Learning Ergonomic Control in Human–Robot Symbiotic Walking

Geoffrey Clark and Heni Ben Amor

***Abstract*—This article presents an imitation learning strategy for extracting ergonomically safe control policies in physical human–robot interaction scenarios. The presented approach seeks to proactively reduce the risk of injuries and musculoskeletal disorders by anticipating the ergonomic effects of a robot's actions on a human partner, e.g., how the ankle angle of a prosthesis affects future knee torques of the user. To this end, we extend ensemble Bayesian interaction primitives to enable the prediction of latent biomechanical variables. This methodology yields a reactive control strategy, which we evaluate in an assisted walking task with a robotic lower limb prosthesis. Building upon the learned interaction primitives, we also present a model-predictive control (MPC) strategy that actively steers the human–robot interaction toward ergonomic and safe movement regimes. We compare the introduced control strategies and highlight the framework's ability to generate ergonomic, biomechanically safe assistive prosthetic control. A rich analysis of constrained MPC shows a 20× reduction in the effects of large perturbations on prosthetic control system. We empirically demonstrate a 16% reduction in vertical knee reaction forces in real-world jumping experiments utilizing our control methodology and examine other optimal control strategies in simulated walking experiments.**

***Index Terms*—Learning from demonstration, optimization and optimal control, physical human–robot interaction, prosthetics and exoskeletons.**

I. Introduction

A CRITICAL challenge in the field of prosthetics is the development of healthy and ergonomic control of assistive devices for those with musculoskeletal conditions. Amputations and other musculoskeletal conditions are especially debilitating when present in lower limbs as human adults walk, on average, several thousand steps per day. Consequently, even minor misalignments or suboptimal control trajectories can substantially increase the risk of further musculoskeletal injuries, such as osteoarthritis, due to increased repetitive joint loads [1]. To ensure the efficacy and long-term safety of such devices, it is, therefore, critical to develop control strategies which are responsive to the users' actions and adaptable to their ergonomic needs.

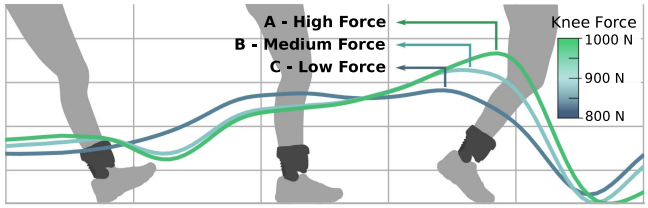

Fig. 1. Introduced model-predictive interaction primitive predicts the future motion and biomechanical state of a human user during walking with a prosthesis. Predictions are conditioned on the robot's actions $\{A, B, C\}$ resulting in variation of the future evolution of the biomechanical state (visualized by three trajectories). The controller then chooses the control action which minimizes an objective function, e.g., minimize the sum of ankle torques over a horizon of future time steps.

In this article, we characterize optimal robotic assistance as a symbiotic interaction [2] in which a robotic device predicts future human states and optimizes control actions to steer the user toward healthy, ergonomic interactions. The fundamental goal is to incorporate the biomechanical well being of the human user into the robot control and decision-making process. To this end, a robotic device needs to be able to anticipate the effects of its actions on a human partner, e.g., how the ankle angle of a prosthesis affects future knee torques of the user, shown in Fig. 1. In addition, such a device needs to leverage generated predictions to *optimize and plan* control trajectories which preemptively avoid high-torque states. To realize the above goals, we focus on three technical aspects of symbiotic interactions with robotic assistive devices: 1) biomechanical prediction, 2) inference of prosthetic controls, and 3) optimization of control trajectories to achieve intended biomechanical and ergonomic outcomes. The presented control algorithm will therefore observe real-time sensor data (as a human user is walking or otherwise interacting with their environment) and use a learned model to estimate an accurate control signal as well as the biomechanical state (such as forces in the knees) of its human partner. Working from the human-centered design perspective, we propose a comprehensive methodology for data collection, augmentation, and learning for human–robot symbiotic walking. Leveraging a data-driven approach, we collect multimodal training

Manuscript received 26 October 2021; revised 30 April 2022; accepted 16 June 2022. Date of publication 5 October 2022; date of current version 8 February 2023. This work was supported in part by the National Science Foundation under the career award under Grant FP00012258, and in part by the Global KAITEKI Center. This article was recommended for publication by Associate Editor C. Yang and Editor E. Yoshida upon evaluation of the reviewers' comments. *(Corresponding author: Geoffrey Clark.)*

The authors are with the School of Computing, Informatics and Decision Systems Engineering, Arizona State University, Tempe, AZ 85281 USA (e-mail: gmclark1@asu.edu; hbenamor@asu.edu).

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the ASU, Institutional Review Board under Application No. 00007521, and performed in line with The Federalwide Assurance (FWA).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TRO.2022.3192779.

Digital Object Identifier 10.1109/TRO.2022.3192779

TABLE I
HIGH-LEVEL OVERVIEW OF POPULAR DYNAMIC MOVEMENT PRIMITIVE INFERENCE METHODS

Method	Model Type	Inference Type	Analysis Type	Planning
Probabilistic Movement Primitives [3] (ProMP)	Probabilistic	Spatial	Exact	No
Interaction Primitives [4] (IP)	Dynamical System	Spatial	Exact	No
Interaction ProMP [5]	Probabilistic	Spatial	Exact	No
Bayesian Interaction Primitives [6] (BIP)	Probabilistic	Spatiotemporal	Exact	No
Ensemble Bayesian Interaction Primitives [7] (enBIP)	Probabilistic	Spatiotemporal	Approximate	No
Model Predictive Interaction Primitives [8] (MPIP)	Probabilistic	Spatiotemporal	Approximate	MPC

demonstrations from able-bodied subjects during human-locomotion tasks. An important feature of this method is the reduction of training trajectories into a low-dimensional latent space using basis function decomposition in a time-independent phase space. We then generate a probabilistic model capable of forecasting the future evolution of the human's full state, including nonobservable biomechanical variables, such as knee reaction force or impulse. Projected estimates are incorporated into control processes to produce an immediate *reactive control* output from the subject's observed state, detailed in Section IV.

However, such a purely reactive control strategy does not adequately take into account the predicted effects of a control policy on the subject's kinematics or kinetics. Going beyond these limitations, we also discuss a model-predictive control (MPC) loop which explicitly minimizes bodily stress over a future time horizon by iteratively querying the learned model (detailed in Section V). To demonstrate our method's predictive and control capabilities, we conducted a variety of simulated and real-world experiments that involve ergonomic control of a lower-limb prosthesis and compare the effectiveness the two control methodologies, namely, reactive control and optimized ergonomic control. This work builds upon our prior work in [7], which introduced the usage of MPC for ergonomic control and provided a demonstration on a jumping task. We substantially expand upon this prior work by 1) providing an expanded and detailed description of the methodology, 2) comparing reactive control to MPC, 3) incorporating comparisons with other motor primitive algorithms, 4) performing a perturbation analysis to investigate robustness, and 5) expanding the evaluation our method to a variety of tasks in both real world and simulation (Experiments 1 and 2).

II. RELATED WORK

Recent years have witnessed substantial progress in the domain of wearable robotic devices, such as lower-limb prosthetics [8]. Early work on prosthetics control has focused on reactive systems utilizing classical control methods, leading to controllers for transtibial as well as transfemoral amputees over varied terrain, such as slopes and stairs [9], [10], [11]. However, a major engineering challenge is to determine how information from healthy humans can be used to estimate the user's locomotive intent, which influences both user adaptation and ergonomy [12].

Machine learning techniques are well positioned to overcome this challenge by modeling healthy human motion in a data-driven fashion and applying the learned features to robotic assistive devices. For example, the work of [13] formulates a deep learning approach to the behavioral cloning and accurate reproduction of a PD controller, with a recurrent neural network. Such approaches utilizing neural networks have shown remarkable predictive power in human intention prediction scenarios [14] including locomotion mode detection [15], terrain recognition [16], or neural prosthetic interfacing [17]. Although deep learning approaches are both popular and powerful, major barriers in the fields of certification, monitoring, and verification for deep neural networks [18], [19] currently prevent their use in safety-critical applications, such as prosthetics and orthotics. It has proven difficult to provide safety criteria necessary to ensure that outputs are both within expected bounds and remain uninfluenced by external forces, such as adversarial attacks [20].

An alternative, and increasingly popular, approach is the use of probabilistic or Bayesian formulations. For example, Thatte et al. [21] leverages sparse Gaussian processes to predict hip angles and hip heights and reactively solves for prosthesis control trajectories with a fast quadratic-program planner for real-time trip avoidance. Locomotion research has leaned heavily on the predictive power of Gaussian processes [22], [23], [24], [25], likely due to their ability to use prior knowledge to generate accurate posterior distributions with uncertainty estimates and the ease with which Gaussian methods are integrated into existing Bayesian filtering and control methods (Kalman or particle filters).

Another probabilistic framework termed interaction primitives (IPs) [4], [6] combines a Bayesian inference scheme with insights from dynamical movement primitives (DMPs) [26]. IPs represent expert trajectories as a weighted superposition of basis functions and then form a probability distribution over the resulting weights. In a similar vein, probabilistic movement primitives (ProMP) [3] also introduced a probabilistic formulation of DMPs that allows for complex inference schemes. As a result, nearly all of the subsequent papers on interaction primitives combine the interaction and phase modeling methodology from [4] with the probabilistic representation and inference scheme introduced [3]. This includes, for example, the work in [27] which introduced mixture models for learning composite behaviors.

Previous applications of IPs illustrate their efficacy in robot-control scenarios requiring human intention prediction, e.g., collaborative assembly [27], catching a thrown ball [35], multiagent handover [28], or engaging in a handshake [29]. IPs have been

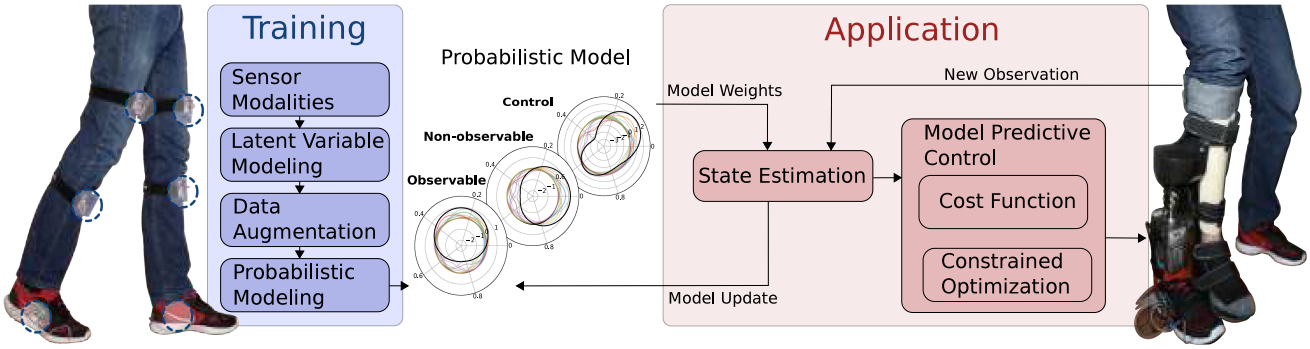

Fig. 2. Overview diagram of our approach depicting both the training and deployment phases. The training phase is composed of the data collection and augmentation via simulation steps, and generates a probabilistic model. The probabilistic model is queried first to generate a new state estimate given observed sensors, and second to formulate an optimal predictive control trajectory to meet biomechanical and ergonomic objectives.

shown to be particularly well suited for predictive modeling of human locomotion and biomechanics [7], termed "predictive biomechanics," through the use of 1) parameterized models for individual agents via basis function decomposition, 2) modeling the joint parameter distributions (covariance of basis weights), and 3) recursive Bayesian filtering (Kalman or particle filtering). Also of note is that although IPs generate locally optimal future state trajectory estimates, they do not take these estimates into account during action generation. Giving IP the ability to incorporate anticipated results of individual actions in an organized way would eliminate the risk of generating unhealthy control actions that lead to undesirable interactions or situations.

