

Learning Predictive Models for Ergonomic Control of Prosthetic Devices

Geoffrey Clark

Interactive Robotics Lab
Arizona State University
gmclark1@asu.edu

Joseph Campbell

Interactive Robotics Lab
Arizona State University
jacampb1@asu.edu

Heni Ben Amor

Interactive Robotics Lab
Arizona State University
hbenamor@asu.edu

Abstract: We present Model-Predictive Interaction Primitives – a robot learning framework for assistive motion in human-machine collaboration tasks which explicitly accounts for biomechanical impact on the human musculoskeletal system. First, we extend Interaction Primitives to enable *predictive biomechanics*: the prediction of future biomechanical states of a human partner conditioned on current observations and intended robot control signals. In turn, we leverage this capability within a model-predictive control strategy to identify the future ergonomic and biomechanical ramifications of potential robot actions. Optimal control trajectories are selected so as to minimize future physical impact on the human musculoskeletal system. We empirically demonstrate that our approach minimizes knee or muscle forces via generated control actions selected according to biomechanical cost functions. Experiments are performed in synthetic and real-world experiments involving powered prosthetic devices.

1 Introduction

Intuitive and safe mobility is a necessary part of everyday life, and a critical challenge to those with musculoskeletal conditions and lower-extremity amputations [1, 2]. According to the United States Bone and Joint Initiative [3], a record number of 126.6 million Americans are affected by a musculoskeletal condition. Out of these, about 2 million are living with a limb loss – a number that is likely to double by the year 2050 [4]. To date, amputees who are longtime prosthesis users are at a high risk for knee osteoarthritis (OA) [5] in the *intact* limb – even when modern microprocessor-controlled prostheses are used. By favoring the intact limb, users may develop asymmetric gaits thereby generating higher knee loads in the healthy leg. Over time, the increased stresses accumulate causing secondary musculoskeletal conditions, further compounding any disability.

Prostheses, orthoses, exoskeletons, and other assistive robotics devices [6] have the potential to substantially improve the quality of life of persons with amputations and other musculoskeletal conditions. However, to ensure the efficacy and long-term safety of such devices, it is critical to develop control strategies that respond to the ergonomic needs of each specific user. More specifically, forces, torques or moments acting on the human body need to be explicitly taken into account during action generation in order to achieve user-specific, ergonomic, and biomechanically-safe control.

In this paper, we identify assisted mobility as a problem of symbiotic interaction between a human and assistive robot. In such a scenario, it is imperative that the assistive device learns models that



Figure 1: Top: Probabilistic modeling predicts upcoming motion and joint stresses. Circles indicate high-impact joints. Bottom: Control signals for our robotic ankle prosthesis are generated in order to minimize joint stress.

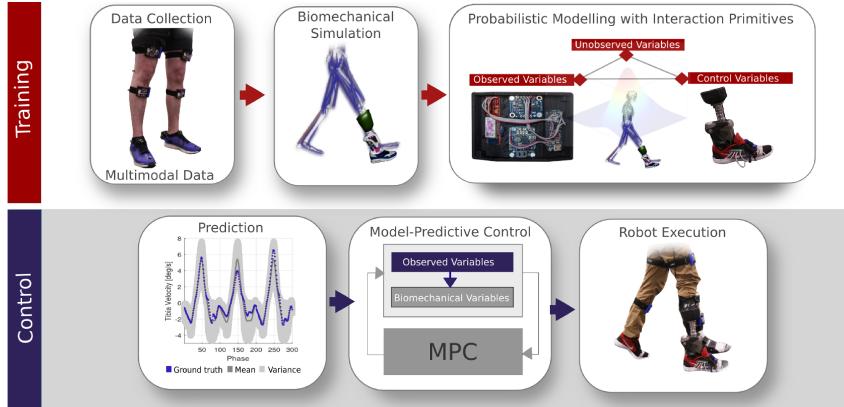


Figure 2: Overview diagram of our approach depicting the training and deployment phases.

can anticipate the future evolution of the interaction and adapt to their specific human partner. To this end, we discuss a unified approach to (a) human state estimation and prediction, (b) inference of unobserved biomechanical variables, and (c) model-predictive control in support of the above ergonomic objectives and constraints. Building upon existing literature and prior work on probabilistic human motion modelling [7, 8], we present the application of interaction primitives for learning predictive models of biomechanics and control in a lower-limb prosthesis. An important feature of this method is the ability to reduce the cooperative system (human and robot) to a latent space using basis function decomposition, in which temporal features are probabilistically related to one another through their covariance. We leverage nonlinear Bayesian filtering, to forecast the future evolution of a multi-modal walking or jumping time-series. At runtime, different types of predictions about the future state of the human user, e.g., the forces acting on different body parts, can be made. Optimal actions are, in turn, chosen in a model predictive control loop so as to meet a specified objective, i.e., minimize knee forces over a finite time horizon. To demonstrate prediction and control abilities of our method, we present a variety of simulated and real-world experiments that involve the real-time, ergonomic control of a commercial lower-limb prosthesis.

