

Probabilistic Multimodal Modeling for Human-Robot Interaction Tasks

Joseph Campbell, Simon Stepputtis, and Heni Ben Amor

School of Computing, Informatics, and Decision Systems Engineering, Arizona State University

{jacampb1, sstepput, hbenamor}@asu.edu

Abstract—Human-robot interaction benefits greatly from multimodal sensor inputs as they enable increased robustness and generalization accuracy. Despite this observation, few HRI methods are capable of efficiently performing inference for multimodal systems. In this work, we introduce a reformulation of Interaction Primitives which allows for learning from demonstration of interaction tasks, while also gracefully handling nonlinearities inherent to multimodal inference in such scenarios. We also empirically show that our method results in more accurate, more robust, and faster inference than standard Interaction Primitives and other common methods in challenging HRI scenarios.

I. INTRODUCTION

Human-robot interaction (HRI) requires constant monitoring of human behavior in conjunction with proactive generation of appropriate robot responses. This decision-making process must often contend with high levels of uncertainty due to partial observability, noisy sensor measurements, visual occlusions, ambiguous human intentions, and a number of other factors. The inclusion of sensor measurements from a variety of separate modalities, e.g., cameras, inertial measurement units, and force sensors, may provide complementary pieces of information regarding the actions and intentions of a human partner, while also increasing the robustness and safety of the interaction. Even in situations in which a complete sensor modality becomes temporarily unavailable, i.e., due to a hardware failure, other available modalities may ensure graceful degradation of the system behavior. Hence, it is critical to support decision-making in HRI with inference and control methods that can deal with a variable number of data sources, each of which may have distinctive numerical and statistical characteristics and limitations.

In this paper, we investigate how multimodal models of human-robot interaction can be efficiently learned from demonstrations and, later, used to perform reasoning, inference, and control from a collection of data sources. Fig. 1 depicts a motivating example – a robot arm catching a ball. In this example, the position of the ball can be continuously tracked using motion capture markers, but is occluded from view while in the human’s hand. Yet, even before the ball is released from the hand, the robot may already intuit the moment of release and travel direction by reading pressure information from a smart shoe, from inertial measurements on the throwing arm, or from human pose data acquired via skeletal tracking.

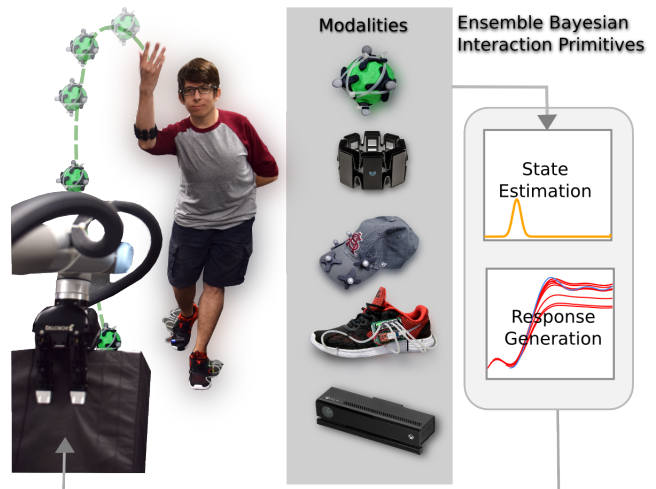


Fig. 1: A robot learning to catch a thrown ball by combining information from different modalities.

Integrating these pieces of information together, we would expect the robot to generate earlier and better predictions of the ball, as well as better estimates of necessary control signals to intercept it.

Few probabilistic inference methods for HRI have examined reasoning across multiple modalities as in the above example, with many instead opting to construct models relating only two modalities, e.g., a single observed modality to a single controlled modality. In the case of Bayesian Interaction Primitives (BIP) [5], demonstrations of two (human) agents interacting with each other are used to form a joint probability distribution among all degrees of freedom (DoF) and all modalities. During inference, this distribution is used as the prior for Bayesian filtering, which is then refined through sensor observations of the observed modalities and subsequently used to infer the controlled DoFs. However, when multiple sensing modalities are employed, several challenges quickly arise. First, in order to maintain computational tractability of the filtering process, limiting assumptions are made both about the form of the joint probability distribution, i.e., unimodal and Gaussian, as well as the linearity of the system as a whole. As the number of sensing modalities increases, each with their own unique statistical characteristics, these assumptions begin to negatively impact inference accuracy. Second, expanding the sensing modalities translates to an increased number of degrees of freedom which magnifies the computational burden and

jeopardizes real-time performance – a vital property in HRI contexts.

In this work, we propose ensemble Bayesian Interaction Primitives (eBIP) for human-robot interaction scenarios. In particular, the following contributions will be made:

- 1) An alternative formulation of interaction primitives that is particularly well-suited for inference and control in the presence of many input modalities, as well as noisy and missing observations.
- 2) An ensemble-based approach to Bayesian inference for HRI, which combines advantages of parametric and non-parametric methods. The approach requires neither an explicit covariance matrix, nor a measurement model. Measurement errors are efficiently calculated in closed-form.
- 3) Our approach allows for inference in nonlinear interactive systems, while avoiding typical inaccuracies due to either linearization errors, the parametric family of the prior, or the underlying dimensionality. Training demonstrations are used to model the non-Gaussian prior distribution of a task. The non-parametric nature of this prior avoids computational overheads and inaccuracies as found, for example, when fitting a mixture model.
- 4) Fast and efficient inference that scales particularly well with increasing dimensionality of the task.

