hand but awkwardly with the robot hand rotated inwards and the elbow pointing out. MILD maintains a human-like posture but falls short of reaching the partner’s hand due to the mismatch in motion retargeting. However, MILD-IK accurately reaches the human’s hand while maintaining a human-like posture.

E. Bimanual Robot-to-Human Handovers

To further test the effectiveness of MILD on a more complex task, we showcase the ability of MILD to reactively predict motions for a Bimanual Robot in a Robot-to-Human handover scenario, wherein the robot has to hand over an object to the human partner. For this handover experiment, we use a Bimanual Franka Emika Panda robot setup, “Kobo” that runs a dual arm cartesian impedance controller to command each arm in the 3D task space. Therefore, we train MILD directly using task space trajectories from the Bimanual Handovers dataset in [49]. We re-scale the hand trajectories of the giver in the training data to fit the robot’s task space limits and train a model using the hand trajectories of both the giver and the receiver. Since we directly predict task space trajectories from Human-Human Interactions and do not need the inverse kinematics adaptation as done with Pepper, we use MILD v1

to learn the robot motions as the giver, with the human as the receiver. We find that MILD generates suitable response motions for various objects, an example of which can be seen in Figure 12. However, some failures still occur, mainly from the lack of incorporating object-specific information, such as the object dimensions or geometry. While the failures could be mitigated with some post-processing to optimize the generated motions [34], this is currently out of the scope of our work. Please refer to the supplementary video for further examples.

V. CONCLUSION AND FUTURE WORK

In this paper, we proposed a system for learning real-world HRI from demonstrations of Human-Human Interactions (HHI). We first learned latent interaction dynamics from the HHI demonstrations in a modular manner using Hidden Markov Models (HMMs) and then demonstrated how the learned dynamics were utilized for learning HRI behaviors. We further showed how the learning of HRI behaviors was improved by incorporating the conditional distribution of the HMM into the training process. This enabled more accurate reactive motion generation during test time which we found is an important aspect of achieving a competitive performance. We then demonstrated the adaptation of the reactively generated robot trajectories during test time with Inverse Kinematics, thereby successfully combining the spatial accuracy of task-space adaptation with the ease of learning joint-space trajectories in a manner that improved the exactness of the learned behaviors. For a contact-rich task like handshaking, we additionally showed how the HMM segment predictions can be used for stiffness modulation to improve the overall perceived quality of the interaction. Through a user study, we found that our method is better perceived by human users in terms of human-likeness, timing, accuracy, and effortlessness. Our user study further validates the effectiveness of our approach in generalizing well to multiple users despite being trained on data from just two interaction partners. We additionally show the effectiveness of MILD in successfully generating Bimanual Robot-to-Human Handovers for different objects in a timely and reactive manner, which shows the usefulness of our approach even on more complex interaction scenarios.

