2018 7th IEEE International Conference on Biomedical
Robotics and Biomechatronics (Biorob)
Enschede, The Netherlands, August 26-29, 2018

Bayesian Optimisation of Exoskeleton Design Parameters

Daniel F. N. Gordon[1,2,*], Takamitsu Matsubara[2,3], Tomoyuki Noda[2], Tatsuya Teramae[2], Jun Morimoto[2], Sethu Vijayakumar[1]

Abstract— **Exoskeletons are currently being developed and used as effective tools for rehabilitation. The ideal location and design of exoskeleton attachment points can vary due to factors such as the physical dimensions of the wearer, which muscles or joints are targeted for rehabilitation or assistance, or the presence of joint misalignment between the human subject and exoskeleton device. In this paper, we propose an approach for identifying the ideal exoskeleton cuff locations based on a human-in-the-loop optimisation process, and present an empirical validation of our method. The muscle activity of a subject was measured while walking with assistance from the XoR exoskeleton (ATR, Japan) over a range of cuff configurations. A Bayesian optimisation process was implemented and tested to identify the optimal configuration of the XoR cuffs which minimised the measured EMG activity. Using this process, the optimal design parameters for the XoR were identified more efficiently than via linear search.**

I. INTRODUCTION

Increasingly, exoskeletons and other assistive robotic devices have been used to great effect for rehabilitation and limb replacement [5][6]. Exoskeletons which fit poorly can cause discomfort and even work against the intended aim of the device [18][21], altering natural muscle activation patterns in undesired ways [10]. As such, the optimal scenario is for exoskeletons to be designed with a particular subject in mind. However, it is not always possible for individuals to obtain these devices on an ad hoc basis. This is especially true in low to middle income countries, where creating personal devices on an individual basis is not feasible due to complexity of the manufacturing process and excessively high cost [15].

Several strategies are being investigated to deal with this problem. One is the creation of modular exoskeleton designs which can be manufactured cheaply and efficiently, in some cases using 3D printable parts [2][3]. A doctor or physiotherapist could request a specific design and have it cheaply manufactured for patients. Another strategy is to include passive degrees of freedom in exoskeleton designs. The inclusion of passive DOFs can in some sense be used to minimise the kinematic inconsistency between exoskeletons and human subjects [12]. Finally, an additional strategy is to develop exoskeletons which are in some way adjustable, so that they can be easily modified to suit he needs of a particular subject. One example of such a robot is the XoR, developed at ATR in Japan, a lower-body exoskeleton which features 4 adjustable cuffs located on the thigh and shank link segments [11]. Robots such as these could in theory be adjusted on the fly according to the needs of an individual, and would suit being used in rehabilitation centres where many patients require treatment.

One difficulty with the use of such exoskeletons is how to decide which cuff positions are optimal for a specific patient and use case. For example, the ideal cuff positions could differ between healthy people due to different limb lengths, heights, muscle masses and other physical characteristics. Studies have shown that the power loss between the torques generated by an exoskeleton and those experienced by a human subject are significant [18][20]. These losses are due in part to interface dynamics involving soft biological tissues, and could in theory be mitigated by a suitable selection of the interface position. Additionally, optimal cuff positions could differ between patients with different disabilities or afflictions - whether the aim is to target the hamstrings or gastrocnemius muscles, for example. If adjustable exoskeletons are to be used effectively in a clinical setting, a means of efficiently identifying the optimal cuff positions for an individual is required. A linear search for the optimal exoskeleton configuration is likely to be too time-expensive to perform, especially as the number of design parameters increases. In this work, we propose the use of Bayesian optimisation as a method to efficiently identify the optimal cuff configuration for an individual. A human-in-the-loop optimisation process was designed and carried out using measurement data obtained from a subject walking with the XoR exoskeleton. An empirical validation of our method is presented, based on recorded EMG data.