Although IPs are particularly powerful at filtering noise in coupled human–robot systems; they require an additional separate step to identify the temporal phase of the interaction. This extra step is required since an interaction involves two separate agents, one of which is strictly observed (human) and the other strictly controlled (robot). Therefore, the controlled agent must infer and adjust for both the temporal and spatial states of the observed agent. As human motion is commonly represented as a low-dimensional subspace of the high-dimensional observation space [30], [31], [32], the proposed solution to temporal prediction in Bayesian interaction primitives (BIP) [6] is to incorporate the system's time dynamics into the inference problem. The joint inference in both time and space sets BIP apart from other dynamic movement primitive methods (DMP) [33], such as ProMP [3] which does not inherently infer the temporal state. Specifying the time dynamics as part of the system state has proven to be a more computationally efficient, accurate, and robust solution than other common methods in human-robot interaction (HRI) scenarios. Additionally, by incorporating a nonlinear observation function through an ensemble Kalman-based approach Ensemble BIP (enBIP) [35] is capable of modeling highly nonlinear physical interactions. Table I contains a high-level overview of these DMP inference methods along with the general model type, inference type, and analysis type for each.

While recent theoretical works have shown that IPs are capable of robust and safe control in highly nonlinear musculoskelatal systems, the control architecture remains singularly reactive in nature [27], [34], [35]. This study goes a step further to enable the planning of optimal trajectories for safer and more fine-grained control of desired of control outputs. Specifically, we illustrate how MPC [36], [37] methods are integrated with BIPs, as an underlying model for biomechanically safe and robust control. The resulting approach is capable of optimal control based on the future nonobservable states of the human–robot system, and is therefore, a powerful tool for generating trajectories which minimize biomechanical stresses within amputees. Although this work singles out the biomechanical regime of human locomotion, the methodology and data-driven framework are applicable to a wide range of HRI scenarios.

III. Learning an Interaction Primitive

As an imitation learning approach, the foundation of our method is a data-driven learning framework, illustrated in Fig. 2. The learning process consists of a) data collection, b) data augmentation, c) latent variable modeling, and d) probabilistic modeling. Typically, the first step in any imitation learning approach is the recording and processing of data, in our case human-locomotion data from a variety of sensors (e.g., IMUs, force plates, motion capture, etc.). It is common to apply a low-pass filter to the raw sensor readings to remove undesirable noise from the data collection process. However, as we will be modeling system time dynamics in a future step, extreme care should be taken to avoid undue influence through the injected phase lag of the filtering process. However, even correctly processed sensor data does not always include all useful or desired variables; in particular, including immeasurable variables, such as knee loads, muscle forces, and prosthetic-control signals, provides critical insights into the biomechanical or ergonomic state of the subject. To this end, we use offline simulations and biomechanics models in order to augment our original dataset.

The high fidelity of demonstration trajectories in time, however, poses a significant challenge with respect to the necessary modeling complexity. We, therefore, project the trajectories into a low-dimensional latent space via basis function decomposition. In turn, the augmented dataset is used to train a probabilistic model of the underlying system and time dynamics, called an enBIP, [6], [29].

The methodology discussed in the following sections extends prior work by reformulating MPC for interaction primitives,

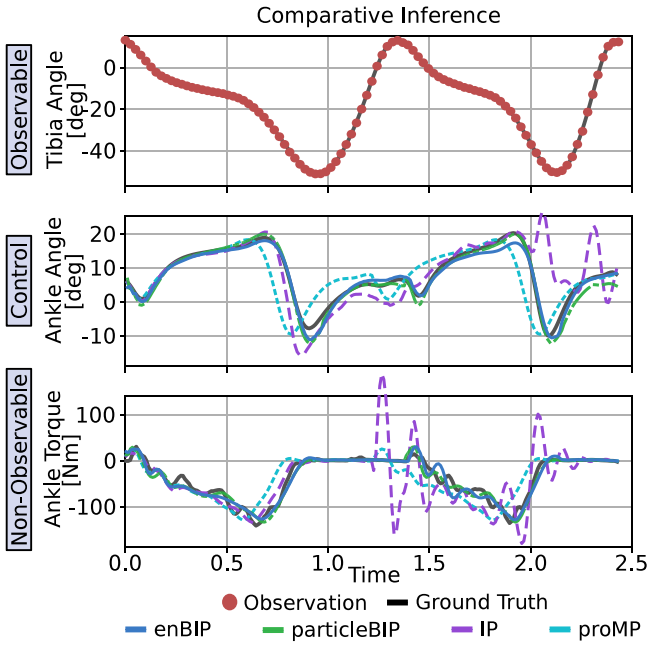

Fig. 3. Comparison of DMP inference methods. We specifically compared the inference capabilities for nonobserved variables using probabilistic movement primitives [3] (proMP), interaction primitives [4] (IP), particle filter Bayesian interaction primitives [7] (particleBIP), and ensemble Bayesian interaction primitives (enBIP).

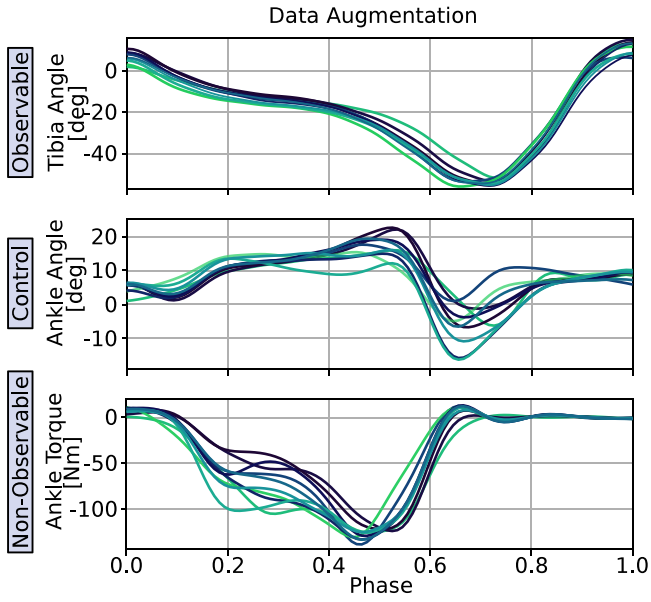

Fig. 4. Walking example with three component variables from human walking data. Observable trajectories (top) are a measurement of the global tibia angle in the sagittal plane, each trajectory is from an independent stride. The augmented control trajectory (center) are collected from the angle of the ankle joint in the sagittal plane. Finally, nonobservable variables are illustrated by ankle torque (bottom), these variables are considered nonobservable since we cannot directly measure them during operation.

while also incorporating constraints. In addition, we substantially expand our experimental section to encapsulate a variety of situations as well as robustness tests. First in Section III-A, our data collection and augmentation process along with the challenges inherent to human walking data and control are discussed. Next in Section III-B, we introduce the latent space formulation and expound on prior work by including a detailed description of von Mises basis functions and their functionality within the context of enBIP. Section IV then mathematically outlines enBIP and illustrates the uses and drawbacks of enBIP in the prediction of human walking data, including a comparative study of enBIP against other DMP-based inference methods, such as: ProMP, IP, and a particle filter inference method (Fig. 3). Finally, Section V presents our MPC scheme and further illustrates both, how to specify constraints in the optimization process, and the benefits of constrained optimization of the systems performance (Fig. 9).

A. Data Collection and Augmentation

To collect human-locomotion demonstrations, we utilize a diverse suite of sensors, including IMUs, smart shoes, motion capture, and an instrumented treadmill. By utilizing high-fidelity state-of-the-art sensor packages we insure that the collected data are of the highest quality possible and that embedded sensors are resilient to external perturbation. Due to the high-quality sensor modalities, we avoid the postprocessing or filtering of raw data, except to downsample higher rate motion-capture data to 100 Hz, which is adequate for human locomotion. However, because these sensors alone do not produce a sufficiently detailed representation of human locomotion, nonobservable biomechanical variables are generated through the use of analytical simulation methods, such as OpenSim [38].

Likewise, control signals can either be recorded via sensor observations, as is the case with a PID controller tracking an observable joint angle, or modeled to generate a more complex impedance control system utilizing joint angles and moments. Regardless of which control method is employed during the data-augmentation process, it should also be adequately implemented and tested on prosthesis hardware to ensure a robust control architecture. Time sequences of both observed and augmented data are processed into N individual strides, $[\mathbf{y}_1, \ldots, \mathbf{y}_{T_n}] \in \mathbb{R}^{D_y \times T_n}$ with D_y signals and T_n time steps for each of the nth demonstrations. As the following section demonstrates, containing all demonstrations in a single matrix simplifies the learning formulation dramatically.

An example is given in Fig. 4, which shows how human subject data are modeled and optimal control trajectories calculated. Our example contains three modalities, (tibia angle, ankle angle, and ankle torque) collected from 100 strides of an able-bodied human subject walking at a fixed speed on an instrumented treadmill. In this example, we will use the observable tibia angle to infer the correct ankle angle and ankle torque values. Since the ankle angle and torque and intrinsically linked we expect to be able to modify the torque by changing the angle.

B. Latent Space Formulation

In order to reduce the dimensionality of the modeling problem to a reasonable level we first project the data from Section III-A into a latent space. The goal of the latent space formulation is

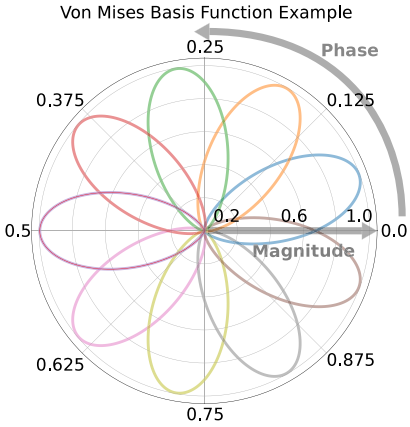

Fig. 5. Example of von Mises basis functions wrapped around a circle, with rotation axis units in phase.

to then estimate a distribution over future observed and nonobserved state variables $\hat{\mathbf{Y}}_{t+1:T}$ given a partial observation of the states and a prior set of N demonstrations $\mathbf{Y}^1_{1:T_1}, \ldots, \mathbf{Y}^N_{1:T_N}$

$$p(\hat{\mathbf{Y}}_{t+1:T} | \mathbf{y}_{1:t}, \mathbf{Y}^1_{1:T_1}, \ldots, \mathbf{Y}^N_{1:T_N}). \tag{1}$$

This formulation, however, requires a generative model over the observed and nonobserved variables as well as a nonlinear transition function, which is a nontrivial challenge. Campbell and Amor [6] proposed to rather transform the full state into a time-invariant representation via basis function decomposition. With basis function decomposition, we define the state space as a linear combination of B^d basis functions $\mathbf{\Phi}_{\phi(t)} \in \mathbb{R}^{B^d}$ with corresponding basis weights $\mathbf{w}^d \in \mathbb{R}^{B^d}$ to approximate each state dimension d as $\mathbf{Y}^d_t = \mathbf{\Phi}^{\intercal}_{\phi(t)} \mathbf{w}^d + \epsilon_y$, with the transition function represented as $\mathbf{y}_t = h(\phi(t), \mathbf{w})$. Note the shift into a relative time measure known as *phase* $\phi(t) \in \mathbb{R}$, where $0 \leq \phi(t) \leq 1$, as well as the approximation error ϵ_y. Basis function selection is usually dependent on application and data; in this work we utilize von Mises basis functions for their ability to represent periodic behavior and cyclical domains

$$\psi_\mu(\phi) = \frac{e^{\kappa \cos((\alpha\phi - \mu))}}{2\pi I_0(\kappa)} \tag{2}$$

where the distribution mean is specified by μ in $0 \leq \mu \leq 2\pi$ and the time scale α is set to 2π to properly scale the phase variable depicted in Fig. 5, as well as Gaussian distributions in noncyclical domains. Finally, κ is the dispersion metric and is analogous to σ^2 in the normal distribution, and I_0 is the zero-order Bessel function. The concatenation of all basis weight vectors $\mathbf{w} = [\mathbf{w}^{0\intercal}, \ldots, \mathbf{w}^{D\intercal}] \in \mathbb{R}^B$ then forms the full state representation, where $B = \sum_d^D B^d$. Intuitively we can reformulate (1) since estimating the time-invariant $\mathbf{w}$ is equivalent to estimating all temporal state values of $\mathbf{Y}$ into $p(\mathbf{w}_t | \mathbf{y}_{1:t}, \mathbf{w}_0)$, where $\mathbf{y}_{1:t}$ utilizes the observation of t time steps. A posterior distribution

$$p(\mathbf{w}_t | \mathbf{Y}_{1:t}, \mathbf{w}_0) \propto p(\mathbf{y}_t | \mathbf{w}_t) p(\mathbf{w}_t | \mathbf{Y}_{1:t-1}, \mathbf{w}_0) \tag{3}$$

is attained with an application of Bayes' rule.

