Bayesian Interaction Primitives: A SLAM Approach to Human-Robot Interaction

Joseph Campbell
CIDSE
Arizona State University, United States
jacampb1@asu.edu

Heni Ben Amor
CIDSE
Arizona State University, United States
hbenamor@asu.edu

Abstract: This paper introduces a fully Bayesian reformulation of Interaction Primitives for human-robot interaction and collaboration. A key insight is that a subset of human-robot interaction is conceptually related to simultaneous localization and mapping techniques. Leveraging this insight we can significantly increase the accuracy of temporal estimation and inferred trajectories while simultaneously reducing the associated computational complexity. We show that this enables more complex human-robot interaction scenarios involving more degrees of freedom.

Keywords: human-robot interaction, Bayesian inference, multimodal perception

1 Introduction

Programming collaborative robots to assist a human partner is a notoriously difficult challenge. The traditional programming paradigm requires a human engineer to foresee all important task parameters, and implement control routines that allow the robot to determine *when* and *how* to engage in the interaction. Unfortunately, even for moderately complex human-robot interaction scenarios this approach becomes intractable, since it involves many possible combinations of states and actions. To overcome this challenge, data-driven approaches for imitation learning in dyadic settings have recently been proposed. Recorded human-human or human-robot demonstrations are leveraged to train machine learning models of the observed dynamics of the interaction. In turn, these models can be used by a robot, at runtime, to engage in similar interactions with a human partner.

A prominent approach to modeling such coupled behavior from example demonstrations has been proposed in [1]. More specifically, the Interaction Primitives method learns probability distributions over joint actions between two interaction partners. Bayesian inference can thereafter be used to generate robot actions, as a function of observed human behavior. A prerequisite for that operation, however, is the identification of the current phase (*i.e.*, time step) within the collaborative task. In Interaction Primitives and related methods, this is solved using Dynamic Time Warping [2] (DTW) or a Monte Carlo sampling strategy, *e.g.*, the Particle Filter [3]. Both strategies are computationally demanding and do not generalize well to new motions with significantly different execution speeds. Most importantly, in Interaction Primitives phase estimation and response generation are separate processes that do not share the underlying computational principles and representations. Hence, inference does not adequately use available data.

In this paper, we propose a fully Bayesian generalization of Interaction Primitives. This new formulation enables us to learn a joint model of both space and time variables and encode them within the same probabilistic representation. Consequently, we can use Bayesian inference to reason about both when and how a robot should engage in a collaborative task. These results leverage a critical, new insight made in this paper: a subset of human-robot interaction shares a theoretical and conceptual relationship to the field of Simultaneous Localization and Mapping (SLAM) [4] – a mature research area with a sound theoretical foundation. We show that this conceptual link between the two fields can immediately be used to leverage the wealth of knowledge regarding SLAM and derive a powerful new formulation of Interaction Primitives. The new, fully Bayesian Interaction Primitives (BIP) (a) support multimodal data sources, (b) infer both phase and action variables, (c) significantly reduce time complexity, and (d) improve inference and prediction accuracy. In our experiments we

1st Conference on Robot Learning (CoRL 2017), Mountain View, United States.

demonstrate that this allows us to effectively perform human-robot interaction in complex scenarios involving multimodal, high-dimensional state spaces.

2 Related Work

Learning from demonstration is an important topic in the field of robotics, in large part due to the complexity of both robots and their operating environments. One such technique, Dynamic Movement Primitives [5], achieves this by fitting observed trajectories to dynamic equations and mapping correlations between the equation parameters. Interaction Primitives [1] are a logical extension of this concept to human-robot interaction, introducing multiple agents and temporal alignment of trajectories. Interaction Primitives were later extended [6] to utilize probabilistic methods [7] as opposed to dynamic equations, improving their generalization ability.

However, Interaction Primitives still require the use of time alignment algorithms such as Dynamic Time Warping to function correctly. This poses several problems: DTW is difficult to use with high-dimensional spaces in real-time [8], and Monte Carlo-based sampling strategies such as Particle Filters scale exponentially with dimension [9] making them impractical.

Ultimately, we would like a framework capable of probabilistic joint models over both phase estimation and parameter estimation that isn't reliant on independent time alignment. Models such as Layered Hidden Markov Models [10] provide some insight, as the output from HMMs are used as inputs to other HMMs and could potentially be used to bridge phase estimation and inference. Similarly, Hierarchical HMMs [11] can fulfill a related role and have been successfully applied to the task of activity recognition [12].

It is the problem of Simultaneous Localization and Mapping, however, that provides the basis for our work. Algorithms such as FastSLAM [13] and EKF SLAM [3] effectively estimate a joint model of a robot's pose with respect to a map while simultaneously generating that same map.