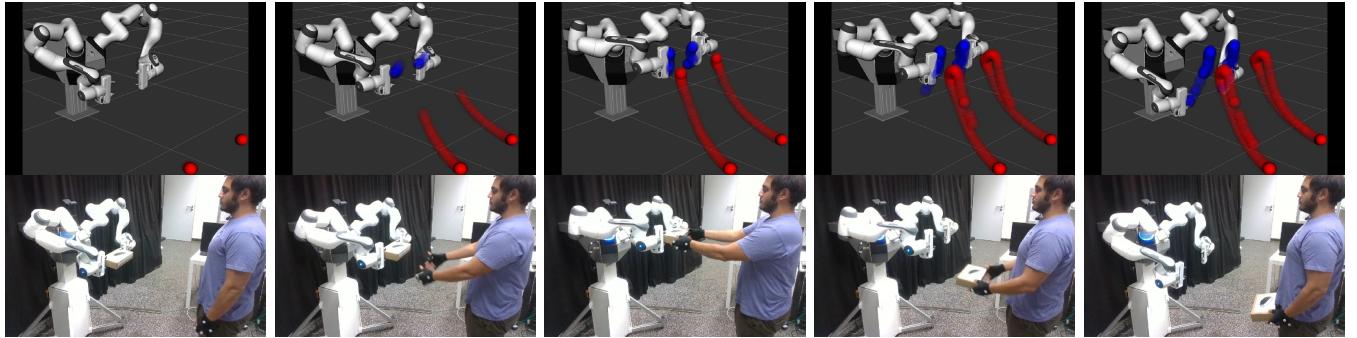


Fig. 12: Example of a Bimanual Robot-to-Human Handover Interaction generated by MILD. The top row shows the observed trajectory of the human in red, and the generated trajectory of the robot in blue. The progression of the interaction can be seen in the bottom row.

A. Limitations

While the current approach generated acceptable and accurate response trajectories, there are still some limitations to our approach which we highlight below. Currently, while training the VAE and HMM together as in the HHI scenario with the conditional loss as in the HRI scenario, we frequently encountered mode collapse of the HMM hidden states. This is why we resort to freezing the learned HMM and Human VAE for training on the HRI scenarios. This could stem from a bad approximation of the forward variable that does not accurately capture the progression of the hidden states but would rather favor just a few or a single state, thereby exacerbating the mode collapse. While this might be solved by propagating gradients through the forward variable, we found that this brings numerical issues arising from the recurrent nature of the forward variable computation, which leads to vanishing/exploding gradients (as in typical RNNs). Mitigating this issue would enable one to incorporate the HMMs into the training process in a more mathematically sound manner.

Although we train the HMMs jointly over the latent spaces of both interacting agents, during testing, we use only the human observations. This can lead to inaccuracies in predicting the state of the interaction when solely using human observations as compared to the joint set of observations of both agents (as highlighted in the appendix). Incorporating the current state of the robot and the relative geometry between the human and the robot into the training process could improve the overall predictive performance of the network.

B. Future Work

On the practical side, coming to the interaction with the robot, we purposely left out auxiliary behaviors such as speech or gaze that make the robot more “alive” as our focus was to evaluate the different interaction algorithms. These auxiliary behaviors could help improve the overall perceived quality of the interaction. The influence of the inherent personality traits also affects the interaction [19]. Further research into quantifying such influences, for example, based on the mental states of the human partner [1] or with underlying emotions [3], [87], to adapt the robot actions and personalize the interactions can subsequently provide a more natural interaction and improve the perception of the robot. For handshaking, we had to incorporate our own mechanism to ensure compliance during the contact-based segments of the interaction. Alternatively, using an underlying Cartesian

Impedance Controller can also help ensure compliance during the interaction. While some works have looked into such compliance in more static scenarios [11], further research is required to adapt this to more contact-rich and dynamic tasks like handshaking for the Pepper robot.

From the learning side, the current bottleneck of our approach, both in terms of training times and in achieving accurate predictions, is the HMM. Currently, we train the HMMs independently from the VAE in a separate step. In this regard, one could look at incorporating this into the variational framework in a more principled manner [5] or training the HMMs and VAEs in a more closely coupled manner by propagating the gradients of the expectation-maximization through the VAE could yield more suitable representations for the task at hand. Alternatively, we plan to look at going beyond HMMs by using neural variants for incorporating the multimodality, such as Mixture Density Networks [9] and additionally handle inputs from multiple sensors, which are important for functional HRI tasks in settings involving shared autonomy with a human. Additionally, incorporating a GAN-like discriminator [94] or diffusion-based losses [62] could improve performance over a VAE for more complex tasks.

ACKNOWLEDGMENTS

This work was supported by the German Research Foundation (DFG) Emmy Noether Programme (CH 2676/1-1), the EU’s Horizon Europe project “ARISE” (Grant no.: 101135959), the German Federal Ministry of Education and Research (BMBF) Projects “IKIDA” (Grant no.: 01IS20045), “KompAKI” (Grant no.: 02L19C150), and “RiG” (Grant no.: 16ME1001), the Förderverein für Marktorientierte Unternehmensführung, Marketing und Personal management e.V., the Leap in Time Stiftung. The authors gratefully acknowledge the compute provided by the high-performance computer Lichtenberg II at TU Darmstadt, funded by BMBF and the State of Hesse. The authors thank Louis Sterker for helping with the experimental setup and Sven Schultze for helping with the NuiSI dataset. The authors thank Snehal Jauhri, S. Phani Teja, and the reviewers for their constructive comments on the paper. We also thank Alap Kshirsagar, Judith Bütepage, Emmanuel Pignat, Pierre Manceron, and Roberto Calandra for open-sourcing their work that helped develop this paper.

