Bayesian Optimization of Soft Exosuits Using a Metabolic Estimator Stopping Process

Myunghee Kim[1,2*], Charles Liu[1*], Jinsoo Kim[1,3], Sangjun Lee[1,3], Adham Meguid[3], Conor J. Walsh[1,3], Scott Kuindersma[1]

***Abstract*— Recent human-in-the-loop (HIL) optimization studies using wearable devices have shown an improved average metabolic reduction by optimizing a small number of control parameters during short-duration walking experiments. However, the slow metabolic dynamics, high measurement noise, and experimental time constraints create challenges for increasing the number of control parameters to be optimized. Prior work applying gradient descent and Bayesian optimization to this problem have decoupled metabolic estimation and control parameter selection using fixed estimation intervals, which imposes a hard limit on the number of parameter evaluations possible in a given time budget. In this work, we take a different approach that couples estimation and parameter selection, allowing the algorithm to spend less time on refining the metabolic estimates for parameters that are unlikely to improve performance over the best observed values. Our approach uses a Kalman filter-based metabolic estimator to formulate an optimal stopping problem during the data acquisition step of standard Bayesian optimization. Performance was analyzed in numerical simulations and in pilot human subject testing with two subjects that involved optimizing six control parameters of a single-joint exosuit and four parameters of a multi-joint exosuit.**

I. Introduction

Wearable robotic devices are intended as a means to augment human economy, strength, and endurance. Over the past decade, a number of devices have been developed for reducing the metabolic cost of walking for able-bodied individuals [1]–[7]. With tethered and portable hardware platforms having advanced considerably [3], [8], [9], it is now possible to accurately control the shape of assistance force profiles [1], [10]–[13], which has led to work focused on how these control parameters influence overall system performance [11], [12].

The control parameters of assistive devices have traditionally been optimized through systematic sweeps using average metabolic response as an objective [4], [14]. Since inter-subject performance variability has been observed when applying the same optimized control strategy to multiple subjects [15], there are significant practical advantages to efficiently identifying optimal parameter settings on an individual basis. Recently, several groups have explored different combinations of devices and optimization algorithms for *human-in-the-loop* (HIL) optimization [16]–[19]. These methods generally use an instantaneous metabolic rate estimate [20] as an objective for optimizing control parameters [16]–[18], [21]. Given the long timescales of metabolic dynamics and noisy observations through breath measurement devices [22], metabolic estimation performance often becomes a bottleneck that limits the number of samples achievable in a fixed experimental time window.

Bayesian optimization is a sequential stochastic optimization strategy that is well suited for optimizing objective functions that are noisy and expensive to evaluate [19], [23]–[25]. The algorithm is typically framed in the context of an iteration budget, abstracting away the actual data sampling process that may actually require multiple measurements to evaluate the objective. The HIL optimization problem, on the other hand, has a limited time constraint rather than an iteration budget since participants can only perform for a fixed interval of time. In previous studies, a fixed observation time window was used for every sample evaluation, matching the time constraint to an iteration constraint. This means that the algorithm can waste valuable time measuring metabolic data for parameter values that are unlikely to improve upon the best known values up until that point in the optimization.

In this work, we propose an estimator stopping process of HIL Bayesian optimization to reduce the total experimental time and/or accommodate more control parameters within a given period of time. A stopping rule is determined after every respiratory measurement to spend less time on unpromising parameter settings. We performed simulations based on previously collected metabolic data to show the time-efficiency of the proposed method compared to conventional approaches. In addition, we conducted preliminary human subject testing with single- and multi-joint soft exosuits to demonstrate the feasibility of the proposed approach.

II. HIL optimization using a stopping process

This section describes our approach and presents the three main components of our algorithm: Bayesian optimization, metabolic cost estimation, and an estimator stopping process.

A. Approach Overview

The goal of the proposed HIL optimization is to minimize the metabolic rate with respect to a set of control parameters in a fixed experimental time period. Note that we will use the terms "metabolic rate" and "metabolic cost" interchangeable throughout this paper. Our approach is based on the insight that for some parameter values, even a very short measurement duration (i.e. a high variance estimate of metabolic

*These authors contributed equally to this work

[1]John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA

[2]Mechanical and Industrial Engineering, University of Illinois at Chicago, Chicago, IL, USA

[3]Wyss Institute for Biologically Inspired Engineering, Harvard University, Cambridge, MA, USA

Exploration points $\mathbb{E}$
Maximum number of measurements N
for e **in** $\mathbb{E}$
    Estimate e for N measurements $\sim \mathcal{N}(\mu_e, \sigma_e^2)$
    Update training set $\mathbb{S} \cup (e, \mu_e)$
end
Determine thresholds from exploration phase $\mathbb{A}$
Initialize $t = 0$
while $t < T$
    Find $x_t = \arg\max_x EI(x|\mathbb{S})$
    Initialize $i = 0$
    Given $x^* = \arg\min_{x \in \mathbb{S}} \mu_x$
    while $i < N$
        Estimate $x_t \sim \mathcal{N}(\mu_{x_t}, \sigma_{x_t}^2)$
        $\mu = \mu_{x_t} - \mu_{x^*}$
        $\sigma^2 = \sigma_{x_t}^2 + \sigma_{x^*}^2$
        if $\mu > \mathbb{A}_i \sigma$
            break
        end
        i+=1, t+=1
    end
    Update training set $\mathbb{S} \cup (x_t, \mu_{x_t})$
end

Algorithm 1: Bayesian Optimization using an Estimator Stopping Process. The algorithm starts with an exploration phase of E sample points using the maximum N measurements to produce a low-variance estimate of metabolic rate, after which it proceeds with the early stopping mechanism until time expires.

cost) can provide enough information to rule them out as being worse than parameters that have already been tried. Stopping the metabolic estimation early in such cases leaves more time in the experimental budget to sample potentially better parameter values. An outline of the algorithm is shown in Algorithm 1.

