

Learning and Blending Robot Hugging Behaviors in Time and Space

Michael Drolet¹, Joseph Campbell², and Heni Ben Amor¹

Abstract—We introduce an imitation learning-based physical human-robot interaction algorithm capable of predicting appropriate robot responses in complex interactions involving a superposition of multiple interactions. Our proposed algorithm, Blending Bayesian Interaction Primitives (B-BIP) allows us to achieve responsive interactions in complex hugging scenarios, capable of reciprocating and adapting to a hug’s motion and timing. We show that this algorithm is a generalization of prior work, for which the original formulation reduces to the particular case of a single interaction, and evaluate our method through both an extensive user study and empirical experiments. Our algorithm yields significantly better quantitative prediction error and more-favorable participant responses with respect to accuracy, responsiveness, and timing, when compared to existing state-of-the-art methods.

I. INTRODUCTION

A hug is a natural embrace and one of the most common forms of social interaction in humans and animals. Beyond a simple salutation, it is an effective means to communicate affection and emotional support [1], and studies have shown hugging to cause physiological responses, thereby resulting in cardiovascular and mental health benefits [2], [3], [4]. With the advent of social robotics, e.g., robots in malls, homes, and theme parks, there is an increasing need for methods that can produce responsive and convincing hugging motions in anthropomorphic agents. However, despite its seemingly simple appearance, hugging is a nuanced and complex process of physical coordination in both time and space.

When execution fails, we are often left feeling uncomfortable, awkward, or embarrassed – the opposite of its intended effect. In turn, implementing such behaviors in robots is an extremely challenging task and is often circumvented in experimental hugging robots by not reciprocating hugs [5], [6] or executing non-adaptive, pre-defined motions [7], [8], [9]. Adaptive hugs require a robot to anticipate the type of hug performed, the current temporal progress, and the upcoming motion. Consequently, it has to generate accurate motor behavior to produce a synchronized motion with the human partner. One challenge in this regard is that hugs can be initiated anytime. Hence, the starting point is not predetermined nor easily identifiable. Another challenge is that hugs are typically fast movements with a duration of

¹M. Drolet and H. Ben Amor are with the School of Computing and Augmented Intelligence at Arizona State University, USA {mdrolet, hbenamor}@asu.edu

²J. Campbell is with the Robotics Institute at Carnegie Mellon University, USA jcampbell@cmu.edu

This work was partially supported by Honda Research Institute, as well as the National Science Foundation under grants CNS-1932068, IIS-1749783, and CNS-1932189.

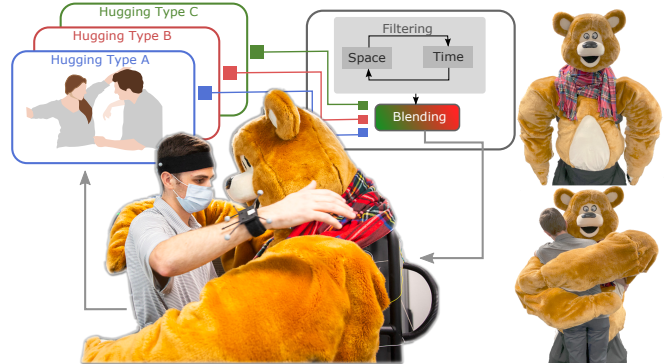


Fig. 1: Robot dressed as a plush bear hugs a human partner.

only a few seconds. Therefore, we need algorithms that allow robots to repeatedly (a) replan their motions in response to (b) the predicted behavior of the human partner. Finally, due to cultural and personal preferences, there may be a number of variations and styles of hugs.

We seek to address the above challenges by introducing a human-robot interaction framework that maintains a list of possible interactions and fluidly transitions between them in response to the interaction partner. Traditionally, HRI methods require us first to classify which interaction is occurring, then proceed with that classification for the remainder of the interaction [10], [11], or to re-classify the interaction at discrete time intervals. However, given the dynamic nature of human interactions, this leads to stilted HRI in which the robot is incapable of smoothly transitioning between discrete actions or even accounting for an interaction that is a blend of multiple actions.

We introduce a generalized version of Bayesian Interaction Primitives [12] in which an interaction may consist of multiple sub-actions. In this context, the original formulation may be considered a particular case of our proposed general form. At each time step, we assess the likelihood of a set of possible interactions based on observations of the human partner. The model associated with each possible interaction is then updated based on the observations, with the magnitude of the update weighted by the likelihood. The intuition here is straightforward: if we are observing a seemingly unlikely interaction, then we do not want to update the model significantly because the observation is unlikely to have been generated from the model. It would only serve to distort it. Aside from being able to responsively transition from one interaction to another at any point in time, our approach has two subtle advantages: a) by updating all likely interactions at every time step, we avoid a sudden discontinuous transition between discrete interactions when switching occurs, and b)

it is possible that interaction is a blend of multiple discrete interactions.