Several recent works have begun to explore using optimisation to improve exoskeletons. Hamaya et al. (2017) used a reinforcement learning framework to learn the assistive strategy of a 1-DoF exoskeleton robot which best reduced human EMG measurements [8]. Zhang et al. (2017) and Ding et al. (2018) used a human-in-the-loop optimisation process to improve the torques applied by exoskeletons in terms of metabolic energy reduction [4][22]. Kim et al. (2017) used Bayesian optimisation to optimize step frequency in unaided human subjects [13]. Our contribution differs from these works in that Bayesian optimisation is used with the intent of optimising exoskeleton design parameters.

II. METHODS

A. Hardware

The exoskeleton which we used to apply assistance is the lower-body exoskeleton XoR, pictured in Figure 1, developed

[1]Department of Informatics, University of Edinburgh, UK.
[2]Department of Brain Robot Interface, ATR-CNS, Kyoto, Japan.
[3]Graduate School of Information Science, Nara Institute of Science and Technology, Nara, Japan.
[*]Corresponding author. Email: daniel.gordon@ed.ac.uk

978-1-5386-8183-1/18/$31.00 ©2018 IEEE

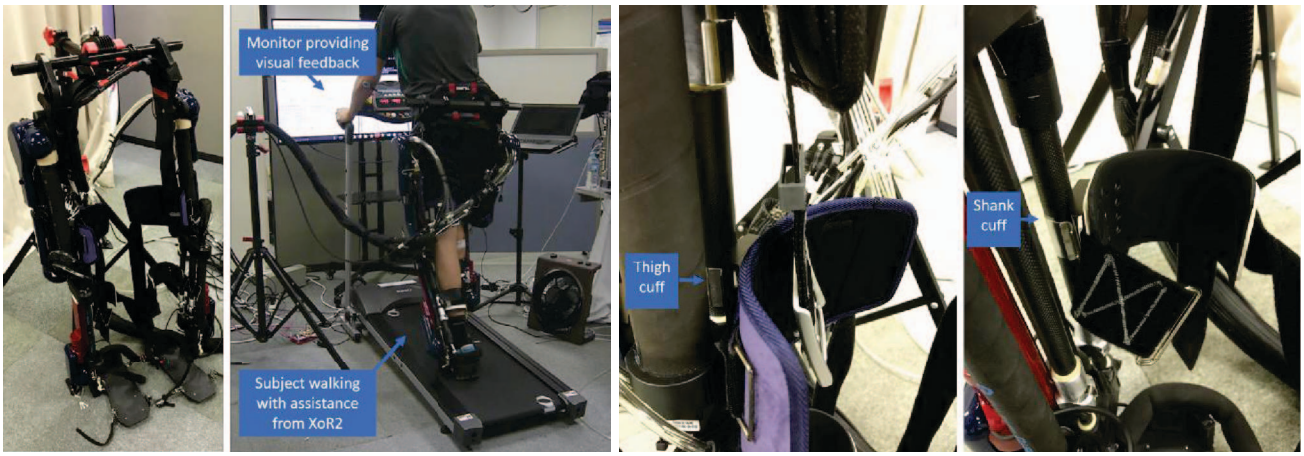

Fig. 1: **Left:** the XoR exoskeleton. **Centre-left:** a snapshot during a data capture. A subject is wearing the XoR while walking on a treadmill, with a monitor providing visual feedback so the subject can attempt to track the reference trajectory. **Centre-right/right:** a close-up view of the thigh and shank cuffs. The cuffs are adjustable vertically on the bar to which they are connected.

by ATR in Kyoto, Japan. The XoR has 14 joints in total, of which 6 are actuated (the hip, knee and ankle flexion degrees of freedom). These joints are powered via a hybrid actuation system involving both DC motors and pneumatic artificial muscle (PAM) actuators [11].

The XoR features 4 adjustable cuffs which attach to the left/right thigh and shank of the user, providing a structural link between the human and robotic systems (see Figure 1). Therefore, the system has four design parameters: x_L, x_R, y_L and y_R, where x and y denote the thigh and shank cuffs and L and R distinguish between the left and right leg, respectively. The 'usable' lengths of the thigh and shank links (meaning the length on which the cuffs could be attached) were measured at 10.3cm and 12.0cm, respectively. The properties of the XoR's design parameters are summarised in Table I.