3 Methodology

In this section, we introduce Bayesian Interaction Primitives in which we model human-robot collaborative tasks as a SLAM problem. In so doing, we couple together the phase estimation and the probabilistic conditioning aspects from the Interaction Primitive framework, improving phase estimation accuracy and significantly reducing its computation overhead. While this concept is applicable to the SLAM problem in general, we will demonstrate its effectiveness through the use of a specific SLAM framework: Extended Kalman Filter (EKF) SLAM. We start by introducing Interaction Primitives and our chosen SLAM framework individually, followed by our new BIP framework which combines elements of both methods.

3.1 Interaction Primitives

Interaction Primitives are a *learning from demonstration* framework for human-robot interaction. Initially conceived as an application of Dynamic Movement Primitives [5] and later extended to Probabilistic Movement Primitives [6], the underlying idea of Interaction Primitives is to take several demonstrations of two agents – consisting of time-dependent sensor trajectories – and compactly model the relationship between the two agents' movements. This trained model is then used to generate motion trajectories of one of the agents when given a partial observation of the other agent.

Formally, we define the combined trajectory of both agents $\boldsymbol{y}$ as consisting of a sequence of T sensor observations over time, $\boldsymbol{y}_{1:T} = \boldsymbol{y}_1, \boldsymbol{y}_2, \dots, \boldsymbol{y}_T$, in which each state $\boldsymbol{y}_t \in \boldsymbol{y}_{1:T}$ consists of sensor measurements for D_c degrees of freedom of the controlled agent (the robot) and D_o degrees of freedom of the observed agent (the human): $\boldsymbol{y}_t = [x^o_1, x^o_2, \dots, x^o_{D_o}, x^c_1, x^c_2, \dots, x^c_{D_c}]$ where $x \in \mathbb{R}$. Each degree of freedom is modeled by a linear combination of time-dependent basis functions with i.i.d. Gaussian white noise: $x_d \approx \Phi_t^T \boldsymbol{w}_d + \epsilon_y$, where $\Phi_t \in \mathbb{R}^{B \times 1}$ is a vector of B Gaussian basis functions and $\boldsymbol{w}_d \in \mathbb{R}^{B \times 1}$. Ultimately, we want to observe a partial trajectory $\boldsymbol{y}^*$ of one agent (the observed agent) and use this to condition $\boldsymbol{w}$ from which a response trajectory for the other agent will be generated:

$$p(\boldsymbol{w}|\boldsymbol{y}^*) \propto p(\boldsymbol{y}^*|\boldsymbol{w})p(\boldsymbol{w}) \tag{1}$$

where $\boldsymbol{w} = [\boldsymbol{w}_0{}^T, \ldots, \boldsymbol{w}_{D_c}{}^T]^T \in \mathbb{R}^{(D_o+D_c)B\times 1}$. Probabilistically, we can represent the posterior of observing a state $\boldsymbol{y}_t$ at time t as,

$$p(\boldsymbol{y}_t|\boldsymbol{w}) = \mathcal{N}\left(\boldsymbol{y}_t \middle| \begin{bmatrix} \Phi_t & \ldots & 0 \\ \vdots & \ddots & \vdots \\ \ldots & \ldots & \Phi_t \end{bmatrix}^T \boldsymbol{w}, \boldsymbol{\Sigma}_y\right). \tag{2}$$

We assume that the weights for a demonstration are drawn from a Gaussian distribution, such that $p(\boldsymbol{w}; \boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w) = \mathcal{N}(\boldsymbol{w}|\boldsymbol{\mu}_w, \boldsymbol{\Sigma}_w)$, $\boldsymbol{\mu}_w \in \mathbb{R}^{(D_o+D_c)B\times 1}$, $\boldsymbol{\Sigma}_w \in \mathbb{R}^{(D_o+D_c)B\times(D_o+D_c)B}$. Multiple demonstrations are recorded so that the parameters $\boldsymbol{\mu}_w$ and $\boldsymbol{\Sigma}_w$ can be learned through maximum likelihood estimation or related parameter estimation methods. However, time is a significant factor when recording demonstrations. Movements which differ in speed will yield different trajectories in absolute time. To account for this, we replace absolute time t with relative time, referred to as phase $\delta \in [0, 1]$, such that the first measurement in a trajectory is 0 and the last is 1.