We propose the following contributions in this work:

- A generalized form of Bayesian Interaction Primitives, which removes the restriction that a single interaction consists of only a single action.
- A probabilistic formulation in which we detect action transitions and update a state approximation.
- An empirical study in which we demonstrate that our formulation can successfully detect and transition between three discrete sub-actions in a physical hugging scenario and compare it to baseline methods.

II. RELATED WORK

A. Social and Physical Human-Robot Interaction

Intimate, social pHRI, such as hugging, has been found to have positive effects on the human emotional state [13], [14], [8], [6]. Subsequently, robotic hugging has been an area of interest as it would be beneficial if robots could confer the same positive emotional benefits as humans. While passive, non-reciprocating huggable robots have been shown to yield emotional, mental, and physical health benefits [15], [6], [5], robots which actively reciprocate hugs have been found to lead to greater interaction duration [16] and are well received if the robot is responsive and comfortable to hug [7].

However, given the challenges involved in developing adaptive hugging algorithms, many robots capable of reciprocating hugs use pre-defined motions in a "one size fits all" setup [8], [17], [9]. Sometimes, limited adaptability is achieved by controlling when the hug is initiated and ended based on visual and haptic feedback [18]. More recently, recognition and reciprocation of intra-hug gestures have been investigated [19], leading to another layer of responsiveness. However, the timing and motion of the hug itself are still largely independent of the user once initiated. It is this latter adaptability for which we propose a solution in this work.

B. Learning Interactions with Multiple Sub-Actions

Methods to combine multiple motion actions, i.e., primitives, into complex movements have been extensively studied, both in terms of sequences [20], [21] and superpositions [22], [23]. However, such approaches have primarily been limited to movement generation for a robot acting independently without coordinating its behavior with a human partner.

Methods targeting human-robot interaction scenarios are often based on graphical models. For example, Tanaka et al. [24] developed a Markov model-based approach to determining which discretized state space region the human partner would be in so the robot could act appropriately. Koppula et al. [25] demonstrate a method based on conditional random fields which can anticipate future human behavior and classify the specific sub-actions. Graphical models are employed by Hawkins et al. [26] to develop a method in which a robot plans its actions to reduce an expected cost associated with a human's predicted sub-action timings.

In the method most similar to ours, Ewerton et al. [11] present a method to accommodate multiple interaction

primitives but do not support a blend or superposition of primitives. Additionally, this method assumes the phase is known/fixed (where observation comes from the last time step of interaction), and then forward control is applied to execute the trajectory. However, we address the problem more generally, in that phase is treated as a random variable. Therefore, our algorithm brings novelty in the sense that we can estimate the correct primitive at an arbitrary phase, which does not need to be provided beforehand; and due to the online, closed-loop nature of the inference algorithm, we can blend between interactions in real-time. In contrast to prior work, we present a fully probabilistic model at the motion planning level, which infers both the current action and the motion and temporal timing into the future.

III. METHODOLOGY

In this work, we model a single physical human-robot interaction as a Bayesian Interaction Primitive [12]. Each primitive explicitly models the relationship between the measured degrees of freedom from a set of training demonstrations, which is then used as a prior for inference during interactions, as depicted in Fig. 2 in the *Training* block.

The goal is for a robot to interact with a human partner in real-time and infer their (a) next actions, and (b) appropriate robot response using this prior knowledge and a sequence of observations of the human, shown in the *Testing* block of Fig. 2. We begin by describing this model in detail, then introduce an improved version, Blending BIP (B-BIP), which allows us to expand a single primitive to encompass multiple interactions while enabling transitions between them.

A. Preliminaries: Bayesian Interaction Primitives

We first train the primitive using a set of demonstrations of the desired interaction. Each demonstration consists of *observed* and *controlled* degrees of freedom, which we model as the matrix $\mathbf{Y} \in \mathbb{R}^{D \times T}$ where $D = |D_c| + |D_o|$ denotes the total number of degrees of freedom (DoFs) in the interaction (having the sets of DoF indices D_c from the controlled agent and D_o from the observed agent) and T samples.

These demonstrations are then transformed to a time invariant latent space \mathbf{w} such that, for a given DoF d , $y_t^d = h^d(\phi(t), \mathbf{w}^d) = \Phi_{\phi(t)}^\top \mathbf{w}^d + \epsilon_y$, where $\Phi_{\phi(t)} \in \mathbb{R}^{1 \times B^d}$ is a row vector of B^d basis functions, $\mathbf{w}^d \in \mathbb{R}^{B^d \times 1}$, and ϵ_y is i.i.d. Gaussian noise. We use the standard Gaussian basis function in this work, although others may be selected if appropriate for the task domain. The full latent state representation for a demonstration \mathbf{w} is obtained by concatenating each degree of freedom together, and may be solved for using standard optimization techniques such as least squares. As is standard [12], the basis functions are dependent on a relative time measure referred to as phase $\phi(t)$.