B. Experimental protocol

Two sets of walking trials were undertaken for a single subject wearing the XoR. Each individual trial consisted of 60s of walking at a steady pace on a treadmill. This 60s period was structured as follows:

1) 10s for initial adjustment and alignment with reference trajectories,
2) 20s of recorded unassisted walking,
3) Engaging assistance, followed by a 10s period for adjustment,
4) 20s of recorded assisted walking.

The assistance provided by the XoR was governed by a PD controller implemented for this experiment. The controller tracked reference kinematic trajectories for the hip, knee and ankle joints, measured while the subject walked with the XoR in passive mode. During each trial, the subject was provided with visual feedback via a screen to aid with alignment to the reference trajectories. The PAM actuators were disabled throughout all walking trials, meaning all joints were actuated only by the DC motors.

Measurements recorded while the subject walked included the encoder readings and motor currents from each active joint of the XoR, as well as EMG signals taken from a sensor placed on the gastrocnemius muscle of the right leg. This location was chosen because it resulted in less noise compared to sensors placed on the deep muscles of the thigh, and also because previous studies have shown that ankle muscle activations can be affected by active assistance at the hip joint [14]. Therefore, the gastrocnemius muscle should be susceptible to changes in location of both the thigh and shank cuffs.

After each trial, the recorded data was segmented in to gait cycles and labelled as assisted or unassisted. The EMG signals were band-pass filtered using a zero-lag Butterworth filter, rectified, then low-pass filtered in order to produce a linear envelope for each full gait cycle of data [17]. The average EMG signal for each gait cycle was calculated by integrating the corresponding envelope and dividing by cycle duration. These values were then averaged by label, and the percentage change in average EMG signal from unassisted to assisted walking was calculated.

In Trial Set 1, the position of the shank cuffs was kept fixed while the thigh cuff position was varied from 0.5cm to 10.0cm in steps of 0.5cm. In Trial Set 2 the thigh cuffs were kept fixed while the shank cuff position was varied from 0.5cm to 11.5cm in steps of 1.0cm[1]. In both trial sets,

[1]The discrepancy in step sizes was due to the process of changing the shank cuff position being more time consuming and intricate than that of the thigh cuff.

Design parameter	Notation	Domain (cm)
Right thigh cuff position	x_R	$[0, 10.3]$
Left thigh cuff position	x_L	$[0, 10.3]$
Right shank cuff position	y_R	$[0, 12.0]$
Left shank cuff position	y_L	$[0, 12.0]$

TABLE I: The design parameters of the XoR exoskeleton, including associated notation and values of the domain.

the left and right cuff positions were coupled, i.e. $x_R = x_L$ and $y_R = y_L$. Measurements of the cuff locations were made using a tape measure. An example of a walking trial in progress is displayed in Figure 1.

After each trial set was complete, a Gaussian process regression (GPR) model was fit to the data using the MATLAB function *fitrgp*. This process resulted in two GPR models, g_t and g_s, representing the ground truth GP posterior for the thigh and shank, respectively. Here, g_t represents the change in activity of the gastrocnemius due to exoskeleton assistance as a function of the thigh cuff position, assuming constant shank cuff position, and similarly for g_s. By collecting EMG measurements in advance and sampling from g_t and g_s offline, the effectiveness of Bayesian optimisation could be investigated without the need for repeatedly taking experimental measurements using the exoskeleton, which was an intricate and time-intensive process.

C. Bayesian optimisation

Consider an objective function f with parameters $\mathbf{x}$, where we do not know the analytical form of f but are free to sample from it. Bayesian Optimisation is a process which ultimately seeks to minimise (or maximise) f by repeatedly drawing samples and updating a posterior expectation of f. Bayesian Optimisation assumes that the function f can be described by a Gaussian process [19]. The sample points are chosen in such a way that an *acquisition function* is maximised, which in some sense defines which regions of f are 'most useful' to sample from i.e. provide the highest utility. The Bayesian Optimisation algorithm proceeds as follows:

1) Given known values $\mathbf{y} = \{f(\mathbf{x}_1), f(\mathbf{x}_2), ...\}$, update the posterior expectation of f using Gaussian Processes.
2) From the posterior expectation of f, find the argmax of the acquisition function to determine x_{new}, the next point to sample from.
3) Compute $f(\mathbf{x}_{\text{new}})$ and add it to the list of known values.