We compare eBIP to other methods on a fast-paced, dynamic human-robot interaction experiment involving multimodal sensor streams. Experiments show that eBIP allows for accurate and rapid inference in high-dimensional spaces.

II. RELATED WORK

In the following section, we will review relevant work on probabilistic modeling of joint actions and multimodal modeling. For a detailed discussion of computational techniques in the HRI domain, see the excellent surveys in [13, 27].

Probabilistic Modeling of Joint Actions: Early work on modeling HRI scenarios using probabilistic representations focused on HMMs [23] as a method of choice, see for instance the works in [15, 26]. The ability of HMMs to perform inference in both time and space, makes them particularly interesting for collaborative and interactive tasks. However, these advantages come at a cost – HMMs require a discretization of the state space and do not scale well in high-dimensional spaces. The concept of Interaction Primitives (IP) was first proposed in [1] as an alternative approach for *learning from demonstration*. Intuitively, an IP models the actions of one or more agents as time dependent trajectories for each measured degree of freedom. The approach has gained popularity in HRI and has been applied to a number of tasks [8, 5, 21, 9, 6, 11].

Most recently, in [5] a fully Bayesian reformulation of IPs called *Bayesian Interaction Primitives* (BIP) was introduced. Most importantly, this work establishes a theoretical link between HRI and joint optimization frameworks as found in the Simultaneous Localization and Mapping (SLAM) literature [28]. The resulting inference framework for BIPs was

shown to produce superior space-time inference performance when compared to previous IP formulations.

Multimodal Modeling and Inference: Multimodal integration, inference and reasoning has been a longstanding and challenging problem of artificial intelligence [7] and signal processing [17]. Following the principles formulated by Piaget [22], many existing multimodal systems separately process incoming data streams of different types, deferring the integration step to later stages of the processing pipeline. In [4], Calinon and colleagues present a fully probabilistic framework in which social cues from the gaze direction and speech patterns of a human partner are incorporated into the robot movement generation process. The multimodal inference process is achieved by modeling such social cues as prior probability distributions. In a similar vein, the work by Dermay et al. [9] uses joint probability distributions over human-robot joint actions, in order to infer robot responses to human visual or physical guidance. However, the approach assumes a fixed user position at training and test time and models the phase variable according to a predetermined relationship to the execution speed. More recently, deep learning approaches for multimodal representations have gained considerable attention. A prominent methodology is to process each data modality with separate sub-networks, which are integrated at a shared final layer [20, 24]. However, such neural network approaches are not well-suited for probabilistic data integration and do not provide an estimate of the uncertainty inherent to observations or outputs. Also, such approaches cannot cope with missing inputs or changing query types, i.e., any change to the number or type of inputs requires a complete retraining of the network.

III. PRELIMINARIES: BAYESIAN INTERACTION PRIMITIVES

The concept of Bayesian Interaction Primitives [5] refers to a human-robot interaction framework which focuses on extracting a model of the interaction dynamics as found in example demonstrations. Given training demonstrations of the interaction task, e.g., a set of human-human interactions, BIPs can be used to capture the observed relationships between the interacting agents. Fig. 2 depicts the training and reproduction process in the BIP framework. After collecting examples for throwing and catching, the training data is represented within a basis function space and encoded as a prior distribution. The distribution is then used during reproduction to perform Bayesian filtering of live human movements, thereby enabling (a) the prediction of the next human movements, and (b) the generation of appropriate robot actions and responses. The basic structure of the above figure applies to both the original BIP formulation and our proposed method. Implementational details, in particular regarding the encoding of multiple modalities, the representation of the joint distribution as an ensemble, as well as the multimodal filtering process differ substantially in our reformulation. Subsequently, we will first provide a discussion of the BIP method as originally proposed in [5]. In particular, we will discuss the basis function decomposition and the Bayesian filtering process in BIP. After that, we introduce

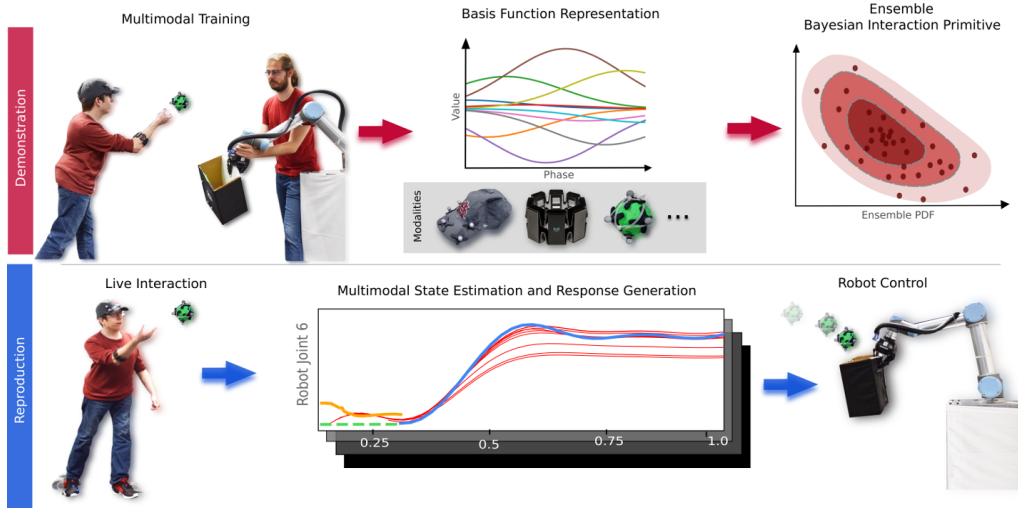


Fig. 2: An overview of eBIP. Top: training demonstrations (left) are decomposed into a latent space (middle) and transformed into an ensemble of samples (right). Bottom: observations are collected during a live interaction (left) which is used to perform filtering with the learned ensemble (middle) and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

our main contribution called ensemble Bayesian Interaction Primitives in Sec. IV.