B. Bayesian Optimization

Bayesian optimization solves problems of the form:

$$x^* = \arg\min_{x \in \mathbb{X}} f(x),$$

where $\mathbb{X} \subset \mathbb{R}^d$, d is dimensionality of the parameter space, and f is an unknown, noisy function that is expensive to evaluate [26]. In the HIL setting, the algorithm proceeds by computing a posterior distribution over metabolic cost functions, $f(x)$, given all previous evaluations and then maximizing an acquisition function on this distribution to globally select the next control parameter value, x, to evaluate. The distribution over f is modeled as a Gaussian process (GP) [27], $\mathbb{G}$,

$$f \sim \mathbb{G}(\mu, \kappa),$$

where $\mu : \mathbb{X} \to \mathbb{R}$ is a mean function typically set to zero, and $\kappa : \mathbb{X} \times \mathbb{X} \to \mathbb{R}$ is a covariance kernel where a squared exponential kernel is a common choice,

$$\kappa(x_i, x_j|\theta) = \sigma_\theta^2 \exp\left(-\sum_{k=1}^{d} \frac{(x_{i,k} - x_{j,k})^2}{2l_k^2}\right)$$

with hyperparameters $\theta = [\sigma_\theta^2, l_1, ..., l_d]$.

Given a set of training samples, $\mathbb{S} = \{(x_i, y_i)\}_{i=1}^n$ where $y_i \sim f(x_i) + \mathcal{N}(0, \sigma_n^2)$, the posterior distribution at x follows a normal distribution with mean $\bar{\mu}(x)$ and variance $\bar{\sigma}^2(x)$ evaluated as in [23]. Note that a small set of exploration points is needed to compute an initial posterior distribution before the optimization process is started.

Given a distribution at point x, an acquisition function captures how attractive that point is to sample next. We used Expected Improvement (EI), which returns the expected reduction in cost over the best parameters observed so far,

$$EI(x|\mathbb{S}) = z\bar{\sigma}(x)\Phi(z) + \bar{\sigma}(x)\phi(z)$$
$$z = \frac{\min(\mathbb{S}) - \bar{\mu}(x) + \xi}{\bar{\sigma}(x)},$$

where ξ is a scaling parameter to adjust the tradeoff between exploration-exploitation, ϕ is the Gaussian PDF, and Φ is the Gaussian CDF [28]. Using EI, the next point is chosen as $x^* = \arg\max_x EI(x|\mathbb{S})$.

C. Metabolic Estimation

For each parameter set, the metabolic cost was continuously estimated by using a first-order dynamical model [20] for up to 45 breaths. We model the metabolic cost, m, as

$$m_t(c_0, c, \tau_0, \tau) = c(1 - e^{\frac{-t}{\tau}}) + c_0 e^{\frac{-t}{\tau_0}},$$

where t is the total measurement time and τ_0, τ are time constants characterizing the rate of change for the initial cost c_0 and the instantaneous energetic cost c, respectively.

An Unscented Kalman Filter (UKF) was used to estimate the parameters of our metabolic cost model using real-time breath measurements [29]. Kalman filters are applied to discrete time systems of the form

$$o(t+1) = F(o(t), v(t), t)$$
$$z(t) = H(o(t), w(t), t),$$

where $o(t)$ represents the unobserved state of the system, $z(t)$ the observed measurement, zero-mean Gaussian vectors $v(t)$ for process disturbances or modeling errors, and $w(t)$ for measurement noise.

As τ_0, τ, c, c_0 are assumed to be constant parameters, the UKF formulation becomes

$$x = [c_0 \ c \ \tau_0 \ \tau]$$
$$F(o(t), v(t), t) = o(t) + v(t)$$
$$H(o(t), w(t), t) = m_t(o(t)) + w(t).$$

Upon subsequent breath measurements, the filter refines the estimate of the metabolic cost function parameters.