REFERENCES

[1] N. Abdulazeem and Y. Hu, "Human factors considerations for quantifiable human states in physical human-robot interaction: A literature review," *Sensors*, 2023.

[2] A. Ajoudani, A. M. Zanchettin, S. Ivaldi, A. Albu-Schäffer, K. Kosuge, and O. Khatib, "Progress and prospects of the human-robot collaboration," *Autonomous Robots*, 2018.

[3] M. Ammi, V. Demulier, S. Caillou, Y. Gaffary, Y. Tsalamalal, J.-C. Martin, and A. Tapus, "Haptic human-robot affective interaction in a handshaking social protocol," in *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2015.

[4] H. B. Amor, G. Neumann, S. Kamthe, O. Kroemer, and J. Peters, "Interaction primitives for human-robot cooperation tasks," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014.

[5] O. Arenz, P. Dahlinger, Z. Ye, M. Volpp, and G. Neumann, "A unified perspective on natural gradient variational inference with gaussian mixture models," *Transactions on Machine Learning Research*, 2023.

[6] S. Balasubramanian, A. Melendez-Calderon, A. Roby-Brami, and E. Burdet, "On the analysis of movement smoothness," *Journal of neuroengineering and rehabilitation*, vol. 12, pp. 1–11, 2015.

[7] C. Bartneck, D. Kulic, E. Croft, and S. Zoghbi, "Godspeed questionnaire series," *International Journal of Social Robotics*, 2008.

[8] P. Becker, H. Pandya, G. Gebhardt, C. Zhao, C. J. Taylor, and G. Neumann, "Recurrent kalman networks: Factorized inference in high-dimensional deep feature spaces," in *International Conference on Machine Learning (ICML)*, 2019.

[9] C. M. Bishop, "Mixture density networks," 1994.

[10] S. Bitzer and S. Vijayakumar, "Latent spaces for dynamic movement primitives," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2009.

[11] A. Bolotnikova, S. Courtois, and A. Kheddar, "Compliant robot motion regulated via proprioceptive sensor based contact observer," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2018.

[12] J. K. Burgoon, "Relational message interpretations of touch, conversational distance, and posture," *Journal of Nonverbal behavior*, 1991.

[13] J. Bütepage, A. Ghadirzadeh, Ö. Ö. Karadag, M. Björkman, and D. Kragic, "Imitating by generating: Deep generative models for imitation of interactive tasks," *Frontiers in Robotics and AI*, 2020.

[14] S. Calinon, "A tutorial on task-parameterized movement learning and retrieval," *Intelligent Service Robotics*, 2016.

[15] S. Calinon, P. Evrard, E. Gribovskaya, A. Billard, and A. Kheddar, "Learning collaborative manipulation tasks by demonstration using a haptic interface," in *International Conference on Advanced Robotics (ICAR)*, 2009.

[16] J. Campbell and H. B. Amor, "Bayesian interaction primitives: A slam approach to human-robot interaction," in *Conference on Robot Learning (CoRL)*, 2017.

[17] J. Campbell, A. Hitzmann, S. Stepputtis, S. Ikemoto, K. Hosoda, and H. B. Amor, "Learning interactive behaviors for musculoskeletal robots using bayesian interaction primitives," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.

[18] P. Capdepuy, S. Bock, W. Benyaala, and J. Laplace, "Improving human-robot physical interaction with inverse kinematics learning," in *International Conference on Social Robotics (ICSR)*. Springer, 2015.