Furthermore, because the Bayesian inference in Eq. 1 requires us to approximate the partially observed trajectory $\boldsymbol{y}^*$ with time-dependent basis functions, we must first estimate the length of $\boldsymbol{y}^*$ with respect to the stored demonstrations, which is referred to as phase estimation. This is accomplished by using a time-alignment algorithm such as Dynamic Time Warping in order to generate a phase alignment from the partially observed trajectory $\boldsymbol{y}^*$ and the approximated mean trajectory of the demonstrations. Inference can now be performed via recursive Bayesian filtering by setting the controlled agent's values in the measurement model to 0 with a suitably high measurement noise:

$$\boldsymbol{H}_t = \begin{bmatrix} \Phi_\delta^o & \ldots & 0 & \ldots & 0 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \ldots & \Phi_\delta^o & \ldots & 0 \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ 0 & \ldots & 0 & \ldots & 0 \end{bmatrix}^T \in \mathbb{R}^{(D_o+D_c)\times(D_o+D_c)B}, \tag{3}$$

$$\boldsymbol{K} = \boldsymbol{\Sigma}_w \boldsymbol{H}_t^T (\boldsymbol{H}_t \boldsymbol{\Sigma}_w \boldsymbol{H}_t^T + \boldsymbol{\Sigma}_z)^{-1}, \tag{4}$$

$$\boldsymbol{\mu}_w^+ = \boldsymbol{\mu}_w + \boldsymbol{K}_t(\boldsymbol{z}_t - \boldsymbol{H}_t \boldsymbol{\mu}_w), \tag{5}$$

$$\boldsymbol{\Sigma}_w^+ = (I - \boldsymbol{K}_t \boldsymbol{H}_t)\boldsymbol{\Sigma}_w. \tag{6}$$

3.2 EKF SLAM

The Simultaneous Localization and Mapping problem is concerned with localizing a robot's state in time $\boldsymbol{s}_t$ with respect to a map m while simultaneously determining the map of the environment. Specifically, we focus on Extended Kalman Filter (EKF) SLAM, in which an Extended Kalman Filter approximates the state $\boldsymbol{s}_t$ as well as the position of landmarks within the map with a Gaussian distribution. Ultimately, we are concerned with finding the following posterior probability of a robot's state at time t given a time-dependent sequence of sensor measurements z of the map landmarks and the control actions which generate the robot's motion u: $p(\boldsymbol{s}_t|\boldsymbol{z}_{1:t}, \boldsymbol{u}_{1:t}) = \mathcal{N}(\boldsymbol{s}_t|\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$. For simplicity, we assume that the correspondence between the sensor measurements and the map landmarks is known.

We define the EKF state as consisting of D degrees of freedom for the robot as well as N map landmarks: $\boldsymbol{s} = [x_1, \ldots, x_D, m_1, \ldots, m_N]^T, x \in \mathbb{R}, m \in \mathbb{R}$. In problems requiring an EKF, either the state transition model $\boldsymbol{s}_t = g(\boldsymbol{\mu}_t, \boldsymbol{s}_{t-1}) + \epsilon_t$ or the measurement model $\boldsymbol{z}_t = h(\boldsymbol{s}_t) + \delta_t$ use nonlinear functions. A Kalman filter is a recursive state estimator and projecting a Gaussian belief through a nonlinear function (g or h) may not yield a Gaussian result. Thus, EKFs approximate the nonlinear functions with linear functions (G and H) evaluated at the current state. The resulting EKF SLAM equations are shown below:

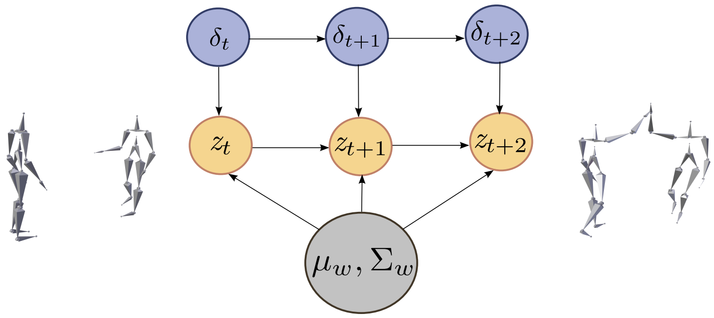

Figure 1: Graphical model of the EKF localization approach to Interaction Primitives.

$$
\begin{aligned}
\boldsymbol{\mu}_t &= g(u_t, \boldsymbol{\mu}_{t-1}), &(7)\\
\boldsymbol{\Sigma}_t &= \boldsymbol{G}_t \boldsymbol{\Sigma}_{t-1} \boldsymbol{G}_t^T + \boldsymbol{R}_t, &(8)\\
\boldsymbol{K}_t &= \boldsymbol{\Sigma}_t \boldsymbol{H}_t^T (\boldsymbol{H}_t \boldsymbol{\Sigma}_t \boldsymbol{H}_t^T + \boldsymbol{Q}_t)^{-1}, &(9)\\
\boldsymbol{\mu}_t^+ &= \boldsymbol{\mu}_t + \boldsymbol{K}_t(\boldsymbol{z}_t - h(\boldsymbol{\mu}_t)), &(10)\\
\boldsymbol{\Sigma}_t^+ &= (I - \boldsymbol{K}_t \boldsymbol{H}_t)\boldsymbol{\Sigma}_t. &(11)
\end{aligned}
$$