2 Related Work

In recent years, we have seen substantial progress in the domain of wearable robotic devices such as prosthetics, orthoses, and exoskeletons [6]. However, user-adaptation and ergonomics still remains a major challenge [9]. Machine learning techniques are well positioned to overcome this challenge by enabling new forms of user-adaptation and customization. The work in [10] builds upon the recent surge in deep learning approaches and describes end-to-end learning of continuous control with a recurrent neural network. Neural network approaches have a remarkable predictive power and are particularly popular for anticipating (discrete) events and walking conditions such as locomotion mode detection [11], human intention inference [12], or environment recognition [13]. Despite recent developments in certification, monitoring, and verification techniques [14, 15], there are still major barriers for the application of deep networks in safety-critical systems such as prosthetics. Chief among these barriers is the difficulty in providing safety certifications – a challenge exacerbated by opportunities for adversarial attacks [16].

Alternatively, a number of approaches for incorporating machine learning into prosthetic control are based on probabilistic and Bayesian formulations. For example, the work in [17] leverages Gaussian processes (GPs) to learn a predictive model for hip angles and hip heights during walking. Gaussian processes have been particularly popular in locomotion research [18, 19, 20, 21], due to their ability to learn from small sample sizes, their ability to produce posterior distributions with uncertainty information, and the relative ease with which they can be integrated into existing (often Kalman-based) filtering and control frameworks.

Another probabilistic framework termed Interaction Primitives (IPs) [22, 23], combines dynamical movement primitives [24] with a Bayesian inference step, to produce an imitation learning approach specialized for modeling the interaction of multiple agents. The IP combination of a.) parameterizing models of two interacting agents via basis function decomposition, followed by b.) jointly

modeling the trajectory distributions (utilizing the basis covariance), and c.) recursive Bayesian filtering (Kalman or Particle filtering); is particularly well-suited for predictive modeling of human motion. Previous applications of IPs for human-robot interaction include, for example, collaborative assembly [25], throwing and catching [7], or electromyography based object hand over [26], show that IPs can infer modeled parameters including robot joint angles given a set of observed partial trajectories. A first attempt at using IPs for robotic prostheses is presented in [8], where control signals were generated in a purely reactive fashion by iteratively conditioning on the sensor readings. The approach is particularly powerful when dealing with noisy sensing and performance gracefully degrades when one or more sensors unexpectedly break down; however it does not take into account the future ramifications of a selected action. Actions of a prosthesis that are locally optimal, may potentially lead to trips, high internal stresses, or other dangerous situations.

In this paper, we build upon previous work on IP, but make critical extensions to allow for planning and biomechanically-safe, predictive control. Going beyond reactive control [25, 8, 7], we show how model-predictive control (MPC) [27, 28] can be implemented with IPs as the underlying data-driven model. The resulting approach has powerful inference capabilities that allow for *predictive biomechanics*: the prediction and control of future, unobserved biomechanical states of the human-robot system from current sensor readings. Furthermore, the approach plans robot control trajectories that minimize biomechanical stress on the human body, and we show that the underlying cost functions specifying biomechanical safety can be modified to realize different ergonomic goals.

3 MPIP: Model-Predictive Interaction Primitives

Our approach, illustrated in Fig. 2, consists of two phases and can be seen as an imitation learning process [29]. In the training phase, we record training data from the process we are interested in modelling, i.e., human locomotion. This step typically involves recording data from a set of input modalities, e.g., accelerometers, force plates, motion capture etc. The raw data set may not yet include variables that are of interest but are not immediately measurable, such as knee loads, muscle forces and the control signals for the prosthesis. These additional variables are acquired through simulations and incorporated into the data set. In turn, the augmented data set is used for training an Interaction Primitive – a probabilistic model of the underlying system dynamics [22, 25, 23]. After training, the real-time controller running on the assistive device uses the observed sensor readings as input to predict future sensor values, biomechanical variables, and reactive control signals. Within a model-predictive control loop, predictions are made on how different robot actions will affect the evolution of these values. Via MPC we optimize an action that best satisfies a given cost function, e.g., the total muscle work, the knee forces or other criteria essential to healthy motion. IPs can effectively be cast as a linear model in which all nonlinearity is solely encoded within a set of basis functions. In the following sections, we will describe each step of our methodology in more detail.