[19] W. F. Chaplin, J. B. Phillips, J. D. Brown, N. R. Clanton, and J. L. Stein, "Handshaking, gender, personality, and first impressions," *Journal of personality and social psychology*, 2000.

[20] M. Chaveroche, A. Malaisé, F. Colas, F. Charpillet, and S. Ivaldi, "A variational time series feature extractor for action prediction," *arXiv preprint arXiv:1807.02350*, 2018.

[21] N. Chen, J. Bayer, S. Urban, and P. Van Der Smagt, "Efficient movement representation by embedding dynamic movement primitives in deep autoencoders," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2015.

[22] N. Chen, M. Karl, and P. Van Der Smagt, "Dynamic movement primitives in latent space of time-dependent variational autoencoders," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2016.

[23] J. Chung, K. Kastner, L. Dinh, K. Goel, A. C. Courville, and Y. Bengio, "A recurrent latent variable model for sequential data," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.

[24] A. Colomé, G. Neumann, J. Peters, and C. Torras, "Dimensionality reduction for probabilistic movement primitives," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2014.

[25] H. Dai, B. Dai, Y.-M. Zhang, S. Li, and L. Song, "Recurrent hidden semi-markov model," in *International Conference on Learning Representations (ICLR)*, 2016.

[26] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the royal statistical society: series B (methodological)*, 1977.

[27] O. Dermy, M. Chaveroche, F. Colas, F. Charpillet, and S. Ivaldi, "Prediction of human whole-body movements with ae-prompts," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2018.

[28] P. Evrard, E. Gribovskaya, S. Calinon, A. Billard, and A. Kheddar, "Teaching physical collaborative tasks: object-lifting case study with a humanoid," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2009.

[29] M. Ewerton, G. Neumann, R. Lioutikov, H. B. Amor, J. Peters, and G. Maeda, "Learning multiple collaborative tasks with a mixture of interaction primitives," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015.

[30] O. Fabius, J. R. van Amersfoort, and D. P. Kingma, "Variational recurrent auto-encoders," in *ICLR (Workshop)*, 2015.

[31] L. Fritzsche, F. Unverzag, J. Peters, and R. Calandra, "First-person tele-operation of a humanoid robot," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2015.

[32] A. Gallace and C. Spence, "The science of interpersonal touch: an overview," *Neuroscience & Biobehavioral Reviews*, vol. 34, no. 2, pp. 246–259, 2010.

[33] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.

[34] Y. Göksu, A. De Almeida Correia, V. Prasad, A. Kshirsagar, D. Koert, J. Peters, and G. Chalvatzaki, "Kinematically constrained human-like bimanual robot-to-human handovers," in *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2024.

[35] S. Gomez-Gonzalez, G. Neumann, B. Schölkopf, and J. Peters, "Adaptation and robust learning of probabilistic movement primitives," *IEEE Transactions on Robotics (T-RO)*, 2020.

[36] F. Hahne, V. Prasad, A. Kshirsagar, D. Koert, R. M. Stock-Homburg, J. Peters, and G. Chalvatzaki, "Transition state clustering for interaction segmentation and learning," in *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2024.

[37] J. Han, M. R. Min, L. Han, L. E. Li, and X. Zhang, "Disentangled recurrent wasserstein autoencoder," in *International Conference on Learning Representations (ICLR)*, 2021.

[38] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "beta-vae: Learning basic visual concepts with a constrained variational framework," in *International Conference on Learning Representations (ICLR)*, 2016.

[39] N. J. Higham, "Computing a nearest symmetric positive semidefinite matrix," *Linear algebra and its applications*, 1988.

[40] S. Hignett and L. McAtamney, "Rapid entire body assessment (reba)," *Applied ergonomics*, vol. 31, no. 2, pp. 201–205, 2000.

[41] E. S. Ho, T. Komura, and C.-L. Tai, "Spatial relationship preserving character motion adaptation," in *ACM Special Interest Group on Computer Graphics (SIGGRAPH)*, 2010.

