

Deep Kernels for Optimizing Locomotion Controllers

Rika Antonova*

Robotics, Perception and Learning, CSC
KTH Royal Institute of Technology
Stockholm, Sweden
antonova@kth.se

Akshara Rai*

Robotics Institute
School of Computer Science
Carnegie Mellon University, USA
arai@andrew.cmu.edu

Christopher G. Atkeson

Robotics Institute
School of Computer Science
Carnegie Mellon University, USA
cga@cs.cmu.edu

Abstract: Sample efficiency is important when optimizing parameters of locomotion controllers, since hardware experiments are time consuming and expensive. Bayesian Optimization, a sample-efficient optimization framework, has recently been widely applied to address this problem, but further improvements in sample efficiency are needed for practical applicability to real-world robots and high-dimensional controllers. To address this, prior work has proposed using domain expertise for constructing custom distance metrics for locomotion. In this work we show how to learn such a distance metric automatically. We use a neural network to learn an informed distance metric from data obtained in high-fidelity simulations. We conduct experiments on two different controllers and robot architectures. First, we demonstrate improvement in sample efficiency when optimizing a 5-dimensional controller on the ATRIAS robot hardware. We then conduct simulation experiments to optimize a 16-dimensional controller for a 7-link robot model and obtain significant improvements even when optimizing in perturbed environments. This demonstrates that our approach is able to enhance sample efficiency for two different controllers, hence is a fitting candidate for further experiments on hardware in the future.

Keywords: Bayesian Optimization, Simulator-to-Robot Transfer, Bipedal Locomotion

1 Introduction

Bayesian Optimization (BO) is rapidly becoming a popular approach for optimizing controllers in robotics. It offers sample-efficient, black-box and gradient-free optimization, well suited for many problems in the field. Recently, some success has also been achieved when optimizing controllers directly on hardware [1, 2, 3]. Hence, this sample-efficient optimization framework has the potential to alleviate the need for manual tuning by experts, to a large extent. However, for high-dimensional controllers and challenging cost functions the performance of conventional BO often degrades. Without an informative prior, the number of data points required could be prohibitively expensive for hardware-only optimization. Hence, it seems ideal to exploit simulation to speed up learning, as proposed in [4] and [3]. These prior approaches, however, need extensive expert domain knowledge to define the problem-specific informed distance metric.

In this work we demonstrate how to construct an informed metric automatically, without relying heavily on domain experts. We propose to learn a distance metric with a neural network, utilizing data obtained from a high-fidelity simulator. This involves first running short simulations of a locomotion controller on a large grid of control parameters and recording the behavior of each set of parameters. The neural network then learns a mapping between input controller parameters and simulation output/behavior. We propose two ways of defining the target to be learned by the network. The first approach is based on the cost function that is to be optimized with BO on hardware, or a

*Both of these authors contributed equally

perturbed simulator. The second is cost-agnostic: learning to reconstruct a summary of robot trajectories obtained from simulation. This provides a useful re-parameterization: controller parameters that produce similar walking trajectory summaries are closer in this re-parameterized space.

In our first set of experiments we optimize a 5-dimensional controller on the ATRIAS robot hardware (Figure 1). We demonstrate that using cost-based kernel obtained with our approach outperforms using an uninformed kernel for BO. The setting we consider for ATRIAS experiments yields a proof-of-concept rather than a large-scale optimization problem. Nonetheless, we believe that its an important step towards optimizing locomotion policies for complex humanoid robots on hardware.

Prior Bayesian Optimization studies often used simpler robots. For example, [5] use snake robots, [3] use a hexapod, which often have statically stable gaits, or spend a significant amount of gait duration in a statically stable state. [1] use a smaller biped with a finite-state-machine controller, which is not widely used. In contrast, ATRIAS is a complex bipedal robot, which cannot be statically stable in single support because of point feet. Hence, it is likely to fall with unstable controllers. Moreover, our control framework is in line with most state-of-the-art robot controllers [6], [7]. Hence, results on our testbed can be transferred to other systems.

Our second set of experiments is on the Neuromuscular model [8]. We optimize a 16-dimensional controller for a 7-link robot model in simulation. Our approach of reconstructing trajectory summaries again yields a significant improvement over using uninformed kernels for BO. This is the case for both a smooth and a challenging non-smooth cost suggested in prior literature [9]. Hence the proposed approach offers a promising way to construct cost-agnostic kernels for BO automatically.



Figure 1: ATRIAS robot.

2 Background

2.1 Optimizing Bipedal Locomotion Controllers

Approaches to optimizing locomotion controllers range from manual tuning to fully automatic optimization. For complex controllers fully manual tuning is sometimes infeasible or excessively time consuming. In such cases, approaches like CMA-ES have been used to find points yielding good performance in simulation first [9]. A domain expert could then use such points as starting points to later manually adjust the parameters such that they are effective on hardware.