3.3 Bayesian Interaction Primitives

Having introduced both Interaction Primitives and EKF SLAM, we now describe how these two methods can be coupled together such that we can perform phase estimation and Bayesian inference in the same step. As presented in Sec. 3.1, phase estimation is required when we have a partial trajectory $\boldsymbol{y}^*$ of observations of the observed agent and wish to perform Bayesian inference on the generated trajectory. In the original IP formulation, phase estimation was performed via a time-alignment algorithm, however, in BIP we choose to view this as a localization problem. In EKF localization, we have a map m consisting of N landmarks and we wish to determine the robot's pose s_t with respect to the landmarks. The phase can be thought of as our robot state, $s_t = \delta_t$: it indicates where we currently are relative to a set of known demonstration trajectories. Thus, we can view the distribution over the demonstration weights, $\boldsymbol{\mu}_w$ and $\boldsymbol{\Sigma}_w$, as our map as shown in Fig. 1. However, there is a weakness with this approach: consider what happens if one or more of the trajectory degrees of freedom is static for a period of time. In this case, if the predicted measurement and the actual measurement do not differ by much – which is likely since the trajectory is static, meaning the phase won't impact the predicted value – then it is not possible to estimate the change in phase. Therefore, we also model the phase velocity $\dot{\delta}$ as a state variable. If we assume a constant velocity model, then static trajectories will be handled properly. The extension to SLAM is trivial; simply add the basis weights to the robot state in place of the landmarks, $\boldsymbol{s}_t = [\delta_t, \dot{\delta}_t, \boldsymbol{w}]^T \in \mathbb{R}^{(D_o+D_c)B+2\times 1}$. There are as many landmarks as there are basis functions for all degrees of freedom.

Accordingly, we define the relevant EKF SLAM equations:

$$
\begin{aligned}
\boldsymbol{s}_t &= [\delta_t, \dot{\delta}_t, \boldsymbol{w}^T]^T, &(12)\\
p(\boldsymbol{s}_t|\boldsymbol{z}_{1:t}) &= \mathcal{N}(\boldsymbol{s}_t|\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t), &(13)\\
\boldsymbol{\mu}_0 &= [0, \beta, \boldsymbol{\mu}_w^T]^T, &(14)\\
\boldsymbol{\Sigma}_0 &= \begin{bmatrix} \boldsymbol{\Sigma}_{\delta,\delta} & \boldsymbol{\Sigma}_{\delta,\Sigma_w} \\ \boldsymbol{\Sigma}_{\Sigma_w,\delta} & \boldsymbol{\Sigma}_{w,w} \end{bmatrix}, &(15)
\end{aligned}
$$

where $\boldsymbol{\mu}_t \in \mathbb{R}^{(D_o+D_c)B+2\times 1}$, $\boldsymbol{\Sigma}_t \in \mathbb{R}^{(D_o+D_c)B+2\times(D_o+D_c)B+2}$, β is the average phase velocity in the stored demonstrations, $\boldsymbol{z}_{1:t} = \boldsymbol{y}_{1:t}*$, and $\boldsymbol{\Sigma}_{w,w} = \boldsymbol{\Sigma}_w$. The motion model is a simple constant velocity model,

$$
\boldsymbol{\mu}_t = \boldsymbol{\mu}_{t-1} + \underbrace{\begin{bmatrix} 1 & \Delta t & \dots & 0 \\ 0 & 1 & \dots & 0 \end{bmatrix}^T}_{\boldsymbol{F}} [1,1]^T + \mathcal{N}(0, \boldsymbol{F}^T \underbrace{\begin{bmatrix} \sigma_{\delta,\delta} & \sigma_{\delta,\dot{\delta}} \\ \sigma_{\dot{\delta},\delta} & \sigma_{\dot{\delta},\dot{\delta}} \end{bmatrix}}_{\boldsymbol{Q}_t} \boldsymbol{F}), \quad \boldsymbol{G}_t = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \qquad (16)
$$

while the measurement model is the same as in standard Interaction Primitives,

$$\boldsymbol{z}_t = \underbrace{\begin{bmatrix} \Phi_{\delta}^T \boldsymbol{w_1} \\ \Phi_{\delta}^T \boldsymbol{w_2} \\ \vdots \\ \Phi_{\delta}^T \boldsymbol{w_{D_c}} \end{bmatrix}}_{h(\boldsymbol{\mu}_t)} + \mathcal{N}(0, \underbrace{\begin{bmatrix} \sigma_1 & 0 & \dots & 0 \\ 0 & \sigma_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \sigma_{D_c} \end{bmatrix}}_{\boldsymbol{R}_t}). \tag{17}$$