During inference, we wish to estimate the latent state \mathbf{w} from which the inferred controlled DoFs can be retrieved, however, it is also necessary to localize both the phase, ϕ , and phase velocity, $\dot{\phi}$, in order to accurately perform inference on demonstrations done at different speeds. Therefore, we augment the state representation to be $\mathbf{s}_t = [\phi_t, \dot{\phi}_t, \mathbf{w}_t]$.




Fig. 2: An overview of B-BIP. Top: training demonstrations (left) are decomposed into a latent space (middle) and transformed into an ensemble of samples (right). Bottom: observations are collected during a live interaction (left) which is used to perform filtering with the learned ensemble (middle) and produce a response trajectory (right).

Given a sequence of measurements, $\mathbf{Y}_{1:t}$, of all modalities, we have the following probabilistic formulation:

$$p(\mathbf{s}_t | \mathbf{Y}_{1:t}, \mathbf{s}_0) \propto p(\mathbf{y}_t | \mathbf{s}_t) p(\mathbf{s}_t | \mathbf{Y}_{1:t-1}, \mathbf{s}_0). \quad (1)$$

As in [27], this posterior density is approximated using a Monte Carlo method in which an ensemble of samples are used to predict the next measurement:

$$\mathbf{x}_{t|t-1}^j = g(\mathbf{x}_{t-1|t-1}^j) + \mathcal{N}(0, \mathbf{Q}). \quad (2)$$

The predicted ensemble is updated based on the error obtained from the actual measurement using a gain coefficient \mathbf{K} , the equations for which are omitted due to space constraints. The interested reader can refer to the standard Ensemble Kalman Filter (EnKF) [28] algorithm for more information.

$$\mathbf{x}_{t|t}^j = \mathbf{x}_{t|t-1}^j + \mathbf{K}(\tilde{\mathbf{y}}_t - h(\mathbf{x}_{t|t-1}^j)). \quad (3)$$

The set of training demonstrations that we start with are used to directly initialize the ensemble members, such that E demonstrations yields E ensemble members.

B. Modeling Multiple Primitives

The BIP framework as described above only supports modeling a single interaction at a time. In order to extend this model to a set of interactions, we first present a probabilistic formulation for what this entails. Suppose we have an ensemble \mathbf{X} in which each of the E ensemble members belongs to a class $c \in \mathcal{C}$, in which the set of classes \mathcal{C} represents different discrete interactions. This allows us to partition \mathbf{X} into $|\mathcal{C}|$ sub-ensembles, such that $\mathbf{X} = \mathbf{X}^{(1)} \cup \mathbf{X}^{(2)} \cup \dots \cup \mathbf{X}^{(|\mathcal{C}|)}$ where a sub-ensemble $\mathbf{X}^{(c)}$ contains $E^{(c)}$ members such that $E = \sum_{c \in \mathcal{C}} E^{(c)}$. We define C as a random variable over the set \mathcal{C} which indicates the class of an interaction. Each sub-ensemble $\mathbf{X}^{(c)}$ is a Monte Carlo approximation of the probability distribution for the c -th interaction, $p(\mathbf{s}_t | \mathbf{Y}_{1:t}, \mathbf{s}_0, C = c)$, for which we can marginalize out C to re-obtain the full posterior distribution:

$$p(\mathbf{s}_t | \mathbf{Y}_{1:t}, \mathbf{s}_0) = \sum_{c \in \mathcal{C}} p(\mathbf{s}_t | \mathbf{Y}_{1:t}, \mathbf{s}_0, C = c) p(C = c | \mathbf{Y}_{1:t}, \mathbf{s}_0). \quad (4)$$

The association of each ensemble member to a class c is static and defined in the prior distribution \mathbf{s}_0 , as demonstrations must be initially collected for each individual interaction and hence we have a mapping from demonstrations

to classes. This allows us to calculate the posterior for a specific class, $p(\mathbf{s}_t | \mathbf{Y}_{1:t}, \mathbf{s}_0, C = c)$, in a similar manner as Eqs. 2-3 but limited to the ensemble members $\mathbf{x} \in \mathbf{X}^{(c)}$; this is covered in Sec. III-D. We do not restrict ourselves to the case that C is fixed to a single value c ; an interaction may transition between multiple classes over time. Therefore, the interaction scenarios examined in previous works are special cases of this formulation and only take on one class value.