D. Estimator Stopping Process: σ-scaled Offset

The decision to continue to the next iteration of Bayesian optimization is framed as a stopping problem using the metabolic estimator's state with the following parameters:

N	Finite horizon
X_t	State at time t
$P(X_t\|X_{t-1})$	State transitions, typically Markovian
λ	Discount Factor $\in (0, 1]$
$r(X)$	Bounded reward function for continuing at state X
$g(X)$	Bounded reward function for stopping at state X

The objective is to find time τ that maximizes

$$V(x) = \max_{\tau:\tau\leq N} \mathbb{E}[\lambda^{\tau} g(X_{\tau}) + \sum_{i=0}^{\tau-1} \lambda^{i} r(X_j)|X_0 = x].$$

The value function can be computed via backward induction by defining

$$J_N(x) = g(x)$$
$$J_n(x) = max\{g(x), r(x) + \lambda \mathbb{E}_{P(y|x)}[J_{n+1}(y)]\},$$

then $V(x) = J_0(x)$ and the optimal stopping time, τ, satisfies [30]

$$\min_{t\geq 0} J_t(X_t) = g(X_{\tau}).$$

Within the context of the metabolic estimator, let N be the maximum measurements allowed at a parameter evaluation, $\hat{x}_t \sim \mathcal{N}(\mu_{x_t}, \sigma^2_{x_t})$ is the energetic cost estimate and associated variance for the current evaluation, and $\hat{x}^* \sim \mathcal{N}(\mu_{x^*}, \sigma^2_{x^*})$ is the energetic cost estimate and variance for the best parameter setting found thus far as defined by the Gaussian process. As Bayesian optimization starts with an exploration phase, $\hat{x}^*$ will always be defined.

The σ-scaled offset model directly embeds $\hat{x}_t$ and $\hat{x}^*$ into the stopping problem. The best estimate of the difference in energetic costs is given by the following Gaussian distribution

$$\hat{x}_t - \hat{x}^* \sim \mathcal{N}(\mu, \sigma^2)$$
$$\mu = \mu_{x_t} - \mu_{x^*}$$
$$\sigma^2 = \sigma^2_{x_t} + \sigma^2_{x^*}.$$

The stopping problem state is then the estimate difference with a non-Markovian transition model of independent draws from the same Gaussian distribution. The lower x_t seems relative to x^*, the higher the reward function incentivizes continuing. In summary,

$$X_t = \mu$$
$$P(X_t|X_{t-1}) = P(X_t) \sim \mathcal{N}(\mu, \sigma^2)$$
$$r(x) = K\sigma - x$$
$$g(x) = 0,$$

where K is a σ-scaled offset to achieve some degree of risk aversion to stopping too early. The optimal stopping point for this formulation can be reduced to the first time $X_t > K\sigma$. As the cost variance decreases over subsequent measurements, σ^2 decreases the offset accordingly.

Having an unfavorable K for scaling the σ-scaled offset can lead to under-sampling potentially favorable settings or reducing to a fixed sampling of the maximum number of measurements. However, since the Bayesian optimization algorithm begins with an exploration phase, these data points can be used to adjust K. A schedule of the trajectory of thresholds is created as if the stopping problem were run on the exploration phase and the most risk-averse settings are used during the optimization. Let $\mathbb{M}$ and Σ be $|\mathbb{E}| \times N$ matrices providing traces where $\mathbb{M}_{ij}, \Sigma_{ij}$ correspond to $\mu_{x_j}, \sigma^2_{x_j}$ for exploration point i and measurement j. For each exploration point, the $\hat{x}^*$ distribution can be the final distribution, $(\mathbb{M}_{iN}, \Sigma_{iN})$, to estimate the magnitude of thresholds when comparing similar cost estimates. Algorithm 2 outlines the process of calculating the adaptive threshold vector $\mathbb{A}$.

Init $|\mathbb{E}| \times N$ **Matrix** $\mathbb{X}$
Init N**-Vector** $\mathbb{A}$
Define Window W
for $i = 1$ **to** $|\mathbb{E}|$
  for $j = 1$ **to** N
    $\mu = \mathbb{M}_{ij} - \mathbb{M}_{iN}$
    $\sigma^2 = \Sigma_{ij} + \Sigma_{iN}$
    $\mathbb{X}_{ij} = \frac{\mu}{\sigma}$
  end
end
for $i = 1$ **to** N
  $\mathbb{A}_i = MAX\{\mathbb{X}_{rc} | 1 \leq r \leq |\mathbb{E}|, i - W \leq c \leq i + W\}$
end
Return $\mathbb{A}$

Algorithm 2: Determining threshold levels at every measurement time based on subject's exploration data. The hyperparameter W specifies a window of time to choose the most risk averse threshold over.

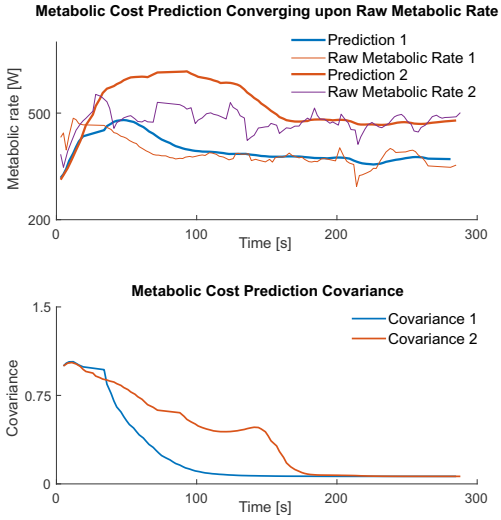

Fig. 1: UKF metabolic estimator applied to two subject trials. The top figure shows the raw measurements and the metabolic cost prediction. The bottom plot displays the normalized covariance associated with the estimate, beginning with a covariance of 1 and converging at different rates depending upon the data.