Steps 1 - 3 are repeated for a set number of iterations or until some termination criterion is satisfied. The algorithm is typically initialised by sampling from a fixed number of randomly generated points.

D. Optimisation protocol

The GP posterior functions g_t and g_s were sampled from as part of an optimisation process, using MATLAB's implementation of Bayesian optimisation *bayesopt*. The process was repeated for three acquisition functions (the MATLAB implementations of Probability of Improvement (PoI), Expected Improvement (EI) and Lower Confidence Bound (LCB) [19]) to determine whether this affected performance.

Since the objective functions of the optimisation (g_t and g_s) were stochastic, a set of 100 optimisations were performed for each combination of acquisition function and objective function. Each optimisation was run for a total of 15 iterations. Once all optimisations were completed, the performance was quantified in terms of the mean and standard deviation of the number of iterations required to converge to within some region of the optimal cuff position.

III. RESULTS

A. Analysis of EMG measurements

The EMG results from the XoR walking trial sets and corresponding ground truth GP posteriors g_t and g_s are shown in Figure 2. These results suggest that the location of the XoR cuffs does have a significant effect on the performance of the exoskeleton. In Trial Set 1, the largest observed reduction in EMG activity due to assistance was a reduction of 38.3%, occurring when the thigh cuff was placed at $x = 2.5$cm. In contrast, the most sub-optimal position for the thigh cuff based on the raw observations was $x = 4.0$cm, where the EMG activity was seen to increase by 25.0%. This represents an observed performance differential of 63.3%. According to the model obtained from GP regression, the optimal cuff location is located at $x = 1.9$cm, where the EMG activity is predicted to be reduced by 29.6% due to XoR assistance. The performance differential according to the GPR model was 39.2%.

In Trial Set 2, the optimal observed cuff placement was $y = 8.5$cm, with an EMG reduction of 42.5%. The most sub-optimal position based on the raw observations was $y = 4.5$cm, where the EMG activity was seen to increase by 12.8%, and therefore the observed performance differential had a value of 55.3%. The GPR model predicted the optimal cuff location to be at $y = 9.3$cm, where the estimated EMG reduction was 35.8%. The GPR model for the Trial Set 2 data predicted the performance differential to be 37.2%.

For reference, the analysis of the EMG measurements and subsequent GPR fitting is summarised in Table II.

B. Bayesian optimisation results

The Bayesian optimisation performance was quantified according to how many iterations were required to converge to a certain minimum value of EMG reduction. The lowest value of reduction analysed was 0%, which corresponds to the cutoff between the XoR assistance hindering and helping the wearer. The performance was then analysed in intervals of 5% up to a maximum of 25%, since the true optimum EMG reduction was in the region of 30% for both the thigh and shank posteriors. To illustrate, a single optimisation process converging to a minimum value of 25% in N iterations would mean that after N iterations the predicted optimum EMG reduction is at least 25% and remains so during the remainder of the optimisation process.

The mean and standard deviation of the number of iterations required to reach each minimum value of EMG reduction was computed for the thigh and shank posteriors and for each of the three acquisition functions considered. The results from this analysis are presented in Figure 3. For Trial Set 1 we see good performance if the EMG reduction required is less than 20%, with the mean number of iterations required being less than 10. For EMG reductions of greater than 20%, the mean number of iterations required sharply