C. Interaction Detection

In this work, we perform Reduced-Rank Linear Discriminant Analysis (LDA) for computing the probability of the interaction class, given the current observations of the human. For each interaction class, $c \in \mathcal{C}$, let $\mathbf{Y}_i^{(c)} \in \mathbb{R}^{D \times T}$ represent the i 'th demonstration from the set of training demonstrations of the corresponding class. Here, $1 \leq i \leq N^{(c)}$, with $N^{(c)}$ denoting the total number of demonstrations for class c . Let the within-class demonstration matrix be defined by $\mathbf{M}^{(c)} = [(\mathbf{Y}_1^{(c)})_{D_o,:}, \dots, (\mathbf{Y}_{N^{(c)}}^{(c)})_{D_o,:}]$ and the between-class demonstration matrix be defined by $\mathbf{M} = [\mathbf{M}^{(1)}, \dots, \mathbf{M}^{(|\mathcal{C}|)}]^\top$. Next, a low-rank representation of \mathbf{M} is found using Multiple Discriminant Analysis [29], [30]. To perform dimensionality reduction, we compute the solution of the Rayleigh coefficient, which is the ratio of the between class scatter to within class scatter. Let \mathbf{S}_W denote the within-class scatter matrix, i.e., the prior-weighted sum of within-class covariance matrices. Here, $\boldsymbol{\mu}^{(c)} \in \mathbb{R}^{D_o}$ is the mean of the columns of $\mathbf{M}^{(c)}$.

The total scatter matrix, \mathbf{S}_T , is the covariance over the dataset, \mathbf{M} , and the between-class scatter matrix is defined as $\mathbf{S}_B = \mathbf{S}_T - \mathbf{S}_W$. We obtain the eigenvectors, \mathbf{w}_i , of the transformation matrix, \mathbf{W} , which maximize the ratio of between-class scatter to within-class scatter by solving $(\mathbf{S}_B - \lambda_i \mathbf{S}_W) \mathbf{w}_i = 0$. Let \mathbf{W}_k be the reduced-rank representation of \mathbf{W} , with eigenvectors corresponding to the $k = |\mathcal{C}| - 1$ largest eigenvalues. The distribution of samples are projected onto the k -dimensional subspace spanned by \mathbf{W}_k . For notational simplicity, let $\mathbf{Z} = \mathbf{W}_k^\top \mathbf{M}$ and $\mathbf{z}_t = \mathbf{W}_k^\top \mathbf{y}_t^{D_o}$. The posterior density is computed as:

$$\begin{aligned} \log p(C = c | \mathbf{Y}_{1:t}, \mathbf{s}_0) &= \log p(\mathbf{Y}_{1:t}, \mathbf{s}_0 | C = c) + \log p(C = c) + \eta \\ &= -\frac{1}{2} (\mathbf{z}_t - \boldsymbol{\mu}_{\mathbf{Z}^{(c)}})^\top \Sigma_{\mathbf{Z}}^{-1} (\mathbf{z}_t - \boldsymbol{\mu}_{\mathbf{Z}^{(c)}}) \\ &\quad + \log p(C = c) + \eta. \end{aligned} \quad (5)$$



Fig. 3: An example of a *left-high* to *right-high* interaction. Left: The participant starts with a left-high interaction. Middle: When switching to the right-high hug, the robot responds accordingly. Right: The participant hugs the robot.

After dropping the quadratic term $\mathbf{z}_t^\top \Sigma_{\mathbf{Z}}^{-1} \mathbf{z}_t$ from 5, which is independent of c , we get the resulting form:

$$\log p(C = c | \mathbf{Y}_{1:t}, \mathbf{s}_0) = \beta_c^\top \mathbf{z}_t + \beta_{c0}. \quad (6)$$

where $\beta_c = \Sigma_{\mathbf{Z}}^{-1} \boldsymbol{\mu}_{\mathbf{Z}(c)}$ and $\beta_{c0} = -\frac{1}{2} \boldsymbol{\mu}_{\mathbf{Z}(c)}^\top \Sigma_{\mathbf{Z}}^{-1} \boldsymbol{\mu}_{\mathbf{Z}(c)} + \log \pi^{(c)}$. The posterior can now be computed by applying the softmax function:

$$p(C = c | \mathbf{Y}_{1:t}, \mathbf{s}_0) = \frac{e^{\beta_c^\top \mathbf{z}_t + \beta_{c0}}}{\sum_{j \in \mathcal{C}} e^{\beta_j^\top \mathbf{z}_t + \beta_{j0}}}. \quad (7)$$

We assume a prior proportional to the number of samples in each training set; i.e., $p(C = c) = \pi^{(c)} = \frac{N^{(c)}}{\sum_{j \in \mathcal{C}} N^{(j)}}$.

D. Interaction Transition

When calculating Eq. 4 we must be careful to weight the magnitude of the ensemble update with the class probability. The intuitive reason is that the standard ensemble member update of Eq. 3 assumes that the observation was generated from a distribution approximated by that ensemble. However, if there is only a small probability that the observation was generated from a particular sub-ensemble and we apply a full-magnitude update then we potentially skew the ensemble members with an out-of-distribution measurement. Thus, unlike Eq. 3, we now weight the gain coefficient with Eq. 7:

$$\mathbf{x}_{t|t}^j = \mathbf{x}_{t|t-1}^j + p(C = c | \mathbf{Y}_{1:t}, \mathbf{s}_0) \mathbf{K}(\tilde{\mathbf{y}}_t - h(\mathbf{x}_{t|t-1}^j)) \quad (8)$$

for all $\mathbf{x}^j \in \mathbf{X}^{(c)}$ and all $c \in \mathcal{C}$.

