An Experimental Evaluation of Bayesian Optimization on Bipedal Locomotion

Roberto Calandra
Intelligent Autonomous Systems Lab
Technische Universität Darmstadt
calandra@ias.tu-darmstadt.de

Jan Peters
Intelligent Autonomous Systems Lab
Technische Universität Darmstadt
mail@jan-peters.net

André Seyfarth
Lauflabor Locomotion Lab
Technische Universität Darmstadt
seyfarth@sport.tu-darmstadt.de

Marc Peter Deisenroth
Department of Computing
Imperial College London
m.deisenroth@imperial.ac.uk

Abstract

The design of gaits and corresponding control policies for bipedal walkers is a key challenge in robot locomotion. Even when a viable controller parametrization already exists, finding near-optimal parameters can be daunting. The use of automatic gait optimization methods greatly reduces the need for human expertise and time-consuming design processes. In this paper, we experimentally evaluate Bayesian optimization for gait optimization of a real bipedal walker. By performing more than 1800 experimental evaluations, we compare Bayesian optimization with various acquisition functions. Additionally, we study the effects of using fixed hyperparameters instead of automatically optimize them.

1   Introduction

Key challenges in bipedal locomotion include balance control, foot placement, and gait optimization. In this paper, we focus on black-box *gait optimization*. Hence, we assume that a suitable controller to generate the desired gait has already been designed, but that appropriate gait parameters for the controller still need to be found. Due to the partially unpredictable effects and correlations among the gait parameters, gait optimization is often an empirical, time-consuming and strongly robot-specific process. As a result, gait optimization may require considerable expert experience, engineering efforts and time-consuming experiments. Additionally, a change in the environment (e.g., different floor surfaces), a variation in the hardware response (e.g., performance decline, substitution of a component or differences in the calibration) or the choice of a performance criterion (e.g., walking speed, energy efficiency, robustness), which differs from the one used during the controller design process, can require searching for new and more appropriate gait parameters. By formulating the search for appropriate gait parameters as an optimization problem, it is possible to automate the gait optimization and reduce the need for engineering expert knowledge. In the context of robotics, and specifically gait optimization, the number of experiments that can be performed on a real system can be extremely limiting. Each experiment can be costly, requires a long time, and inevitably leads to wear and tear of the robot's hardware. As a result, it is crucial for the chosen optimization method to limit the number of experiments to perform while reliably finding near-optimal parameters. In this paper, we propose to use Bayesian optimization in the context of gait optimization. To evaluate its applicability, we perform an experimental comparison of different variants of Bayesian optimization.

2 Automatic Gait Optimization

The search for appropriate parameters of a controller can be formulated as an optimization problem, such as the minimization

$$\underset{\boldsymbol{\theta} \in \mathbb{R}^D}{\text{minimize}} \; f\left(\boldsymbol{\theta}\right) \qquad (1)$$

of an *objective function* $f\left(\cdot\right)$ with respect to the parameters $\boldsymbol{\theta}$. In the context of gait optimization, this optimization problem is characterized as a global optimization of a zero-order stochastic objective function. Therefore, the use of Bayesian optimization well suits this challenging optimization task.

Bayesian Optimization   Bayesian optimization is an iterative model-based global optimization method [7, 5, 11, 1]. In Bayesian optimization, for each iteration (i.e., evaluation of the objective function f), a GP model $\boldsymbol{\theta} \mapsto f\left(\boldsymbol{\theta}\right)$ is learned from the data set $\mathbb{T} = \{\boldsymbol{\theta}, f\left(\boldsymbol{\theta}\right)\}$ composed by the past parameters $\boldsymbol{\theta}$ and the corresponding measurements $f\left(\boldsymbol{\theta}\right)$ of the objective function. This model is used to predict the response surface $\hat{f}$ and the corresponding ac-

Figure 1: The bio-inspired dynamical bipedal walker *Fox* used for the experimental evaluation.

quisition surface $\alpha\left(\boldsymbol{\theta}\right)$[1], once the response surface $\hat{f}\left(\cdot\right)$ is mapped through the acquisition function $\alpha\left(\cdot\right)$. Using a global optimizer, the minimum $\boldsymbol{\theta}^*$ of the acquisition surface $\alpha\left(\boldsymbol{\theta}\right)$ is computed without any evaluation of the objective function, e.g., no robot interaction. The current minimum $\boldsymbol{\theta}^*$ is evaluated on the robot and, together with the resulting measurement $f\left(\boldsymbol{\theta}^*\right)$, added to the dataset $\mathbb{T}$.