Notation: we define an interaction \mathbf{Y} as a time series of D -dimensional sensor observations over time, $\mathbf{Y}_{1:T} = [\mathbf{y}_1, \dots, \mathbf{y}_T] \in \mathbb{R}^{D \times T}$. Of the D dimensions, D_o of them represent *observed* DoFs from one agent (the human) and D_c of them represent the *controlled* DoFs from the other agent (the robot), such that $D = D_c + D_o$.

A. Basis Function Decomposition

Working with the time series directly is impractical due to the fact that the state space dimension would be proportional to the number of observations, so we transform the interaction \mathbf{Y} into a latent space via basis function decomposition. Each dimension $d \in D$ of \mathbf{Y} is approximated with a weighted linear combination of time-dependent basis functions: $[y_1^d, \dots, y_t^d] = [\Phi_{\phi(1)}^d \mathbf{w}^d + \epsilon_y, \dots, \Phi_{\phi(t)}^d \mathbf{w}^d + \epsilon_y]$, where $\Phi_{\phi(t)}^d \in \mathbb{R}^{1 \times B^d}$ 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. As this is a linear system with a closed-form solution, the weights \mathbf{w}^d can be found through simple linear regression, i.e., least squares. The full latent model is composed of the aggregated weights from each dimension, $\mathbf{w} = [\mathbf{w}^{1\top}, \dots, \mathbf{w}^{D\top}] \in \mathbb{R}^{1 \times B}$ where $B = \sum_d B^d$ and we denote the basis transformation as $\mathbf{y}_t = h(\phi(t), \mathbf{w})$.

We note that the time-dependence of the basis functions is not on the absolute time t , but rather on a relative phase value $\phi(t)$. Consider the basis function decompositions for a motion performed at slow speeds and fast speeds with a fixed measurement rate. If the time-dependence is based on the absolute time t , then the decompositions will be different despite the motion being spatially identical. Thus, we substitute the absolute time t with a linearly interpolated relative phase value, $\phi(t)$, such that $\phi(0) = 0$ and $\phi(T) = 1$. For notational simplicity, from here on we refer to $\phi(t)$ as simply ϕ .

B. Bayesian Filtering in Time and Space

Given t observations of an interaction, $\mathbf{Y}_{1:t}$, the objective in BIP is to infer the underlying latent model \mathbf{w} while taking into account a prior model \mathbf{w}_0 . We assume that the t observations made so far are of a partial interaction, i.e., $\phi(t) < 1$, and that T is unknown. This requires the simultaneous estimation of the phase, as well as the phase velocity, i.e., how fast we are proceeding through the interaction, alongside the latent model. This joint estimation process is possible since the uncertainty estimates of each weight in the latent model are correlated due to a shared error in the phase estimate. In other words, if we mis-estimate where we are in the interaction in a temporal sense, we will mis-estimate where we are in a physical sense as well. Probabilistically, we represent this insight with the augmented state vector $\mathbf{s} = [\phi, \dot{\phi}, \mathbf{w}]$ and the following definition:

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

The posterior density in Eq. 1 is computed with a recursive linear state space filter, i.e., an extended Kalman filter [28]. Such filters are composed of two steps performed recursively: state prediction in which the state is propagated forward in time according to the system dynamics $p(\mathbf{s}_t | \mathbf{Y}_{1:t-1}, \mathbf{s}_0)$, and measurement update in which the latest sensor observation is incorporated in the predicted state $p(\mathbf{y}_t | \mathbf{s}_t)$. Applying Markov assumptions, the state prediction density can be defined as:

$$\begin{aligned} p(\mathbf{s}_t | \mathbf{Y}_{1:t-1}, \mathbf{s}_0) \\ = \int p(\mathbf{s}_t | \mathbf{s}_{t-1}) p(\mathbf{s}_{t-1} | \mathbf{Y}_{1:t-1}, \mathbf{s}_0) d\mathbf{s}_{t-1}. \end{aligned} \quad (2)$$

As with all Kalman filters, we assume that all error estimates produced during recursion are normally distributed, i.e., $p(\mathbf{s}_t | \mathbf{Y}_{1:t}, \mathbf{s}_0) = \mathcal{N}(\boldsymbol{\mu}_{t|t}, \boldsymbol{\Sigma}_{t|t})$ and $p(\mathbf{s}_t | \mathbf{Y}_{1:t-1}, \mathbf{s}_0) = \mathcal{N}(\boldsymbol{\mu}_{t|t-1}, \boldsymbol{\Sigma}_{t|t-1})$. The state evolves according to a linear constant velocity model:

$$\boldsymbol{\mu}_{t|t-1} = \underbrace{\begin{bmatrix} 1 & \Delta t & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}}_{\mathbf{G}} \boldsymbol{\mu}_{t-1|t-1}, \quad (3)$$