IV. EXPERIMENTAL DESIGN

In order to empirically evaluate our algorithm, we conduct an IRB-approved participant study consisting of a hugging scenario between a robot and a human.

A. Training Data Collection

Motion capture data is collected at 120 Hz from participants wearing a hat and a wrist band on each hand, totaling three observed modalities. 15 demonstrations of each interaction type (*left-high*, *middle*, and *right-high*) are collected from 15 different participants, totaling 225 demonstration hugs per interaction class and 675 in total. The term *left-high* is used to indicate an interaction where the left hand of the robot and the human are raised such that the human's left hand approaches the robot over its right shoulder; the same symmetry holds for the *right-high* interaction, and *middle* is used to denote a hug where the robot hugs with arms at the same height

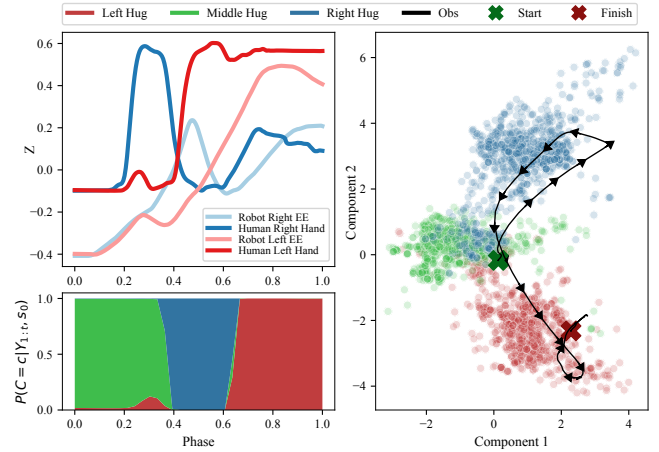


Fig. 4: Top left: The observed z positions of the participant's hands and the robot's end effectors during a right-high to left-high interaction. Bottom left: The corresponding interaction class weights for each hug type, with vertical area equal to the class probability. Right: The trajectory of the observed DoF (black arrows) projected to the reduced-rank LDA space, overlaid on the distributions (circles) for each hug type.

(see Fig. 3). We applied a response elicitation technique [31], where the robot actuates according to an open loop control policy and the human responds accordingly. Outlier training demonstrations, which have any DOF value outside of four standard deviations from the distribution of demonstrations at the same point in time, are removed from the dataset. After which, an 80 – 20 percent train-test split is used to train and validate the methods for mean-squared error (MSE) comparison. This results in 439 training demonstrations and 110 validation demonstrations.

B. Prediction Methods

The goal of the study is to compare 1) B-BIP, 2) BIP, 3) Probabilistic Movement Primitives (ProMP) [11], [32], [23], and 4) a LSTM network, all of which are evaluated on *non-switching* interactions (left-high, middle, and right-high hugs) as well as on *switching* interactions (transitions from either left-high to right-high or right-high to left-high hugs). The LSTM architecture contains: 28 hidden units (twice the number of DOFs from the robot), a dropout layer with rate of 0.2, a batch normalization layer, a fully-connected layer with 28 units, and a fully-connected layer with 14 units at the output. In the case of ProMP, phase estimation is performed via Dynamic Time Warping (DTW) as described in [32]. All methods are trained on the 439 demonstrations previously described; the alternative methods treat each interaction as one class, while Blending takes into account the interaction class labels.

C. Experimental Hypotheses

We identified three important factors to be evaluated when comparing the B-BIP to the alternative methods. After every hug, we ask the following questions:

- 1) On a scale of 1 to 5, how good was the timing of the robot during the interaction?




Fig. 5: Distribution of scores for the three questions- which are used for hypothesis tests- after switching interactions.

- 2) On a scale of 1 to 5, how well did the robot match your type of hug during the interaction?
- 3) On a scale of 1 to 5, how responsive was the robot to your motion?

We state our main hypotheses – which are applied to both non-switching and switching interactions – as follows:

- H_1 : Proposed method better matches the hug type than baseline methods.
- H_2 : Proposed method has better timing than baseline methods.
- H_3 : Proposed method elicits more responsive behavior than baseline methods.

D. Participant Study

The participant study was conducted on 22 new participants whose data was not used to train the methods, including 4 female and 18 male participants between the ages of 18 and 53. A preliminary power analysis for the Wilcoxon test (which we use to conduct our hypotheses, see Section V) indicates that a sample size of $n = 20$ is sufficient to achieve a power level of 0.9, hence the number of participants. Every participant performed 24 hugs such that each method was deployed on the three non-switching interactions (totaling 12) and on switching interactions for the remaining 12. For every hug, the participant is assigned an interaction and the robot is assigned a method, both of which are randomly drawn without replacement. In the switching interactions, participants performed left-high to right-high hugs, right-high to left-high hugs, and a randomly chosen option from the former two interactions. This setup ensures that both switching and non-switching interactions have the same number of samples for every hypothesis test.