Recently there has been significant interest in developing methods for automatic parameter optimization. Bayesian Optimization has been suggested as one of the promising approaches due its sample efficiency. However, it still can take 30-40 samples to optimize a 4 dimensional controller [1]. To enhance kernel flexibility [10] suggests supervised learning of a feature transform during regression. However, this approach does not directly support incorporating a very large amount of data from simulation. Even if it is extended to pre-train on simulated data, it is not clear whether further joint optimization would be desirable: 10-100 hardware samples might not be enough to meaningfully affect the transform built from hundreds of thousands of points from simulation.

Recent works proposed using simulation to aid learning on hardware, for example [2], [3], [4]. [2] propose adding noisy evaluations from simulation to BO posterior directly. The limitation is the need to carefully balance the influence of the samples obtained from simulation versus hardware. [3] tabulate best performing points versus their average score on a behavioural metric – average contact time of their hexapod system in simulation. This metric guides trial-and-error learning to quickly find behaviours that compensate for damage of the robot. The search is done in behaviour space, and limited to pre-selected “successful” points from simulation. This helps make their search faster and potentially safer. However, if an optimal point was not pre-selected, BO cannot sample it during optimization. “Best points” are cost-specific (the map needs to be re-generated for each cost) and problem specific, so expert-knowledge is needed to apply the method to other systems. [4] propose a new distance metric using domain knowledge about bipedal locomotion. Short simulations are used to compute this metric for a large number of points (sets of control parameters),

thus distinguishing points based on their behaviour in simulation, rather than the Euclidean distance between them. The method generalizes to different costs and locomotion controllers. However, the distance metric is specifically designed for bipedal locomotion. Further domain-specific expertise would be needed to adapt this approach to other settings.

Another recent direction for learning locomotion controllers utilized deep neural networks. [11] formulates the problem of learning locomotion gaits as actor-critic Reinforcement Learning with neural networks as function approximators for policy and value functions. However, it is not straightforward to make such approaches data-efficient enough for real hardware ([11] uses 10 million state-action transitions for training). So in our work we are interested in combining sample efficiency of an approach like Bayesian Optimization with the flexibility and scalability of deep neural networks.

2.2 Background on Bayesian Optimization

Bayesian Optimization is a framework for sample-efficient global search ([12] gives a recent overview). The goal is to find \mathbf{x}^* that optimizes a given objective function $f(\mathbf{x})$, while executing as few evaluations of f as possible. In order to select the most promising points to evaluate next, an “acquisition” function is defined. One example is Expected Improvement (EI) function that selects \mathbf{x} to maximize expected improvement over the value of the best result obtained so far [13]. EI requires defining the prior/posterior mean and variance of f , and Gaussian Process is frequently used for this:

$$f(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}_i, \mathbf{x}_j))$$

Here μ is a mean function and k defines a kernel. $k(\mathbf{x}_i, \mathbf{x}_j)$ encodes the similarity of two inputs $\mathbf{x}_i, \mathbf{x}_j$. The value of $f(\mathbf{x}_i)$ has a significant influence on the posterior value of $f(\mathbf{x}_j)$ if $\mathbf{x}_i, \mathbf{x}_j$ have high similarity according to the kernel. Squared Exponential kernel is widely used:

$$k_{SE}(\mathbf{x}_i, \mathbf{x}_j) = \sigma_k^2 \exp\left(-\frac{1}{2}(\mathbf{x}_i - \mathbf{x}_j)^T \text{diag}(\boldsymbol{\ell})^{-2}(\mathbf{x}_i - \mathbf{x}_j)\right),$$

where hyperparameters: σ_k^2 , $\boldsymbol{\ell}$ are signal variance and a vector of length scales respectively. It is customary to adjust these automatically during optimization to learn the overall variance and how quickly f varies in each input dimension.

Gaussian Process conditioned on evidence represents a posterior distribution for f . After evaluating f at points $\mathbf{x}_1, \dots, \mathbf{x}_t$ the predictive posterior $P(f_{t+1} | \mathbf{x}_{1:t}, \mathbf{y}, \mathbf{x}_{t+1}) \sim \mathcal{N}(\mu_t(\mathbf{x}_{t+1}), \text{cov}_t(\mathbf{x}_{t+1}))$ can be computed in closed form with mean and covariance:

$$\mu_t(\mathbf{x}_{t+1}) = \mathbf{k}^T [\mathbf{K} + \sigma_{noise}^2 \mathbf{I}]^{-1} \mathbf{y} \quad \text{cov}_t(\mathbf{x}_{t+1}) = k(\mathbf{x}_{t+1}, \mathbf{x}_{t+1}) - \mathbf{k}^T [\mathbf{K} + \sigma_{noise}^2 \mathbf{I}]^{-1} \mathbf{k},$$

where $\mathbf{k} \in \mathbb{R}^t$, with $\mathbf{k}_i = k(\mathbf{x}_{t+1}, \mathbf{x}_i)$; $\mathbf{K} \in \mathbb{R}^{t \times t}$ with $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$; \mathbf{I} is an identity $\in \mathbb{R}^{t \times t}$, and \mathbf{y} is a vector of values obtained after evaluating $f(\mathbf{x}_1), \dots, f(\mathbf{x}_t)$, assuming Gaussian noise with variance σ_{noise}^2 : $y_i = f(\mathbf{x}_i) + \epsilon_{\mathcal{N}(0, \sigma_{noise}^2)}$. More details can be found in [14].