III. Simulation Results

1) Metabolic Estimation: We evaluated the online estimation method using subject data from a prior study. The results from two representative subjects are shown in Fig. 1. The accuracy of this estimator was tested in percentage error by comparing the prediction values at certain time instants (5 min, 2 min, 1.5 min, 20 breaths, and 30 breaths) to the "ground truth," which we take to be the average of the last two minutes of the five minute data. The errors were 1%, 1%, 1%, 3%, and 6%, respectively.

2) Optimization using an Estimator Stopping Process: Four standard functions (Table I) were tested using the early estimator stopping process, compared to a fixed measure-

Function	Domain	Min	$\tilde{\mu}$	$\tilde{\sigma}$
Hartmann-6	$[0,1]^6$	-3.322	-0.259	0.383
Ackley	$[-32.768, 32.768]^4$	0	20.882	1.027
Levy	$[-10,10]^4$	0	42.544	27.939
Branin	$([-5,10],[0,15])$	0.398	54.452	51.129

TABLE I: Validation data set. Four different objective functions are used for simulation validation.

ment interval. These functions exhibit different behavior: the Hartmann-6 and Ackley functions have larger domain spaces with lower standard deviations, generally making them more difficult to enter local minima. While the Levy and Branin functions have large regions of local minima, and in Branin's case three global minima, much higher noise is injected due to the standard deviations of the respective functions.

A Gaussian state estimate, $\hat{x}_t$, was initialized to $(\tilde{\mu}, \tilde{\sigma}^2)$ of the respective function. At each timestep a measurement, z_t, was taken with Gaussian noise proportional to the standard deviation of the function. The updated state estimate is the product of Gaussian PDFs of the current state estimate, $\hat{x}_t$, with the measurement estimate, (z_t, σ_z), where $\sigma_z = [\frac{\tilde{\sigma}^2}{10}, \tilde{\sigma}^2, 10\tilde{\sigma}^2]$. Each function used the same exploration phase of 6 points, a window for the adaptive thresholds of 2, and the maximum number of timesteps was set to 45 measurements. The optimization methods then had 630 remaining measurements to find an optimal point, which for the fixed measurement interval implies 14 iterations of standard Bayesian Optimization.

For time t, let x_t be the GP's estimated minimum point. Define

$$y_t = \frac{f(x_t) - f(x^*)}{f(x_0) - f(x^*)},$$

where f is the true, noiseless function and $f(x^*)$ is the function minimum. Fig. 2 plots y_t averaged over 100 trials. As the exploration phase is the same for each method, it is omitted from the charts for brevity.

The early stopping mechanism allowed for greater exploration of the domain space as evidenced by the smoother trajectory over time as compared to more of a step ladder trajectory with the fixed measurement interval. The scenario where measurement noise is underestimated may lead to the possibility where promising points are stopped too early, while overestimating measurement noise may reduce the algorithm to fixed measurements at the maximum number of timesteps. However in our simulations, regardless of the inaccuracy in the measurement noise the early stopping method produced results closer to the optimum as compared to the fixed measurement interval method (Fig. 2).

IV. Human Walking Studies with Soft Exosuits

We conducted preliminary human walking experiments using single- and multi-joint soft exosuits to test the proposed algorithm. The Harvard Longwood Medical Area Institutional Review Board approved the studies, and all participants provided written informed consent.