Gaussian Process Model for Objective Function   To create the model that maps $\boldsymbol{\theta} \mapsto f\left(\boldsymbol{\theta}\right)$, we make use of Bayesian non-parametric Gaussian Process regression [12]. Such a GP is a distribution over functions

$$f(\boldsymbol{\theta}) \sim \text{GP}\left(m_f, k(\boldsymbol{\theta}_p, \boldsymbol{\theta}_q)\right) , \qquad (2)$$

fully defined by a prior mean m_f and a covariance function $k(\boldsymbol{\theta}_p, \boldsymbol{\theta}_q)$. As prior mean, we choose $m_f \equiv \mathbf{0}$, while the chosen covariance function $k(\boldsymbol{\theta}_p, \boldsymbol{\theta}_q)$ is the squared exponential with automatic relevance determination and Gaussian noise

$$k(\boldsymbol{\theta}_p, \boldsymbol{\theta}_q) = \sigma_f^2 \exp(-\tfrac{1}{2}(\boldsymbol{\theta}_p - \boldsymbol{\theta}_q)^T \mathbf{\Lambda}^{-1}(\boldsymbol{\theta}_p - \boldsymbol{\theta}_q)) + \sigma_w^2 \delta_{pq} \qquad (3)$$

with $\mathbf{\Lambda} = \text{diag}([l_1^2, ..., l_D^2])$. Here, l_i are the characteristic length-scales, σ_f^2 is the variance of the latent function $f(\cdot)$ and σ_w^2 the noise variance. A practical issue, for both GP modeling and Bayesian optimization, is the choice of the hyperparameters of the GP model, such as the characteristic length-scales l_i, the variance of the latent function σ_f^2 and the noise variance σ_w^2. In gait optimization, these hyperparameters are often fixed a priori [9]. There are suggestions [8] that fixing the hyperparameters can considerably speed up the convergence of Bayesian optimization. However, manually choosing the value of the hyperparameters requires extensive expert knowledge about the system that we want to optimize, which is often an unrealistic assumption. An alternative common approach is to automatically select the hyper-parameters by optimizing with respect to the marginal likelihood [12].

Acquisition Function   A number of acquisition functions $\alpha\left(\cdot\right)$ exist, such as Probability of Improvement [7], Expected Improvement [10], Upper Confidence Bound [3] and Entropy-Based Improvements [4]. Experimental results [4] suggest that Expected Improvement on specific families of artificial functions performs better on average than Probability of Improvement and Upper Confidence Bound. However, these results do not necessarily hold true for real-world problems such as gait optimization, where the objective functions are more complex to model. Both Probability of improvement [9] and Expected Improvement [14] have been previously employed in gait optimization. In our experiments, we evaluate Probability of Improvement, Expected Improvement and Upper Confidence Bound.

---

[1]The correct notation would be $\alpha\left(\hat{f}\left(\boldsymbol{\theta}\right)\right)$, but we use $\alpha\left(\boldsymbol{\theta}\right)$ for notational convenience.
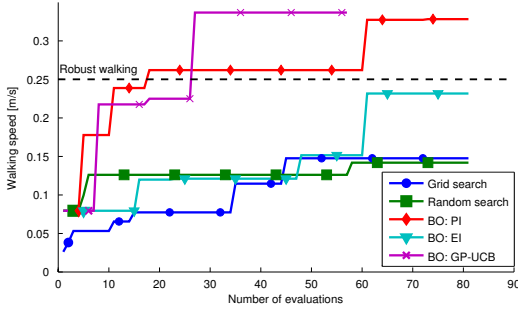
2

Figure 2: Maximum walking speed of *Fox* evaluated during the gait optimization process. For Bayesian optimization different acquisition functions, with automatically determined hyperparameters, are shown.
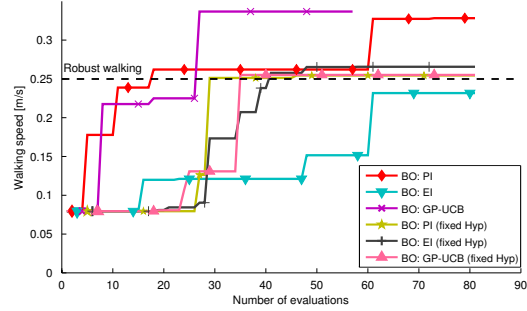
Figure 3: Maximum walking speed during the gait optimization process of *Fox*. Bayesian optimization using different acquisition functions is evaluated for both automatically optimized and fixed hyperparameters.