3 Problem Formulation

In this work we aim to automatically optimize parameters of controllers for bipedal locomotion with respect to some commonly used cost functions. We assume that for a d -dimensional controller there is a bounded region of interest (a hypercube) defined by low/high limits on the values of controller parameters: $\mathbf{x} \in [\mathbf{x}_{low}, \mathbf{x}_{high}] \subset \mathbb{R}^n$. Some parts of this region contain points corresponding to parameter sets of the controller that yield the desired walking behavior. Such regions might comprise a large part of the space with numerous local optima, or might comprise only a small part of the space (e.g. less than 1%). In other words: we do not impose any overly restrictive assumptions on the space of controller parameters, local/global optima, or structure and properties of the cost functions of interest.

The first setting we consider is the case of optimizing 5-dimensional parameters for Raibert locomotion controller of the ATRIAS robot similar to [15], [16] and [17]. This controller has a Raibert-like foot placement policy [18]. It uses a linear feedback law operating on horizontal speed and displacement of the center of mass (CoM) to determine a desired foot touch down point:

$$x_p = k(v - v_{tgt}) + C \cdot d + 0.5 \cdot v \cdot T$$

Here, x_p is the desired location for the end of swing; v is the current speed of the CoM; k is a feedback term that regulates v towards the target speed v_{tgt} ; C is a constant and d is the measured

distance between the stance leg and the CoM; T is the step time. The term $0.5 \cdot v \cdot T$ is a feedforward term, similar to [18]. The swing foot trajectory is defined as a 5th order spline ending at x_p .

In stance, we regulate both the torso pitch and CoM height to maintain constant desired values:

$$F_x = K_{pt}(\theta_{des} - \theta) + K_{dt}(\dot{\theta}_{des} - \dot{\theta}) \quad F_z = K_{pz}(z_{des} - z) + K_{dz}(\dot{z}_{des} - \dot{z})$$

These desired forces are sent to an inverse dynamics solver to return the corresponding joint torques that produce these desired ground reaction forces. Our 5-dimensional controller consists of $[k, C, T, K_{pt}, K_{dt}]$. Other parameters can be included, but the performance is not sensitive to them. This controller does not specify a target CoM trajectory. Instead it tries to maintain a constant height and torso angle in stance. The foot-placement strategy determines the resulting motion and speed.

To demonstrate applicability to a challenging setting with a higher-dimensional controller we also experiment with a Neuromuscular model for control [9]. Since it has not yet been fully adapted to work on ATRIAS in hardware, for this setting we evaluate our work on a 7-link planar model [19]. To facilitate comparison of our results with prior work in [4], we optimize over a 16-dimensional subspace of controller parameters. Description of the Neuromuscular controller and detailed information about the 16 parameters that are optimized can be found in Section III of [4]. We collect training data from simulations on flat ground. We conduct the evaluation of our approaches on perturbed models to create a simulated mismatch between simulation and hardware. We generate a set of model disturbances for each link of the robot, perturbing the mass, inertia and center of mass location up to $\pm 15\%$ of the original value. In addition, instead of walking on flat ground, we use a set of randomly generated rough ground profiles with step height of up to ± 6 cm.

4 Proposed Approach

The aim of our approach is to automatically learn an informed kernel for optimizing bipedal locomotion controllers with Bayesian Optimization. An uninformed kernel, like Squared Exponential, operates with vectors that represent controller parameters directly. In contrast, we learn a re-parameterization that incorporates information from simulation. We run short simulations for a range of parameter sets and record the resulting costs from the same cost function as that used in Bayesian Optimization. Costs obtained during short simulations serve as approximate indicators of the quality of the controller parameters. Our idea is to train a neural network to reconstruct the cost landscape of short simulations while focusing on the more promising parts of the space. Section 4.1 describes how this approach yields an informed kernel that helps focus the search on the well-performing regions. We also develop a cost-agnostic approach of reconstructing trajectory summaries instead of cost landscape from short simulations. This is described in Section 4.2.