3.1 Data Collection and Augmentation

The presented method is inherently an imitation learning technique with planning capabilities. As such, it must be trained with example demonstrations of the behavior to be learned. For data collection, a human user is equipped with a diverse set of sensors. These may include accelerometers, gyros, pressure-sensitive shoes, and other sensors that provide information about the state of the user. The necessary sensor data representing the subject’s state is collected from N individual demonstrations of an action, $[\mathbf{y}_1, \dots, \mathbf{y}_{T_n}] \in \mathbb{R}^{D_y \times T_n}$, such that D_y is the number of sensor signals and T_n is the number of time steps in the n^{th} demonstration. In order to cover a variety of settings, demonstrations should ideally cover multiple gait cycles, speeds and ground inclinations. However, these values do not yield a sufficiently detailed picture of the walking process. In particular, non-observable biomechanical variables are missing. Examples are joint forces, muscle effort and mechanical work. These can be calculated via analytical methods for which a number of musculoskeletal simulation libraries are available, e.g., OpenSim [30]. In our specific scenario, we use pressure-sensing shoes to measure the ground reaction forces at four points and, in turn, estimate joint reaction forces through the Newton-Euler formulation of 3D inverse dynamics as described in [31]. External forces acting on the foot are reduced to a single equivalent force which represents the sum of all effective forces acting on the body. Consequently, we project the one-dimensional foot force from the smart shoe into a 3D force by assuming that the force originates at the center of pressure (COP) in the shoe and traverses a straight line through the divergence point located just

above the subject’s center mass (CM). Newton-Euler equations are then used to calculate the force and moment at the proximal end of each segment up to the CM.

We follow best practices to ensure sufficiently accurate approximation of dynamics. In particular, we calibrate with accurate physiological measurements from the user (see supplemental material), incorporate sanity checks for expected range of values, and use the best practices in Hicks et al. [32]. Most importantly, we cross-compare the results from OpenSim with results generated from motion capture and force plate technology (Accugait, AMTI, Watertown, MA, USA) to ensure a match between the results ($\leq 5\%$ error). While such force plate technology is the de facto standard in force calculation and used for clinical gait analysis, it is static and cannot be built into the shoes. The work in Nejad et al. [33] noticed “that current generic musculoskeletal modelling techniques are able to reproduce the in vivo conditions (i.e. ground truth) measured during walking” which is one of the main targets of the paper. Please also note, that while not standard today, technology exists to measure such forces in vivo [34]. Our approach is agnostic to the type of sensor and can also be used with these in vivo sensors instead of simulation data.

The final step in data augmentation is to provide robot control signals. These signals can either be recorded from within the demonstration, i.e., setting the robot’s target angle to the human ankle ankle; or can be developed through a modeling process such as a complex impedance control system. The objective of this step is to provide an approximate control signal which is later refined through model-predictive control. Using the above methodologies, trajectories of latent variables (biomechanical variables and control signals) $[\mathbf{m}_1, \dots, \mathbf{m}_{T_n}] \in \mathbb{R}^{D_m \times T_n}$ are incorporated into the dataset where D_m is the number of latent variables. The full data matrix for a demonstration then consists of $\mathbf{Y}_{1:t} = [\mathbf{y}_t, \mathbf{m}_t] \in \mathbb{R}^D$, where $D = D_y + D_m$.

3.2 Preliminaries: Latent Space Formulation and Inference

Latent Space Formulation: We approach imitation learning from a probabilistic perspective, with the high-level goal of estimating a distribution over future observed and unobserved state variables $\hat{\mathbf{Y}}_{t+1:T}$ given previously observed states $\mathbf{y}_{1:t}$ and a prior set of N demonstrations $\mathbf{Y}_{1:T_1}^1, \dots, \mathbf{Y}_{1:T_N}^N$:

$$p(\hat{\mathbf{Y}}_{t+1:T} | \mathbf{y}_{1:t}, \mathbf{Y}_{1:T_1}^1, \dots, \mathbf{Y}_{1:T_N}^N). \quad (1)$$

This formulation, however, would require a nonlinear transition function for observed and unobserved variables, as well as the definition of a generative distribution over them, which is a non-trivial challenge. An existing strategy [23] is to transform the state representation into a time-invariant latent space via basis function decomposition. In this method, basis functions capture the evolution of the state variables over time while the relationship between basis functions is modeled with a joint probability distribution. We define the basis decomposition as a linear combination of weighted basis functions which approximates each state dimension d as $\mathbf{Y}_t^d = \Phi_{\phi(t)}^T \mathbf{w}^d + \epsilon_y$ where $\Phi_{\phi(t)} \in \mathbb{R}^{B^d}$

is a vector of B^d basis functions, $\mathbf{w}^d \in \mathbb{R}^{B^d}$ is a vector of corresponding weights and acts as our latent state variable, $\phi(t) \in \mathbb{R}$ such that $0 \leq \phi(t) \leq 1$ is a relative measure of time known as *phase* (hereafter denoted as ϕ), and ϵ_y is the approximation error. The choice of basis function depends on the data and in this work we opt to utilize a Gaussian function, however both von Mises as well as polynomial basis function are viable given the domain. The full latent state representation is given by the concatenation of all latent state vectors, $\mathbf{w} = [\mathbf{w}^0, \dots, \mathbf{w}^D] \in \mathbb{R}^B$ where $B = \sum_d B^d$, together with the phase ϕ and phase velocity $\dot{\phi}$ such that $\mathbf{s} = [\phi, \dot{\phi}, \mathbf{w}]$ [23]. Since \mathbf{w} is time-invariant, estimating this quantity is equivalent to estimating all past, present, and future state values of \mathbf{Y} ; thus we can reformulate Eq. 1 as $p(\mathbf{s}_t | \mathbf{y}_{1:t}, \mathbf{s}_0)$. After application of Bayes’ rule, this yields the following posterior distribution:

$$p(\mathbf{s}_t | \mathbf{y}_{1:t}, \mathbf{s}_0) \propto p(\mathbf{y}_t | \mathbf{s}_t) p(\mathbf{s}_t | \mathbf{y}_{1:t-1}, \mathbf{s}_0). \quad (2)$$