3   Comparative Evaluation on the *Fox* Robot

Experimental Set-up   To validate our Bayesian gait optimization approach we used the 2-D dynamic walker *Fox*, shown in Figure 1. This robot consists of a rudimentary trunk, two legs, made of a rigid segment connected by a knee joint to a telescopic leg spring, and two spherical feet with touch sensors [13]. *Fox* is equipped with four actuated degrees of freedom at both hip and knee joints. Moreover, there are six sensors on the robot: two on the hip joints, two on the knee joints, and one under each foot. The sensors on the hip and knee joints return voltage measurements corresponding to angular positions of the leg segments. The touch sensor under each foot returns binary ground contact signals. The walker is mounted on a boom that enforces planar, circular motion. An additional sensor in the boom measures the angular position of the walker on the circle. The controller of the walker is a *Finite State Machine (FSM)* that controls the four actuated joints. For the optimization process, we identified four parameters of the controller crucial for the resulting gait.The objective function f to be minimized was defined as $f(\boldsymbol{\theta}) = -\frac{1}{N}\sum_{i=1}^{N}\bar{v}_i(\boldsymbol{\theta})$, i.e., the negative mean of the average walking velocity $\bar{v}$ over $N = 3$ experiments on the robot. Minimizing this performance criterion does not only guarantee a fast walking gait, but also reliability since the gait must be robust to noise and initial configurations across multiple experiments. The chosen parameter space is sufficiently large that only a small percentage of the possible parameter values can achieve stable walking, while for most of the configurations the robot falls down after one or two steps. Each experiment was initialized from similar initial configurations, and each experiment consisted of 12 seconds starting from the moment when the foot of the robot first touched the ground. As a baseline to compare Bayesian optimization we also evaluate grid search and pure random search. For grid search optimization, we used 3 evaluations along each of the four dimensions for a total of 81 evaluations. For comparability, we performed for all other methods the same number of evaluations. In our experiments, as global optimizer of the acquisition surface, we used DIRECT [6] to find the approximate global minimum, followed by L-BFGS [2] to refine it. To initialize Bayesian optimization, we used the first three evaluations from pure random search (i.e., uniformly randomly sampled sets of parameters), thus, leaving 78 evaluations to be selected.

Experimental Results   The maximum walking speed of *Fox* evaluated during the gait optimization process for the different methods is shown in Figure 2. The optimization process of GP-UCB is limited to 57 evaluations due to a mechanical failure that forcefully interrupted the experiment. Values of the objective function below 0.1 m/s indicate that the robot fell down after a single step. Values between 0.1 and 0.15 m/s indicate that the robot was capable of executing multiple steps but showed systematic falls. Between 0.15 m/s and 0.25 m/s occasional falls occurred. Above 0.25 m/s the achieved gait was robust and did not manifested any fall. From the results, we see that both grid search and random search performed poorly, finding a maximum that can only barely walk. We can notice how Bayesian optimization, using any acquisition functions, performed considerably better. Bayesian optimization using PI and GP-UCB achieved robust gaits with similar walking speed, while GP-UCB being slightly faster in finding the maximum. On the other hand, Bayesian optimization using EI was incapable of achieving robust gaits.

3

This result is unexpected as EI is considered a versatile acquisition function, and there are experimental results [4], which suggest that EI on specific families of artificial functions performs better than GP-UCB and PI. The reason of this result were the inaccuracies of the model of the underlying objective function. The automatically selected hyperparameters had overly long length-scales (see Equation (3)) which converted into an inappropriate model. In turn, we observed that EI behaved excessively greedily, exploring the parameter space insufficiently. In a spiral, this insufficient exploration resulted in overly long length-scales and so on. We hypothesize that in case of complex real-world objective functions, such as the one we optimized, the use of the EI acquisition function might not perform as well as in the artificial functions. As second comparison, we stud-

Figure 4: Intensity map of a GP model of the parameter space, along two of the dimensions.