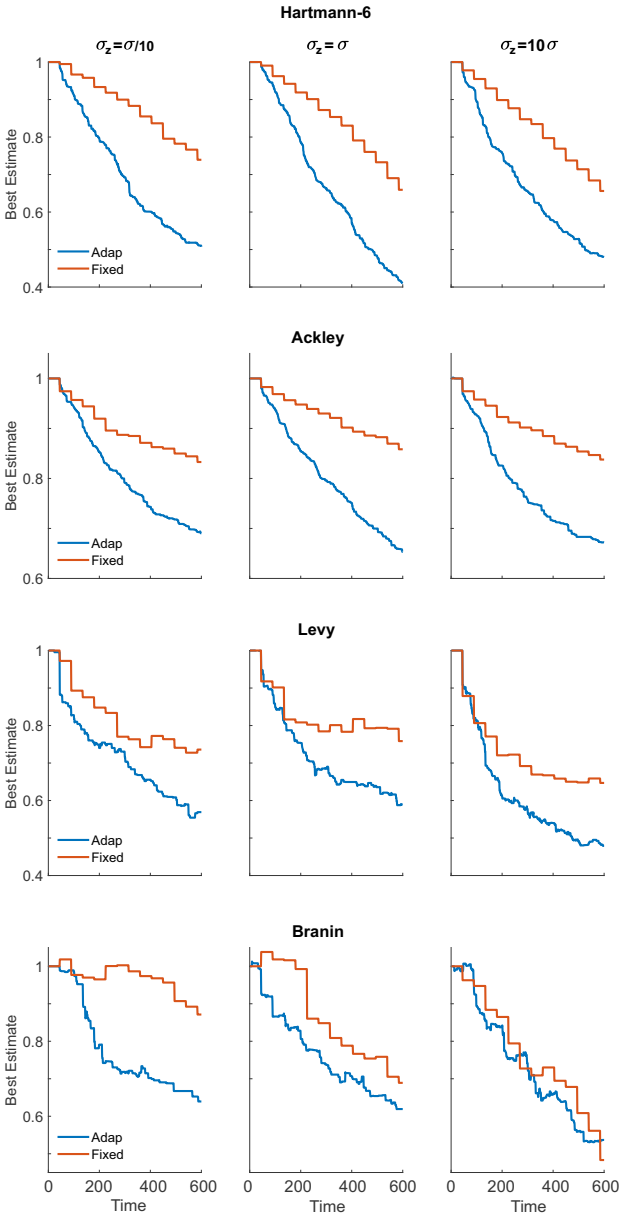

Fig. 2: Performance comparison between an estimator stopping method (Adap) and a fixed measurement interval method (Fixed). Four different objective functions with three different levels of measurement variation were tested. The exploration phase is the same between methods and is omitted for brevity.

A. Single-joint exosuit optimization

We optimized parameters for hip extension assistance using a soft exosuit, a textile based wearable device [31] with a non-portable, tethered system consisting of two independent actuators to deliver assistance through the hip extension load path of both legs. Further details about the soft exosuit and off-board system can be found in [7].

1) Controls overview: The electronics hardware has a three-layer configuration for the control system architecture: real-time target machine (top layer), microprocessor (middle layer), and servomotor driver (bottom layer). In the top layer, the real-time target machine (Speedgoat, Switzerland) runs a Simulink model (MathWorks, Natick, MA, USA) in a host

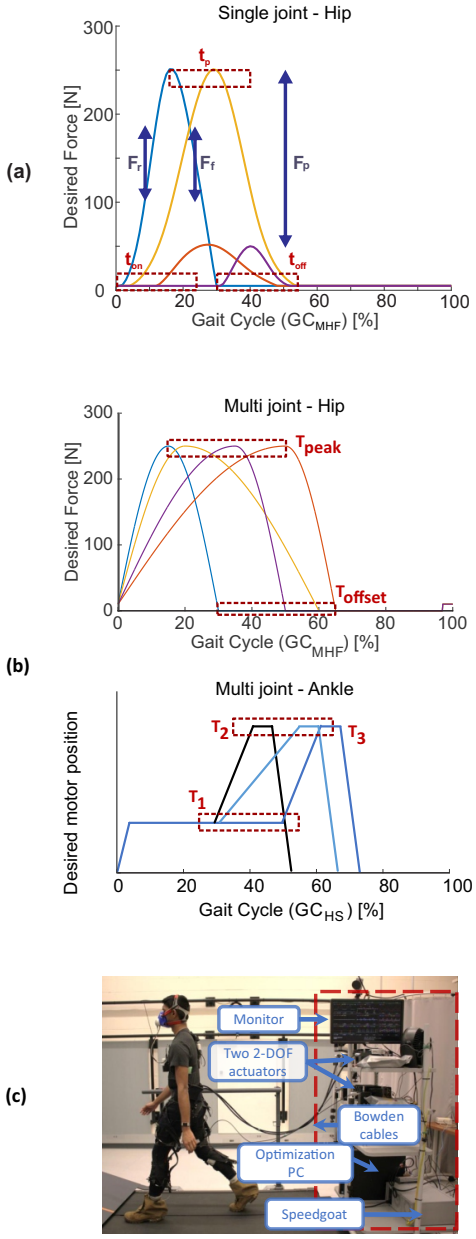

Fig. 3: (a) Single joint, Hip extension force profiles (6 parameters, 1: onset timing, 2: peak timing, 3: offset timing, 4: peak magnitude of force, 5: rising force node, 6: falling force node) (b) Multi-joint, Hip extension force profiles (2 parameters, 1: peak timing, 2: offset timing) and ankle plantarflexion position profiles (2 parameters, 1: onset timing, 2: peak timing) (c) Experimental setup

laptop. The Atmel 32-bit microprocessor (ATSAME70N21, Atmel Corporation, CA, USA) in the middle layer acts as a signal hub between the target machine and the servomotor driver. In the bottom layer, Gold Twitter (Elmo Motion Control Ltd, Israel) is selected as the current servomotor driver. The CAN communication between layers closes the control loop at a 1kHz loop rate [7].

The algorithm PC collects the subject's metabolic data, runs the optimization code, and sends the assistance profile parameters that will be explored in the next condition to the host laptop. Ethernet communication is used between a real-time target machine, a host laptop, and the algorithm PC. Further details can be found in [18].

2) Control parameter optimization: We used a six dimensional parameter space (increased from two parameters used in prior work [18]). We chose the following parameters to optimize a force profile as a function of gait cycles which consist of two cubic splines: the magnitude of peak force, F_p, three timings (onset t_{on}, peak t_p, offset t_{off}), and two node points (rising F_r and falling F_f) at the middle of the rising and falling duration as described in Fig. 3 (a). The gait cycle ($\%GC_{MHF}$) was calculated using a maximum hip flexion gait event [18].