Inference: We employ an approximate inference scheme found in ensemble Bayesian Interaction Primitives [35] solve the posterior distribution in Eq. 2. Assuming that a) the Markov property holds and b) errors in the state estimate are normally distributed (that is, we ignore higher-order statistical moments), the posterior defined by Eq. 2 can be solved with a recursive two-step Bayesian filter in which the distributions are approximated with an ensemble of state members. We define an ensemble \mathbf{X} consisting of E members such that $\mathbf{X} = [\mathbf{x}^1, \dots, \mathbf{x}^E]$. The initial ensemble, \mathbf{X}_0 , is sampled directly from the prior: $\mathbf{x}_0 \sim p(\mathbf{s}_0)$ for all $\mathbf{x}_0 \in \mathbf{X}_0$. In practice, since we do not know

$p(\mathbf{s}_0)$ it is common to set $E = N$ and utilize the demonstration latent states directly as samples. The first step approximates $p(\mathbf{s}_t|\mathbf{y}_{1:t-1}, \mathbf{s}_0)$ by propagating each ensemble member forward in time with

$$\mathbf{x}_{t|t-1}^j = g(\mathbf{x}_{t-1|t-1}^j) + \epsilon_x, \quad 1 \leq j \leq E, \quad (3)$$

where $g(\cdot)$ is a constant-velocity transition operator and ϵ_x is a noise perturbation. The second step updates the members with the observation, leveraging the nonlinear observation operator $h(\cdot)$:

$$\begin{aligned} \mathbf{H}_t \mathbf{X}_{t|t-1} &= \left[h(\mathbf{x}_{t|t-1}^1), \dots, h(\mathbf{x}_{t|t-1}^E) \right]^\top, \\ \mathbf{H}_t \mathbf{A}_t &= \mathbf{H}_t \mathbf{X}_{t|t-1} - \left[\frac{1}{E} \sum_{j=1}^E h(\mathbf{x}_{t|t-1}^j), \dots, \frac{1}{E} \sum_{j=1}^E h(\mathbf{x}_{t|t-1}^j) \right]. \end{aligned} \quad (4)$$

As a result, no linearization error is introduced as would be the case in analytical formulations such as the EKF. The innovation covariance can now be found with

$$\mathbf{S}_t = \frac{1}{E-1} (\mathbf{H}_t \mathbf{A}_t) (\mathbf{H}_t \mathbf{A}_t)^\top + \mathbf{R}, \quad (5)$$

where \mathbf{R} is the observation noise matrix. This is then used to compute the Kalman gain as

$$\mathbf{A}_t = \mathbf{X}_{t|t-1} - \frac{1}{E} \sum_{j=1}^E \mathbf{x}_{t|t-1}^j, \quad \mathbf{K}_t = \frac{1}{E-1} \mathbf{A}_t (\mathbf{H}_t \mathbf{A}_t)^\top \mathbf{S}_t^{-1}. \quad (6)$$

Finally, the ensemble is updated to incorporate the new measurement perturbed by stochastic noise:

$$\tilde{\mathbf{y}}_t = [\mathbf{y}_t + \epsilon_y^1, \dots, \mathbf{y}_t + \epsilon_y^E], \quad \mathbf{X}_{t|t} = \mathbf{X}_{t|t-1} + \mathbf{K}_t (\tilde{\mathbf{y}}_t - \mathbf{H}_t \mathbf{X}_{t|t-1}). \quad (7)$$

It is important to note that during inference observations can be partial. In our case, we leverage this ability to generate a posterior over biomechanical variables despite the fact that they cannot be immediately observed. This follows from the data augmentation strategy employed in training.

3.3 Model-Predictive Control with Interaction Primitives

Given the trained probabilistic model which can predict the temporal evolution of the system, we can construct an efficient model predictive control (MPC) framework. MPC [28] computes an open-loop optimal control strategy over a given time horizon. The loop is, then, closed by instantiating the first control value and periodically recomputing the control strategy from an updated state. As opposed to the traditional state-based sampling method, which predicts the time horizon trajectory by iteratively calling the state transition function, our probabilistic model inherently contains the future state prediction for observed as well as latent signals. We start with a cost function formulation to find the optimal control strategy \mathbf{u}^* :

$$J(\mathbf{u}^*) = \int_t^{H_x} \|\mathbf{x}^u\|^2 dt + \int_t^{H_u} \|\Delta \mathbf{u}\|^2 dt, \quad (8)$$