ied the effects of manually fixing hyperparameters, based on our expert knowledge. Figure 3 shows the comparison between the various acquisition functions, when using the manually fixed hyperparameters. From these results, it can be observed that all the different acquisition functions, when using the fixed hyperparameters, found similar sub-optimal solutions. The reason is that for all three acquisition functions with fixed hyperparameters, one parameter reached only a sub-optimal value. This observation suggests that, at least for this one parameter, the chosen length-scales were sufficiently wrong to prevent the creation of an accurate model and, therefore, the optimization process was hindered. A confirmation of this hypothesis derived from training a GP model and automatically selected the hyperparameters using the marginal likelihood, using all the evaluations performed. Both GP-UCB and PI using fixed hyperparameters performed worst than the respective cases with automatic hyperparameters selection. In contrast, for EI the use of fixed hyperparameters was beneficial. The hyperparameters of the GP model directly influence the amount of exploration performed by the acquisition functions. Hence, fixing the hyperparameters using expert knowledge can be an attractive choice, since forcing the right amount of exploration can speed up the optimization process. However, the presented experimental results also show that a poor choice of hyperparameters can potentially harm the optimization process by limiting the exploration and leading to sub-optimal solution. A visual representation between two of the parameters of the parameter space as predicted using the data collected from all the over 1800 evaluations is shown in Figure 4. This space is complex and non-convex, and, therefore, motivate the use of global optimization methods, such as Bayesian optimization.
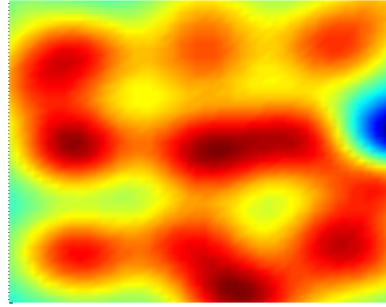
4 Discussion & Conclusion

Gait optimization is a key research topic in order to obtain efficient bipedal locomotion. Bayesian optimization is a promising method for efficient optimization, especially in fields like locomotion where only few evaluations can be performed before wearing out the hardware. To compare the performances of Bayesian optimization in different configurations, we performed over 1800 evaluations on a real bipedal walker. We firstly compared different acquisition functions. While GP-UCB had the best performances, EI performed poorly. Secondly, we compared the manually fixing hyperparameters against automatically selecting them. The results showed that manually fixing the hyperparameters can strongly influence the outcome of the optimization process. The GP modeling capabilities are often overlooked when evaluating Bayesian optimization's performances, with most of the emphasis on the use of different acquisition functions. Following from the results of our experimental evaluation, we conclude that the GP modeling capabilities are equally important to the use of different acquisition functions, when evaluating Bayesian optimization's performances. We speculate that for complex objective functions exist a strict and yet unclear connection between the exploration properties of the acquisition function and the capabilities of the GP modeling. The performance of an acquisition function depends on the capabilities of properly modeling the function. On the other hand accurate modeling takes place only when the acquisition function evaluates relevant parameters.

References

[1] E. Brochu, V. M. Cora, and N. De Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.

[2] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995.

[3] D. D. Cox and S. John. SDO: A statistical method for global optimization. *Multidisciplinary design optimization: state of the art*, pages 315–329, 1997.

[4] P. Hennig and C. J. Schuler. Entropy search for information-efficient global optimization. *Journal of Machine Learning Research (JMLR)*, 13:1809–1837, 2012.

[5] D. R. Jones. A taxonomy of global optimization methods based on response surfaces. *Journal of Global Optimization*, 21(4):345–383, 2001.

[6] D. R. Jones, C. D. Perttunen, and B. E. Stuckman. Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications*, 79(1):157–181, 1993.

[7] H. J. Kushner. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Basic Engineering*, 86:97, 1964.

[8] D. J. Lizotte, R. Greiner, and D. Schuurmans. An experimental methodology for response surface optimization methods. *Journal of Global Optimization*, 53(4):699–736, 2012.

[9] D. J. Lizotte, T. Wang, M. Bowling, and D. Schuurmans. Automatic gait optimization with Gaussian process regression. In *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 944–949, 2007.

[10] J. Mockus, V. Tiesis, and A. Zilinskas. The application of Bayesian methods for seeking the extremum. *Towards Global Optimization*, 2:117–129, 1978.

[11] M. A. Osborne, R. Garnett, and S. J. Roberts. Gaussian processes for global optimization. In *3rd International Conference on Learning and Intelligent Optimization (LION3)*, pages 1–15, 2009.

[12] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.

[13] D. Renjewski. *An engineering contribution to human gait biomechanics*. PhD thesis, TU Ilmenau, 2012.

[14] 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.