$$\boldsymbol{\Sigma}_{t|t-1} = \mathbf{G} \boldsymbol{\Sigma}_{t-1|t-1} \mathbf{G}^\top + \underbrace{\begin{bmatrix} \Sigma_{\phi,\phi} & \Sigma_{\phi,\dot{\phi}} & \dots & 0 \\ \Sigma_{\dot{\phi},\phi} & \Sigma_{\dot{\phi},\dot{\phi}} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix}}_{\mathbf{Q}_t}, \quad (4)$$

where \mathbf{Q} is the process noise associated with the state transition, e.g. discrete white noise. The observation function $h(\cdot)$ is nonlinear with respect to ϕ and must be linearized via Taylor expansion:

$$\mathbf{H}_t = \frac{\partial h(\mathbf{s}_t)}{\partial \mathbf{s}_t} = \begin{bmatrix} \frac{\partial \Phi_\phi^\top \mathbf{w}^1}{\partial \phi} & 0 & \Phi_\phi^1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \Phi_\phi^\top \mathbf{w}^D}{\partial \phi} & 0 & 0 & \dots & \Phi_\phi^D \end{bmatrix}. \quad (5)$$

This yields the measurement update

$$\mathbf{K}_t = \boldsymbol{\Sigma}_{t|t-1} \mathbf{H}_t^\top (\mathbf{H}_t \boldsymbol{\Sigma}_{t|t-1} \mathbf{H}_t^\top + \mathbf{R}_t)^{-1}, \quad (6)$$

$$\boldsymbol{\mu}_{t|t} = \boldsymbol{\mu}_{t|t-1} + \mathbf{K}_t (\mathbf{y}_t - h(\boldsymbol{\mu}_{t|t-1})), \quad (7)$$

$$\boldsymbol{\Sigma}_{t|t} = (\mathbf{I} - \mathbf{K}_t \mathbf{H}_t) \boldsymbol{\Sigma}_{t|t-1}, \quad (8)$$

where \mathbf{R}_t is the Gaussian measurement noise associated with the sensor observation \mathbf{y}_t .

The prior model $\mathbf{s}_0 = [\phi_0, \dot{\phi}_0, \mathbf{w}_0]$ is computed from a set of initial demonstrations. That is, given the latent models for N demonstrations, $\mathbf{W} = [\mathbf{w}_1^\top, \dots, \mathbf{w}_N^\top]$, we define \mathbf{w}_0 as simply the arithmetic mean of each DoF. The initial phase ϕ_0 is set to 0 under the assumption that all interactions start from the beginning. The initial phase velocity $\dot{\phi}_0$ is the arithmetic mean of the phase velocity of each demonstration (reciprocal length $1/T$). The prior density is defined as $p(\mathbf{s}_0) = \mathcal{N}(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ where

$$\boldsymbol{\mu}_0 = \mathbf{s}_0, \quad (9)$$

$$\boldsymbol{\Sigma}_0 = \begin{bmatrix} \Sigma_{\phi,\phi} & 0 \\ 0 & \Sigma_{\mathbf{w},\mathbf{w}} \end{bmatrix}, \quad (10)$$

and $\Sigma_{\phi,\phi}$ is the variance in the phases and phase velocities of the demonstrations, with no initial correlations between them.

IV. ENSEMBLE BAYESIAN INTERACTION PRIMITIVES

The introduction of additional sensor modalities is intended to increase the robustness of the inference process for the latent model, as defined in Eq. 1, by revealing additional information about the true state of the environment. However, naively increasing the number of observed degrees of freedom often harms the inference process. This is due to three reasons: 1) the

approximation errors in the prior distribution may increase, 2) the linearization errors may increase, and 3) the state dimension increases. This motivates the form of our proposed method as we seek to explicitly address these three issues.

Non-Gaussian Uncertainties: In general, the extended Kalman filter employed for recursive filtering in BIP relies on the assumption that uncertainty in the state prediction is approximately Gaussian. When this is not the case, the estimated state can diverge rapidly from the true state. One potential source of non-normality in the uncertainty is the nonlinear state transition or observation function in the dynamical system. The original formulation of BIP addresses this challenge by linearizing these functions about the estimated state via first-order Taylor approximation, which is performed in Eq. 5 for the nonlinear observation function $h(\cdot)$. Unfortunately, this produces linearization errors resulting from the loss of information related to higher-order moments. In strongly nonlinear systems this can result in poor state estimates and in the worst case cause divergence from the true state [19].

As we add additional degrees of freedom from modalities with their own unique numerical fingerprint (we do not make assumptions about statistical independence, however), we potentially increase the nonlinearity of the observation model. We follow an ensemble-based filtering methodology [10] which avoids the Taylor series approximation and hence the associated linearization errors. Fundamentally, we approximate the state prediction with a Monte Carlo approximation where the sample mean of the ensemble models the mean $\boldsymbol{\mu}$ and the sample covariance models the covariance $\boldsymbol{\Sigma}$. Thus, rather than calculating these values explicitly during state prediction at time t as in Eq. 8, we instead start with an ensemble of E members sampled from the prior distribution $\mathcal{N}(\boldsymbol{\mu}_{t-1|t-1}, \boldsymbol{\Sigma}_{t-1|t-1})$ such that $\mathbf{X}_{t-1|t-1} = [\mathbf{x}^1, \dots, \mathbf{x}^E]$. Each member is propagated forward in time using the state evolution model with an additional perturbation sampled from the process noise,

$$\mathbf{x}_{t|t-1}^j = \mathbf{G} \mathbf{x}_{t-1|t-1}^j + \mathcal{N}(0, \mathbf{Q}_t), \quad 1 \leq j \leq E. \quad (11)$$

As E approaches infinity, the ensemble effectively models the full covariance calculated in Eq. 4 [10]. We note that in BIP the state transition function is linear, however, when this is not the case the nonlinear function $g(\cdot)$ is used directly.