[42] K. S. Hone and R. Graham, "Towards a tool for the subjective assessment of speech system interfaces (sassi)," *Natural Language Engineering*, 2000.

[43] M. Karl, M. Soelch, J. Bayer, and P. Van der Smagt, "Deep variational bayes filters: Unsupervised learning of state space models from raw data," in *International Conference on Learning Representations (ICLR)*, 2017.

[44] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *International Conference on Learning Representations (ICLR)*, 2014.

[45] D. Koert, S. Trick, M. Ewerton, M. Lutter, and J. Peters, "Online learning of an open-ended skill library for collaborative tasks," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2018.

[46] R. Krishnan, U. Shalit, and D. Sontag, "Structured inference networks for nonlinear state space models," in *AAAI Conference on Artificial Intelligence (AAAI)*, 2017.

[47] S. Krishnan, A. Garg, S. Patil, C. Lea, G. Hager, P. Abbeel, and K. Goldberg, "Transition state clustering: Unsupervised surgical trajectory segmentation for robot learning," *The International Journal of Robotics Research (IJRR)*, 2017.

[48] O. Kroemer, C. Daniel, G. Neumann, H. Van Hoof, and J. Peters, "Towards learning hierarchical skills for multi-phase manipulation tasks," in *IEEE international conference on robotics and automation (ICRA)*. IEEE, 2015.

[49] A. Kshirsagar, R. Fortuna, Z. Xie, and G. Hoffman, "Dataset of bimanual human-to-human object handovers," *Data in Brief*, 2023.

[50] R. Lioutikov, G. Neumann, G. Maeda, and J. Peters, "Learning movement primitive libraries through probabilistic segmentation," *The International Journal of Robotics Research (IJRR)*, 2017.

[51] D. Liu, A. Honoré, S. Chatterjee, and L. K. Rasmussen, "Powering hidden markov model by neural network based generative models," in *European Conference on Artificial Intelligence (ECAI)*. IOS Press, 2020.

[52] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations (ICLR)*, 2018.

[53] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *ICML Workshop on Deep Learning for Audio, Speech and Language Processing*, 2013.

[54] G. Maeda, M. Ewerton, R. Lioutikov, H. B. Amor, J. Peters, and G. Neumann, "Learning interaction for collaborative tasks with probabilistic movement primitives," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2014.

[55] G. Maeda, M. Ewerton, T. Osa, B. Busch, and J. Peters, "Active incremental learning of robot movement primitives," in *Conference on Robot Learning (CoRL)*, 2017.

[56] P. Manceron, "Ikpy," 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.6551158>

[57] K. L. Marsh, M. J. Richardson, and R. C. Schmidt, "Social connection through joint action and interpersonal coordination," *Topics in cognitive science*, 2009.

[58] L. McAtamney and N. Corlett, "Rapid upper limb assessment (rula)," in *Handbook of human factors and ergonomics methods*. CRC Press, 2004.

[59] T. Mu and H. Su, "Boosting reinforcement learning and planning with demonstrations: A survey," *arXiv preprint arXiv:2303.13489*, 2023.

[60] M. Nagano, T. Nakamura, T. Nagai, D. Mochihashi, I. Kobayashi, and W. Takano, "Hvgh: unsupervised segmentation for high-dimensional time series using deep neural compression and statistical generative model," *Frontiers in Robotics and AI*, 2019.

[61] S. Nasiriany, T. Gao, A. Mandlikar, and Y. Zhu, "Learning and retrieval from prior data for skill-based imitation learning," in *Conference on Robot Learning (CoRL)*, 2023.

[62] E. Ng, Z. Liu, and M. Kennedy, "Diffusion co-policy for synergistic human-robot collaborative tasks," *IEEE Robotics and Automation Letters*, 2023.

[63] S. Niekum, S. Osentoski, G. Konidaris, and A. G. Barto, "Learning and generalization of complex tasks from unstructured demonstrations," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012.

[64] O. S. Oguz, W. Rampeltshammer, S. Paillan, and D. Wollherr, "An ontology for human-human interactions and learning interaction behavior policies," *ACM Transactions on Human-Robot Interaction (THRI)*, 2019.