Considering the hip extension moment and hardware limitations, the following constraints were applied.

$$0\%GC_{MHF} \leq t_{on} \leq t_p - 15\%GC_{MHF}$$
$$t_p + 15\%GC_{MHF} \leq t_{off} \leq 55\%GC_{MHF}$$
$$50N \leq F_p \leq 250N$$
$$F_p \times 0.3 \leq F_r \leq F_p \times 0.7$$
$$F_p \times 0.3 \leq F_f \leq F_p \times 0.7.$$

3) Experimental protocol: Two healthy male adults with prior experience of walking with the exosuit's assistance participated in this study (S1: 48 yrs, 85 kg, 178 cm; S2: 27 yrs, 74 kg, 177 cm). The study consisted of a single-day experimental session involving an exosuit parameter optimization process during treadmill walking at 1.25 m/s. The optimization process included two separate walking bouts: a 20-min (approx.) bout initializing the algorithm with an evenly distributed pseudo-random set of parameters and a 40-min (approx.) bout to run actual Bayesian optimization, making the entire optimization process less than 60 min including 3-min warm-up periods in the beginning of each bout. The metabolic cost of walking 5-min with and without the exosuit were then measured using (i) individually optimized parameters (OPT) and (ii) fixed generic parameters (FIX) from the literature (i.e. ESLP condition in [32]) for comparison.

4) Result: The optimization procedure ran 37 iterations with 17 early estimator stops (subject 1) and 27 iterations with 8 early estimator stops (subject 2). Participants reduced their metabolic cost by 35% (subject 1) and 7% (subject 2) for the optimal condition (OPT) and 29% (subject 1) and -4% (subject 2) for the fixed generic condition (FIX), compared to no-exosuit condition. The optimized force profile is shown in Fig. 4. The soft exosuit followed the desired trajectory within 5% error relative to the peak force.

B. Multi-joint exosuit optimization

We further evaluated the performance of the algorithm by optimizing parameters for both hip and ankle assistance using a multi-joint soft exosuit [9]. An off-board actuation system was used, similar to the one in the single-joint optimization study but with two sets of 2-DOF actuators. Further details about the textile architecture can be found in [9].

1) Controls overview: Similar to the single-joint optimization study, a Simulink-based real-time control architecture (Speedgoat) was used to regulate the exosuit assistance.

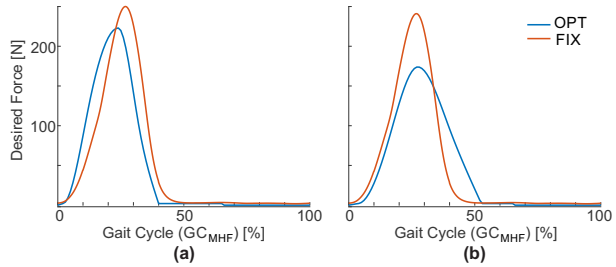

Fig. 4: Hip extension force profiles of the two subjects: the optimal parameters were (a) 2%, 24%, 40%, $223N$, $0.6\times223N$, and $0.45\times223N$ (subject 1) and (b) 3%, 27%, 53%, $174N$, $0.44\times174N$, and $0.54\times174N$ (subject 2), for onset, peak, offset timings, peak force, rising node, and falling node, respectively.

For the hip assistance, we used a combination of piece-wise sinusoidal force profiles as a function of gait cycles ($\%GC_{MHF}$) (Fig. 3 (b), top). For the ankle assistance, on the other hand, the controller detected heel strike by using the foot IMUs to define gait cycles ($\%GC_{HS}$). A force-based position control was performed following a piece-wise linear position profile as a function of $\%GC_{HS}$ (Fig. 3 (b), bottom) [9], [13].

2) Control parameter optimization: In this study, four control parameters across different lower-limb joints (two for the hip extension load path and the other two for the ankle plantarflexion load path) were simultaneously optimized for each individual using the proposed algorithm. For the hip extension assistance, the peak (T_{peak} in Fig. 3 (b)) and the offset (T_{offset}) timings of the force profile were optimized, while the peak force magnitude was maintained at 40% of the wearer's body weight. Parameter constraints were defined as follows:

$$15\%GC_{MHF} \leq T_{peak} \leq 50\%GC_{MHF}$$
$$30\%GC_{MHF} \leq T_{offset} \leq 65\%GC_{MHF}$$
$$T_{offset} + 15\%GC_{MHF} \leq T_{peak}.$$

For the ankle assistance, the cable retraction onset timing (T_1 in Fig. 3 (b)) and the completion timing (T_2) were optimized, while the cable release onset timing (T_3) was set relative to T_2. The target peak force magnitude was set to 60% of the wearer's body weight. Parameter constraints were defined as follows:

$$30\%GC_{HS} \leq T_1 \leq 40\%GC_{HS}$$
$$50\%GC_{HS} \leq T_2 \leq 60\%GC_{HS}$$
$$T_1 + 10\%GC_{HS} \leq T_2$$
$$T_3 = T_2 + 10\%GC_{HS}.$$

3) Experimental protocol: Two healthy male adults with prior experience of walking with the exosuit assistance participated in this study (S1: 48 yrs, 85 kg, 178 cm; S2: 23 yrs, 75 kg, 183 cm). The experimental procedure was identical to the procedure in the single-joint optimization study except for an additional comparison with a slack (powered-off) condition. As for the fixed generic parameters for comparison, averaged parameters from [18] were used for the hip extension assistance, while averaged parameters from [7] were used for the ankle assistance.