During the measurement update step, we calculate the innovation covariance \mathbf{S} and the Kalman gain \mathbf{K} directly from the ensemble, with no need to specifically maintain a covariance matrix. We begin by calculating the transformation of the ensemble to the measurement space, via the nonlinear observation function $h(\cdot)$, along with the deviation of each ensemble member from the sample mean:

$$\mathbf{H}_t \mathbf{X}_{t|t-1} = [h(\mathbf{x}_{t|t-1}^1), \dots, h(\mathbf{x}_{t|t-1}^E)], \quad (12)$$

$$\mathbf{H}_t \mathbf{A}_t = \mathbf{H}_t \mathbf{X}_{t|t-1} \quad (13)$$

$$- \left[\frac{1}{E} \sum_{j=1}^E h(\mathbf{x}_{t|t-1}^j), \dots, \frac{1}{E} \sum_{j=1}^E h(\mathbf{x}_{t|t-1}^j) \right].$$

The innovation covariance can now be found with

$$\mathbf{S}_t = \frac{1}{E-1} (\mathbf{H}_t \mathbf{A}_t) (\mathbf{H}_t \mathbf{A}_t)^\top + \mathbf{R}_t, \quad (14)$$

which is then used to compute the Kalman gain as

$$\mathbf{A}_t = \mathbf{X}_{t|t-1} - \frac{1}{E} \sum_{j=1}^E \mathbf{x}_{t|t-1}^j, \quad (15)$$

$$\mathbf{K}_t = \frac{1}{E-1} \mathbf{A}_t (\mathbf{H}_t \mathbf{A}_t)^\top \mathbf{S}_t^{-1}. \quad (16)$$

With this information, the ensemble can be updated to incorporate the new measurement perturbed by stochastic noise:

$$\begin{aligned} \tilde{\mathbf{y}}_t &= [\mathbf{y}_t + \epsilon_y^1, \dots, \mathbf{y}_t + \epsilon_y^E], \\ \mathbf{X}_{t|t} &= \mathbf{X}_{t|t-1} + \mathbf{K}(\tilde{\mathbf{y}}_t - \mathbf{H}_t \mathbf{X}_{t|t-1}). \end{aligned} \quad (17)$$

It has been shown that when $\epsilon_y \sim \mathcal{N}(0, \mathbf{R}_t)$, the measurements are treated as random variables and the ensemble accurately reflects the error covariance of the best state estimate [3]. The measurement noise \mathbf{R}_t can be calculated with the following closed-form solution:

$$\mathbf{R}_t = \frac{1}{N} \sum_i^N \frac{1}{T_i} \sum_t^{T_i} (\mathbf{y}_t - h([\phi(t), \mathbf{w}_i]))^2. \quad (18)$$

This value is equivalent to the mean squared error of the regression fit for our basis functions over every demonstration. Intuitively, this represents the variance of the data around the regression and captures both the approximation error and the sensor noise associated with the observations.

One of the advantages of this algorithm is the elimination of linearization errors through the use of the nonlinear functions. While this introduces non-normality into the state uncertainties, it has been shown that the stochastic noise added to the measurements pushes the updated ensemble towards normality, thereby reducing the effects of higher-order moments [14, 16] and improving robustness in nonlinear scenarios.

Non-Gaussian Prior: Another source of non-Gaussian uncertainty is from the initial estimate (the prior) itself. In BIP, our prior is given by a set of demonstrations which indicate where we believe a successful interaction would lie in the state space. As we have yet to assimilate any observations of a new interaction, the (unknown) true distribution from which the demonstrations are sampled represents our best initial estimate of what it may be. However, given that these are real-world demonstrations, they are highly unlikely to be normally distributed. As such, two options are available in this case: we can either use the demonstrations directly as samples from the non-Gaussian prior distribution or approximate the true distribution with a Gaussian and sample from it. The latter approach is used by BIP in Eq. 9 and Eq. 10, however, this comes with its own risks since a poor initial estimate can lead to poor state estimates [12]. Given that the ensemble-based filtering proposed here provides a degree of robustness to non-Gaussian uncertainties, we choose to use samples from the non-Gaussian prior directly in eBIP, with the knowledge that the ensemble will be pushed towards normality. If the

Ensemble Bayesian Interaction Primitives

Input: $\mathbf{W} = [\mathbf{w}_1^\top, \dots, \mathbf{w}_N^\top] \in \mathbb{R}^{B \times N}$: set of B basis weights corresponding to N demonstrations, $\mathbf{l} = [\frac{1}{T_1}, \dots, \frac{1}{T_N}] \in \mathbb{R}^{1 \times N}$: reciprocal lengths of demonstrations, $\mathbf{y}_t \in \mathbb{R}^{D \times 1}$: sensor observation at time t .

Output: $\hat{\mathbf{y}}_t \in \mathbb{R}^{D \times 1}$: the inferred trajectory at time t .

- 1) Create the initial ensemble \mathbf{X}_0 such that

$$\mathbf{x}_0^j = [0, \phi^j, \mathbf{w}^j], \quad 1 \leq j \leq E,$$

eBIP⁻: $\mathbf{w}^j \sim \sum_k^K \alpha_k \mathcal{N}(\mu_k, \Sigma_k)$ where μ_k, Σ_k , and α_k are found via EM over \mathbf{W} , $\phi^j \sim \mathcal{N}(\mu_t, \sigma_t^2)$

eBIP: $i \sim \mathcal{U}\{1, N\}$, $\mathbf{w}^j = \mathbf{w}_i$, $\phi^j = \frac{1}{T_i}$.