E. Quantitative Experimental Design

In addition to the participant study, offline experiments are conducted to quantitatively evaluate the performance of our proposed algorithm against the baseline methods in terms of MSE. All methods predicted a response for the demonstrations in the validation set as well as for demonstrations from a set of unseen switching interactions. In order to get the ground truth

Switching	H_1 : p -value	H_2 : p -value	H_3 : p -value
B-BIP vs. BIP	3.62×10^{-7}	1.94×10^{-6}	6.56×10^{-8}
B-BIP vs. ProMP	6.16×10^{-7}	1.25×10^{-7}	2.27×10^{-7}
B-BIP vs. LSTM	6.64×10^{-4}	9.85×10^{-2}	1.82×10^{-4}
Supported	Yes	Partially	Yes
Non-Switching	H_1 : p -value	H_2 : p -value	H_3 : p -value
B-BIP vs. BIP	1.00	1.21×10^{-1}	1.00
B-BIP vs. ProMP	2.29×10^{-1}	3.08×10^{-3}	1.59×10^{-2}
B-BIP vs. LSTM	4.80×10^{-4}	1.40×10^{-4}	5.70×10^{-5}
Supported	Partially	Partially	Partially

TABLE I: Top: p -values for Switching Interactions. Bottom: p -values for Non-Switching Interactions. Grey cells indicate a p -value greater than $\alpha = 0.05$.

dataset for the switching interactions, we manually design hug trajectories where the robot transitions from left-high to right-high and from right-high to left-high. We collect 25 demonstrations of the left-to-right interaction and 25 demonstrations of the right-to-left interaction. Here, each of the 5 participants partake in 5 demonstrations per switching interaction, where the robot executes the designed trajectories. After which, outliers are removed in the same manner as the validation set, resulting in 39 demonstrations in total.

V. RESULTS AND DISCUSSION

In the following section, we discuss the performance of our proposed method by analyzing participant responses and MSE prediction values with regard to the given hypotheses. Figure 3 shows an example an interaction where the human switches from a left-high to a right-high hug and the robot reciprocates appropriately. Additionally, Figure 4 displays predictions from Blending BIP during a right-high to left-high interaction with a test participant. The top left plot shows the trajectories from the human and robot DoFs during the interaction. The bottom left plot shows the inferred weights for each interaction, indicating which interaction Blending BIP model estimates is active at every time step.

A. Survey Responses

Participant survey responses are visualized in Fig. 5 as a histogram, where we can qualitatively observe a difference in distributions. Notably, B-BIP yields the largest number of survey responses with a maximum score of 5 for each question. When performing hypothesis tests, we cannot make the assumption that responses across different treatments are independent; namely, every participant partakes in the same set of methods. Additionally, given that the variable of interest, *score*, takes on an integer value from 1 through 5 (as in the Likert scale) and that the scores do not appear to be normally distributed, we opt to use a paired non-parametric hypothesis test. A one-way repeated measures analysis of variance is conducted to test for differences in methods across participant responses by using the Friedman test. After obtaining a p -value of $p < 10^{-5}$, we perform post-hoc analysis by applying the two-sided Wilcoxon signed-rank test to every baseline method paired with our proposed method. To account for




Fig. 6: Top: The average of the time lags (in seconds) which maximize the correlation coefficients (below) between human and robot trajectories over all switching interactions from the participant study (see Section V-B).

multiple comparisons, a Bonferroni correction is applied. This procedure is performed separately for non-switching and switching interaction scores. For a given hypothesis test, there are 66 pairs of responses (i.e., 22 participants who perform 3 switching/non-switching interactions), and each pair contains the participant’s scores to the question of interest for the methods being compared.

The resulting significance results are shown in Table I. For switching hugs, we find that participants strongly preferred B-BIP over all baseline methods with respect to hugging type, timing, and responsiveness. Hypotheses H_1 and H_3 were fully supported with participants reporting B-BIP to offer improved timing and responsiveness over BIP, ProMP, and LSTM. Hypothesis H_2 was partially supported, with participants finding B-BIP to yield better timing than BIP and ProMP with statistical significance, but not to LSTM. Responses are more mixed for non-switching hugs, as participants did not prefer B-BIP over BIP in any category, while preferring B-BIP over ProMP only in terms of timing (H_2) and responsiveness (H_3). These results are not unexpected, as any given single interaction falls within the prior distribution modeled by BIP, which means it is not unreasonable for B-BIP to perform similarly. B-BIP was, however, preferred over the LSTM for all categories, indicating that LSTM struggled to generalize with the small number of training demonstrations.