4.1 Regression with Implicitly Asymmetric Loss

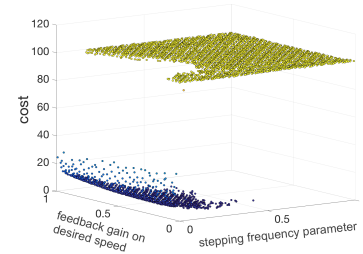
We consider a cost function focused on matching the desired walking speed and heavily penalizing falls:

$$cost_{atrias} = \begin{cases} 100 - x_{fall}, & \text{if fall} \\ 10 \cdot \|\mathbf{v}_{tgt} - \mathbf{v}_{actual}\|^2, & \text{if walk} \end{cases} \quad (1)$$

where x_{fall} is the distance travelled before falling, \mathbf{v}_{tgt} is the target velocity and \mathbf{v}_{actual} is the vector containing actual velocities of the robot. This kind of cost function is of interest because it helps easily distinguish points that walk from points that fall. Similar costs have been considered in prior work when optimizing locomotion controllers [9, 4].

Figure 2 shows a scatter plot of applying cost from Equation 1 to simulations of Raibert controller for the ATRIAS robot as introduced in Section 3. For visualization we restrict attention to a 2-dimensional subspace of the parameter space. We pick a well-performing set of parameters (in 5D), then vary the first two dimensions to obtain a 2D subspace. The challenge comes from the fact that the boundary between the well-performing (blue) and poorly performing (yellow) parameters is discontinuous. This is a typical landscape for bipedal systems, where a controller that makes the robot fall is much worse than one than walks, and the boundary is extremely sharp. Fitting such cost function with regression could

Figure 2: 2D slice of cost landscape.



be difficult. Learning to reconstruct the boundary exactly using the training set might result in overfitting and poor performance on the test set. Applying regularization is likely to yield high loss and uncertainty about points close to the boundary. This is particularly problematic if poorly performing points lie close to the most promising regions of the parameter space, which is the case in our setting.

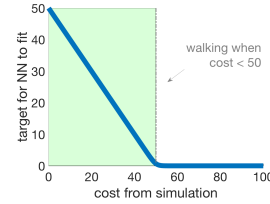
We propose to use a transformation of the cost as the target for the supervised learning. Our approach is to train a deep neural network to reconstruct a reflected shifted softplus function of the cost:

$$score_{NN} = \zeta(c_{walk} - f_{sim}(\mathbf{x})) \quad (2)$$

Here ζ is a softplus function: $\zeta(a) = \ln(1 + e^a)$, c_{walk} is the average cost for the parameter sets that walk during short simulations, $f_{sim}(\mathbf{x})$ is the cost computed by the simulator for vector \mathbf{x} of controller parameter values. Using this transformation yields a “score” function such that parameter sets which produce poor results in simulation are mapped to values close to zero. With this, the differences in scores of the poorly performing parameter sets become small or zero. In contrast, the differences in scores of the parameter sets yielding potentially promising results remain proportional to the difference in the corresponding costs. Figure 3 gives a visualization of this transformation.

Cost transformation in Equation 2 serves to essentially create an asymmetric loss for neural network training. This loss is minimized when the promising (low-cost, high-score) points are reconstructed correctly. For the poorly performing (high-cost, low-score) points, it only matters that the output of the neural network is close to zero. Such asymmetric loss can be interpreted as implementing a hybrid of regression and “soft” classification. The regression aspect aims to fit the promising points which correspond to walking behaviors. The “soft” classification aspect gives an increase in the loss only if a poorly performing point is “mis-classified” as well-performing.

Figure 3: Cost transform.



When training the neural network we apply L1 loss instead of the usual L2 loss. With this, errors in reconstructing points on the boundary contribute only linearly to the overall loss. This helps achieve a better fit of the stable parts of the parameter space, instead of focusing on the boundary.

We utilize the reconstructed transformed costs to define *asymNN* kernel for Bayesian Optimization:

$$k_{asymNN}(\mathbf{x}_i, \mathbf{x}_j) = \sigma_k^2 \exp\left(-\frac{1}{2\ell^2} |score_{NN}(\mathbf{x}_i) - score_{NN}(\mathbf{x}_j)|^2\right) \quad (3)$$

with hyperparameters σ_k^2, ℓ as described in Section 2.2. The proposed approach is able to clearly separate the unpromising part of the parameter space. Under the resulting distance metric poorly performing sets of parameters are close together and far from well-performing ones.

4.2 Reconstructing Cost-agnostic Trajectory Summaries

Utilizing costs obtained from short simulations provides a way to build an informed kernel without specifying any additional domain knowledge. If simulations are computationally expensive it is desirable to minimize the need to repeat data collection. Often, the cost function needs to be modified to accommodate different objectives, hence, there is a need for a cost-agnostic approach. For such cases we propose to train a neural network to reconstruct summaries of trajectories that are cost-agnostic, then utilize these trajectory summaries for constructing kernel distance metric.