- 2) For time step t , propagate the ensemble forward in time as in Eq. 11:

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

- 3) If a measurement \mathbf{y}_t is available, perform the measurement update step from Eq. 17:

$$\mathbf{X}_{t|t} = \mathbf{X}_{t|t-1} + \mathbf{K}(\tilde{\mathbf{y}}_t - \mathbf{H}_t \mathbf{X}_{t|t-1})$$

- 4) Extract the estimated state and uncertainty from the ensemble:

$$\boldsymbol{\mu}_{t|t} = \frac{1}{E} \sum_{j=1}^E \mathbf{x}_{t|t-1}^j, \quad \boldsymbol{\Sigma}_{t|t} = \frac{1}{E-1} \mathbf{A}_t \mathbf{A}_t^\top$$

- 5) **Output** the trajectory for each controlled DoF:

$$\hat{\mathbf{y}}_t = h(\boldsymbol{\mu}_{t|t})$$

- 6) Repeat steps 2-5 until the interaction is concluded.

Fig. 3: Ensemble Bayesian Interaction Primitives

number of ensemble members is greater than the number of available demonstrations, then the density of the true interaction distribution will need to be estimated given the observed demonstrations. This can be accomplished using any density estimation technique, e.g., a Gaussian mixture model, and we denote this as the alternative formulation eBIP⁻.

Computational Performance: The increased state dimension resulting from the introduction of additional sensor modalities leads to undesirable increases in computation times in the BIP algorithm. This is due to the necessary covariance matrix updates defined in Eq. 8, which causes BIP to yield an asymptotic computational complexity of approximately $O(n^3)$ with respect to the state dimension n [28]; we ignore terms related to the measurement dimension as it is significantly smaller than the state dimension. However, as eBIP is ensemble-based, we no longer explicitly maintain a covariance matrix; this information is implicitly captured by the ensemble. As a result, the computational complexity for eBIP is approximately $O(E^2 n)$, where E is the ensemble size and n is the state




Fig. 4: A sequence of images from three live interactions. The robot is already reacting to the human by the second image in each sequence and catches the ball in different poses due to the different ball trajectories.

dimension [18]. Since the ensemble size is typically much smaller than the state dimension, this results in a performance increase when compared to BIP. Furthermore, the formulation presented in this work also obviates the need to explicitly construct the observation matrix H . The creation of the observation matrix introduces an additional overhead for BIP as it must be initialized at each time step due to the phase-dependence, a process which is unnecessary in eBIP.

In addition, we also benefit from the computational performance-accuracy trade off inherent to all sample-based methods. Inference accuracy can be sacrificed for computational performance by lowering the number of ensemble members when called for. While this is also true for particle filters, they generally scale poorly to higher state dimensions due to sample degeneracy. In particle filtering, ensemble members are re-sampled according to their weight in a scheme known as importance sampling. However, in large state spaces it is likely that only a small number of ensemble members will have high weights, thus eventually causing all members to be re-sampled from only a few. In our proposed method this is not the case, as all members are treated equally, thus lending itself well to high-dimensional state spaces.

Algorithm: Putting together all of the components, our full proposed algorithm is shown in Fig. 3.

V. EXPERIMENTS AND RESULTS

We show the effectiveness of our proposed algorithm in the multimodal HRI scenario described in Sec. I (Fig. 1). In this scenario, a human participant outfitted with a variety of sensors tosses a ball which is caught by a UR5 [25] arm. The sensors can be broadly grouped into two categories: modalities that observe the human and modalities that observe the ball. Observations of the ball are unavailable while it is grasped by the human, due to occlusions, and do not become available until the ball is thrown. In empirical tests, we have observed that it is not possible for the robot to catch the ball using a purely reactive strategy given the limited time to react, kinematic constraints, phase lag, etc. Hence, we leverage the observation modalities of the human to predict how the robot should react while the human is still in the preparatory phase, i.e., the "wind up" for the throw. This strategy is fundamentally built upon a

predictive approach – we can begin reacting as early as possible and refine our predictions as more detailed observations become available.

A. Experimental Setup

The experiment is designed to emphasize the advantages of a multimodal observation set by having different sensors reveal different information about the true environment state at different points in time. However, throwing and catching are fast-paced actions requiring a high frequency observation rate and appropriately low computation times for inference; without these properties an HRI algorithm will likely fail at catching the ball in real experiments. We utilize sensor observations of 8 objects from 5 modalities: the positions of the human participant's hands and feet, inertial measurements of the throwing arm, pressure measurements from the soles of both left and right feet, the orientation of the head, the position of the ball being thrown, and the joint positions of the robot. The observations were synchronized and collected at a frequency of 60Hz. The basis decomposition for each sensor object was chosen from a set of candidate basis spaces comprised of Polynomial, Gaussian, and Sigmoid functions – standard choices in this type of application [2] – using the Bayesian Information Criterion, yielding a total state space dimension of 559 dimensions.