Sbj	Cond.	Optimized parameter value				Metabolic reduction	
		Hip peak	Hip offset	Ankle onset	Ankle peak	Ref: Slack	Ref: No-suit
1	OPT	21.3	45.3	44.9	55.9	36%	39%
	FIX	24.9	42.4	42.5	57.5	35%	38%
2	OPT	26.4	41.4	50.0	60.0	36%	33%
	FIX	24.9	42.4	42.5	57.5	36%	32%

TABLE II: Multi-joint optimization results

4) Result: The optimization ran 26 iterations within 50 minutes with 11 early estimator stops (subject 1) and 20 iterations within 32 minutes with 8 early estimator stops (subject 2), excluding warm-up periods. A summary of the metabolic reduction can be found in Table II.

V. Conclusion

In this work, we introduced an early stopping process to address the limited time budget for human-in-the-loop optimization. By coupling the metabolic cost estimation with the optimization process, the optimization allowed up to 37 iterations in 50 minutes, which would have required 80 minutes using Bayesian optimization with fixed estimation time budget [10], [18]. When the optimized parameters were used, participants reduced their metabolic cost to below, or approximately equal to, the generic condition. Our preliminary results lend support to the hypothesis that an estimator stopping process can make more efficient use of experimental time compared to approaches that use fixed estimation intervals while sampling, opening the door to explore richer parameters spaces in human subject experiments.

We note that while our approach may support modest increases to the parameter space dimensionality, performing HIL optimization in high-dimensional (i.e. dozens or more) spaces remains intractable in the absence of prior information that could otherwise bias search in promising directions. Also, while our results demonstrate significant reductions in metabolic cost for both devices, our results are based on $N = 2$ subjects and experimental protocol does not rule out possible mitigating factors such as human adaptation over the duration of the experiment. In future work, we aim to collect a larger set of human subject data and add explicit mechanisms that model the nonstationarity of metabolic cost over time.

Acknowledgements

The authors would like to thank Nikos Karavas, Chih-Kang Chang, Asa Eckert-Erdheim, Maria Athanassiu, Brice Mikala Iwangou, Nicolas Menard, and Sarah Sullivan for their contributions to this work. J. Kim and S. Lee also appreciate the financial support from the Samsung Scholarship. This work was also supported in part by the Technology for Equitable and Accessible Medicine (TEAM) initiative at Harvard.

REFERENCES

[1] P. Malcolm, W. Derave, S. Galle, and D. De Clercq, "A simple exoskeleton that assists plantarflexion can reduce the metabolic cost of human walking," *PLoS one*, vol. 8, no. 2, p. e56137, 2013.

[2] S. H. Collins, M. B. Wiggin, and G. S. Sawicki, "Reducing the energy cost of human walking using an unpowered exoskeleton," *Nature*, vol. 522, pp. 212–215, 2015.

[3] L. M. Mooney and H. M. Herr., "Biomechanical walking mechanisms underlying the metabolic reduction caused by an autonomous exoskeleton." *J. NeuroEng. Rehabil.*, vol. 13, no. 1, p. 4, 2016.

[4] B. Quinlivan, S. Lee, P. Malcolm, D. Rossi, M. Grimmer, C. Siviy, N. Karavas, D. Wagner, A. Asbeck, I. Galiana, and C. Walsh, "Assistance magnitude versus metabolic cost reductions for a tethered multiarticular soft exosuit." *Science Robotics,*, vol. 2, no. 2, p. eaah4416, 2017.

[5] F. A. Panizzolo, I. Galiana, A. T. Asbeck, C. Siviy, K. Schmidt, K. G. Holt, and C. J. Walsh, "A biologically-inspired multi-joint soft exosuit that can reduce the energy cost of loaded walking." *J. NeuroEng. Rehabil.*, vol. 13, no. 1, p. 43, 2016.

[6] S. Lee, J. Kim, L. Baker, A. Long, N. Karavas, N. Menard, I. Galiana, and C. J. Walsh, "Autonomous multi-joint soft exosuit with augmentation-power-based control parameter tuning reduces energy cost of loaded walking," *J. NeuroEng. Rehabil.*, 2018.

[7] G. Lee, Y. Ding, I. G. Bujanda, N. Karavas, Y. M. Zhou, and C. J. Walsh, "Improved assistive profile tracking of soft exosuits for walking and jogging with off-board actuation," in *Proc. Int. Conf. Intel. Rob. Sys.* IEEE, 2017, pp. 1699–1706.

[8] J. M. Caputo and S. H. C. Collins, "A universal ankle-foot prosthesis emulator for experiments during human locomotion," *J. Biomech. Eng.*, vol. 136, p. 035002, 2014.

[9] S. Lee, N. Karavas, B. T. Quinlivan, D. L. Ryan, D. Perry, A. E. Erdheim, P. Murphy, T. G. Goldy, N. Menard, M. Athanassiu, J. Kim, G. Lee, I. Galiana, and C. J. Walsh, "Autonomous multi-joint soft exosuit for assistance with walking overground," *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2018.