[65] P. Oikonomou, A. Dometos, M. Khamassi, and C. S. Tzafestas, "Reproduction of human demonstrations with a soft-robotic arm based on a library of learned probabilistic movement primitives," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022.

[66] A. K. Pandey and R. Gelin, "A mass-produced sociable humanoid robot: Pepper: The first machine of its kind," *IEEE Robotics & Automation Magazine*, 2018.

[67] A. Paraschos, C. Daniel, J. Peters, and G. Neumann, "Using probabilistic movement primitives in robotics," *Autonomous Robots*, 2018.

[68] A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann, "Probabilistic movement primitives," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2013.

[69] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

[70] E. Pignat and S. Calinon, "Learning adaptive dressing assistance from human demonstration," *Robotics and Autonomous Systems (RAS)*, 2017.

[71] V. Prasad, D. Koert, R. Stock-Homburg, J. Peters, and G. Chalvatzaki, "Mild: Multimodal interactive latent dynamics for learning human-robot interaction," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2022.

[72] V. Prasad, R. Stock-Homburg, and J. Peters, "Learning human-like hand reaching for human-robot handshaking," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

[73] B. Rammstedt and O. P. John, "Measuring personality in one minute or less: A 10-item short version of the big five inventory in english and german," *Journal of research in Personality*, 2007.

[74] D. Rao, F. Sadeghi, L. Hasenclever, M. Wulfmeier, M. Zambelli, G. Vezzani, D. Tirumala, Y. Aytar, J. Merel, N. Heess, et al., "Learning transferable motor skills with hierarchical latent mixture policies," in *International Conference on Learning Representations (ICLR)*, 2021.

[75] D. J. Rezende, S. Mohamed, and D. Wierstra, "Stochastic backpropagation and approximate inference in deep generative models," in *International Conference on Machine Learning (ICML)*, 2014.

[76] L. Rozo, J. Silverio, S. Calinon, and D. G. Caldwell, "Learning controllers for reactive and proactive behaviors in human-robot collaboration," *Frontiers in Robotics and AI*, 2016.

[77] A. Saarinen, V. Harjunen, I. Jasinskaja-Lahti, I. P. Jääskeläinen, and N. Ravaja, "Social touch experience in different contexts: A review," *Neuroscience & Biobehavioral Reviews*, vol. 131, pp. 360–372, 2021.

[78] S. Schaal, "Dynamic movement primitives-a framework for motor control in humans and humanoid robotics," in *Adaptive motion of animals and machines*, 2006.

[79] N. Sebanz, H. Bekkering, and G. Knoblich, "Joint action: bodies and minds moving together," *Trends in cognitive sciences*, 2006.

[80] N. Sebanz and G. Knoblich, "Prediction in joint action: What, when, and where," *Topics in cognitive science*, 2009.

[81] F. Semeraro, A. Griffiths, and A. Cangelosi, "Human-robot collaboration and machine learning: A systematic review of recent research," *Robotics and Computer-Integrated Manufacturing*, 2023.

[82] A. Shahroudy, J. Liu, T.-T. Ng, and G. Wang, "Ntu rgb+ d: A large scale dataset for 3d human activity analysis," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[83] T. Shankar and A. Gupta, "Learning robot skills with temporal variational inference," in *International Conference on Machine Learning (ICML)*, 2020.

[84] T. Shu, X. Gao, M. S. Ryoo, and S.-C. Zhu, "Learning social affordance grammar from videos: Transferring human interactions to human-robot interactions," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

[85] T. Shu, M. S. Ryoo, and S.-C. Zhu, "Learning social affordance for human-robot interaction," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

[86] T. Shu, M. Ryoo, and S.-C. Zhu, "Learning social affordance for human-robot interaction," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.

[87] R. Stock-Homburg, "Survey of emotions in human-robot interactions: Perspectives from robotic psychology on 20 years of research," *International Journal of Social Robotics*, 2022.