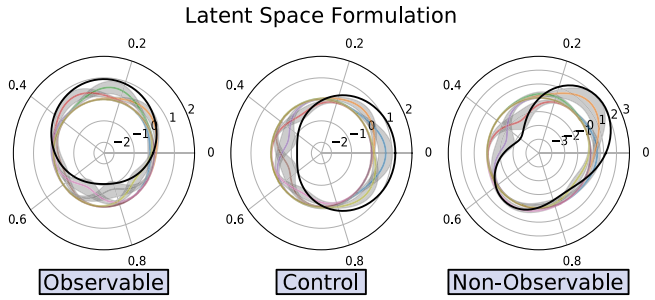

Fig. 6. Example with three component variables (black) decomposed into von Mises basis functions (color) with uncertainty estimates (gray).

An example of our methodology for the latent space formulation is illustrated in Fig. 6. A single observation for each of the state variables is decomposed into a latent set of von Mises basis functions; the cyclical nature of the data is evident as the phase is wrapped around a circle. Gaussian and polynomial basis functions tend to perform poorly on cyclical data as there is a discontinuity between the end of one cycle and the start of the next. Each basis function component is shown in color with the uncertainty estimates places as bounds on each component.

By reducing complex nonlinear trajectories into a linear combination of basis components, we enable a model of sufficient complexity to observe a small subset of the component weights and predict the remainder. One such model, enBIPs, has a number of notable features, namely, 1) prediction of future system states and, 2) inference of nonobserved variables from observed ones, 3) providing uncertainty estimates for inferred components, and 4) robustness against observation and process noise as well as missing values. The functionality provided by enBIPs, enables its use in complex human–robot interactions in which one or more estimates of control variables are needed. Additionally, through the data-augmentation step, enBIPs will generate optimal inference of nonobservable biomechanical variables, such as joint forces, moments, or muscle forces.

Continuing to Section IV we show how the latent formulation is integrated into a Bayesian filtering process to predict the nonobservable augmented variables from a novel set of observed sensor data. By predicting nonobservable control variables, e.g., the ankle angle of the prosthesis, we enable *reactive control*: the generation of control values based on the immediate sensory context. However, the ability to perform predicted biomechanics can be utilized in a more complex way, as shown in Section V, to generate optimized control trajectories over a prediction horizon, thereby maximizing user ergonomics or comfort.

IV. Inference and Reactive Control With Ensemble Bayesian Interaction Primitives

Our objective is to utilize the latent space formulation of human locomotion to generate robot control parameters in human–robot symbiotic tasks. In order to achieve this objective, we need the ability to perform real-time inference of the current and future state in latent space. To this end, observed sensor readings are passed through a Bayesian filter to generate optimal

predictions of sensor values, nonobserved biomechanical variables, and augmented control signals. We achieve approximate inferences regarding the posterior distribution in (3) through the ensemble Bayesian estimation method found in [29]. Fundamentally, the posterior in (3) can be solved by approximating the system state with an ensemble Bayesian filter, as long as 1) the Markov property [39] holds and 2) we can ignore higher order statistical moments in the state error [40]. More specifically, we must first generate an ensemble of latent models which we select at random from the set of example demonstrations, including both state and time dynamics. During execution, as new sensor data becomes available, we then propagate each member of the ensemble one step forward with the state transition function. Observed sensor data are then used to perform the measurement update, perturbed with stochastic noise across the ensemble members. The mean and covariance of the ensemble can easily be calculated from the ensemble members and further projected into the trajectory space by applying the linear combination of basis functions, also known as the nonlinear observation function. The desired reactive control values at the current interaction phase are finally applied to the robot, while the interaction continues. Because we extract full trajectories from the interaction along with estimated time dynamics, this method allows for approximate high-frequency control with periodic low-frequency updates of the model. Which is convenient if either the system is very complex or control is required at a very high rate.

We start by defining an ensemble $\mathbf{X}$ of E members shown by $\mathbf{X} = [\mathbf{x}^1, \ldots, \mathbf{x}^E]$. Optimally we want to sample the initial ensemble $\mathbf{X}_0$ directly from the prior $\mathbf{x}_0 \sim p(\mathbf{w}_0)$ for all $\mathbf{x}_0 \in \mathbf{X}_0$; however, since we do not have direct access to $p(\mathbf{w}_0)$, as a data-driven method, it is standard to instead sample from observed training demonstrations. Random selection on ensemble members is reasonable as the ensemble-based filtering approach provides robustness against possible non-Gausian uncertainties, provided the number of ensemble members is not less than the number of example demonstrations $E \leq N$. As a two-step Bayesian estimation method, our first step approximates $p(\mathbf{w}_t|\mathbf{y}_{1:t-1}, \mathbf{w}_0)$ by propagating each ensemble member forward one time step with

$$\mathbf{x}^j_{t|t-1} = g(\mathbf{x}^j_{t-1|t-1}) + \epsilon_x, \quad 1 \leq j \leq E \tag{4}$$

with constant-velocity state transition operator $g(\cdot)$, and noise error ϵ_x. Next, the ensemble members are updated from the observation and the nonlinear observation operator $h(\cdot)$ via

$$\mathbf{H}_t\mathbf{X}_{t|t-1} = \left[h(\mathbf{x}^1_{t|t-1}), \ldots, h(\mathbf{x}^E_{t|t-1})\right]^\intercal \tag{5}$$

$$\mathbf{H}_t\mathbf{A}_t = \mathbf{H}_t\mathbf{X}_{t|t-1} - \left[\frac{1}{E}\sum_{j=1}^{E} h(\mathbf{x}^j_{t|t-1}), \ldots, \frac{1}{E}\sum_{j=1}^{E} h(\mathbf{x}^j_{t|t-1})\right]. \tag{6}$$

Both the state transition and nonlinear observation operators in ensemble filters differ from classical Kalman filter methods in that they can represent any function. A notable improvement in this formulation over previous work utilizing analytical formulations, such as extended Kalman filter (EKF), is the lack of an introduced linearization error. The deviation of each ensemble member from the sample mean $\mathbf{H}_t\mathbf{A}_t$ and the observation noise matrix $\mathbf{R}$ are then used to compute the innovation covariance with

$$\mathbf{w}_t = \frac{1}{E-1}(\mathbf{H}_t\mathbf{A}_t)(\mathbf{H}_t\mathbf{A}_t)^\intercal + \mathbf{R}. \tag{7}$$

The Kalman gain is likewise calculated directly from the ensemble, with no need to specify an explicit covariance matrix, with

$$\mathbf{A}_t = \mathbf{X}_{t|t-1} - \frac{1}{E}\sum_{j=1}^{E}\mathbf{x}^j_{t|t-1} \tag{8}$$

$$\mathbf{K}_t = \frac{1}{E-1}\mathbf{A}_t(\mathbf{H}_t\mathbf{A}_t)^\intercal\mathbf{w}_t^{-1}. \tag{9}$$

Lastly, the ensemble update occurs by applying the Kalman gain to the difference between actual and expected observations

Algorithm 1: Ensemble Bayesian Interaction Primitives.

Inputs:
- $\mathbf{W} = [\mathbf{w}_1^\intercal, \ldots, \mathbf{w}_N^\intercal] \in \mathbb{R}^{B\times N}$: N demos of B weights.
- $\mathbf{l} = [\frac{1}{T_1}, \ldots, \frac{1}{T_N}] \in \mathbb{R}^N$: Reciprocal demo lengths.
- $\mathbf{y}_t \in \mathbb{R}^D$: State observation at time t.
- $E \in \mathbb{Z}^+$: Number of desired ensemble members.

Outputs:
- $\hat{\mathbf{y}}_t \in \mathbb{R}^D$: State trajectory inference at time t.
- $\hat{\mathbf{v}}_t \in \mathbb{R}^D$: State trajectory variance at time t.

1) Generate initial ensemble $\mathbf{X}_0$:
$$\mathbf{x}^j_0 = [0, \dot{\phi}^j, \mathbf{w}^j], \quad 1 \leq j \leq E$$
$$\dot{\phi}^j = \frac{1}{T_i}, \quad \mathbf{w}^j = \mathbf{w}_i, \quad i \sim \mathcal{U}(1, N)$$
2) Propagate each member of the ensemble one step forward:
$$\mathbf{x}^j_{t|t-1} = g(\mathbf{x}^j_{t-1|t-1}) + \mathcal{N}(0, \mathbf{Q}_t), \quad 1 \leq j \leq E$$
3) Perturb observation with stochastic noise for each ensemble member:
$$\tilde{\mathbf{y}}_t = \left[\mathbf{y}_t + \epsilon^1_y, \ldots, \mathbf{y}_t + \epsilon^E_y\right]$$
4) Perform measurement update:
$$\mathbf{X}_{t|t} = \mathbf{X}_{t|t-1} + \mathbf{K}(\tilde{\mathbf{y}}_t - \mathbf{H}_t\mathbf{X}_{t|t-1})$$
5) Precipitate inferred state and variance from ensemble:
$$\boldsymbol{\mu}_{t|t} = \frac{1}{E}\sum_{j=1}^{E}\mathbf{x}^j_{t|t-1}, \quad \boldsymbol{\Sigma}_{t|t} = \frac{1}{E-1}\mathbf{A}_t\mathbf{A}_t^\intercal$$
6) **Output** reactive trajectory for each controlled DoF:
$$\hat{\mathbf{y}}_t = h(\boldsymbol{\mu}_{t|t}), \quad \hat{\mathbf{v}}_t = h(\boldsymbol{\Sigma}_{t|t})$$
7) Repeat steps 2–6.

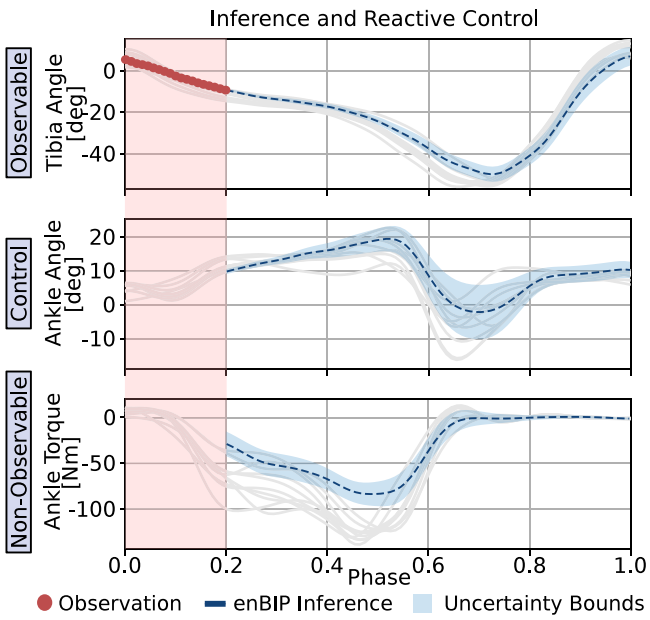

Fig. 7. Inference and reactive control for walking example. Illustration of how the latent formulation infers future system states and reactive control (blue) with uncertainty estimates from a partial state observation (red). The high uncertainty in nonobservable and control variables is due to the subject making small adjustments to their stride which do not effect the observation.

$$\tilde{\mathbf{y}}_t = \left[\mathbf{y}_t + \epsilon_y^1, \ldots, \mathbf{y}_t + \epsilon_y^E\right] \tag{10}$$

$$\mathbf{X}_{t|t} = \mathbf{X}_{t|t-1} + \mathbf{K}(\tilde{\mathbf{y}}_t - \mathbf{H}_t\mathbf{X}_{t|t-1}). \tag{11}$$

As is typical in Kalman filtering, partial observations are sufficient to optimally estimate the full state, which we leverage to generate a posterior over unobservable latent variables, i.e., biomechanics and controls. In the case of partial observations, nonobservable variables are set to zero or random Gaussian noise, while the observation noise is artificially inflated to insure the filtering process rejects the nonobserved input values. A full algorithmic description of the reactive control and inference process is available in Algorithm1.