During training, the ball was thrown from a distance of approx. 3.7m and was caught within a box grasped by the robot (the end effector used in this experiment actuates too slowly for in-hand interception). An initial set of 221 demonstrations was provided via kinesthetic teaching in which the robot was manually operated by a human (top left of Fig. 2) in order to catch the ball while joint positions were recorded. These demonstrations provided the only source of prior knowledge for the interaction (for both state estimation and control); no inverse kinematics or other models were employed at any point in time. We compare our algorithm, to the original BIP formulation, as well as particle filtering (PF). In all cases, the PF model used the same number of ensemble members as in eBIP and employed a systematic resampling scheme when the effective number of members was less than $E/2$.

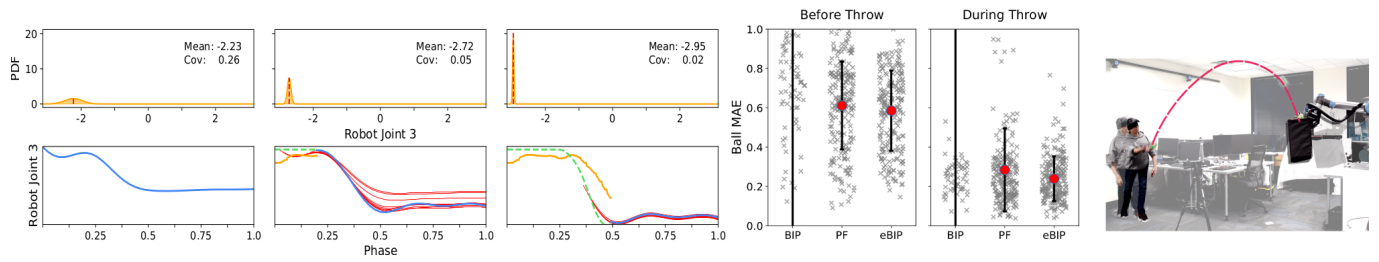


Fig. 5: **(Left)** A sequence of frames from different time points during an interaction. Top: the PDF of the third robot joint; the initial uncertainty is high and decreases over time. Bottom: the inferred trajectory for the robot joint. The blue line indicates the current prediction while the red lines indicate the predictions for the past 10 time steps. The yellow line is the actual response from the robot while it attempts to follow the inferred trajectory, and the dashed green line indicates the expected trajectory from the demonstration. **(Center)** The ball MAE of the {All} subset. While overall error decreases for both PF and eBIP over time, only eBIP experiences a reduced variance. **(Right)** A blindfolded user throws a ball which is, in turn, caught by the robot.




Fig. 6: Top: computation time required for filtering observations of varying lengths. Bottom: the computation time-accuracy-ensemble size trade-off for eBIP.

the inference error of the final ball position by a factor of 2. However, we also gain the ability to initially predict the ball position much earlier in the interaction, before the ball is visible. This process is visualized in Fig. 5 through the uncertainty in the inferred joint trajectories. The width of the catchable region is approximately 180cm. Hence, the fact that we can predict the interception point to within about 60cm, or 1/3 of our operating region, *before the ball is visible* is quite significant. By the time the ball is in the air (82% of the trajectory is observed), we have further narrowed the prediction down to 23cm with additional modalities. While this error is still significantly higher than the 13cm error produced by incorporating only the ball, we observed empirically that a 23cm error still results in a catch in most cases and justifies the inclusion of additional modalities. Given that the radius of the ball itself is 4.5cm and the width of the box is 32cm, this amount of error is adequate and is offset by the proactive behavior of the robot in this setting. Still, the above results suggests that there may be substantial benefit to the ability of the inference process to switch modalities on and off in real-time according to the context.

Real-time interactions: Results of the real-time interaction and reproduction with the robot can be seen in Fig. 4. To minimize the computation time required for inference, the number of ensemble members was limited to 80. This setup ensured a reasonable trade-off between accuracy and computational performance as shown in Fig. 6. Three observations can be made from the image sequences in Fig. 4: the throws all have

significantly different trajectories (the top throw has a low apex and fast velocity while the bottom throw has a high apex and low velocity), the robot is already moving into position before the ball is thrown (second image in each sequence), and the robot catches the ball in a different pose for each throw (last image in each sequence). Recordings of a variety of throwing experiments can be found in the accompanying video. To avoid habituation or any unconscious effort to throw the ball directly at the robot, we also performed a set of experiments in which the user threw the ball while in a blindfolded condition, see Fig. 5 (right) for an example. Even under this condition, the ball was successfully caught 12 out of 20 times, for a success rate of 60%. We noticed, however, that in this condition the user frequently threw the ball outside the robot’s reach.

VI. CONCLUSIONS

In this paper, we introduced ensemble Bayesian Interaction Primitives and discussed their application to state estimation, inference, and control in challenging, fast-paced HRI tasks with many data sources. We discussed an ensemble-based approach to Bayesian inference in eBIP, which requires neither an explicit formation of a covariance matrix, nor a measurement model, resulting in significant computational speed-ups. The approach allows for fast inference from high-dimensional, probabilistic models and avoids typical sources for inaccuracies, e.g., linearization and Gaussian priors. In our real-robot experiments, a relatively small number of ensemble members produced a reasonable trade-off between accuracy and computational performance. However, our results also indicate that the uncontrolled inclusion of many data sources is not always beneficial. Some modalities may introduce spurious correlations or significant amounts of noise into the filtering process, thereby harming the accuracy of predictions. These challenges may be overcome by incorporating feature selection mechanisms, or by switching individual modalities on and off according to context.

ACKNOWLEDGMENTS

This work was supported by a grant from the Honda Research Institute and by the National Science Foundation under Grant No. (IIS-1749783).