where the cost function J takes the integral of a variable-of-interest time vector \mathbf{x}^u with respect to the control output \mathbf{u} , as well as the change in the control output vector $\Delta \mathbf{u}^* = \mathbf{u} - \hat{\mathbf{u}}$, shown as the difference between the control signal to be optimized \mathbf{u} and the estimated control signal from the IP process $\hat{\mathbf{u}}$. Moreover H_x and H_u indicate the signal and control horizons respectively. To make use of the probabilistic model, we first transform the cost function into the phase domain by integrating from the current phase to the time horizons in phase $H_{x\phi}$ and $H_{u\phi}$. The benefit of utilizing a basis function model is that, once in the phase domain, we can easily extract the weight vector \mathbf{w} out of the integrals, since we represent the full trajectory $\tilde{\mathbf{Y}}_{1:P}^d = [\Phi^\top \hat{\mathbf{w}}_t^d]$ as the inner product of basis vector and the weight vector, resulting in:

$$J(\mathbf{u}^*) = \int_{\phi}^{H_{x\phi}} \|\Phi^\top\| d\phi \|\mathbf{w}_{x^u}\|^2 + \int_{\phi}^{H_{u\phi}} \|\Phi^\top\| d\phi \|\Delta \mathbf{w}_u\|^2, \quad (9)$$

The weight vector for the variable-of-interest (e.g. knee forces, muscle forces) is simply the possible control strategy times the learned prior model, namely the covariance between the weight vectors of the control and variable-of-interest $\mathbf{w}_{x^u} = \text{Cov}(\mathbf{w}_x, \mathbf{w}_x) \cdot \mathbf{w}_{u^*}$

For computational speed-ups, the analytical integral $\Psi_{\phi:H}$ can be pre-calculated to further reduce the number of computations. In the case of the Gaussian basis function, integration reduces to:

$$\Phi(\phi_{0:1}) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(\frac{\phi-\mu}{2\sigma})^2} \quad (10)$$

$$\Psi_{\phi:H} = \int_{\phi}^H \|\Phi(\phi_{0:1})\|^2 = \sqrt{\frac{2}{\pi}} e^{\frac{\mu-\phi}{2\sigma}} \quad (11)$$

resulting in a simple and easy to pre-compute closed-form integral that does not require stochastic sampling. Building on this result, the formulation of the optimization problem for generating controls thus becomes:

$$\mathbf{u}^* = \operatorname{argmin}(\Psi_{\phi:H_{x_\phi}}^\top \|\mathbf{w}_{x^u}\|^2 + \Psi_{\phi:H_{u_\phi}}^\top \|\Delta\mathbf{w}_u\|^2). \quad (12)$$

Solutions to the above optimization problem can be found using modern constrained optimization solvers such as Sequential Least Squares Programming (SLSQP) [36]. We impose constraints on the controller by using the min and max values of each basis function from the training demonstrations. Alternatively, one can also use a multiobjective MPC solver to compute a Pareto-optimal solution, see Bemporad and de la Peña [37].

4 Experiments

In this section, we present a set of experiments that highlight the capabilities of the introduced method. We our method to the target domain of assisted mobility in both simulation and real-world.

4.1 Experiment 1: Simulation

In the first experiment, we apply our MPIP method to the problem of prosthesis control for assisted mobility. To this end, we build upon the NeurIPS “AI for Prosthetics” competition in 2018. The competition featured a passive prosthesis (not controlled) and focused on controlling the virtual human through muscle activations. The problem that is at the center of our paper, however, is the challenge of actively controlling a prosthesis so as to best adapt to a human user (not controlled). Accordingly, we modified the simulation environment by adding a controllable degree-of-freedom for powered actuation at the ankle of the prosthesis. The NeurIPS competition simulator uses OpenSim in order to provide a physiologically-accurate musculoskeletal model, along with a physics-based environment, and various data sets of gait kinematics. For training, we used the winning submission of 2018 and incorporated low-frequency random control signals to the prosthesis. More specifically, we adopt a quasi-active control scheme for the prosthesis – a popular approach for modern powered-ankle prosthetic devices. In this control strategy, the offset of the spring attached to the prosthesis is set to mimick the quasi-static stiffness curve of an intact ankle during walking [38]. We used the above controller, along with biomechanical information including: a.) Vertical Knee Reaction Force - sum on all external forces from simulation including body kinetics and muscle forces in the vertical direction, and b.) Sum of Muscle Forces - summation of the right leg muscle groups available in the simulation environment including: hamstring, biceps femoris, vasti muscles, rectus femoris, iliopsoas, and gluteus maximus; gathered from OpenSim in order to train an IP of the underlying walking behavior. After training, we performed walking experiments in which the virtual agent repeatedly performs 60 strides with 3 different modes of our controller. In the first mode, we only use the inference mechanisms of IPs to generate reactive control values. This can be seen as a behavioral cloning approach and corresponds to the method introduced in [8]. In addition we evaluated our MPIP variant with a cost function that minimizes a.) the vertical knee reaction force and b.) the sum of muscle forces. The results of this experiment can be seen in Tab. 1. The passive prosthesis corresponds to the winning NeurIPS submission with a powered ankle described above. We can observe that the reactive IP approach behaves similar to the passive prosthesis in regards to

Table 1: Results of walking experiment in simulation.