The inference and reactive control process is further illustrated in Fig. 7, where from a partial observation of a single state variable, the ensemble-based probabilistic model estimates the nonobservable variables forward in phase via the latent formulation. A critical insight at this stage is that although the uncertainty estimate for the observable variable is exceedingly low, showing a well-modeled system, the observable variable does not contain information regarding the relationship between nonobservable and control variables. While the uncertainty is, therefore, much larger for nonobservable and control variables, the model does, however, accurately model the expected or mean trajectory, shown by the blue dashed line. Although a high level of uncertainty in the control estimate is problematic, it is not necessarily a failure of the model, as no model can generate accurate estimates without necessary information. In fact, we can take the models uncertainty about optimal the control trajectory as an opportunity to select specific control trajectories from the latent formulation which correspond to desirable outcomes for other observable or nonobservable variables. For example we could seek to minimize or maximize the ankle torque by modulating the ankle angle.

While there is an entire family of DMP methods which use similar motor primitive schemes with basis functions, we ultimately settled on the enBIP inference method. EnBIP was selected as it has the lowest error of those we tested over long time horizons as it explicitly models the time dynamics as part of the state. The particle filter-based BIP was a close runner up in terms of accuracy. However, since it has not been thoroughly evaluated in other contexts yet, we did not feel confident applying it to human subjects. Both the vanilla Kalman filter-based IP and ProMP models performed poorly as they model the temporal state as a linear process, which progresses as a function of some time value δt. Additionally while the IP method includes an optimal filtering process, ProMP does not and, therefore, executed the estimated mean trajectory given the predetermined phase. If a subject is walking at a consistently slower or faster rate than the ProMP or vanilla IP models expect, we get catastrophic failure of the models over time.

V. Model-Predictive Control With Interaction Primitives

The primary concern we address in this work is generating control signals which conform to requirements of long-term comfort and safety. Because enBIP models the latent formulation between nonobservable biomechanics and prosthesis control variables, but cannot select an optimal control from the range of possible solutions, we construct an efficient MPC [37] framework to optimize the open-loop prosthesis control strategy for more favorable biomechanical outcomes given the learned probabilistic model. The control loop is closed by instantiating only the first control value, and periodically recomputing the control strategy from new observations. Of seminal importance is the cost function structure, in which we integrate across the latent formulation to achieve long-term estimates of the subject's comfort and safety as they relate to the myriad of control trajectories. Optimization, via cost function minimization, incorporates the biomechanical and control expectations from enBIP and effectively produces control solutions able to maximize the health, comfort, and safety of human partners.

MPC formulations are usually noted as the optimal control strategy $\mathbf{u}^*$ which minimizes a specified cost function $J(\mathbf{u})$. The classical MPC method recursively modifies the control vector $\mathbf{u}$ and then iteratively calls a state transition function to compute the sum of its effects on some reference variables $\mathbf{r}$ out to the specified time horizon H_r

$$\mathbf{u}^* = \text{argmin}(J(\mathbf{u})) \tag{12}$$

$$J(\mathbf{u}) = \sum_{i=1}^{H_r} ||r||^2 + \sum_{i=1}^{H_u} ||\Delta\mathbf{u}||^2. \tag{13}$$

A common addition to the cost function as seen in (13) is a cost related to the deviation of the new control strategy from the previous control strategy $\Delta\mathbf{u} = \mathbf{u}_{t-1} - \mathbf{u}_t$ out to the control horizon H_u. Minimizing the control deviation helps to

ensure stability and rejection of perturbations. Building upon (13), the latent formulation provides the notable advantage of an analytical time-invariant composition due to the basis function construction of the system $\mathbf{Y}_t^d = \mathbf{\Phi}_{\phi(t)}^\intercal \mathbf{w}^d$. We can, therefore, simply compute the cost over the time horizon as an integration of weighted basis functions

$$J(\mathbf{w}^{\mathbf{u}}) = \int_{\phi}^{H_r} ||\mathbf{\Phi}_{\phi(t)}^\intercal \mathbf{w}^{r|u}||^2 \, d\phi + \int_{\phi}^{H_u} ||\mathbf{\Phi}_{\phi(t)}^\intercal \Delta \mathbf{w}^u||^2 \, d\phi. \tag{14}$$

It follows that we now have a cost over the reference variable weights $\mathbf{w}^{r|u} = \mathbf{w}^u Cov(\mathbf{w}^u, \mathbf{w}^r)$ (calculated using the covariance model from enBIP) and control trajectory change $\Delta \mathbf{w}^u = \mathbf{w}^{u_{t-1}} - \mathbf{w}^{u_t}$ (the difference between the last and current control weights). To facilitate greater efficiency, we pull the weight vectors out of the integration

$$J(\mathbf{w}^{\mathbf{u}}) \leq \int_{\phi}^{H_r} ||\mathbf{\Phi}_{\phi(t)}^\intercal||^2 d\phi \, ||\mathbf{w}^{r|u}||^2 + \int_{\phi}^{H_u} ||\mathbf{\Phi}_{\phi(t)}^\intercal||^2 d\phi \, ||\Delta \mathbf{w}^u||^2 \tag{15}$$

since the analytical integrals $\mathbf{\Psi}_{\phi:H} = \int_{\phi}^{H} ||\mathbf{\Phi}_{\phi(t)}^\intercal||^2 d\phi$ can be precomputed and stored in a hash table trading memory for speed. Incorporating the new cost function into the optimization problem then resolves to

$$\mathbf{u}^* = \text{argmin}(\mathbf{\Psi}_{\phi:H_r}^\intercal ||\mathbf{w}^{r|u}||^2 + \mathbf{\Psi}_{\phi:H_u}^\intercal ||\Delta \mathbf{w}^u||^2) \tag{16}$$

where in the case of the Gaussian basis function, the integration resolves to

$$\mathbf{\Phi}(\phi_{0:1}) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(\frac{\phi-\mu}{2\sigma})^2} \tag{17}$$

$$\mathbf{\Psi}_{\phi:H} = \int_{\phi}^{H} ||\mathbf{\Phi}(\phi_{0:1})||^2 = \sqrt{\frac{2}{\pi}} e^{\frac{\mu-\phi}{2\sigma}}. \tag{18}$$

Application Remark 1: It is important to note that the von Mises basis function in (2) does not have an analytical integral as it requires the integration of an infinite series of Bessel functions. The indefinite integral is therefore difficult to attain, and we instead compute, via numerical quadrature integration [41], the definite integral; a reasonable approximation as the upper and lower bounds of the integration are naturally constrained in the von Mises distribution to $\pm\pi$.

Application Remark 2: Modern constrained optimization solvers, such as sequential least squares quadratic programming (SLSQP) [42], provide both fast and accurate solutions to the optimization problem in (16). Alternately, for more robust solutions multiobjective MPC solvers compute a Pareto-optimal solution, seen in [43]. Regardless of the particular solver, upper and lower bounds can be imposed on the control trajectory and selected from the maximum and minimum values for each basis function during training.

Application Remark 3: Using enBIP we are able to condition on past events and predict future control states with a probabilistic model. However, when applying this same model to the MPC process to optimize the control trajectory for specific results we may not want to optimize future control states based on past executed controls. Recent theoretical insights [44] into learning models for control show that conditioning on prior actions can cause "self-delusions," i.e., the model takes its own actions as evidence about the world, thereby slowly corrupting the inference process. To avoid the occurence of such self-delusions, we zero out all weights below the diagonal of the weight matrix $\mathbf{w}^{r|u}$ from (15), which represent the delusional effects of future controls on past actions.

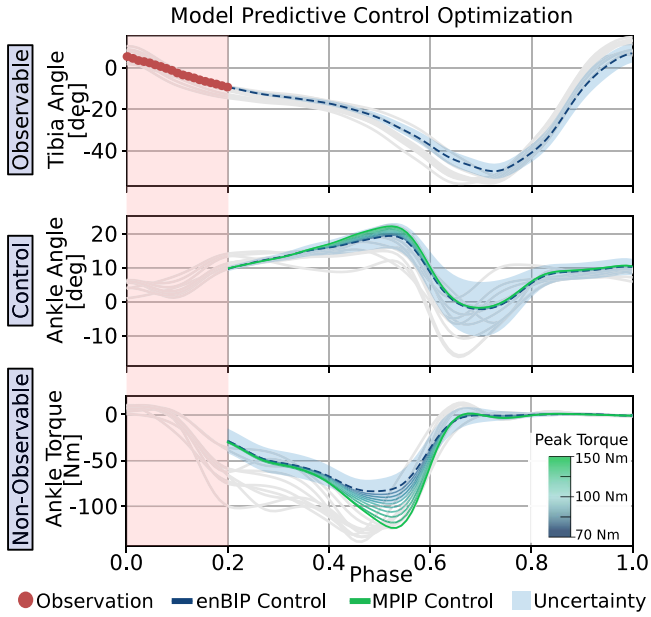

Fig. 8. Reactive versus model predictive control for walking example. Inference of the reactive control and nonobservable trajectories (blue) with uncertainty estimates from a partial state observation. Using MPIP we optimize (green) the ankle angle control to maximize the magnitude of the nonobservable ankle torque. Intermediate optimization trajectories are shown between the inferred and optimal trajectories.

An example model predictive interaction primitives (MPIP) solution is shown in Fig. 8 for the walking problem, and compared with the reactive inference and control values. In this example, we selected a control trajectory through optimization to maximize the magnitude of the nonobservable (ankle torque) trajectory. Small changes in the ankle angle during pushoff, where the peak force occurs, produce a drastic effect on the expected torque. The intermediate lines illustrate the path the optimization process took to maximize the ankle torque in this example. To complete the MPC process we actuate only on the current temporal control signal and reoptimize the entire trajectory every time-step, thus closing the control loop.

In the walking example shown in Fig. 8 we included the control difference from (15), in which we minimize the difference between the current and last control values $\mathbf{u}_{t-1} - \mathbf{u}_t$ for the time horizon $H_u = 0.1$, which is approximately 10% of the phase. The control delta minimization keeps the control from fluctuating too quickly. The control time horizon, therefore, influences the robustness of the controller by filtering large changes in the control. We illustrate the effectiveness of different control horizons in disturbance rejection in Fig. 9. We compare the baseline (no MPC, blue) against MPIP (red) with progressively increasing control horizons, while applying a large

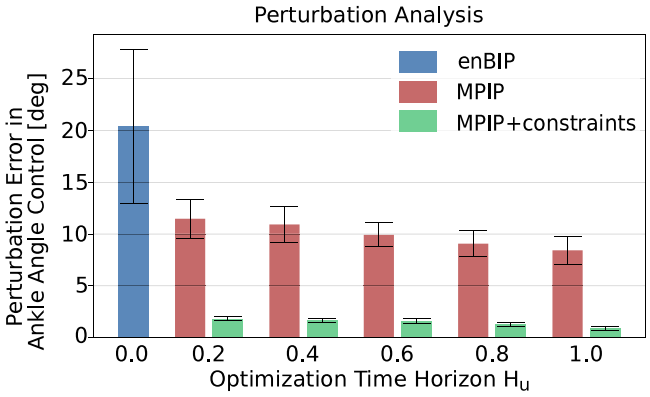

Fig. 9. Perturbation analysis of MPIP, without the MPC portion (blue), without constraints (yellow), and our version of MPIP with both MPC and constraints implemented. A large Dirac delta perturbation was added to the observation, to analyze its effects on the control output.

impulse type perturbation to the observation. The large impulse is partially filtered by the Kalman filter, however, there remains a significant change in the control signal, on the order of 20 degrees change. MPIP, however, cuts the effects in half and further reduces as the time horizon approaches 1. Furthermore, Fig. 9 shows that the application of constrained optimization (green) in the MPC process significantly increases the rejection of perturbations from output control signals by approximately $20\times$. More specifically we incorporate the error covariance from the Kalman filter as bounds for the constrained optimization.