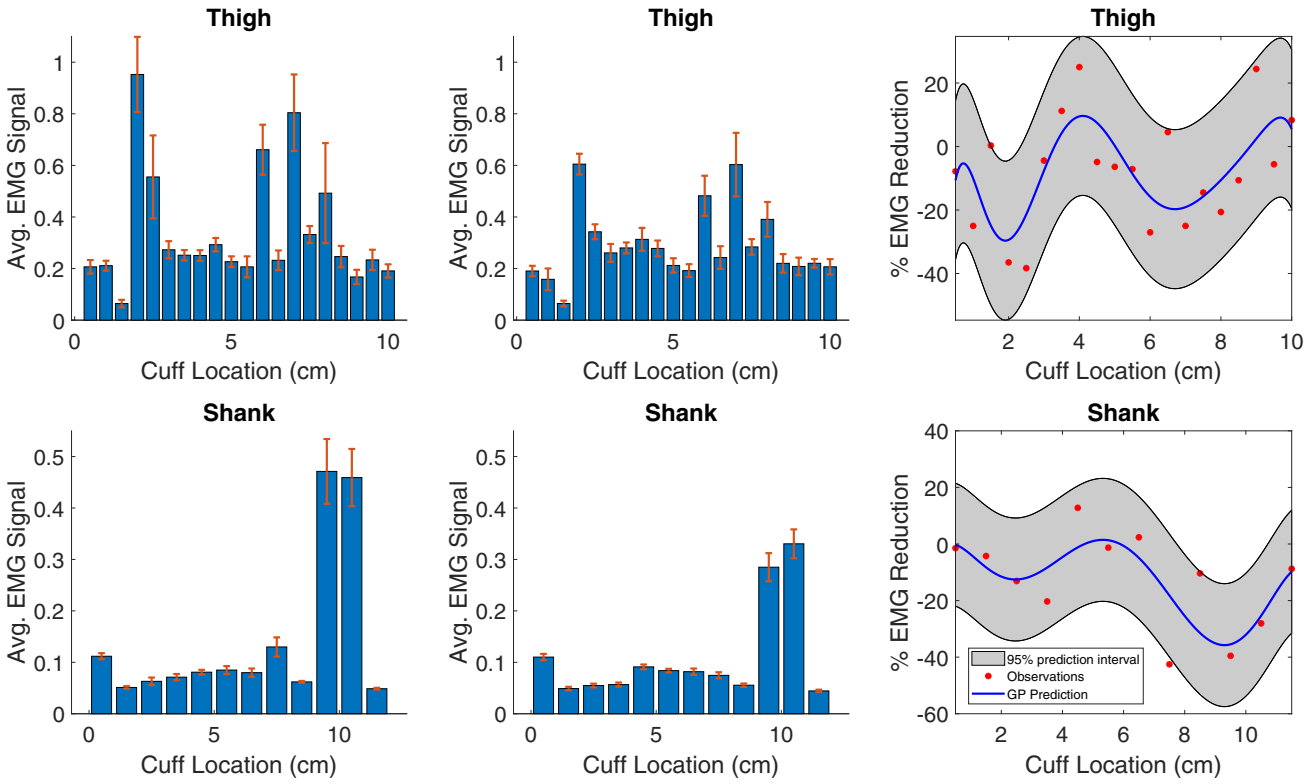

Fig. 2: **Top row:** the average total EMG signal measured per gait cycle during unassisted walking (left) and assisted walking (centre) as well as the GPR model of the relative change in muscle activation from unassisted to assisted (right) for Trial Set 1. **Bottom row:** analogous results for Trial Set 2. Note that the GPR model fits shown in the rightmost figure of the top and bottom row, respectively, are equivalent to g_t and g_s, discussed in Section II-B.

rises to more than 14. This is less than required for a linear search, but is subject to variability as evidenced by the error bars. For Trial Set 2 we see good performance in terms of mean iterations even up to the highest minimum EMG reduction of 25%.

For Trial Set 1 the performance of each of the tested acquisition functions is comparable in both mean and standard deviation across all values of EMG reduction. For Trial Set 2, the PoI acquisition function is slightly more effective than EI and LCB for minimum EMG reductions of greater than 15%, particularly in terms of standard deviation.

The Bayesian optimisation results are summarised for reference in Table III, which gives the optimum acquisition function as well as mean and standard deviation of the number of iterations for each value of minimum EMG reduction.

IV. DISCUSSION

By collecting EMG data of a subject walking with assistance from the XoR exoskeleton over a range of cuff positions, we observed that changing the attachment point locations meaningfully affected the EMG signals received from the subject, and therefore how their muscles were being assisted by the exoskeleton. Notably, for some configurations of the exoskeleton cuffs, the EMG activity of the subject was increased by active assistance, rather than decreased. In particular, GPR models fit to the EMG data estimated the performance differential, measured in terms of reduction of EMG signals, between the optimal and most suboptimal cuff configurations to be approximately 40%. These results underline the need to optimise exoskeleton cuff placement to ensure assistive forces are applied correctly.