B. Interaction Responsiveness

In order to further assess the responsiveness, we evaluate how well the robot matched switching interactions from the group of test participants. We calculate the translation (backwards in time, in seconds) which, when applied to the robot end effector trajectories, maximizes the sum of the Pearson correlation coefficients from: 1) the z position of the human’s right hand with the z position of the robot’s right end effector, and 2) the z position of the human’s left hand with the z position of the robot’s left end effector. Namely, samples from these matching modalities (from the recorded test interactions) are paired to compute the correlation coefficients, and the average is taken over all switching interactions for every method. These calculations

	Switching	Non-Switching
BIP	0.127 ± 0.008	0.040 ± 0.002
ProMP	0.128 ± 0.005	0.036 ± 0.001
LSTM	0.107 ± 0.005	0.102 ± 0.004
B-BIP	0.062 ± 0.005	0.018 ± 0.000

TABLE II: MSE values of the controlled DoFs (predicted at each time step) compared to the ground-truth response, using all validation demonstrations. Green cells indicate the method with the smallest mean values. Tukey’s HSD Test indicates statistical significance for B-BIP in all cases, having $p < 10^{-5}$ in all cases.

take advantage of the symmetrical nature of the interaction, and the intuition here is that preferable methods should exhibit strong correlations with the participant within a small time lag, especially when the switching occurs- with a perfect response having zero time delay. Figure 6, shows that B-BIP has a correlation which is maximized at the smallest time delay of all methods, and B-BIP produces the highest total correlation (i.e., is the sum of the correlations coefficients from the left and right matching modalities), with a value of 1.31 out of a maximum value of 2.0. These findings provide further support toward the hypothesis that B-BIP exhibits better responsiveness during switching interactions.

C. Quantitative Analysis

The MSE values for all controlled DoFs in the quantitative offline experiments are shown in Table II. We find that B-BIP yields significantly lower prediction errors than all baseline methods for both switching and non-switching hug types. From the significance in switching hug types, we can conclude that the interaction detection mechanism works as intended and yields more accurate inference than BIP. The significance for non-switching hugs is somewhat surprising given the lack of statistical significance in the survey responses for the non-switching hugs in Sec. V-A, and we conjecture that error values of this magnitude do not always consistently result in noticeable behavioral differences. It is clear, however, that despite non-switching hugs falling within BIP’s prior distribution, the wider distribution results in significantly larger spatio-temporal errors when compared to the per-interaction priors of B-BIP.

VI. CONCLUSION

In this paper we present a method for learning and blending human-robot interactions from demonstration. A carefully-designed user study is conducted to validate whether our method produces a more responsive, timely, and suitably-matching behavior, all of which are supported through qualitative and quantitative analysis. Notably, Blending BIP achieves 1) statistically significant participant responses on switching interactions for most hypotheses, 2) a nearly twofold reduction in mean-squared prediction error compared to the second-best methods, and 3) the highest correlation and lowest response lag with test participants.