Control Methods	Simulation Experiment 60 Stride Average				
	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 [8]	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



Figure 3: Assisted mobility with model-predictive interaction primitives (MPIP) running in real-time on a SpringActive Ruggedized Odyssey prosthesis. All control is handled by our method.

the internal stresses (e.g. knee impulse). These results are to be expected since the IP effectively clones the controller behavior. Interestingly, however, an improvement in the Local Dynamic Stability (LDS) can be observed. Empirically, we noticed that this is likely due to the fact that the typical IP approach executes the *mean* control signal of the generated posterior. As a result, some stability “buffer” is induced by the variance around the mean prediction. Tab. 1 also shows the results for our MPIP method, when optimizing knee force. A significant reduction in the medial and vertical knee impulse can be observed. In addition, a significant improvement to the LDS Lyapunov exponent can be observed, which corresponds to higher stability and reduced risk of falls. The LDS metric has been shown to be a reliable measure for assessing dynamic stability of locomotion in pathological and healthy populations [39]. Our experiment also shows a significant reduction in the muscle impulse when optimizing for muscle force using MPC. In this case, however, this impulse reduction did not also translate to an improvement in the LDS. Rather, the optimization of muscle force lead to a significant reduction in stability. It can be argued that higher muscle forces are required for resilient responses to external perturbations.

4.2 Experiment 2: Real-World Prosthetic Experiments

Next, we performed real-world experiments using the “Ruggedized Odyssey” an ankle prosthesis by SpringActive ¹ which incorporates a high power motor and parallel-elastic mechanism. To this end, the different control strategies discussed in this paper were implemented on the embedded controller of the prosthesis. An able-bodied human subject was asked to perform 20 jumps, 10 examples for soft landing and 10 examples for hard landing with more impact on the ground, for training (a bypass connected to the prosthesis was used to attach it to the participant’s leg). After training, the same participant was asked to repeat 10 trials with different control strategies running on the prosthesis. Tab. 2 presents the results of this experiment. Besides the reactive IP inference, we also evaluated our MPIP approach with a cost function for a.) reducing knee forces (minimization), b.) increasing knee forces (maximization), and c.) increasing symmetry of the ankle angles between both legs. We can observe, that the lowest knee force is achieved when using the MPIP Reduction strategy. Fig. 4 depicts the MPC optimization process and visualizes a number of candidate solutions and the resulting minimal force profile. Interestingly, the peak toe forces increase, indicating a shift in force distribution towards the metatarsal area. The MPIP Symmetry strategy also implicitly leads to a reduction in knee forces – an interesting observation that confirms the importance of maintaining symmetric gaits in order to reduce the risks of osteoarthritis.

An in-depth visualization of the effects of the different strategies can be seen in Fig. 5. The top row depicts the knee forces for the different strategies. MPIP Increase leads to a substantial increase in the knee force. In contrast to that both MPIP Symmetry and MPIP Reduction maintain a low force profile. In the bottom row we see effects of the strategies on the ankle angle. Substantial differences can be observed between ankle angle trajectories for the knee force reduction and the knee force increase. In the case of symmetry matching (bottom right) we can observe that the generated control trajectory closely matches the ankle angles of the intact leg. While the resulting trajectory substantially differs from the control trajectory for the knee force reduction strategy in flight, both of these options produce a low knee force during

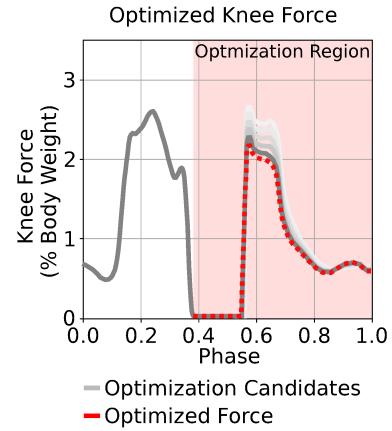


Figure 4: MPC resulting in reduced knee forces.

¹<https://springactive.com/>

Table 2: 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	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

landing, i.e., the force profiles in the top row resemble each other, since the subject is able to control both ankles effectively. Finally, to demonstrate the real-time capabilities of our method, we tested the approach outside of laboratory conditions. In particular, we asked the participant to walk with the prosthesis in a typical outdoor environment, i.e., a street side walk. The MPIP control loop runs at a rate of approx. 15 Hz and generates smooth walking trajectories for the lower-limb prosthesis. Fig. 3 shows an image sequence of the resulting walking gait. Hence, the presented method is readily applicable on commercial devices and does not require any changes to the hardware or the computational resources.