Collecting the experimental data required 20 measurements in the case of the thigh cuffs and 11 measurements in the case of the shank cuffs. Using Bayesian Optimisation, we were able to identify cuff configurations which performed well in terms of EMG reduction in a far lower number of measurements, as summarised in Table III. Specifically, our analysis suggests that overall EMG reductions of more than 25% could be achieved in a mean of just 6-7 iterations when modifying the shank cuff position, and overall EMG

	Trial Set 1		Trial Set 2	
	Observed	Modelled	Observed	Modelled
Position (cm)	2.5	1.9	8.5	9.3
Value (%)	−38.3	−29.6	−42.5	−35.8
Differential (%)	63.3	39.2	55.3	37.2

TABLE II: The optimal cuff position, optimal percentage EMG reduction, and performance differential listed for both Trial Set 1 and Trial Set 2. For comparison purposes values are presented based on both the raw EMG results and the GPR models.

Minimum	Trial Set 1			Trial Set 2		
EMG Reduction	Best Acquisition Function	Mean Iterations	STD	Best Acquisition Function	Mean Iterations	STD
$> 0\%$	LCB	2.39	2.70	LCB	1.12	0.50
$> 5\%$	PoI	2.98	3.31	PoI	2.09	2.21
$> 10\%$	PoI	5.77	5.76	PoI	3.25	3.41
$> 15\%$	PoI	9.09	6.96	PoI	4.43	4.38
$> 20\%$	LCB	14.7	7.09	PoI	5.43	4.97
$> 25\%$	PoI	16.03	5.76	PoI	6.31	5.14

TABLE III: A summary of the Bayesian optimisation results. The best acquisition function, mean and STD results are presented for Trial Set 1 and Trial Set 2, in terms of iterations required to converge to the minimum EMG reduction listed in the left-most column.

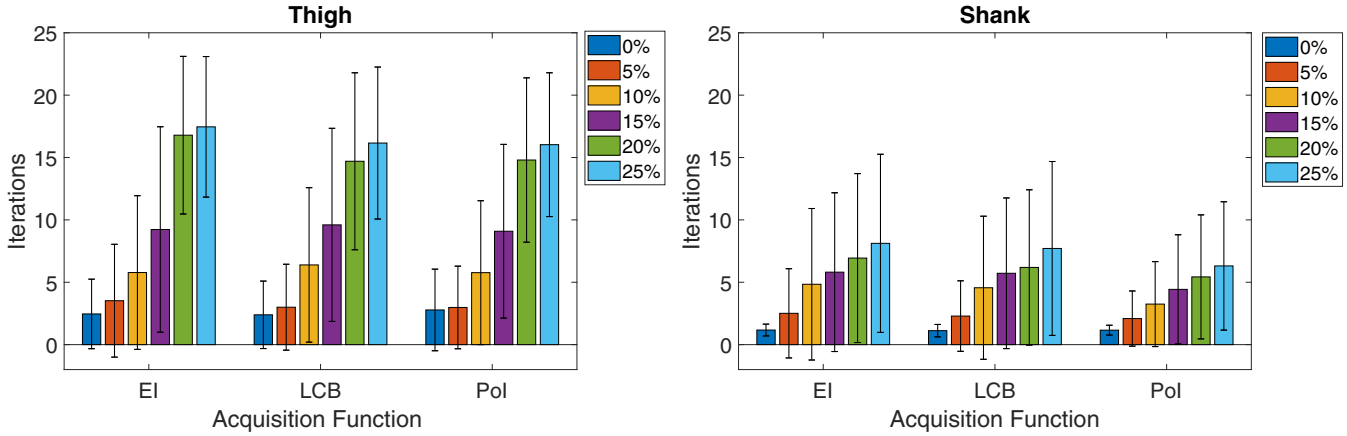

Fig. 3: The mean and standard deviation of the number of iterations required to converge to a certain minimum EMG reduction for the thigh (**left**) and shank (**right**) GP posteriors. The bars denote the mean number of iterations required to ensure that the point found by the optimisation results in at least a reduction in EMG activity corresponding to some minimum level. Note the legend which illustrates the ordering of minimum EMG reduction zones.