However, when calculating the Jacobian we must now account for both the phase and phase velocity:

$$\begin{aligned} \boldsymbol{H}_t = \frac{\partial h(\boldsymbol{\mu}_t)}{\partial x_t} &= \begin{bmatrix} \frac{\partial \Phi_{\delta}^T \boldsymbol{w}_1}{\partial \delta} & \frac{\partial \Phi_{\delta}^T \boldsymbol{w}_1}{\partial \dot{\delta}} & \frac{\partial \Phi_{\delta}^T \boldsymbol{w}_1}{\partial w_1} & \dots & \frac{\partial \Phi_{\delta}^T \boldsymbol{w}_1}{\partial w_{D_c}} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \Phi_{\delta}^T \boldsymbol{w}_{D_c}}{\partial \delta} & \frac{\partial \Phi_{\delta}^T \boldsymbol{w}_{D_c}}{\partial \dot{\delta}} & \frac{\partial \Phi_{\delta}^T \boldsymbol{w}_{D_c}}{\partial w_1} & \dots & \frac{\partial \Phi_{\delta}^T \boldsymbol{w}_{D_c}}{\partial w_{D_c}} \end{bmatrix}, \\ &= \begin{bmatrix} \frac{\partial \Phi_{\delta}^T \boldsymbol{w}_1}{\partial \delta} & 0 & \Phi_{\delta}^T & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \frac{\partial \Phi_{\delta}^T \boldsymbol{w}_{D_c}}{\partial \delta} & 0 & 0 & \dots & \Phi_{\delta}^T \end{bmatrix}. \end{aligned} \tag{18}$$

The standard EKF SLAM equations are used from this point forward. Our approach has several advantages over the standard Interaction Primitive framework. The first and foremost is that the phase estimation is now coupled to conditioning in recursive state estimation. With the DTW approach, phase estimation is either performed repeatedly against the conditioned weights with the entire observed trajectory (thus being computationally inefficient), or performed iteratively with only the new observations and thus suffering from an aggregation of errors. By incorporating the phase and phase velocity into the Bayesian inference we can avoid these issues entirely. Second, we now get phase estimation "for free", computationally speaking, in that the weights must be conditioned regardless. However, adding two dimensions to the EKF state that we are already conditioning over is relatively negligible compared to performing DTW every time step.

4 Experiments and Results

In order to analyze the effectiveness of Bayesian Interaction Primitives[1], we test the accuracy, robustness, and computational efficiency against standard Interaction Primitives in three scenarios: one involving synthetic benchmark data and the other two real data from a complex, multimodal, human-robot interaction scenario.

4.1 Experimental Setup

The first experiment utilizes a sample of two-dimensional handwriting data, as shown in Fig. 2. The original trajectory consists of 100 measured states, $y_t = [x, y]^T$, from which new trajectories are created via translation, interpolation, and truncation. Accuracy is measured in terms of the mean absolute error of both dimensions over all time steps.

The effectiveness of BIP is first compared to IP in terms of temporal robustness: 50 training trajectories are generated in which each state $y_t \in y$ is transformed by a linear translation drawn from the Gaussian distribution $\mathcal{N}(0, 5.0)$ with i.i.d. Gaussian noise $\mathcal{N}(0, 0.1)$ applied to each state. Test trajectories are generated with varying lengths by sampling from the trained linear basis model and applying a linear transformation drawn from the same distribution as applied to the training demonstrations. The sampled trajectories vary from 25 to 400 samples (20 trajectories for each domain value), equivalent to a four-fold decrease in temporal speed up to a four-fold increase, assuming a constant sampling frequency. Next we analyze the spatial robustness of BIP by applying different magnitudes of linear translations. The training and test trajectories are held constant at 100 samples, however, the translation is now sampled from a uniform distribution $\mathcal{U}(-a, a)$ where $a \in [1, 15]$. Lastly, we analyze the robustness of BIP to partially observed trajectories. Training and test trajectories are generated via translations drawn from the Gaussian distribution $\mathcal{N}(0, 5.0)$ with i.i.d.

[1] Library source code available at: http://interactive-robotics.engineering.asu.edu/interaction-primitives

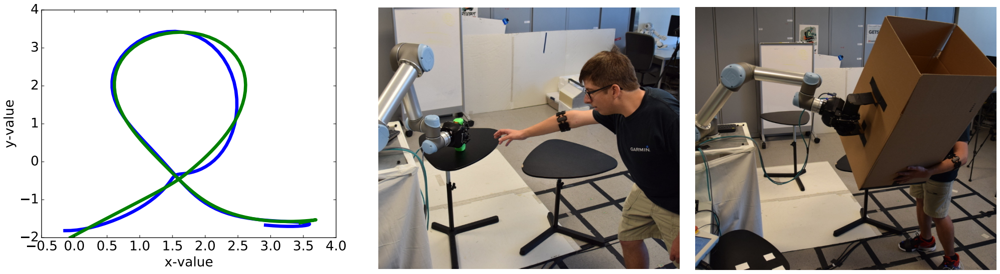

Figure 2: Experimental setup for 1 (left), 2 (middle), and 3 (right).