REFERENCES

- [1] Heni Ben Amor, Gerhard Neumann, Sanket Kamthe, Oliver Kroemer, and Jan Peters. Interaction primitives for human-robot cooperation tasks. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 2831–2837. IEEE, 2014.
- [2] Christopher M Bishop. Pattern recognition and machine learning. 2006.
- [3] Gerrit Burgers, Peter Jan van Leeuwen, and Geir Evensen. Analysis scheme in the ensemble kalman filter. *Monthly weather review*, 126(6):1719–1724, 1998.
- [4] Sylvain Calinon and Aude Billard. A framework integrating statistical and social cues to teach a humanoid robot new skills. In *Proc. IEEE Intl Conf. on Robotics and Automation (ICRA), Workshop on Social Interaction with Intelligent Indoor Robots*, May 2008.
- [5] Joseph Campbell and Heni Ben Amor. Bayesian interaction primitives: A slam approach to human-robot interaction. In *Conference on Robot Learning*, pages 379–387, 2017.
- [6] L. Chen, H. Wu, S. Duan, Y. Guan, and J. Rojas. Learning human-robot collaboration insights through the integration of muscle activity in interaction motion models. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, pages 491–496, Nov 2017.
- [7] Michael H. Coen. Multimodal integration: A biological view. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 2, IJCAI'01*, pages 1417–1424, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc.
- [8] Yunduan Cui, James Poon, Jaime Valls Miro, Kimitoshi Yamazaki, Kenji Sugimoto, and Takamitsu Matsubara. Environment-adaptive interaction primitives through visual context for human-robot motor skill learning. *Autonomous Robots*, Aug 2018.
- [9] Oriane Dermy, François Charpillet, and Serena Ivaldi. Multi-modal intention prediction with probabilistic movement primitives. In *Human Friendly Robotics*, pages 181–196. Springer, 2019.
- [10] Geir Evensen. The ensemble kalman filter: Theoretical formulation and practical implementation. *Ocean dynamics*, 53(4):343–367, 2003.
- [11] Marco Ewerton, Gerhard Neumann, Rudolf Lioutikov, Heni Ben Amor, Jan Peters, and Guilherme Maeda. Learning multiple collaborative tasks with a mixture of interaction primitives. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1535–1542, May 2015. doi: 10.1109/ICRA.2015.7139393.
- [12] Eric L Haseltine and James B Rawlings. Critical evaluation of extended kalman filtering and moving-horizon estimation. *Industrial & engineering chemistry research*, 44(8):2451–2460, 2005.
- [13] Tariq Iqbal and Laurel D Riek. Human-robot teaming: Approaches from joint action and dynamical systems. *Humanoid Robotics: A Reference*, pages 2293–2312, 2019.
- [14] W Gregory Lawson and James A Hansen. Implications of stochastic and deterministic filters as ensemble-based data assimilation methods in varying regimes of error growth. *Monthly weather review*, 132(8):1966–1981, 2004.
- [15] Dongheui Lee and Yoshihiko Nakamura. Mimesis model from partial observations for a humanoid robot. *The International Journal of Robotics Research*, 29(1):60–80, 2010.
- [16] Jing Lei, Peter Bickel, and Chris Snyder. Comparison of ensemble kalman filters under non-gaussianity. *Monthly Weather Review*, 138(4):1293–1306, 2010.
- [17] Christian Lundquist, Zoran Sjanic, and Fredrik Gustafsson. *Statistical Sensor Fusion: Exercises*. Studentlitteratur AB, Sweden, 2015.
- [18] Jan Mandel. *Efficient implementation of the ensemble Kalman filter*. University of Colorado at Denver and Health Sciences Center, Center for Computational Mathematics, 2006.
- [19] Robert N Miller, Michael Ghil, and Francois Gauthiez. Advanced data assimilation in strongly nonlinear dynamical systems. *Journal of the atmospheric sciences*, 51(8):1037–1056, 1994.
- [20] Kuniaki Noda, Hiroaki Arie, Yuki Suga, and Tetsuya Ogata. Multimodal integration learning of robot behavior using deep neural networks. *Robotics and Autonomous Systems*, 62(6):721–736, 2014.
- [21] Ozgur S. Oguz, Zhehua Zhou, and Dirk Wollherr. A hybrid framework for understanding and predicting human reaching motions. *Frontiers in Robotics and AI*, 5:27, 2018.
- [22] Jean Piaget. *The Child's Construction of Reality*. Routledge & Paul, 1955.
- [23] L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, Feb 1989.
- [24] Dushyant Rao, Mark De Deuge, Navid Nourani-Vatani, Stefan B Williams, and Oscar Pizarro. Multimodal learning and inference from visual and remotely sensed data. *The International Journal of Robotics Research*, 36(1):24–43, 2017.
- [25] Universal Robots. UR5 Robot Arm. <https://www.universal-robots.com/products/ur5-robot/>. [Online; accessed May 15, 2019].
- [26] Leonel Rozo, Joao Silvério, Sylvain Calinon, and Darwin G. Caldwell. Learning controllers for reactive and proactive behaviors in human-robot collaboration. *Frontiers in Robotics and AI*, 3(30):1–11, June 2016. Specialty Section Robotic Control Systems.
- [27] Andrea Thomaz, Guy Hoffman, Maya Cakmak, et al. Computational human-robot interaction. *Foundations and Trends® in Robotics*, 4(2-3):105–223, 2016.
- [28] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic robotics*. MIT press, 2005.