We summarize trajectories by recording fairly generic aspects of locomotion: walking time (time before falling), energy used during walking, position of the torso, angle of the torso, coordinates of the center of mass at the end of the short simulation runs. These summaries of simulated trajectories are collected for a range of controller parameters and comprise the training set for the neural network to fit (input: \mathbf{x} – a set of control parameters; output: $\mathbf{traj}_{\mathbf{x}}$ – the corresponding trajectory summary obtained from simulation). The outputs of the (trained) neural network offer the reconstructed/approximate trajectory summaries: $f_{NN}(\mathbf{x}) = \widehat{\mathbf{traj}}_{\mathbf{x}}$, where \mathbf{x} is the input controller parameters, and $\widehat{\mathbf{traj}}_{\mathbf{x}}$ is the corresponding reconstructed trajectory summary. These are then used to construct an informed cost-agnostic kernel for Bayesian Optimization:

$$k_{trajNN}(\mathbf{x}_i, \mathbf{x}_j) = \sigma_k^2 \exp\left(-\frac{1}{2} \mathbf{t}_{ij}^T \text{diag}(\ell)^{-2} \mathbf{t}_{ij}\right), \quad \mathbf{t}_{ij} = f_{NN}(\mathbf{x}_i) - f_{NN}(\mathbf{x}_j) \quad (4)$$

The general concept of utilizing trajectory data to improve sample efficiency of BO has been proposed before, for example in [20]. However, prior work assumed obtaining trajectory data is possible

every time kernel values $k(\mathbf{x}_i, \mathbf{x}_j)$ need to be evaluated. This is not the case in our setting. Trajectory summaries are initially obtained via costly high-fidelity simulations, and it would be infeasible to compute trajectory information via simulation during BO. Hence, our approach is to train a neural network to learn reconstructing trajectory summaries first. Running a forward pass of the neural network is a relatively inexpensive operation, hence reconstructed/approximate trajectory summaries can be quickly obtained during BO whenever $k_{trajNN}(\mathbf{x}_i, \mathbf{x}_j)$ needs to be computed.

When defining trajectory summaries we did not focus on carefully selecting what aspects to include/exclude. Our goal was an approach that could be quickly adapted to other domains. When applying this approach to a new domain, the strategy could be simply to include trajectory information used to compute cost functions that are of interest/relevance in the domain. For example, for a manipulator, the coordinates of end-effector(s) could be recorded at relevant points. In principle, our approach also could utilize domain- and task-specific ‘descriptors’, like those proposed in [4, 3].

5 Experimental Results

In this section we describe our experiments with cost-based and trajectory-based kernels. We first consider the setting of optimizing a 5-dimensional controller for the ATRIAS robot. We show that the cost-based kernel is able to improve sample efficiency over standard Bayesian Optimization. We present hardware experiments to demonstrate that our kernel allows obtaining a set of parameters close to optimal on the second trial. We then discuss simulation experiments with a 16 dimensional controller that utilizes a Neuromuscular model [9]. These experiments show that our trajectory-based kernel is able to significantly outperform standard Bayesian Optimization for a higher-dimensional controller even when a sharply discontinuous cost is used during optimization.

5.1 Experiments with Raibert controller on the ATRIAS robot

For our experiments on the ATRIAS robot we used a high-fidelity ATRIAS simulator [17] to generate the kernel. We did an initial analysis of the performance of our approach in simulation, followed by hardware experiments. As described in Section 4.1, we trained a neural network to reconstruct cost information obtained from short simulations. We created a sobol grid on the input parameter space with 20K points and ran short 3.5 second simulations on each of the corresponding 20K parameter sets to compute the costs. We then trained a fully connected network with 4 hidden layers (128, 64, 16, 4 units) to reconstruct $score_{NN}$ (the transformation of the cost described in Section 4.1).

In Figure 4 we first compare the performance of BO that used our neural network-based kernel (*asymNN*) versus using a standard Squared Exponential kernel (*SE*) in simulation. For these experiments we used the cost from Equation 1, Section 4.1 with target velocity of $1m/s$. Simulations with cost less than 50 yielded walking behavior, those with cost less than 20 yielded stable walking close to the target speed. BO with *asymNN* kernel reliably found stable walking points after only 8 trials. In contrast, BO with *SE* kernel did not find stable walking solutions in the first 20 trials reliably. We also compare with a recently proposed Determinants of Gait (*DoG*) kernel [4]. *DoG* utilized more specific domain knowledge to construct an informed kernel for BO of locomotion controllers for a fixed set of points. *asymNN* was able to closely match the performance of *DoG* in this setting after 8 trials.