REFERENCES

- [1] J. Holt-Lunstad, W. A. Birmingham, and K. C. Light, "Influence of a "warm touch" support enhancement intervention among married couples on ambulatory blood pressure, oxytocin, alpha amylase, and cortisol," *Psychosomatic medicine*, vol. 70, no. 9, pp. 976–985, 2008.
- [2] K. C. Light, K. M. Grewen, and J. A. Amico, "More frequent partner hugs and higher oxytocin levels are linked to lower blood pressure and heart rate in premenopausal women," *Biological psychology*, vol. 69, no. 1, pp. 5–21, 2005.
- [3] K. M. Grewen, B. J. Anderson, S. S. Girdler, and K. C. Light, "Warm partner contact is related to lower cardiovascular reactivity," *Behavioral medicine*, vol. 29, no. 3, pp. 123–130, 2003.
- [4] S. Cohen, D. Janicki-Deverts, R. B. Turner, and W. J. Doyle, "Does hugging provide stress-buffering social support? a study of susceptibility to upper respiratory infection and illness," *Psychological science*, vol. 26, no. 2, pp. 135–147, 2015.
- [5] R. Yamazaki, L. Christensen, K. Skov, C.-C. Chang, M. F. Damholdt, H. Sumioka, S. Nishio, and H. Ishiguro, "Intimacy in phone conversations: anxiety reduction for danish seniors with hugvie," *Frontiers in psychology*, vol. 7, p. 537, 2016.
- [6] H. Sumioka, A. Nakae, R. Kanai, and H. Ishiguro, "Huggable communication medium decreases cortisol levels," *Scientific reports*, vol. 3, no. 1, pp. 1–6, 2013.
- [7] A. E. Block and K. J. Kuchenbecker, "Softness, warmth, and responsiveness improve robot hugs," *International Journal of Social Robotics*, vol. 11, no. 1, pp. 49–64, 2019.
- [8] M. Shiomi, A. Nakata, M. Kanbara, and N. Hagita, "A robot that encourages self-disclosure by hug," in *International Conference on Social Robotics*. Springer, 2017, pp. 324–333.
- [9] H. Hedayati, S. Bhaduri, T. Sumner, D. Szafer, and M. D. Gross, "Hugbot: a soft robot designed to give human-like hugs," in *Proceedings of the 18th ACM International Conference on Interaction Design and Children*, 2019, pp. 556–561.
- [10] H. B. Amor, G. Neumann, S. Kamthe, O. Kroemer, and J. Peters, "Interaction primitives for human-robot cooperation tasks," in *Robotics and Automation (ICRA), 2014 IEEE International Conference on*. IEEE, 2014, pp. 2831–2837.
- [11] 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 *Robotics and Automation (ICRA), 2015 IEEE International Conference on*. IEEE, 2015, pp. 1535–1542.
- [12] J. Campbell and H. Ben Amor, "Bayesian interaction primitives: A slam approach to human-robot interaction," in *Conference on Robot Learning*, 2017, pp. 379–387.
- [13] K. Wada, T. Shibata, T. Saito, and K. Tanie, "Effects of robot-assisted activity for elderly people and nurses at a day service center," *Proceedings of the IEEE*, vol. 92, no. 11, pp. 1780–1788, 2004.
- [14] D. Rus and M. T. Tolley, "Design, fabrication and control of soft robots," *Nature*, vol. 521, no. 7553, p. 467, 2015.
- [15] C. DiSalvo, F. Gemperle, J. Forlizzi, and E. Montgomery, "The hug: an exploration of robotic form for intimate communication," in *The 12th IEEE International Workshop on Robot and Human Interactive Communication, 2003. Proceedings. ROMAN 2003*. IEEE, 2003, pp. 403–408.
- [16] M. Shiomi, A. Nakata, M. Kanbara, and N. Hagita, "Robot reciprocation of hugs increases both interacting times and self-disclosures,"
- [20] T. Kulvicius, K. Ning, M. Tamosiunaite, and F. Wörgötter, "Joining movement sequences: Modified dynamic movement primitives for robotics applications exemplified on handwriting," *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 145–157, 2011.
- International Journal of Social Robotics*, vol. 13, no. 2, pp. 353–361, 2021.
- [17] M. Gaitán-Padilla, J. C. Maldonado-Mejía, L. Fonseca, M. J. Pinto-Bernal, D. Casas, M. Múnera, and C. A. Cifuentes, "Physical human-robot interaction through hugs with castor robot," in *International Conference on Social Robotics*. Springer, 2021, pp. 814–818.
- [18] A. E. Block, S. Christen, R. Gassert, O. Hilliges, and K. J. Kuchenbecker, "The six hug commandments: design and evaluation of a human-sized hugging robot with visual and haptic perception," in *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, 2021, pp. 380–388.
- [19] A. E. Block, H. Seifi, O. Hilliges, R. Gassert, and K. J. Kuchenbecker, "In the arms of a robot: Designing autonomous hugging robots with intra-hug gestures," *arXiv preprint arXiv:2202.09935*, 2022.
- [21] J. Kober, M. Gienger, and J. J. Steil, "Learning movement primitives for force interaction tasks," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3192–3199.
- [22] B. Williams, M. Toussaint, and A. J. Storkey, "Modelling motion primitives and their timing in biologically executed movements," *Advances in neural information processing systems*, vol. 20, 2007.
- [23] A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann, "Probabilistic movement primitives," *Advances in neural information processing systems*, vol. 26, 2013.
- [24] Y. Tanaka, J. Kinugawa, Y. Sugahara, and K. Kosuge, "Motion planning with worker's trajectory prediction for assembly task partner robot," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 1525–1532.
- [25] H. S. Koppula and A. Saxena, "Anticipating human activities using object affordances for reactive robotic response," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 1, pp. 14–29, 2015.
- [26] K. P. Hawkins, N. Vo, S. Bansal, and A. F. Bobick, "Probabilistic human action prediction and wait-sensitive planning for responsive human-robot collaboration," in *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)*. IEEE, 2013, pp. 499–506.
- [27] J. Campbell, S. Stepputtis, and H. Ben Amor, "Probabilistic multimodal modeling for human-robot interaction tasks," in *Robotics: Science and Systems*, 2019.
- [28] G. Evensen, "The ensemble kalman filter: Theoretical formulation and practical implementation," *Ocean dynamics*, vol. 53, no. 4, pp. 343–367, 2003.
- [29] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. New York: Wiley, 2001.
- [30] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [31] J. Campbell, A. Hitzmann, S. Stepputtis, S. Ikemoto, K. Hosoda, and H. Ben Amor, "Learning interactive behaviors for musculoskeletal robots using bayesian interaction primitives," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019.
- [32] G. J. Maeda, G. Neumann, M. Ewerton, R. Lioutikov, O. Kroemer, and J. Peters, "Probabilistic movement primitives for coordination of multiple human–robot collaborative tasks," *Autonomous Robots*, vol. 41, no. 3, pp. 593–612, 2017.