[88] D. S. Syrdal, K. Dautenhahn, K. L. Koay, and M. L. Walters, "The negative attitudes towards robots scale and reactions to robot behaviour in a live human-robot interaction study," *Adaptive and emergent behaviour and complex systems*, 2009.

[89] D. Tanneberg, K. Ploeger, E. Rueckert, and J. Peters, "Skid raw: Skill discovery from raw trajectories," *IEEE Robotics and Automation Letters (RA-L)*, 2021.

[90] M. Tavassoli, S. Katyara, M. Pozzi, N. Deshpande, D. G. Caldwell, and D. Prattichizzo, "Learning skills from demonstrations: A trend from motion primitives to experience abstraction," *IEEE Transactions on Cognitive and Developmental Systems*, 2023.

[91] C. Van Gemeren, R. Poppe, and R. C. Veltkamp, "Spatio-temporal detection of fine-grained dyadic human interactions," in *Human Behavior Understanding: 7th International Workshop, HBU 2016, Amsterdam, The Netherlands, October 16, 2016, Proceedings 7*. Springer, 2016, pp. 116–133.

[92] P. Vinayavekhin, M. Tatsubori, D. Kimura, Y. Huang, G. D. Magistris, A. Munawar, and R. Tachibana, "Human-like hand reaching by motion prediction using long short-term memory," in *International Conference on Social Robotics (ICSR)*. Springer, 2017.

[93] D. Vogt, S. Stepputis, S. Grehl, B. Jung, and H. B. Amor, "A system for learning continuous human-robot interactions from human-human demonstrations," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

[94] C. Wang, C. Pérez-D'Arpino, D. Xu, L. Fei-Fei, K. Liu, and S. Savarese, "Co-gail: Learning diverse strategies for human-robot collaboration," in *Conference on Robot Learning*. PMLR, 2022.

[95] X. Zhao, S. Chumkamon, S. Duan, J. Rojas, and J. Pan, "Collaborative human-robot motion generation using lstm-rnn," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2018.

[96] Y. Zhou, J. Gao, and T. Asfour, "Learning via-point movement primitives with inter-and extrapolation capabilities," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.

TABLE A.1: p -values for pairwise MSE comparisons of the results reported in Table II. For $p < 0.001$, we report zeros.

		Version of MILD				
		v1	v2.1	v2.2	v3.1	v3.2
HRI-Yumi [13]	[13]	0.	0.	0.	0.	0.
	MILD v1	–	0.	0.	0.	0.
	MILD v2.1	–	–	0.	0.284	0.337
	MILD v2.2	–	–	–	0.	0.
	MILD v3.1	–	–	–	–	0.675
	[13]	0.	0.	0.397	0.	0.
	MILD v1	–	0.	0.	0.	0.
	MILD v2.1	–	–	0.018	0.	0.
	MILD v2.2	–	–	–	0.	0.
	MILD v3.1	–	–	–	–	0.008
HRI-Pepper [13]	[13]	0.	0.	0.	0.976	0.999
	MILD v1	–	0.	0.	0.	0.
	MILD v2.1	–	–	0.	0.	0.
	MILD v2.2	–	–	–	0.	0.
	MILD v3.1	–	–	–	–	0.909
	[13]	0.	0.136	0.098	0.	0.
	MILD v1	–	0.	0.	0.055	0.
	MILD v2.1	–	–	0.785	0.	0.
	MILD v2.2	–	–	–	0.	0.
	MILD v3.1	–	–	–	–	0.003
HRI-Pepper (NuiSI)	[13]	0.	0.	0.	0.	0.
	MILD v1	–	0.	0.	0.	0.
	MILD v2.1	–	–	0.	0.	0.
	MILD v2.2	–	–	–	0.	0.
	MILD v3.1	–	–	–	–	0.119
	[13]	0.	0.	0.	0.	0.
	MILD v1	–	0.	0.	0.	0.
	MILD v2.1	–	–	0.	0.	0.
	MILD v2.2	–	–	–	0.	0.
	MILD v3.1	–	–	–	–	0.041
Parachute Fistbump	[13]	0.	0.291	0.838	0.	0.
	MILD v1	–	0.	0.	0.	0.
	MILD v2.1	–	–	0.	0.	0.
	MILD v2.2	–	–	–	0.	0.
	MILD v3.1	–	–	–	–	0.
	[13]	0.	0.	0.	0.	0.
	MILD v1	–	0.	0.	0.	0.
	MILD v2.1	–	–	0.016	0.	0.
	MILD v2.2	–	–	–	0.	0.
	MILD v3.1	–	–	–	–	0.018
Parachute Fistbump	[13]	0.	0.	0.139	0.158	0.309
	MILD v1	–	0.	0.	0.	0.
	MILD v2.1	–	–	0.	0.	0.
	MILD v2.2	–	–	–	0.004	0.530
	MILD v3.1	–	–	–	–	0.014
	[13]	0.	0.	0.	0.	0.
	MILD v1	–	0.	0.	0.	0.
	MILD v2.1	–	–	0.862	0.095	0.003
	MILD v2.2	–	–	–	0.104	0.007
	MILD v3.1	–	–	–	–	0.065
Parachute Fistbump	[13]	0.	0.	0.	0.	0.
	MILD v1	–	0.	0.	0.	0.
	MILD v2.1	–	–	0.036	0.	0.
	MILD v2.2	–	–	–	0.050	0.023
	MILD v3.1	–	–	–	–	0.685
	[13]	0.	0.	0.	0.	0.
	MILD v1	–	0.	0.	0.	0.
	MILD v2.1	–	–	0.026	0.017	0.
	MILD v2.2	–	–	–	0.808	0.018
	MILD v3.1	–	–	–	–	0.023