After experiments in simulation suggested that *asymNN* kernel can yield a significant improvement in sample efficiency of BO, we conducted a set of experiments on the ATRIAS robot. We completed 6 sets of runs of BO: 3 using *asymNN* kernel and 3 using a standard *SE* kernel with 10 trials each, leading to a total of 60 hardware experiments.

Since ATRIAS walks around a rather short boom in 2D, walking at high speeds needs a lot of torque from the robot motors. This means higher lateral forces between the robot and the boom, which

Figure 4: Initial tests of Bayesian Optimization for 5-dimensional controller in simulation. The plot shows mean of best cost so far over 30 runs for each kernel, error bars are 95% confidence intervals.

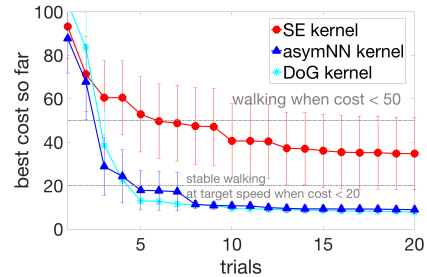
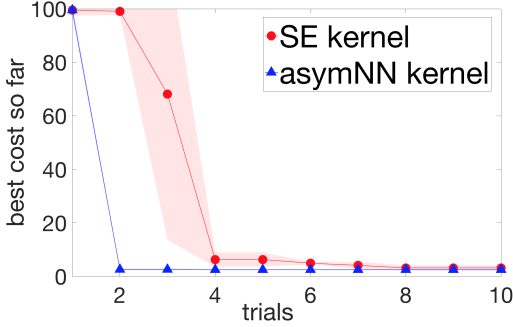
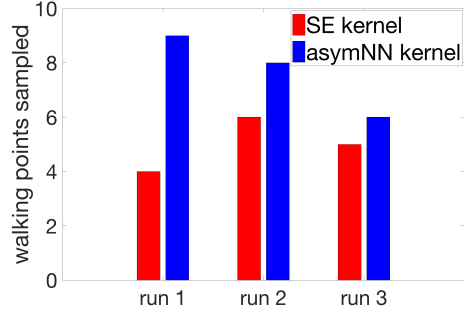


Figure 6: Hardware experiments on the ATRIAS robot.

(a) Best cost so far during BO (mean over 3 runs, shaded region indicates \pm one standard deviation).



(b) Number of “walking” points sampled (out of 10 trials in each run).



do not affect our direction of motion but can lead to a lot of internal forces, eventually breaking the robot. So, in our first attempt, we tried to start with lower speeds of $0.4m/s$ so that we could do hardware experiments and analyze the validity of our approach on hardware without breaking the robot too often. In this setting with low target speed, stable walking points comprised $\approx \frac{1}{6}$ of the parameter space. We anticipated it would be challenging to improve over BO with *SE* kernel, since it was able to find stable walking solutions after only 3-4 trials.

Figure 5: ATRIAS during BO with *asymNN* kernel.

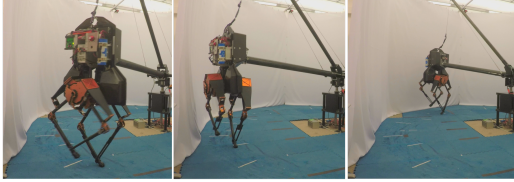


Figure 6a shows the performance of BO with *SE* versus *asymNN* kernel. *SE* obtains a stable walking solution on the 3rd trial in one run, and on the 4th trial in the two other runs. *asymNN* kernel is able to find the best-performing set of parameters on the second trial in each of the 3 runs. This confirms that using *asymNN* kernel offers an improvement over using *SE* kernel in this setting. We suggest that *asymNN* reliably selects an excellent point on the 2nd trial because such points lie far from poorly performing subspace of parameters (under the distance metric constructed with *asymNN*). *asymNN* kernel also helped sampling more walking points overall (Figure 6b). This is desirable as stable points are less likely to break the robot.