Utilizing MPC with enBIPs produces smooth output control trajectories which effortlessly transition the reactive control output to minimize the latent formulation. Constraints are imposed on the MPC solution via the uncertainty bounds to give an added measure of safety by ensuring that the MPC solution does not deviate from the latent distribution. The full algorithmic description of utilizing MPC with enBIP is shown in Algorithm 2.

Algorithm 2: Model Predictive Interaction Primitives.

Inputs:

- $\mathbf{W} = [\mathbf{w}_1^\intercal, \ldots, \mathbf{w}_N^\intercal] \in \mathbb{R}^{B\times N}$: N demos of B weights.
- $\mathbf{l} = [\frac{1}{T_1}, \ldots, \frac{1}{T_N}] \in \mathbb{R}^N$: Reciprocal demo lengths.
- $\mathbf{y}_t \in \mathbb{R}^D$: State observation at time t.
- $E \in \mathbb{Z}^+$: Number of desired ensemble members.
- J: Cost function for biomechanical optimization.

Outputs:

- $\hat{\mathbf{y}}_t \in \mathbb{R}^D$: State trajectory inference at time t.
- $\hat{\mathbf{v}}_t \in \mathbb{R}^D$: State trajectory variance at time t.

1) Generate initial ensemble $\mathbf{X}_0$:

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

$$\dot{\phi}^j = \frac{1}{T_i}, \quad \mathbf{w}^j = \mathbf{w}_i, \quad i \sim \mathcal{U}(1, N)$$

2) Perform steps 2-5 of Ensemble Bayesian Interaction Primitives (Algorithm 1)

3) Calculate basis integral for necessary time horizons:

$$\mathbf{\Psi}_{\boldsymbol{\mu}^1_{t|t}:H} = \int_{\boldsymbol{\mu}^1_{t|t}}^{H} \qquad ||\mathbf{\Phi}(\boldsymbol{\mu}^1_{t|t})||^2 = \sqrt{\frac{2}{\pi}} e^{\frac{\mu - \boldsymbol{\mu}^1_{t|t}}{2\sigma}},$$

$$1 \le b \le B$$

4) Optimize control weights for cost function:

$$\mathbf{u}_t^* = \text{argmin}(J(\boldsymbol{\mu}^u_{t|t}, \mathbf{\Psi}_{\phi:H}))$$

5) **Output** reactive trajectory for each controlled DoF:

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

6) Repeat steps 2-8.

VI. Experimental Setup

Real-World Experiments: Our experiments heavily utilize human–subject data to train our predictive models. The following section details our data collection and augmentation methodology, including sensor and algorithm descriptions for Experiments 1 and 4. Real-world experiments were performed in two stages: 1) data collection and augmentation and 2) model evaluation. In the data collection phase, we performed multimodal data collection on participants, fitted with inertial measurement units (IMUs) on each lower-limb segment and retroreflective markers for a VICON motion capture system, which included a split-belt instrumented treadmill to capture participant kinematics. Additionally, ground reaction force sensors in the form of pressure sensitive shoes, measuring the subject's force at four points in the shoe (across the toes, center of the heel, first metatarsal, and fourth metatarsal) were fitted to each participant to capture kinetic data in environments a treadmill cannot simulate. Each participant was then asked to perform a specific locomotive task, such as walking or jumping.

In addition to the IMU and shoe force data, we processed the motion capture and ground reaction force data with the commercially available OpenSim [38] to extract accurate kinematics (angles/velocities) and kinetics (reflected forces and moments) for data augmentation purposes. Because OpenSim, or any human biomechanics simulation software for that matter, has no scientific consensus regarding its accuracy, specific quantitative estimates may be of questionable value. We instead followed the advice from the OpenSim best practice [45] and model, in Section III-B, the qualitative relationships between the simulated forces and observable sensor values as (conditional) distributions. It follows empirically that individual subject's joint loads may differ dramatically due to biological difference (i.e., sex, weight, or age); however, we seek to model general trends between variables, allowing for MPC to, later on, maximize or minimize trajectories without focusing on quantitative values. To ensure that data are of the highest quality, we go a step further and incorporate sanity checks with data range, calibrate to accurate physiological measurements, and cross-compare the results from OpenSim with results generated from the Vicon plug-in gait model [46] ($\leq$5% error). Overall, nonobservable variables include: ankle angle, velocity, and moment, as well as biomechanical variables strongly linked to osteoarthritis in amputees, that is, the vertical knee-reaction force and medial

knee moment. These variables were, in turn, augmented to the recorded sensor dataset.

Although simulation and modeling strategies are the de facto standard in the biomechanical analysis of clinical gaits, technology does exist to measure joint forces and moments *in vivo* [47]. It has been reported in [48] "that current generic musculoskeletal modeling techniques are able to reproduce the in vivo conditions (i.e., ground truth) measured during walking,". We would, therefore, like to highlight that our methodology is agnostic to the type of approach used to collect the nonobserved, biomechanical data and that it can be used with either simulations or *in vivo* measurements.

In the model-evaluation phase, control methods were tested on able-bodied subjects utilizing a powered prosthesis for the same locomotive task. A large variety or powered prosthetics exist, we chose to utilize the SpringActive "Ruggedized Odyssey Ankle" [49] because its parallel elastic mechanism allows both the prosthesis and subject to drive the interaction. The high level of synergy between human and robot in our scenario necessitates a *cooperative* approach to physical interaction. Subjects participating in the real-world tests were fitted with an ankle bypass; a carbon fiber structure molded the the foot and lower limb and was constructed such that a prosthetic ankle can be fitted to allow the able-bodied subject to walk on the prosthesis. Fitting and alignment of a prosthesis with the ankle bypass is done in much the same way as a traditional prosthesis and per device recommendations.

All experiments with human subjects were conducted in accordance with Arizona State University Guidelines. Written informed consent was obtained under an appropriate institutional review board (IRB) review.

Simulation Experiments: To evaluate the modeling accuracy with respect to biomechanical prediction, a simulation environment was used. Since internal biomechanical features can be calculated directly without external sensor noise, this approach allows for more accurate testing conditions. Simulation experiments were performed in two stages 1) data collection and augmentation and 2) model evaluation.

As a simulation environment, we utilized the winning submission from the NeurIPS AI for Prosthetics competition. The competition focused on controlling a biomechanically realistic musculoskeletal simulation through muscle activations and it notably featured a passive prosthesis (spring/damper) mounted to the right lower limb. Furthermore, the NeurIPS competition simulator uses OpenSim in a physics-based environment to provide a physiologically accurate musculoskeletal dynamics. The NeurIPS competition provides an excellent comparison model since it enables access to traditionally nonobservable variables. We performed data collection and augmentation by modifying the simulator such that kinematic and kinetic variables, such as joint angles and forces, are output while the winning submission controls the model at a range of speeds from 0.5 to 1.2 m/s to be in line with the locomotion speeds of Experiment 1. As traditionally nonobservable biomechanical forces were the focus of this experiment, we collected data from a multitude of biomechanical regimes including 1) vertical knee-reaction force: the reflected force across the knee joint as the sum of all external forces in the vertical world frame, 2) medial knee moment: moment about the knee joint in the medial direction, and 3) sum of muscle forces: aggregate of forces in right leg simulated muscle groups, including the hamstring, biceps femoris, vasti muscles, rectus femoris, iliopsoas, and gluteus maximus.

Model evaluation utilized an identical simulation environment, except for the addition of a controllable series elastic actuator at the ankle joint and an accompanying PID controller. In this control strategy, the series elastic spring is set to mimic the quasi-static stiffness curve of the ankle during walking, whereas the controlled actuator modifies the shape of the curve between strides. Quasi-active control further allows for the ankle control values to be easily tuned for different speeds or subject conditions.

VII. Experiments

In this section, we present experiments in the domain of assisted mobility that highlight the capabilities of our methodology and compare them to other state-of-the-art methods. We target two phases of the MPIP methodology and highlight results in both simulated as well as real-world environments. Experiment 1-Section VII-A and experiment 2-Section VII-B highlight the performance of reactive control and inference of enBIPs in human locomotion tasks. Experiment 3-Section VII-C and experiment 4-Section VII-D add a comparison of the purely reactive control method to MPIP by evaluating safety, comfort, and biomechanical stresses placed on the human body.

A. Experiment 1: Reactive Control on Robotic Prosthesis

In the first experiment, we apply enBIPs to evaluate the inference capabilities on real-world and simulated locomotion data when predicting future reactive control and biomechanical trajectories. Each participant was instructed to walk on the instrumented treadmill at five uniformly distributed speeds between 0.5 and 1.3 m/s.

Training incorporates both sensors and augmented biomechanics from OpenSim into a single model and can be seen as a behavioral cloning approach, whereas only the inference mechanism of the enBIP is used to generate reactive control values from recorded joint angles. As a generic model of human locomotion is desired, data from all five participants were combined in training to form a single model. In this way, we hope to capture as general a model as possible without imposing a single individual's distinct gait or posture dynamics on other subjects. The joint subject model was tested on reserved sensor data from each participant in a test of the model's inference capabilities for both observed and nonobserved variables.

Table II shows the mean absolute error (MAE) of predictions averaged across all five subjects. Our method tracks the observed kinematics closely, with an average error less than 1.5 degrees. However, although the inference error of nonobserved ankle kinematics is also very low, at approximately 1.45 degrees, we see significantly higher errors across other biomechanical variables including ankle moment and vertical knee force. As suggested by [50], frequent errors in biomechanical variables are likely driven by two factors 1) imperfect biomechanical

Fig. 10. Assisted mobility with model predictive interaction primitives running in real-time on a SpringActive Ruggedized Odyssey prosthesis.

TABLE II
PREDICTION ERRORS ON ALL VARIABLES

Real-World Observed Inference Error

		MAE	% Error	Range
GRF Vertical	N	23.44	2.74	882.74
Tibia Angle	deg	1.50	2.09	68.54
Tibia Velocity	deg/s	10.68	2.16	484.41
Femur Angle	deg	0.93	2.23	44.34
Femur Velocity	deg/s	9.31	3.27	284.88
Knee Angle	deg	1.82	2.66	67.17
Knee Velocity	deg/s	18.79	2.72	685.35

Real-World Nonobserved Inference Error

		MAE	% Error	Range
Ankle Angle	deg	1.45	5.19	24.69
Ankle Velocity	deg/s	30.11	6.92	320.44
Ankle Moment	Nm	8.14	4.22	172.41
Vertical Knee Force	N	43.90	9.53	458.16

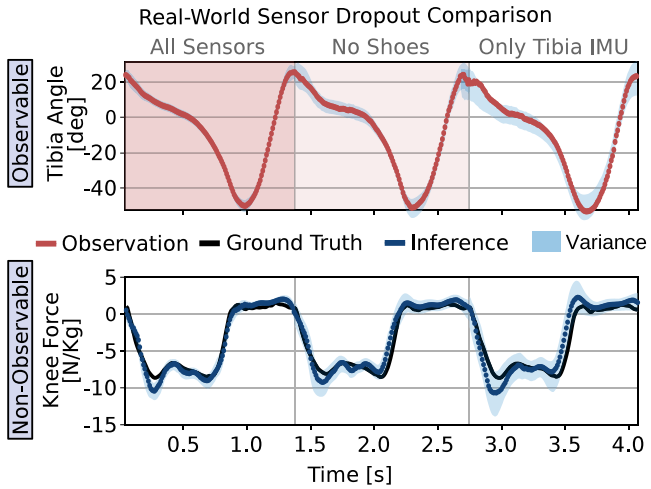

Fig. 11. Inference of the observable tibia angle in degrees (red) and nonobservable vertical knee-reaction force (blue) under three conditions (1) all sensors are functioning properly, (2) shoe force sensors have failed, and (3) all sensors failed except for the tibia IMU.