reductions of more than 15% could be achieved in a mean of 8-9 iterations when modifying the thigh cuff position. Using Bayesian Optimisation as an alternative to sampling evenly over the entire domain can therefore potentially offer significant increases in efficiency, which could be particularly useful in a clinical setting such as a rehabilitation centre, where robotic exoskeletons such as the XoR could efficiently be adapted to provide effective assistance to multiple patients.

The Bayesian optimisation process was less successful in identifying the optimal thigh cuff position compared to the shank cuff position. A potential explanation for the difficulty in finding the true optimum of the posterior for Trial Set 1 is that, referring to Figure 2, the thigh posterior contains two peaks in EMG reduction (at approximately 2cm and 6.5cm, respectively) which are close in magnitude. Therefore, the Bayesian optimisation process is more susceptible to becoming trapped in the local minima near the $x = 6.5$cm cuff position. This is not true for the shank posterior, which has a marked global minimum near a cuff position of 9cm.

The Bayesian optimisation process was not hugely affected by the choice of acquisition function, but overall PoI gave the best results for both of our Trial Sets.

A potential source of error in the collected data came from the relatively strenuous data collection process, which introduced subject sweating and fatigue, exacerbating the problem of noisy EMG signals. Additionally, it was difficult to ensure that the position of the sensor was unchanged between walking trials, due to interaction between the exoskeleton and the EMG sensor. To offset these issues, EMG measurements were taken during both passive and assisted walking, and the relative improvement is reported rather than the absolute EMG signals. A potential alternative to using EMG measurements would be to use a respiratory device to measure metabolic rate, similarly to Kim et al. (2017) [13]. Alternatively, kinematic and kinetic data could be collected during walking trials to allow for human effort to be computed using a human-exoskeleton musculoskeletal model [7][9][16].

A further test of this technique would be how it performs on higher dimensional problems. In the case of the XoR, considering the 2D problem of optimising the thigh and shank cuff locations simultaneously would require more than 500 measurements using linear search (assuming a resolution of 0.5cm). Bayesian Optimisation could potentially be used as an alternative which could actually be applied in practice. Re-applying this investigation to the 2D case is a potential source of further work.

Overall, we have shown the successful application of Bayesian optimisation to improve the performance of a lower-body exoskeleton by modifying freely adjustable design parameters. This result, combined with recent previous works on human-in-the-loop optimisation for exoskeleton control [4][8][13][22], illustrates the potential usefulness of Bayesian optimisation as a tool for both the control and design of exoskeletons.

ACKNOWLEDGMENT

This research was supported by the Commissioned Research of NICT; the ImPACT Program of the Council for Science, Technology, and Innovation (Cabinet Office, Government of Japan); AMED SRPBS; JSPS KAKENHI JP16H06565; and NEDO as part of employment at ATR-CNS, Kyoto, Japan, and the Engineering and Physical Sciences Research Council (EPSRC) as part of the CDT in Robotics and Autonomous Systems at Heriot-Watt University and the Univeristy of Edinburgh.

REFERENCES

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

[2] Cardoso, R. and Silva, M.T., 2017. Design, Analysis and Simulation of a Novel Device for Locomotion Support. In *Converging Clinical and Engineering Research on Neurorehabilitation II* (pp. 833-837). Springer, Cham.

[3] Cui, L., Phan, A. and Allison, G., 2015, August. Design and fabrication of a three dimensional printable non-assembly articulated hand exoskeleton for rehabilitation. In *Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE* (pp. 4627-4630). IEEE.

[4] Ding, Ye, Myunghee Kim, Scott Kuindersma, and Conor J. Walsh. "Human-in-the-loop optimization of hip assistance with a soft exosuit during walking." *Science Robotics* 3, no. 15 (2018): eaar5438.