Gaussian noise $\mathcal{N}(0, 0.1)$ with 100 samples per trajectory. However, now the test trajectories are truncated from the end such that only the first $10b$ samples remain and the phase value for the last observed measurement is b, with $b \in [0.1, 1.0]$.

The second and third experiments consist of a human-robot interaction scenario between a human outfitted with wearable sensors of different modalities and a Universal Robotics UR5 6-DOF robotic arm equipped with a three-finger Robotiq Adaptive Gripper. In both experiments, the trajectories consist of 6 separate sensor modalities for a total of 74 DOF. These include: Kinect skeleton tracking for 13 joints, 6-axis inertial measurement unit (IMU) sensor readings on each forearm, electromyography (EMG) readings from each forearm, pressure sensor readings from each foot, and joint angles for each joint in the UR5 and the gripper. Sensor readings are synchronized online and collected at a frequency of 30 Hz. Training trajectories are obtained via kinesthetic teaching, *i.e.*, the robot is manually manipulated in tandem with the human to compose demonstrations. Testing is performed via leave-one-out cross validation, where the UR5 trajectories are intentionally excluded such that we can measure the error between the generated UR5 trajectory and the expected trajectory. In order to test temporal robustness, test trajectories are truncated and expanded via non-linear deletion and duplication of measurement states drawn from a uniform distribution.

In the second experiment, the human attempts to reach for an object placed on a table out of reach, as shown in Fig. 2. The robot, moving in response to the human, reaches for the object, grasps it, and performs a hand-over maneuver such that the trajectory ends with the human holding the object. It should be noted that no visual detection is used in this experiment; the location of the object is inferred purely from the human's trajectory. The purpose of this experiment is to demonstrate a robotic response to human ergonomics and safety.

The third experiment involves the human and UR5 cooperatively manipulating an object, as shown in Fig. 2. The human grasps the box from one side and the UR5 from the other and together they cooperatively lift the box such that it is placed on a nearby table. In this scenario, we demonstrate the importance of utilizing multiple sensor modalities. The object being manipulated is a large box which obstructs visual skeleton tracking by the Kinect, rendering it all but useless for trajectory prediction. This is an extremely common scenario in everyday life and it is important for human-robot interaction algorithms to be capable of accommodating it.

In each experiment, the BIP algorithm presented in this paper is compared against a standard IP implementation utilizing DTW for phase estimation. Our IP implementation performs a DTW search every 5 measurement updates so as to improve computational efficiency with the reference trajectory projected from the most recently conditioned weights.

4.2 Results and Discussion

The results from the first experiment indicate clear advantages in the BIP framework over standard IP. We first note that while BIP performed slightly worse than IP in terms of temporal robustness (Fig. 4) for very slow and very fast motions ($< 50\%$ or $> 300\%$ movement speed), the generated trajectories in between these extremes are more accurate as well as more consistent, *i.e.*, a smaller standard deviation. Particularly noteworthy is the excellent inference accuracy when the test trajectory is close to the same temporal speed as the training demonstrations; this is a common scenario so accuracy is vital. At the same time, BIP significantly outperformed IP in the spatial robustness

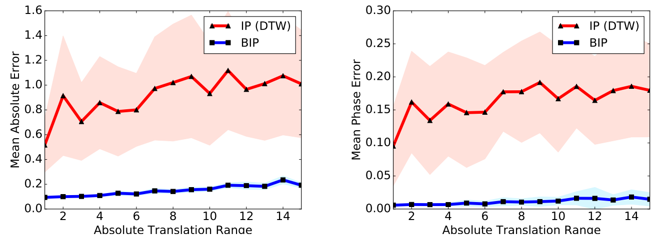

Figure 3: The results for the spatial robustness test for Experiment 1. The x-axis indicates the value of a for the uniform distribution from which translations were drawn.