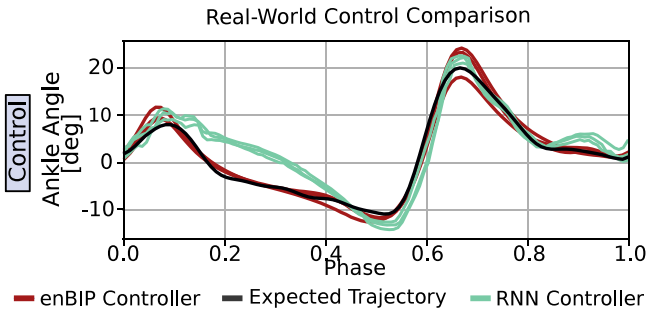

Fig. 12. Comparison of able-bodied ankle angle (black) to a state-of-the-art adaptive RNN controller [13] (green) and the enBIP inferred control (red).

modeling and 2) motion capture errors from the complacency of soft tissues.

Knee forces are known to be especially susceptible to simulation errors due to the small movement of the knee relative to its compliance in the motion capture model and the difficulty in modeling the complex multidimensional structure of the biomechanical joint. The idea that data collection and biomechanical modeling errors compound to influence the modeling abilities of the enBIP is further corroborated by the high noise estimate seen in the latent variables associated with biomechanical forces as compared to the direct kinematic measurement.

Of critical importance when evaluating machine learning methods for powered prosthetic devices is *safety*. One aspect of safety, aside from model accuracy, is the minimal deterioration of projected inference due to external noise or sensor dropout, known as graceful degradation. Graceful degradation as exhibited during inference is shown in Fig. 11, where sensors are iteratively removed until a single inertial sensor remains. It can be seen that as sensors are removed, the uncertainty of the prediction increases and inference accuracy decreases, as expected from a loss of sensing modalities. However, the system remains stable and does not exhibit significant or unexpected adverse effects from the sensor dropout.

Finally, the reactive control output was used to actively control a robotic prosthetic ankle in a real-world locomotion task, shown in Fig. 10. The subject was asked to walk, with the assistive device, first on a treadmill under controlled conditions and later on a sidewalk environment using the joint model from all five test subjects. Interaction between the subject and the powered prosthesis was successful even when performed in out-of-laboratory conditions. Fig. 12 shows our reactive control method (in red) and a state-of-the-art adaptive controller (in green) compared against the expected trajectory from data collection (in black). The adaptive recurrant neural network (RNN) controller was trained based on the methodology of [13], while using the SpringActive robotic prosthetic device. The adaptive RNN controller utilizes the angular position and velocities from the femur and tibia in order to compute the motor current necessary to drive the ankle controller. The RNN controller was trained with

TABLE III
PREDICTION ERRORS ON ALL VARIABLES

Simulation Observed Inference Error

		MAE	% Error	Range
GRF Vertical	N	41.13	3.67	1119.24
Tibia Angle	deg	1.79	3.40	52.70
Tibia Velocity	deg/s	50.37	7.31	689.16
Femur Angle	deg	2.62	6.27	41.79
Femur Velocity	deg/s	25.92	5.61	461.88
Knee Angle	deg	1.64	3.37	48.51
Knee Velocity	deg/s	54.96	7.81	703.40

Simulation Nonobserved Inference Error

		MAE	% Error	Range
Ankle Angle	deg	0.73	3.17	23.02
Ankle Velocity	deg/s	40.50	5.51	734.52
Vertical Knee Force	N	166.84	4.51	3700.71
Medial Knee Moment	Nm	15.56	6.75	230.48
Sum Muscle Forces	N	352.48	7.33	4809.41

early stopping to avoid overfitting and attained an mse of 7.16 in the real-world control task. The significantly lower testing mse of 1.81 for the enBIP model, Fig. 12 illustrates that the RNN controller does not generalize well to unseen data. Since we train on a limited dataset, the control output of the RNN controller significantly deviates from the expected trajectory during the stance phase. This insight indicates that the PID portion of the RNN controller was not correctly trained. In comparison, the enBIP controller accurately matched the ankle angle during the stance phase, which further illustrates the power of our method to correctly learn a controller from limited expert data. Furthermore, because the enBIP implementation included multimodal sensor data, the controller is able to distinguish between the stance and swing phases and apply the necessary control outputs accordingly. In the final swing phase, the adaptive RNN controller has a small output oscillation due to the controller not correctly differentiating the end of one stride from the beginning of the next. While this amounts to a minor deviation, it shows that the adaptive RNN controller is susceptible to fluctuactions which may lead to dangerous situations.

B. Experiment 2: Reactive Inference in Simulated Environment

An important motivation behind the use of Bayesian inference is the ability to leverage multiple low-cost and low-quality sensors to infer high-quality variables which would otherwise be too (computationally or economically) costly to generate. The underlying rationale is that such high-quality data, such as information on human biomechanics dynamics, are available at training time and can be used to train our models. The training process, in turn, reveals the hidden relationships and covariances between the high-quality data and our low-cost sensors. After training, these relationships can be used to infer the biomechanics from low-cost sensors alone.

To further evaluate the efficacy of inference for nonobservable biomechanical factors, we ran a comparative analysis of locomotion with a passive prosthesis in the simulation environment described in Section VI. Data collection in the simulation environment focused on the variables listed in Table III, while training following the enBIP methodology was performed and the resulting model was tested on a new simulation run at 0.8 m/s. The results of this experiment are shown in Table III.

When comparing Tables III to II, it is immediately clear that inference of direct kinematics is slightly less accurate in the simulation environment. We postulate that the increased modeling error derives from the lower local dynamic stability (LDS) the simulated subject exhibits, with a Lyapunov exponent of 1.78 in comparison to the human subject's. However, even with a lower LDS, inference error on biomechanical variables is reduced by utilizing the simulated environment data. Furthermore, the simulation environment enables collection of novel and interesting biomechanical values, such as muscle forces, providing a detailed look into how different control strategies may effect ergonomic quantities, such as muscle fatigue and energy expenditure.

C. Experiment 3: Optimal Control With MPIP in Simulation

Although reactive control can be very efficient in many scenarios, it does not take into account the future bodily ramifications of generated robot actions. By contrast, the MPIP algorithm is able to generate control trajectories which 1) minimize vertical force at the knee, 2) minimize knee moment in the medial direction, and 3) minimize muscle forces in the right leg.

To adequately train our model, we again use the winning submission from the 2018 NeurIPS competition [51] to collect kinematic as well as biomechanical features of locomotion, while incorporating low-frequency pseudorandom control of the lower-limb prosthesis. After training, we performed locomotion experiments in the virtual environment for four different modes 1) passive prosthesis as experimental control, 2) reactive control generated via the inference step of enBIP, 3) MPIP variant in which the cost function minimizes vertical knee-reaction force, and 4) MPIP variant in which the cost function minimizes muscle forces of the amputated limb. Each cost function specified time horizons covering the entire phase as well as equally long control horizons to prevent high-frequency changes in the prosthesis control from individual perturbations.

Results averaged across 60 strides from each mode are listed in Table IV. The first, reactive controller, mode generates approximately equal internal stresses (i.e., knee forces) as compared to the passive prosthesis, however, it exhibits a significant increase in the LDS calculated via the Lyapunov exponent. We speculate that empirically due to the tight coupling of the system, actuating a *mean* or average expected trajectory from the reactive controller posterior adds stability to the human robot system by constantly nudging the subject toward a more dynamically stable regime. For both healthy and pathological locomotion analysis the LDS has been shown to be a reliable indicator of dynamic stability corresponding to risk factors of injuries and falls. When applying MPC in tandem with an enBIP the cost function pertaining to a reduction in the vertical knee-reaction force performed the best. It results in significant reductions in the vertical knee impulse and medial knee impulse, both of which are highly associated with increased risk of osteoarthritis [52]. Additionally, minimization of the vertical knee force produces

TABLE IV
RESULTS OF WALKING EXPERIMENT IN SIMULATION

Simulation Experiment 60 Stride Average					
Control Methods	Medial Knee Impulse (Nm· s)	Vertical Knee Impulse (N· s)	Muscle Impulse (N· s)	Prosthetic Impulse (N· s)	Dynamic Stability Lyapunov Exponent
Passive Prosthesis	44.6 ± 16.3	1114 ± 209	1880 ± 337	25.1 ± 4.01	1.78
Interaction Primitives (enBIP)	44.9 ± 23.9	1160 ± 275	1880 ± 445	22.3 ± 3.31	-0.617
MPIP (Knee Force Reduction)	**29.4** ± 12.5	**913** ± 245	1900 ± 238	21.3 ± 4.82	**-0.226**
MPIP (Muscle Force Reduction)	47.6 ± 19.2	1200 ± 186	**1680** ± 469	23.9 ± 2.45	3.31

The best (lowest) values are highlighted for each optimized variable.

TABLE V
RESULTS OF REAL-WORLD JUMPING EXPERIMENT WITH A SPRINGACTIVE PROSTHESIS

Prosthesis Experiment 10 Trials					
Control Methods	Peak Heel Force (Body Weight)	Peak Toe Force (Body Weight)	Peak Knee Force (Body Weight)	Total Knee Impulse (Body Weight· s)	Impact Ankle Angle (Deg)
AB Soft Landing Training	0.211 ± 0.041	0.374 ± 0.093	2.36 ± 0.355	118 ± 16.6	-22.6 ± 3.96
AB Hard Landing Training	0.660 ± 0.045	0.406 ± 0.086	3.90 ± 1.25	158 ± 27.1	7.10 ± 5.22
Interaction Primitives (enBIP)	0.333 ± 0.087	0.342 ± 0.098	2.68 ± 0.687	130 ± 19.6	-11.6 ± 4.81
MPIP Reduction	0.191 ± 0.049	0.384 ± 0.095	**2.26** ± 0.387	**109** ± 13.4	**-28.6** ± 2.00
MPIP Increase	0.650 ± 0.048	0.321 ± 0.096	3.49 ± 0.989	146 ± 22.7	9.03 ± 0.942
MPIP Symmetry	0.511 ± 0.059	3.98 ± 0.090	2.41 ± 0.764	122 ± 26.0	-14.8 ± 8.67

The best (lowest) values are highlighted for each optimized variable.

an increase in LDS and does not significantly alter the sum of muscle forces in any way. The final MPIP cost function, produces a reduction in muscle impulse, however, does so at the expense of LDS. Empirically, it follows that higher muscle forces are necessary for sufficient resilience to external perturbations, which occur cyclically during locomotion with heel-strike and toe-off.However, this control application also serves as a warning that the model predictive controller is not a catchall solution without a robust and effective cost function.

D. Experiment 4: Model Predictive Control on Real-World Prosthesis

Lastly, we implemented multiple MPC control strategies for assisted jumping with MPIP on a real-world robotic prosthetic ankle. A subject was tasked with jumping in different ways to generate an accurate model of joint-reaction loads in reference to the ankle angle. We then examined the effect of different MPIP strategies on the kinematics and biomechanics of the jump and on the interaction between human subject and robotic prosthetic device.

During training, an able-bodied subject performed ten vertical jumps while landing softly and ten jumps while landing hard, for a total of 20 jumps. To augment the sensor data estimates, joint-reaction forces were constructed from the shoe sensors through the Newton–Euler formulation of 3-D inverse dynamics, by means of [53]. In our 3-D inverse dynamics formulation, the four external forces acting on the foot are consolidated to a singe equivalent force, and projected from the center of pressure at the foot through a point just above the subjects center mass (CM), known as the divergence point. Consequently, Newton–Euler equations provide the 3-D forces and moments as reaction forces of each body segment. Of particular note is the use of the Gaussian basis function in this experiment, as opposed to the von Mises basis functions used throughout the rest of the article, resulting in greater independence between the beginning and end of the interaction because they are not directly connected by any single basis function. After training, the subject performed ten trials for each control strategy, with a carbon fiber ankle bypass which allows an able-bodied subject to interact with a prosthesis without utilizing their own foot. We evaluate four control strategies, with results presented in Table V, 1) reactive control from enBIPs, 2) vertical knee reaction force (minimization), 3) vertical knee-reaction force (maximization), and 4) maximization of ankle angle symmetry between left and right lower limbs. In each trial, the time horizon was set from the current phase to the end of the interaction, with no control time horizon.