[5] Dollar, A.M. and Herr, H., 2008. Lower extremity exoskeletons and active orthoses: challenges and state-of-the-art. *IEEE Transactions on robotics*, 24(1), pp.144-158.

[6] Jimenez-Fabian, R. and Verlinden, O., 2012. Review of control algorithms for robotic ankle systems in lower-limb orthoses, prostheses, and exoskeletons. *Medical engineering & physics*, *34*(4), pp.397-408.

[7] Gordon, D., Henderson, G. and Vijayakumar, S., 2018. Effectively quantifying the performance of lower-limb exoskeletons over a range of walking conditions. *Frontiers in Robotics and AI*, *5*, p.61.

[8] Hamaya, M., Matsubara, T., Noda, T., Teramae, T. and Morimoto, J., 2017. Learning assistive strategies for exoskeleton robots from user-robot physical interaction. *Pattern Recognition Letters*, *99*, pp.67-76.

[9] Henderson, G., Gordon, D. and Vijayakumar, S., 2017, December. Identifying invariant gait metrics for exoskeleton assistance. In *Robotics and Biomimetics (ROBIO), 2017 IEEE International Conference on* (pp. 520-526). IEEE.

[10] Hidler, J.M. and Wall, A.E., 2005. Alterations in muscle activation patterns during robotic-assisted walking. *Clinical Biomechanics*, *20*(2), pp.184-193.

[11] Hyon, S.H., Hayashi, T., Yagi, A., Noda, T. and Morimoto, J., 2013, November. Design of hybrid drive exoskeleton robot XoR2. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on* (pp. 4642-4648). IEEE.

[12] Jarrasse, N. and Morel, G., 2012. Connecting a human limb to an exoskeleton. *IEEE Transactions on Robotics*, *28*(3), pp.697-709.

[13] Kim, M., Ding, Y., Malcolm, P., Speeckaert, J., Siviy, C.J., Walsh, C.J. and Kuindersma, S., 2017. Human-in-the-loop Bayesian optimization of wearable device parameters. *PloS one*, *12*(9), p.e0184054.

[14] Lenzi, T., Carrozza, M.C. and Agrawal, S.K., 2013. Powered hip exoskeletons can reduce the user's hip and ankle muscle activations during walking. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *21*(6), pp.938-948.

[15] C. Penchak, M. Thompson, Disability and Rehabilitation in Developing Countries, Global Health Education Consortium, Dallas, Texas, USA, 2007.

[16] Rajagopal, A., Dembia, C.L., DeMers, M.S., Delp, D.D., Hicks, J.L. and Delp, S.L., 2016. Full-Body Musculoskeletal Model for Muscle-Driven Simulation of Human Gait. *IEEE Trans. Biomed. Engineering*, *63*(10), pp.2068-2079.

[17] Rose, W., 2011. Electromyogram analysis. *Online course material. University of Delaware. Retrieved July, 5*, p.2016.

[18] Schiele, A., 2008, May. An explicit model to predict and interpret constraint force creation in phri with exoskeletons. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on* (pp. 1324-1330). IEEE.

[19] Snoek, J., Larochelle, H. and Adams, R.P., 2012. Practical bayesian optimization of machine learning algorithms. In *Advances in neural information processing systems* (pp. 2951-2959).

[20] Yandell, M.B., Quinlivan, B.T., Popov, D., Walsh, C. and Zelik, K.E., 2017. Physical interface dynamics alter how robotic exosuits augment human movement: implications for optimizing wearable assistive devices. *Journal of neuroengineering and rehabilitation*, *14*(1), p.40.

[21] Zanotto, D., Akiyama, Y., Stegall, P. and Agrawal, S.K., 2015. Knee joint misalignment in exoskeletons for the lower extremities: Effects on user's gait. *IEEE Transactions on Robotics*, *31*(4), pp.978-987.

[22] Zhang, J., Fiers, P., Witte, K.A., Jackson, R.W., Poggensee, K.L., Atkeson, C.G. and Collins, S.H., 2017. Human-in-the-loop optimization of exoskeleton assistance during walking. *Science*, *356*(6344), pp.1280-1284.