Vignesh Prasad is a Postdoctoral Researcher at the Interactive Robot Perception and Learning (PEARL) in the Computer Science Department of the Technical University of Darmstadt, Germany, working on on Robot Learning, Computer Vision, and Human-Robot Interaction. He pursued a Master's degree in Computer Science with a focus on robotics at the Robotic Research Center (RRC), IIIT Hyderabad, India. He obtained a Ph.D. from the Intelligent Autonomous Systems Group (IAS) at the Technical University of Darmstadt, Germany, in 2023.



Lea Heitlinger received her Ph.D. at the Chair of Marketing and Human Resource Management at the Technical University of Darmstadt, Germany, in 2024. Her research focuses on the psychological aspects of the integration of service robots in our everyday lives from both customer and employee perspectives. Prior to her doctoral studies, she received a Bachelor's and Master's degree in Psychology from the University of Heidelberg, Germany, with a focus on Organizational Behavior & Adaptive Cognition.



Dorothea Koert is an Independent Research Group Leader of the interdisciplinary BMBF junior research group IKIDA. She has Master's degrees in Autonomous Systems and Computational Engineering with a focus on Robotics and completed her Ph.D. in 2020 with the Intelligent Autonomous Systems Group (IAS) at the Technical University of Darmstadt. Her research focus is at the intersection of Interactive Machine Learning and Human-Robot Interaction. In 2019 she received the AI-Newcomer award of the German Informatics Society.



received several best paper awards from the American Marketing Association.



her Ph.D. in 2019 at the Electrical and Computer Engineering School of the National Technical University of Athens, Greece.



Jan Peters is a full professor (W3) for Intelligent Autonomous Systems at the Computer Science Department of the Technical University of Darmstadt, dept head of the research department on Systems AI for Robot Learning (SAIROL) at the German Research Center for Artificial Intelligence (Deutsches Forschungszentrum für Künstliche Intelligenz, DFKI), and a founding research faculty member of The Hessian Center for Artificial Intelligence. He has received the Dick Volz Best 2007 US Ph.D. Thesis Runner-Up Award, RSS - Early Career Spotlight, INNS Young Investigator Award, and IEEE Robotics & Automation Society's Early Career Award, as well as numerous best paper awards. He received an ERC Starting Grant and was appointed an IEEE fellow, AIAA fellow and ELLIS fellow.