The optimization strategies for knee force minimization (blue), knee force maximization (red), and ankle angle symmetry (green) are visualized in Fig. 14. In each subplot the ankle angle control parameter is on the top and the calculated knee force reaction on the bottom. Furthermore, the shaded areas of each plot illustrate the difference of the optimized solution and the mean action. Image sequences of the force minimization and maximization solutions are presented in Fig. 13 and match up with the plots in Fig. 14, respectively. Comparing the three strategies by looking first at the control actions, the force reduction strategy (Fig. 14, blue) pushes the toe down throughout the entire flight phase of the jump. Therefore, when the toe touches down the internal compliant mechanism absorbs the force over a longer time reducing the peak force substantially. By contrast, the force increase method (Fig. 14, red) optimizes the control trajectory to pull the toe up as fast as possible such that little to none of the force is absorbed by the internal compliant mechanism and is instead applied directly to the knee joint. This is the expected result, and illustrates the algorithm's ability to exploit the design and construction of compliant prosthetics. To show the wide range of optimization strategies possible with MPIP, we applied an additional optimization

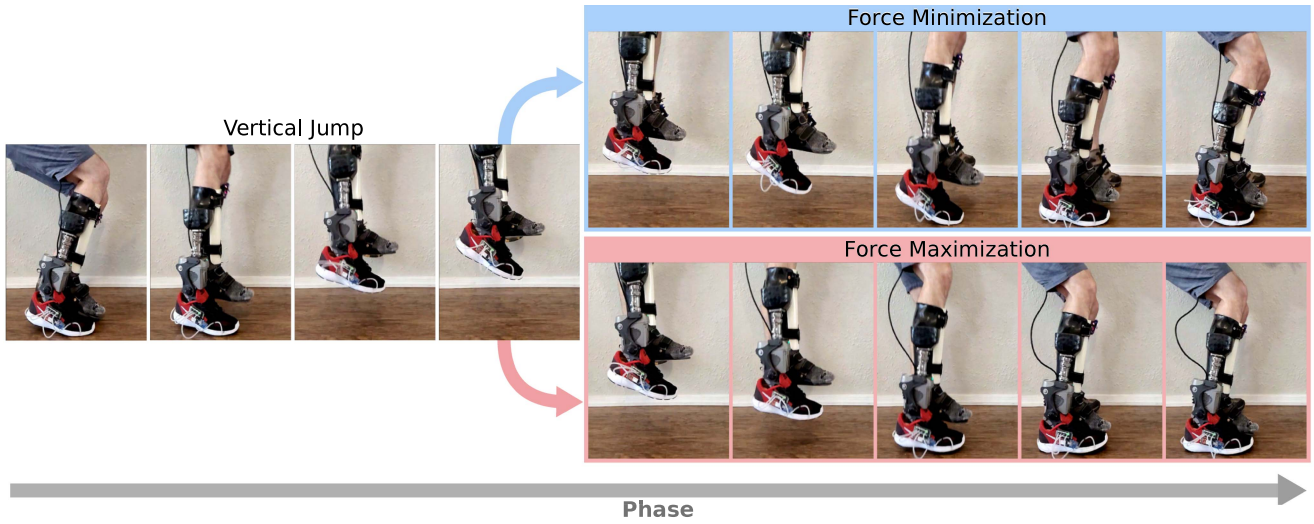

Fig. 13. Depiction of possible variations in the jumping controller. After the initial vertical jump the optimization can yield many control solutions for landing, such as force minimization with the toe down (top) or force maximization with the toe up (bottom).

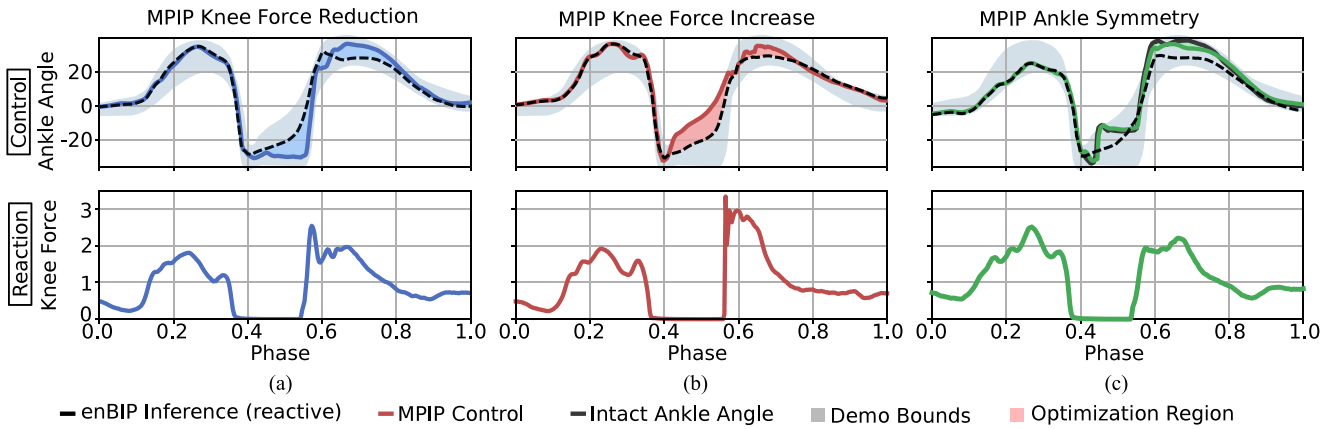

Fig. 14. Three separate MPIP optimization strategies for vertical jumping with a powered prosthesis. In each experiment a subject jumps with the powered prosthesis, while the subject is in the air the MPIP optimizes the reference angle of the prosthetic ankle in attempts to (a) reduce the internal vertical knee force, (b) increase the internal vertical knee force, and (c) enforce symmetry between the prosthesis and intact ankles. For each optimization strategy we compare the inferred ankle angle reference trajectory from enBIP, to the optimized MPIP ankle angle reference trajectory. By actualizing the optimized ankle angle references trajectories with a PID controller, we compare the actual knee forces experienced by the subject with each optimization strategy.

strategy geared toward maximizing the symmetry between the prosthetic and intact ankles (Fig. 14, green). The symmetry-based optimization strategy pushes the toe down somewhat but reduces the vertical knee impulse through mimicking the intact ankle.

Although the knee force reduction (Fig. 14, blue) and symmetry (Fig. 14, green) control trajectories differ significantly during the jump, both methods produce decreased vertical knee reaction forces during landing, showing that multiple mechanisms can be exploited to reduce the risk of osteoarthritis. However, when assessing the knee force it is clear that the knee force maximization strategy (Fig. 14, red) results in a 50% greater impact on the subject's knee. These optimization strategies show that MPIP is capable of incorporating expected future outcomes into the decision-making process.

VIII. Conclusion

In this article, we present a data-driven methodology for learning robot control policies that take ergonomic considerations into account. Specifically, we present an imitation learning strategy for learning MPIPs for human–robot interaction scenarios. Once trained, these policies can be used to generate robot control signals that are both functional and biomechanically safe. An important element of this approach is the ability to perform *predictive biomechanics*: the prediction of future nonobserved biomechanical variables of a human partner from multimodal, low-fidelity sensor readings of the human–robot system. In turn, these predictions are used within an efficient MPC loop to optimize robotic actions, so as to minimize future stress to the human body.

Results of Current Study: We show how IPs, in a general sense, are a useful tool for learning predictive models that bridge the divide between the state of the controlled system (prosthesis), the biomechanical state of the human user, and the required control signals. Our experiments show that we can use these models to produce accurate estimates of human biomechanics over a range of simulated and real-world locomotion tasks. We specifically tested multiple inference models, including ProMP, IP, particleBIP, and enBIP on human subject data to show that the nonlinear BIP-based models outperform all others. At training time, a comprehensive dataset consisting of both system states and high-fidelity biomechanics data was collected. However, once a model was learned, multimodal low-cost and fidelity sensors were sufficient to infer the human biomechanical features. Likewise, learned predictive models can integrate unique and interesting cost functions into the control loop, to meet user-specific ergonomic requirements, that is, minimization of knee loads, muscle forces, or symmetrical constraints. In our experiments we showed, that the resulting MPIP approach successfully reduced vertical knee-reaction forces by approximately 18% in simulated locomotion trials and by approximately 16% in real-world prosthesis-assisted jumping trials.

Safety is a critical aspect of HRI. As such, our framework introduced safe functionality in a number or ways. 1) During operation, our framework provided an estimate of the error covariance, i.e., how certain the framework is about the control outputs generated by the ensemble Kalman filter. In the case of a large error, a default controller can be used or retraining can be requested. 2) We incorporated ergonomic and safety-critical features, such as joint forces to explicitly lower the risks of secondary musculoskeletal conditions, such as osteoarthritis. Our framework is extensible and allows for other features and cost functions to be used should new insights in biology or biomechanics warrant that. However, as shown in Table IV, optimizing for one feature might effect others, such as stability. In such a case, a multiobjective solver may be needed. 3) To ensure that the control output is robust and stable in the presence of perturbations, we utilized a constrained MPC framework to restrict the control actions to the bounds of the demonstrations. Fig. 9 indicates that adding constrained optimization that minimizes the change in control output reduces the effects of a perturbation by approximately $20\times$.

Limitations and Future Work: In this article, we restricted the data collection process to focus on isolated behaviors, which is an unrealistic expectation for real-world locomotion scenarios, where users are expected to navigate dynamic environments, including varying ground types, slopes, and stairs. Furthermore, because subjects, especially those with amputations, can vary dramatically from generic human locomotion models, we cannot say with certainty that our models will provide safe or healthy control for all groups. We, therefore, intend to explore calibration of model parameters during deployment with active learning methods to compensate for uncertainties due to abnormal gaits. Further tailoring of predictive models to individual users would both increase inference/control capabilities and give critical insights into optimal fitment of prosthetic devices.

To attain longer, more comprehensive interactions, we will further investigate Bayesian formulations which can incorporate a diverse set of interactions in a single dexterous model. As nonlinear features will become more important as the models we utilize become more complex, we expect to incorporate more powerful modeling formulations, such as neural networks to better handle nonlinear system dynamics. Finally, because safety is our highest priority, we expect to incorporate novel safety features, such as barrier functions [54] or more complex planning algorithms [55].

REFERENCES

[1] D. C. Morgenroth, A. C. Gellhorn, and P. Suri, "Osteoarthritis in the disabled population: A mechanical perspective," *PM R*, vol. 4, no. 5, pp. S20–S27, 2012.

[2] G. Jacucci, A. Spagnolli, J. Freeman, and L. Gamberini, "Symbiotic interaction: A critical definition and comparison to other human-computer paradigms," in *Proc. Int. Workshop Symbiotic Interact.*, 2015, pp. 3–20.

[3] A. Paraschos et al., "Probabilistic movement primitives," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 1–6.

[4] H. B. Amor, G. Neumann, S. Kamthe, O. Kroemer, and J. Peters, "Interaction primitives for human-robot cooperation tasks," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2014, pp. 2831–2837.

[5] G. Maeda, M. Ewerton, R. Lioutikov, H. Ben Amor, J. Peters, and G. Neumann, "Learning interaction for collaborative tasks with probabilistic movement primitives," in *Proc. IEEE-RAS Int. Conf. Humanoid Robots*, 2014, pp. 527–534.

[6] J. Campbell and H. B. Amor, "Bayesian interaction primitives: A SLAM approach to human-robot interaction," in *Proc. Conf. Robot Learn.*, 2017, pp. 379–387.