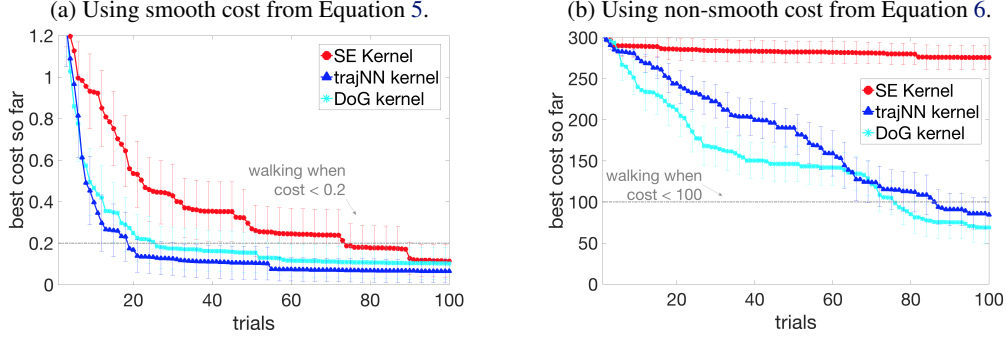
While in our hardware setup most methods are likely to sample walking points within 10 trials, we believe our experimentation is an important step towards optimizing locomotion policies for complex humanoid robots. BO studies in the past also used real robot hardware (e.g. [1, 3, 5]). However, [5, 21, 3] used robots which are statically stable for significant parts of their gait, making discontinuities in the cost function landscape less likely and in turn making the optimization easier. In contrast, ATRIAS is a complex bipedal system which is likely to fall with unstable controllers due to point feet. [1] use a walking robot similar to ours. However, their controller parametrization is very different, and not widely used, unlike our inverse dynamics and force-based controller which is more modern and state-of-the-art [7], [22], [6]. While with our hardware setting it might be hard to show improvement over simpler approaches, we believe the setting is adequate, because our testbed is fairly complex and our problem formulation is widely applicable.

5.2 Experiments with the Neuromuscular Model

16-dimensional controller of the Neuromuscular model (described in Section 3) yielded a challenging optimization setting: walking points comprised less than 2% of the parameter space in simulation. Here we describe our experiments with cost-agnostic approach for constructing an informed kernel introduced in Section 4.2. We created a grid of 100K points in the input parameter space and ran short 5 second simulations on each of the corresponding 100K parameter sets to collect the trajectory summaries. We then trained a fully connected network with 3 hidden layers (512, 128, 32 units) with L1 loss to reconstruct 8-dimensional trajectory summaries (as described in Section 4.2). All experiments were done on perturbed models, as described in Section 3.

Figure 7 compares using *trajNN* versus *SE* kernel for BO with two different costs from prior literature. The first cost promotes walking further and longer before falling, while penalizing deviations

Figure 7: Bayesian Optimization for the Neuromuscular model controller in simulation. *trajNN* and *DoG* kernels were constructed with undisturbed model on flat ground. BO is run with mass/inertia disturbances on different rough ground profiles to simulate mismatch. Plots show means over 50 runs, 95% confidence intervals.



from the target speed [4]:

$$cost_{smooth} = 1/(1+t) + 0.3/(1+d) + 0.01(s - s_{tgt}), \quad (5)$$

where t is seconds walked, d is the final hip position, s is mean speed and s_{tgt} is the desired walking speed (1.3m/s in our case). The second cost function is a simplified version of the cost used in [9]. It penalizes falls explicitly, and encourages walking at desired speed and with lower cost of transport:

$$cost_{non-smooth} = \begin{cases} 300 - x_{fall}, & \text{if fall} \\ 100||v_{avg} - v_{tgt}|| + c_{tr}, & \text{if walk} \end{cases} \quad (6)$$

where x_{fall} is the distance covered before falling, v_{avg} is the average speed of walking, v_{tgt} is the target velocity, and c_{tr} captures the cost of transport.

Figure 7a shows that *trajNN* offers a significant improvement in sample efficiency when using $cost_{smooth}$. Points with cost less than 0.2 correspond to robust walking behavior. With *trajNN*, more than 90% of runs obtain walking solutions after only 25 trials. In contrast, using *SE* requires more than 90 trials for such success rate. The performance of *trajNN* matches that of a *DoG* kernel from prior work [4]. This is notable, since *trajNN* is learned automatically, whereas *DoG* kernel is constructed using domain expertise. Figure 7b shows that *trajNN* also provides a significant improvement when using the second cost. Points with cost less than 100 correspond to walking. With *trajNN*, 70% of the runs find walking solutions after 100 trials. In contrast, optimizing non-smooth cost is very challenging for BO with *SE* kernel: a walking solution is found only in 1 out of 50 runs after 100 trials. The difference in performance on the two costs is due to the nature of the costs. For example, if a point walks some distance d , Equation 5 includes a term $\frac{1}{d}$ and Equation 6 includes $-d$. A sharper fall in the first cost causes BO to exploit around points that walk some distance. It then quickly finds points that walk forever, while BO with the second cost continues to explore.

6 Conclusion and Future Work

In this work we proposed learning informed kernels for Bayesian Optimization of locomotion controllers without relying heavily on domain experts. We optimized a 5-dimensional controller on the ATRIAS robot and showed that our cost-based kernel offered an improvement over using an uninformed kernel. We also proposed a cost-agnostic alternative. Experiments with a 16-dimensional Neuromuscular controller in simulation showed a significant improvement with different costs.