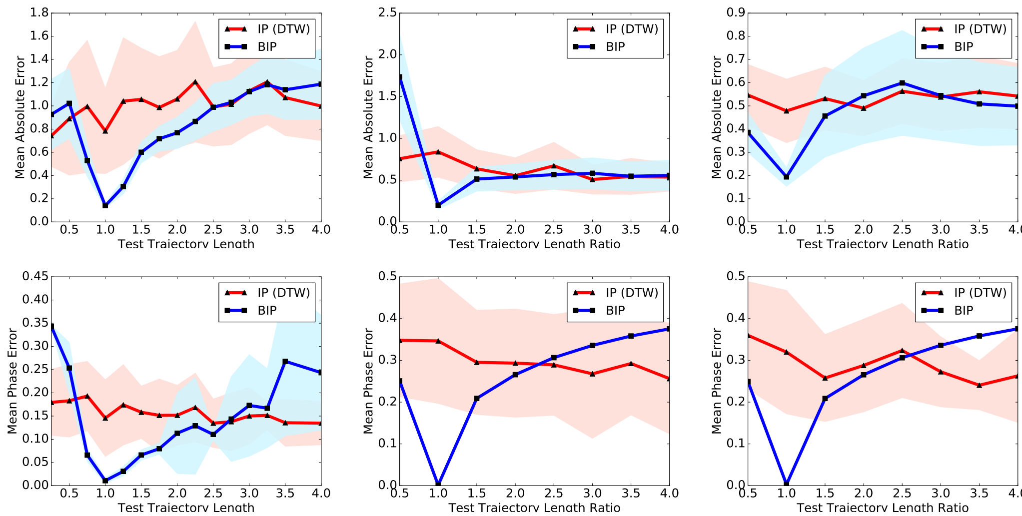

Figure 4: The mean absolute error (top) and mean phase error (bottom) for the temporal robustness test for Experiment 1 (left), Experiment 2 (middle), and Experiment 3 (right). A trajectory length of 1.0 is normal speed, 2.0 twice as fast, etc. The lines indicate the mean result and the shaded regions indicate the standard deviation.

test which is shown in Fig. 3, indicating that BIP is quite robust with regard to spatial uncertainty. In terms of the partial visibility robustness results shown in Fig. 5, the results between the two methods are closer; however, BIP again outperforms IP in both mean absolute error and mean phase error at nearly every length trajectory, only lagging 20% or less of the test trajectory is observed.

These same trends only become more evident for the more complicated scenarios in Experiments 2 and 3. It is clear from Fig. 4 and Fig. 5 that the standard IP algorithm struggles to estimate the correct phase; even with 100% of the trajectory visible there is still a mean phase error of 0.34 and 0.32 in Experiments 2 and 3 respectively. As a result, the inference accuracy does not display any noticeable improvement as more of the observed trajectory is made available. BIP, by contrast, exhibits a constant increase in phase estimation accuracy which leads to a significant improvement in the predicted trajectory error. This is the same type of behavior that is seen in the partial visibility test of Experiment 1. We also observe that the choice of process and measurement noise values becomes extremely important with BIP; a poor choice of noise can cause severe effects in the accuracy of the Bayesian inference.

What makes these results all the more striking is the fact that BIP enjoys a significant computational advantage over IP in every experiment, as shown in Fig. 6. While computational times are difficult to estimate and highly dependent on algorithm implementation, there is the simple fact that the standard IP framework must spend upwards of $O(N^2)$ computational cycles performing DTW at each measurement update step. This is an overhead that BIP simply does not have, and as a result achieves upwards of a 66% reduction in computation time for fully observed trajectories regardless of the dimensionality of the state vector.

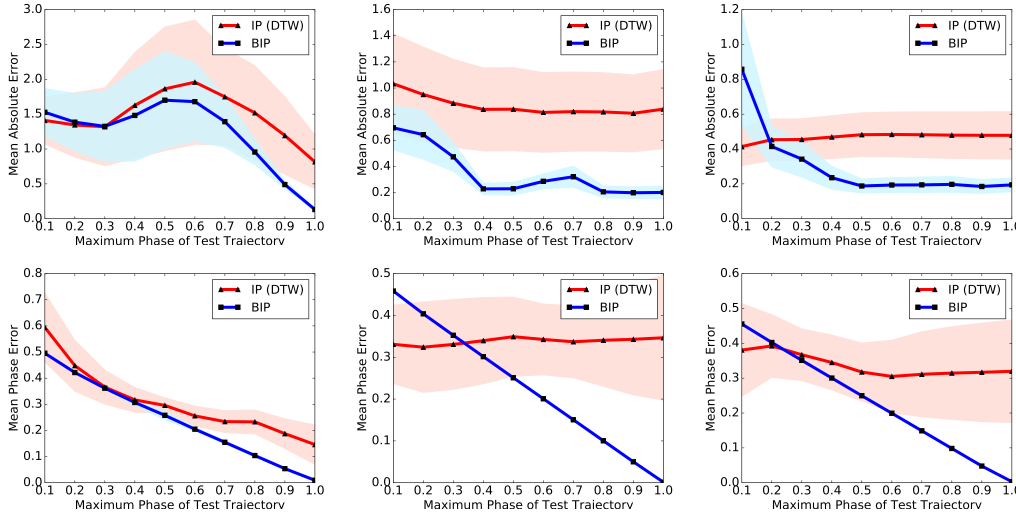

Figure 5: The mean absolute error (top) and mean phase error (bottom) for the partial visibility robustness test for Experiment 1 (left), Experiment 2 (middle), and Experiment 3 (right). The length of the test trajectories are given as a ratio of the observed trajectory to the full trajectory along the x-axis.