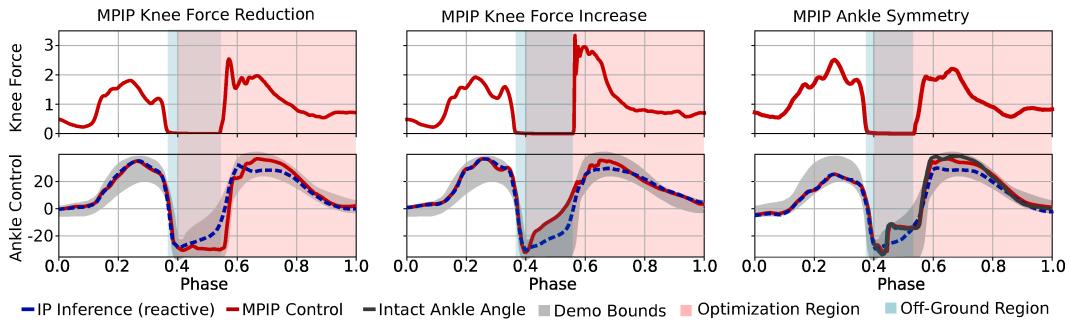


Figure 5: Different control strategies based on changing cost functions. Control signals are chosen to minimize knee loads, maximize knee loads or generate ankle angles for both legs that are symmetric.

5 Conclusion

We present a method for generating ergonomic, biomechanically safe assistive motion in tasks that require physical human-robot collaboration. An important novelty of this approach is the ability to perform predictive biomechanics – predicting future biomechanical states of the human user given current observations and intended robot actions. In turn, this feature is used within an efficient model-predictive control loop to generate control signals for a commercially available prosthetic to meet user-specific ergonomic requirements, i.e., minimize internal knee loads.

Limitations and Future Work: The accuracy of our prediction depends on the quality of the data set. We aim to investigate active learning techniques that improve the model by collecting data after deployment. The learned models in this paper focused on a single behavior and may not perform well in transitions, e.g., walking, running, staircases. We will investigate Bayesian hierarchical methods in order to cover a wide range of locomotive skills within a single model. To guarantee safety at all times, we would like to investigate learning and incorporating barrier functions [40] into the MPC process. We recognize that the lower dimensional relationships between observable, control, and latent variables is not a catch-all solution, and individuals with amputations may not fit the general learned representation. However, this can be partially remedied by re-training the approach with data from the specific user, by recording data from people with amputations, or by collecting data from high functioning amputees. We aim to further investigate this aspect going forward.

6 Acknowledgments

This work was supported by the National Science Foundation under the career award grant FP00012258, as well as by the Global KAITEKI Center.

References

- [1] R. Sinha, W. J. van den Heuvel, and P. Arokiasamy. Factors affecting quality of life in lower limb amputees. *Prosthetics and orthotics international*, 35(1):90–96, 2011.
- [2] S. A. Deans, A. K. McFadyen, and P. J. Rowe. Physical activity and quality of life: A study of a lower-limb amputee population. *Prosthetics and orthotics international*, 32(2):186–200, 2008.
- [3] USBJI. *United States Bone and Joint Initiative*, 2017. Available at <https://www.usbji.org>.
- [4] T. P. Sheehan and G. C. Gondo. Impact of limb loss in the united states. *Physical Medicine and Rehabilitation Clinics*, 25(1):9–28, 2014.
- [5] D. C. Morgenroth, A. C. Gellhorn, and P. Suri. Osteoarthritis in the disabled population: A mechanical perspective. *PM&R*, 4(5):S20 – S27, 2012. ISSN 1934-1482. Osteoarthritis.
- [6] H. Herr. Exoskeletons and orthoses: classification, design challenges and future directions. *Journal of neuroengineering and rehabilitation*, 6(1):21, 2009.
- [7] J. Campbell, S. Stepputis, and H. B. Amor. Probabilistic multimodal modeling for human-robot interaction tasks. In *Proceedings of Robotics: Science and Systems*, FreiburgimBreisgau, Germany, June 2019. doi:10.15607/RSS.2019.XV.047.
- [8] 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 *2020 IEEE International Conference on Robotics and Automation*, 2020.
- [9] M. R. Tucker, J. Olivier, A. Pagel, H. Bleuler, M. Bouri, O. Lambery, J. del R Millán, R. Riener, H. Vallery, and R. Gassert. Control strategies for active lower extremity prosthetics and orthotics: a review. *Journal of neuroengineering and rehabilitation*, 12(1):1, 2015.
- [10] 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. *arXiv preprint arXiv:2002.03197*, 2020.
- [11] A. S. Alharthi, S. U. Yunas, and K. B. Ozanyan. Deep learning for monitoring of human gait: a review. *IEEE Sensors Journal*, 19(21):9575–9591, 2019.
- [12] 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 Transactions on Neural Systems and Rehabilitation Engineering*, 27(5):1032–1042, 2019.
- [13] G. Khademi and D. Simon. Convolutional neural networks for environmentally aware locomotion mode recognition of lower-limb amputees. In *Dynamic Systems and Control Conference*, volume 59148, page V001T07A005. American Society of Mechanical Engineers, 2019.
- [14] G. Katz, C. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer. Reluplex: An efficient smt solver for verifying deep neural networks. In *International Conference on Computer Aided Verification*, pages 97–117. Springer, 2017.
- [15] A. Dokhanchi, H. B. Amor, J. V. Deshmukh, and G. Fainekos. Evaluating perception systems for autonomous vehicles using quality temporal logic. In *International Conference on Runtime Verification*, pages 409–416. Springer, 2018.
- [16] J. Su, D. V. Vargas, and K. Sakurai. One pixel attack for fooling deep neural networks. *IEEE Transactions on Evolutionary Computation*, 23(5):828–841, 2019.
- [17] N. Thatte, N. Srinivasan, and H. Geyer. Real-time reactive trip avoidance for powered trans-femoral prostheses. In *Proceedings of Robotics Systems and Sciences (RSS 2019)*, June 2019.
- [18] D. J. Lizotte, T. Wang, M. H. Bowling, and D. Schuurmans. Automatic gait optimization with gaussian process regression. In *IJCAI*, volume 7, pages 944–949, 2007.