[7] G. Clark, J. Campbell, and H. B. Amor, "Learning predictive models for ergonomic control of prosthetic devices," in *Proc. Mach. Learn. Res.*, 2020, pp. 1–11.

[8] R. Versluys, P. Beyl, M. Van Damme, A. Desomer, R. Van Ham, and D. Lefeber, "Prosthetic feet: State-of-the-art review and the importance of mimicking human ankle-foot biomechanics," *Disabil. Rehabil.: Assistive Technol.*, vol. 4, no. 2, pp. 65–75, 2009.

[9] S. K. Au, J. Weber, and H. Herr, "Powered ankle-foot prosthesis improves walking metabolic economy," *IEEE Trans. Robot.*, vol. 25, no. 1, pp. 51–66, Feb. 2009.

[10] F. Sup, H. A. Varol, J. Mitchell, T. J. Withrow, and M. Goldfarb, "Preliminary evaluations of a self-contained anthropomorphic transfemoral prosthesis," *IEEE/ASME Trans. Mechatronics*, vol. 14, no. 6, pp. 667–676, Dec. 2009.

[11] F. Sup, H. A. Varol, and M. Goldfarb, "Upslope walking with a powered knee and ankle prosthesis: Initial results with an amputee subject," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 19, no. 1, pp. 71–78, Feb. 2011.

[12] M. R. Tucker et al., "Control strategies for active lower extremity prosthetics and orthotics: A review," *J. Neuroengineering Rehabil.*, vol. 12, no. 1, 2015, Art. no. 1.

[13] C. Gao, R. Gehlhar, A. D. Ames, S.-C. Liu, and T. Delbruck, "Recurrent neural network control of a hybrid dynamic transfemoral prosthesis with edgeDRNN accelerator," *IEEE Int. Conf. Robot. Automat. (ICRA)*, 2020.

[14] B.-Y. Su, J. Wang, S.-Q. Liu, M. Sheng, J. Jiang, and K. Xiang, "A CNN-based method for intent recognition using inertial measurement units and intelligent lower limb prosthesis," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 27, no. 5, pp. 1032–1042, May 2019.

[15] A. S. Alharthi, S. U. Yunas, and K. B. Ozanyan, "Deep learning for monitoring of human gait: A review," *IEEE Sensors J.*, vol. 19, no. 21, pp. 9575–9591, Nov. 2019.

[16] G. Khademi and D. Simon, "Convolutional neural networks for environmentally aware locomotion mode recognition of lower-limb amputees," in *Proc. Dyn. Syst. Control Conf.*, vol. 59148, 2019, Art. no. V001T07A005.

[17] T. C. Noel and B. R. Snider, "Utilizing deep neural networks for brain-computer interface-based prosthesis control," *Fac. Publications–Dept. Elect. Eng. Comput. Sci.*, 2019, Art. no. 28.

[18] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, "Reluplex: An efficient SMT solver for verifying deep neural networks," in *Proc. Int. Conf. Comput. Aided Verification*, 2017, pp. 97–117.

[19] A. Dokhanchi, H. B. Amor, J. V. Deshmukh, and G. Fainekos, "Evaluating perception systems for autonomous vehicles using quality temporal logic," in *Proc. Int. Conf. Runtime Verification*, 2018, pp. 409–416.
[20] J. Su, D. V. Vargas, and K. Sakurai, "One pixel attack for fooling deep neural networks," *IEEE Trans. Evol. Comput.*, vol. 23, no. 5, pp. 828–841, Oct. 2019.
[21] N. Thatte, N. Srinivasan, and H. Geyer, "Real-time reactive trip avoidance for powered transfemoral prostheses," in *Proc. Robot. Syst. Sci.*, 2019, pp. 1–9.
[22] D. J. Lizotte, T. Wang, M. H. Bowling, and D. Schuurmans, "Automatic gait optimization with Gaussian process regression," in *Proc. Int. Joint Conf. Artif. Intell.*, vol. 7, 2007, pp. 944–949.
[23] R. Calandra, A. Seyfarth, J. Peters, and M. P. Deisenroth, "Bayesian optimization for learning gaits under uncertainty," *Ann. Math. Artif. Intell.*, vol. 76, no. 1, pp. 5–23, 2016.
[24] R. Antonova, A. Rai, and C. G. Atkeson, "Sample efficient optimization for learning controllers for bipedal locomotion," in *Proc. IEEE-RAS 16th Int. Conf. Humanoid Robots*, 2016, pp. 22–28.
[25] N. Dhir, H. Dallali, E. M. Ficanha, G. A. Ribeiro, and M. Rastgaar, "Locomotion envelopes for adaptive control of powered ankle prostheses," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 1488–1495.
[26] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, "Dynamical movement primitives: Learning attractor models for motor behaviors," *Neural Comput.*, vol. 25, no. 2, pp. 328–373, 2013.
[27] 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 *Proc. IEEE Int. Conf. Robot. Automat.*, 2015, pp. 1535–1542.
[28] 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 *Proc. IEEE-RAS 17th Int. Conf. Humanoid Robot.*, 2017, pp. 491–496.
[29] J. Campbell, A. Hitzmann, S. Stepputtis, S. Ikemoto, K. Hosoda, and H. B. Amor, "Learning interactive behaviors for musculoskeletal robots using Bayesian interaction primitives," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, 2019, pp. 5071–5078.
[30] R. D. Gregg, E. J. Rouse, L. J. Hargrove, and J. W. Sensinger, "Evidence for a time-invariant phase variable in human ankle control," *PLoS One*, vol. 9, no. 2, 2014, Art. no. e89163.
[31] S. Srinivasan, I. Raptis, and E. R. Westervelt, "Low-dimensional sagittal plane model of normal human walking," *J. Biomechanical Eng.*, vol. 130, no. 5, 2008, Art. no. 051017.
[32] M. Castelán and G. Arechavaleta, "Approximating the reachable space of human walking paths: A low dimensional linear approach," in *Proc. 9th IEEE-RAS Int. Conf. Humanoid Robots*, 2009, pp. 81–86.
[33] S. Schaal, "Dynamic movement primitives—A framework for motor control in humans and humanoid robotics," in *Adaptive Motion of Animals and Machines*. Berlin, Germany: Springer, 2006, pp. 261–280.
[34] G. Clark, J. Campbell, S. M. R. Sorkhabadi, W. Zhang, and H. B. Amor, "Predictive modeling of periodic behavior for human-robot symbiotic walking," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 7599–7605.
[35] J. Campbell, S. Stepputtis, and H. B. Amor, "Probabilistic multimodal modeling for human-robot interaction tasks," in *Proc. Robot.: Sci. Syst.*, FreiburgimBreisgau, Germany, 2019, pp. 1–9.
[36] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, vol. 50, no. 12, pp. 2967–2986, 2014.
[37] S. J. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control Eng. Pract.*, vol. 11, no. 7, pp. 733–764, 2003.
[38] S. L. Delp et al., "Opensim: Open-source software to create and analyze dynamic simulations of movement," *IEEE Trans. Biomed. Eng.*, vol. 54, no. 11, pp. 1940–1950, Nov. 2007.
[39] Y.-F. Zhang, Q.-F. Zhang, and R.-H. Yu, "Markov property of Markov chains and its test," in *Proc. Int. Conf. Mach. Learn. Cybern.*, vol. 4, 2010, pp. 1864–1867.
[40] W. Sacher and P. Bartello, "Sampling errors in ensemble Kalman filtering. Part I: Theory," *Monthly Weather Rev.*, vol. 136, no. 8, pp. 3035–3049, 2008.
[41] R. Piessens, E. de Doncker-Kapenga, C. W. Überhuber, and D. K. Kahaner, *Quadpack: A Subroutine Package for Automatic Integration*, vol. 1. Berlin, Germany: Springer, 2012.
[42] D. Kraft, "A software package for sequential quadratic programming," Deutsche Forschungs- und Versuchsanstalt fur Luft- und Raumfahrt, Cologne, Germany, DFVLR-FB 88-28, 1988.
[43] A. Bemporad and D. M. de la Peña, "Multiobjective model predictive control," *Automatica*, vol. 45, no. 12, pp. 2823–2830, 2009.
[44] P. A. Ortega et al., "Shaking the foundations: Delusions in sequence models for interaction and control," 2021, *arXiv:2110.10819*.
[45] J. Hicks et al., "Simulation with Opensim-best practices," Nat. Center Simul. Rehabil. Res., Stanford Univ., Stanford, CA, USA, 2015.
[46] Vicon, Oxford, U.K., "Plug-in gait model details." [Online]. Available: https://www.vicon.com/downloads/documentation/plug-in-gait-model-details
[47] D. D. D'Lima, B. J. Fregly, and C. W. Colwell, "Implantable sensor technology: Measuring bone and joint biomechanics of daily life in vivo," *Arthritis Res. Ther.*, vol. 15, no. 1, 2013, Art. no. 203.
[48] Z. I. Nejad et al., "The capacity of generic musculoskeletal simulations to predict knee joint loading using the CAMS-knee datasets," *Ann. Biomed. Eng.*, vol. 48, no. 4, pp. 1430–1440, 2020.
[49] J. Ward, K. Schroeder, D. Vehon, R. Holgate, A. Boehler, and M. Grimmer, "A rugged microprocessor controlled ankle-foot prosthesis for running," in *Proc. 39th Annu. Meeting Amer. Soc. Biomech.*, 2015.
[50] B. M. Potvin, M. S. Shourijeh, K. B. Smale, and D. L. Benoit, "A practical solution to reduce soft tissue artifact error at the knee using adaptive kinematic constraints," *J. Biomech.*, vol. 62, pp. 124–131, 2017.
[51] Ł. Kidzinski et al., "Artificial intelligence for prosthetics: Challenge solutions," in *Competition: From Machine Learning to Intelligent Conversations*. Cham, Switzerland: Springer, 2019, ch. 69.
[52] A. Ekizos, A. Santuz, A. Schroll, and A. Arampatzis, "The maximum Lyapunov exponent during walking and running: Reliability assessment of different marker-sets," *Front. Physiol.*, vol. 9, 2018, Art. no. 1101.
[53] G. E. Robertson, G. E. Caldwell, J. Hamill, G. Kamen, and S. Whittlesey, *Research Methods in Biomechanics*. Champaign, IL, USA: Human Kinetics, 2013.
[54] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. 18th Eur.Control Conf.*, 2019, pp. 3420–3431.
[55] T. Schoels, P. Rutquist, L. Palmieri, A. Zanelli, K. O. Arras, and M. Diehl, "CIAO: MPC-based safe motion planning in predictable dynamic environments," *IFAC-PapersOnLine*, vol. 3, no. 2, pp. 6555–6562, 2020.

Geoffrey Clark received the master's degree in electrical engineering focusing on machine learning for ergonomic human–robot interaction in 2019 from Arizona State University, Tempe, AZ, USA, where he is currently working toward the Ph.D. degree in electrical engineering.

He is a Ph.D. Research Associate with the Interactive Robotics Laboratory, Arizona State University, Tempe, AZ, USA. In addition to his Ph.D. research, he has worked for 10 years on the research and development of robotic prosthetic devices and exoskeletons. His research interest include the intersection of machine learning and controls theory.

Heni Ben Amor received the Ph.D. degree in computer science from the Technical University Freiberg, Freiberg, Germany, in 2010, focusing on artificial intelligence and machine learning.

He is an Associate Professor of Robotics with Arizona State University, Tempe, AZ, USA. He is the Director of the ASU Interactive Robotics Laboratory. He was a Research Scientist with Georgia Tech, Atlanta, GA, USA, a Postdoctoral Researcher with the Technical University Darmstadt, Darmstadt, Germany, and a Visiting Research Scientist with the Intelligent Robotics Lab, University of Osaka, Osaka, Japan. His research interests include machine learning, robotics, human–robot interaction, and virtual reality.

Dr. Amor was the recipient of the NSF CAREER Award, the Fulton Outstanding Assistant Professor Award, as well as the Daimler-and-Benz Fellowship.