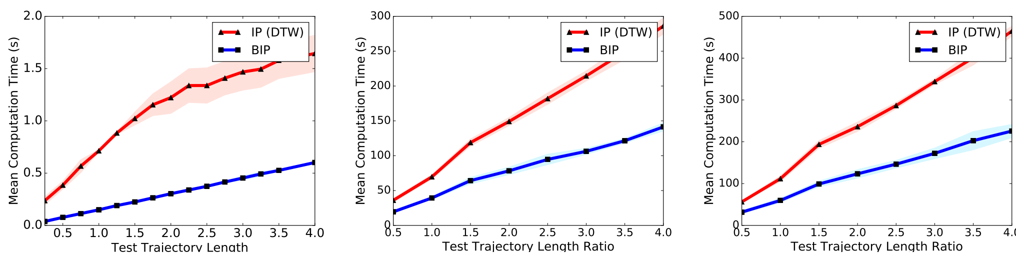

Figure 6: The computational time required to process trajectories of varying lengths for Experiment 1 (left), Experiment 2 (middle), and Experiment 3 (right).

5 Conclusion

In this work we have introduced a fully Bayesian generalization of Interaction Primitives in which phase estimation and trajectory conditioning are coupled together in a joint model of both time and space. We have shown that this enables greater phase estimation accuracy, and consequently inference accuracy, along with reduced computational complexity in human-robot interaction scenarios. Most importantly, this allows us to tackle far more complex scenarios than is possible with normal IP. Despite interacting with a robot in a multimodal, high-dimensional environment, Bayesian Interaction Primitives are capable of accurately performing trajectory inference, thus opening the door to further work in such complex scenarios. Furthermore, we have established a connection between SLAM and human-robot interaction; this is advantageous in that the vast body of knowledge regarding localization and mapping may now be applicable to interaction.

References

[1] 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*, pages 2831–2837. IEEE, 2014.

[2] D. J. Berndt and J. Clifford. Using dynamic time warping to find patterns in time series. In *KDD workshop*, volume 10, pages 359–370. Seattle, WA, 1994.

[3] S. Thrun, W. Burgard, and D. Fox. *Probabilistic robotics*. MIT press, 2005.

[4] H. Durrant-Whyte and T. Bailey. Simultaneous localization and mapping: part i. *IEEE robotics & automation magazine*, 13(2):99–110, 2006.

[5] S. Schaal. Dynamic movement primitives-a framework for motor control in humans and humanoid robotics. In *Adaptive motion of animals and machines*, pages 261–280. Springer, 2006.

[6] G. Maeda, M. Ewerton, R. Lioutikov, H. B. Amor, J. Peters, and G. Neumann. Learning interaction for collaborative tasks with probabilistic movement primitives. In *Humanoid Robots (Humanoids), 2014 14th IEEE-RAS International Conference on*, pages 527–534. IEEE, 2014.

[7] A. Paraschos, C. Daniel, J. R. Peters, and G. Neumann. Probabilistic movement primitives. In *Advances in neural information processing systems*, pages 2616–2624, 2013.

[8] M. Shokoohi-Yekta, J. Wang, and E. Keogh. On the non-trivial generalization of dynamic time warping to the multi-dimensional case. In *Proceedings of the 2015 SIAM International Conference on Data Mining*, pages 289–297. SIAM, 2015.

[9] P. Rebeschini, R. Van Handel, et al. Can local particle filters beat the curse of dimensionality? *The Annals of Applied Probability*, 25(5):2809–2866, 2015.

[10] N. Oliver, E. Horvitz, and A. Garg. Layered representations for human activity recognition. In *Proceedings of the 4th IEEE International Conference on Multimodal Interfaces*, page 3. IEEE Computer Society, 2002.

[11] S. Fine, Y. Singer, and N. Tishby. The hierarchical hidden markov model: Analysis and applications. *Machine learning*, 32(1):41–62, 1998.

[12] N. T. Nguyen, D. Q. Phung, S. Venkatesh, and H. Bui. Learning and detecting activities from movement trajectories using the hierarchical hidden markov model. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 2, pages 955–960. IEEE, 2005.

[13] M. Montemerlo and S. Thrun. Fastslam 2.0. *FastSLAM: A scalable method for the simultaneous localization and mapping problem in robotics*, pages 63–90, 2007.