[10] Y. Ding, I. Galiana, A. T. Asbeck, S. M. M. D. Rossi, J. Bae, T. R. T. Santos, V. L. de Araujo, S. Lee, K. G. Holt, and C. Walsh, "Biomechanical and physiological evaluation of multi-joint assistance with soft exosuits." *Trans. Neural Syst. Rehabil. Eng.*, vol. 25, no. 2, pp. 119–130, 2017.

[11] M. Kim and S. H. Collins, "Once-per-step control of ankle-foot prosthesis push-off work reduces effort associated with balance during walking," *J. NeuroEng. Rehabil.*, vol. 12, no. 1, p. 43, 2015.

[12] ——, "Step-to-step ankle inversion/eversion torque modulation can reduce effort associated with balance," *Front. Neurorobotics*, vol. 11, p. 62, 2017.

[13] S. Lee, S. Crea, P. Malcolm, I. Galiana, A. Asbeck, and C. Walsh, "Controlling negative and positive power at the ankle with a soft exosuit," in *Proc. Int. Conf. Rob. Autom.*, 2016, pp. 3509–3515.

[14] J. M. Caputo, P. G. Adamczyk, and S. H. Collins, "Informing ankle-foot prosthesis prescription through haptic emulation of candidate devices," in *Proc. Int. Conf. Rob. Autom.*, 2015, pp. 6445 – 6450.

[15] R. E. Quesada, J. M. Caputo, and S. H. Collins, "Increasing ankle push-off work with a powered prosthesis does not necessarily reduce metabolic rate for transtibial amputees," *J. Biomech.*, vol. 49, no. 14, pp. 3452–3459, 2016.

[16] J. R. Koller, D. H. Gates, D. P. Ferris, and C. D. Remy, "'body-in-the-loop optimization of assistive robotic devices: A validation study," in *Proc. Rob: Sci. Sys. Conf.*, 2016.

[17] J. Zhang, P. Fiers, K. A. Witte, R. W. Jackson, K. L. Poggensee, C. G. Atkeson, and S. H. Collins, "Human-in-the-loop optimization of exoskeleton assistance during walking." *Science*, vol. 356, no. 6344, pp. 1280–1284, 2017.

[18] Y. Ding, M. Kim, S. Kuindersma, and C. J. Walsh, "Human-in-the-loop optimization of hip assistance with a soft exosuit during walking," *Sci. Robot.*, vol. 3, no. 15, 2018.

[19] M. Kim, Y. Ding, P. Malcolm, J. Speeckaert, C. J. Siviy, C. J. Walsh, and S. Kuindersma, "Human-in-the-loop bayesian optimization of wearable device parameters," *PLoS One*, vol. 12, no. 9, p. e0184054, 2017.

[20] J. C. Selinger and J. M. Donelan., "Estimating instantaneous energetic cost during non-steady-state gait," *J. Appl. Physiol.*, vol. 117, no. 11, pp. 1406–1415, 2014.

[21] W. Felt, J. C. Selinger, J. M. Donelan, and C. D. Remy, ""Body-In-The-Loop": Optimizing Device Parameters Using Measures of Instantaneous Energetic Cost," *PLoS One*, vol. 10, no. 8, p. e0135342, 2015.

[22] T. Wai and T. Langer, "Mitochondrial dynamics and metabolic regulation." *Transmission Electron Microscope*, vol. 27, no. 2, pp. 105–117, 2016.

[23] V. M. C. Brochu, Eric and N. De Freitas, "A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning." *Technical Report*, pp. TR–2009–23, 2010.

[24] H. J. Kushner, "A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise." *J. Basic. Eng.*, vol. 86, no. 1, pp. 97–106, 1964.

[25] V. T. Mockus, J. and A. Zilinskas, "The application of bayesian methods for seeking the extremum." in *Toward Global Optimization*, vol. 2. Elsevier, 1978, p. 117128.

[26] 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*, vol. 104, no. 1, pp. 148–175, 2016.

[27] C. E. Rasmussen, "Gaussian processes in machine learning," in *Summer School on Machine Learning*. Springer, 2003, pp. 63–71.

[28] D. J. Lizotte, "Practical bayesian optimization," PhD dissertation, University of Alberta, 2008.

[29] S. J. Julier and J. K. Uhlmann, "New extension of the kalman filter to nonlinear systems," in *Signal processing, sensor fusion, and target recognition VI*, vol. 3068. International Society for Optics and Photonics, 1997, pp. 182–194.

[30] T. Ferguson, *Optimal Stopping and Applications*, 2006. [Online]. Available: https://www.math.ucla.edu/ tom/Stopping/Contents.html

[31] G. Lee, J. Kim, F. A. Panizzolo, Y. M. Zhou, L. M. Baker, P. Malcolm, and C. J. Walsh, "Reducing the metabolic cost of running with a tethered soft exosuit," *Sci.Robot.*, vol. 2, eaan6708, 2017.

[32] Y. Ding, F. A. Panizzolo, C. Siviy, P. Malcolm, I. Galiana, K. G. Holt, and C. J. Walsh, "Effect of timing of hip extension assistance during loaded walking with a soft exosuit," *J. NeuroEng. Rehabil.*, vol. 13, no. 1, p. 87, 2016.