[19] R. Calandra, A. Seyfarth, J. Peters, and M. P. Deisenroth. Bayesian optimization for learning gaits under uncertainty. *Annals of Mathematics and Artificial Intelligence*, 76(1):5–23, 2016.

[20] R. Antonova, A. Rai, and C. G. Atkeson. Sample efficient optimization for learning controllers for bipedal locomotion. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, pages 22–28, 2016.

[21] N. Dhir, H. Dallali, E. M. Ficanha, G. A. Ribeiro, and M. Rastgaard. Locomotion envelopes for adaptive control of powered ankle prostheses. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1488–1495, 2018.

[22] H. B. Amor, G. Neumann, S. Kamthe, O. Kroemer, and J. Peters. Interaction primitives for human-robot cooperation tasks. In *2014 IEEE international conference on robotics and automation (ICRA)*, pages 2831–2837. IEEE, 2014.

[23] J. Campbell and H. B. Amor. Bayesian interaction primitives: A slam approach to human-robot interaction. In *Conference on Robot Learning*, pages 379–387, 2017.

[24] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal. Dynamical movement primitives: learning attractor models for motor behaviors. *Neural computation*, 25(2):328–373, 2013.

[25] 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 *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1535–1542, May 2015.

[26] L. Chen, H. Wu, S. Duan, Y. Guan, and J. Rojas. Learning human-robot collaboration insights through the integration of muscle activity in interaction motion models. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, pages 491–496. IEEE, 2017.

[27] D. Q. Mayne. Model predictive control: Recent developments and future promise. *Automatica*, 50(12):2967–2986, 2014.

[28] S. J. Qin and T. A. Badgwell. A survey of industrial model predictive control technology. *Control engineering practice*, 11(7):733–764, 2003.

[29] S. Schaal. Is imitation learning the route to humanoid robots? *Trends in cognitive sciences*, 3(6):233–242, 1999.

[30] S. L. Delp, F. C. Anderson, A. S. Arnold, P. Loan, A. Habib, C. T. John, E. Guendelman, and D. G. Thelen. Opensim: open-source software to create and analyze dynamic simulations of movement. *IEEE transactions on biomedical engineering*, 54(11):1940–1950, 2007.

[31] G. E. Robertson, G. E. Caldwell, J. Hamill, G. Kamen, and S. Whittlesey. *Research methods in biomechanics*. Human kinetics, 2013.

[32] J. Hicks, A. Seth, S. Hamner, M. Demers, et al. Simulation with opensim-best practices, 2015.

[33] Z. I. Nejad, K. Khalili, S. H. H. Nasab, P. Schütz, P. Damm, A. Trepczynski, W. R. Taylor, and C. R. Smith. The capacity of generic musculoskeletal simulations to predict knee joint loading using the cams-knee datasets. *Annals of biomedical engineering*, pages 1–11, 2020.

[34] 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 research & therapy*, 15(1):1–8, 2013.

[35] J. Campbell, A. Hitzmann, S. Stepputtis, S. Ikemoto, K. Hosoda, and H. B. Amor. Learning interactive behaviors for musculoskeletal robots using bayesian interaction primitives. *arXiv preprint arXiv:1908.05552*, 2019.

[36] D. Kraft. A software package for sequential quadratic programming. *Forschungsbericht-Deutsche Forschungs- und Versuchsanstalt fur Luft- und Raumfahrt*, 1988.

[37] A. Bemporad and D. M. de la Peña. Multiobjective model predictive control. *Automatica*, 45(12):2823–2830, 2009.

- [38] S. K. Au, H. Herr, J. Weber, and E. C. Martinez-Villalpando. Powered ankle-foot prosthesis for the improvement of amputee ambulation. In *2007 29th annual international conference of the IEEE engineering in medicine and biology society*, pages 3020–3026. IEEE, 2007.
- [39] A. Ekizos, A. Santuz, A. Schroll, and A. Arampatzis. The maximum lyapunov exponent during walking and running: Reliability assessment of different marker-sets. *Frontiers in Physiology*, 9:1101, 2018. ISSN 1664-042X.
- [40] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada. Control barrier functions: Theory and applications. In *2019 18th European Control Conference (ECC)*, pages 3420–3431. IEEE, 2019.