In future work it would be interesting to further analyze various approaches that enhance sample efficiency of BO. Approaches that embed simulation-based information into the kernel (like those we proposed) can enhance sample efficiency dramatically by focusing BO on regions that look promising in simulation. Approaches that use simulation-based samples in BO posterior mean directly (e.g. [2]) could be more robust to simulation-based inaccuracies after collecting a larger amount of data from hardware experiments. However, they can only incorporate cost-based information (e.g. no way to add trajectory information directly to the posterior mean). Perhaps there is an effective way to combine the two directions. Another promising line for future work is learning flexible models of simulation-vs-hardware mismatch. Such models could help decrease the influence of distortion from incorrect simulations and could help enhance both ‘kernel-based’ and ‘mean-based’ methods.

Acknowledgments

This research was supported in part by National Science Foundation grant IIS-1563807, the Max-Planck-Society, & the Knut and Alice Wallenberg Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding organizations.

References

- [1] R. Calandra, A. Seyfarth, J. Peters, and M. P. Deisenroth. Bayesian Optimization for Learning Gaits Under Uncertainty. *Annals of Mathematics and Artificial Intelligence*, 76(1-2):5–23, 2016.
- [2] A. Marco, F. Berkenkamp, P. Hennig, A. P. Schoellig, A. Krause, S. Schaal, and S. Trimpe. Virtual vs. real: Trading off simulations and physical experiments in reinforcement learning with bayesian optimization. *arXiv preprint arXiv:1703.01250*, 2017.
- [3] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret. Robots that can adapt like animals. *Nature*, 521(7553):503–507, 2015.
- [4] R. Antonova, A. Rai, and C. G. Atkeson. Sample efficient optimization for learning controllers for bipedal locomotion. In *Humanoid Robots (Humanoids), 2016 IEEE-RAS 16th International Conference on*, pages 22–28. IEEE, 2016.
- [5] M. Tesch, J. Schneider, and H. Choset. Using response surfaces and expected improvement to optimize snake robot gait parameters. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 1069–1074. IEEE, 2011.
- [6] S. Feng, E. Whitman, X. Xinjilefu, and C. G. Atkeson. Optimization-based full body control for the darpa robotics challenge. *Journal of Field Robotics*, 32(2):293–312, 2015.
- [7] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake. Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot. *Autonomous Robots*, 40(3):429–455, 2016.
- [8] H. Geyer and H. Herr. A Muscle-reflex Model that Encodes Principles of Legged Mechanics Produces Human Walking Dynamics and Muscle Activities. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 18(3):263–273, 2010.
- [9] S. Song and H. Geyer. A Neural Circuitry that Emphasizes Spinal Feedback Generates Diverse Behaviours of Human Locomotion. *The Journal of Physiology*, 593(16):3493–3511, 2015.
- [10] R. Calandra, J. Peters, C. E. Rasmussen, and M. P. Deisenroth. Manifold gaussian processes for regression. In *Neural Networks (IJCNN), 2016 International Joint Conference on*, pages 3338–3345. IEEE, 2016.
- [11] X. B. Peng, G. Berseth, and M. van de Panne. Terrain-adaptive locomotion skills using deep reinforcement learning. *ACM Transactions on Graphics (TOG)*, 35(4):81, 2016.
- [12] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas. Taking the Human Out of the Loop: A Review of Bayesian Optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- [13] J. Mockus, V. Tiesis, and A. Zilinskas. Toward Global Optimization, volume 2, chapter Bayesian Methods for Seeking the Extremum. 1978.
- [14] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005. ISBN 026218253X.
- [15] W. C. Martin, A. Wu, and H. Geyer. Experimental evaluation of deadbeat running on the atrias biped. *IEEE Robotics and Automation Letters*, 2(2):1085–1092, 2017.

- [16] C. Hubicki, A. Abate, P. Clary, S. Rezazadeh, M. Jones, A. Peekema, J. Van Why, R. Domres, A. Wu, W. Martin, et al. Walking and running with passive compliance: Lessons from engineering a live demonstration of the atrias biped. *IEEE Robotics and Automation Magazine*, 2(4.1):4–1, 2016.
- [17] W. C. Martin, A. Wu, and H. Geyer. Robust spring mass model running for a physical bipedal robot. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pages 6307–6312. IEEE, 2015.
- [18] M. H. Raibert. *Legged robots that balance*. MIT press, 1986.
- [19] 16D Simulator for Neuromuscular Models for Biped Locomotion. *Code available from <https://github.com/nthatte/Neuromuscular-Transfemoral-Prosthesis-Model>*.
- [20] A. Wilson, A. Fern, and P. Tadepalli. Using Trajectory Data to Improve Bayesian Optimization for Reinforcement Learning. *The Journal of Machine Learning Research*, 15(1):253–282, 2014.
- [21] 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.
- [22] A. Herzog, N. Rotella, S. Mason, F. Grimmering, S. Schaal, and L. Righetti. Momentum control with hierarchical inverse dynamics on a torque-controlled humanoid. *Autonomous Robots*, 40(3):473–491, 2016.