

MEDICAL ROBOTS

Human-in-the-loop optimization of hip assistance with a soft exosuit during walking

Ye Ding,^{1,2*} Myunghee Kim,^{1,2*} Scott Kuindersma,^{1†} Conor J. Walsh^{1,2†}

Wearable robotic devices have been shown to substantially reduce the energy expenditure of human walking. However, response variance between participants for fixed control strategies can be high, leading to the hypothesis that individualized controllers could further improve walking economy. Recent studies on human-in-the-loop (HIL) control optimization have elucidated several practical challenges, such as long experimental protocols and low signal-to-noise ratios. Here, we used Bayesian optimization—an algorithm well suited to optimizing noisy performance signals with very limited data—to identify the peak and offset timing of hip extension assistance that minimizes the energy expenditure of walking with a textile-based wearable device. Optimal peak and offset timing were found over an average of 21.4 ± 1.0 min and reduced metabolic cost by $17.4 \pm 3.2\%$ compared with walking without the device (mean \pm SEM), which represents an improvement of more than 60% on metabolic reduction compared with state-of-the-art devices that only assist hip extension. In addition, our results provide evidence for participant-specific metabolic distributions with respect to peak and offset timing and metabolic landscapes, lending support to the hypothesis that individualized control strategies can offer substantial benefits over fixed control strategies. These results also suggest that this method could have practical impact on improving the performance of wearable robotic devices.

INTRODUCTION

Wearable robotic devices have demonstrated the potential to enhance human economy and endurance (1–3). Recent breakthroughs in wearable robotics have substantially reduced energy expenditure in human walking using both passive (1) and active autonomous (2–4) or tethered (5–11) devices. In particular, advances in active devices provided flexibility to regulate assistance control parameters related to timing (5, 7, 10), magnitude (6, 12), or delivered power (12, 13). Studies have shown that control strategies can significantly affect performance (4–7, 10, 12), which raises questions about how to reliably and efficiently design optimal controllers. Assistive strategies have commonly been derived from simulations (14, 15) and biomechanical measurements (10, 16) or tuned manually based on average responses (11). Specifically, there is a growing interest in designing control strategies using musculoskeletal simulations (14, 17), and recently, this approach has shown promise in guiding assistive profiles for running (14, 18). However, physiological and neurological differences between individuals can cause divergent responses to an identical controller, that is, one participant's optimal control strategy may perform poorly on another (5–7, 10, 19). Thus, although generic musculoskeletal simulations may provide general guidelines on assistance, participant-specific models may be required when considering how to find optimal system parameters for individualized assistance.

Conventionally, discrete step (1, 4–7) and continuous sweep (20, 21) protocols have been used to investigate a participant's performance and to explore the landscape of control parameter settings for wearable robotic devices. With these approaches, metabolic cost is measured by varying a control parameter in either a discrete or a continuous manner. A curve fitting process is then followed to identify the optimal parameter value that results in the maximum metabolic benefit (1, 5–7, 20, 21). Unfortunately, both continuous and discrete protocols involve a lengthy evaluation process, and the time required increases exponentially with

added control parameter dimensions. Long walking times in protocols may affect the accuracy of metabolic measurements due to high exertion or fatigue, which in turn leads to cardiopulmonary drift (22), especially for clinical populations who may not be able to sustain long walking bouts (23).

Human-in-the-loop (HIL) optimization aims to address the aforementioned challenges in protocol length by adjusting control parameters based on real-time measurements of human physiological signals, such as metabolic cost. This optimization is inspired by an observation of humans continuously adjusting their coordination pattern to minimize the metabolic cost of walking (24) and expands the concept to wearable devices. Some promising efforts in this domain have recently demonstrated the ability to optimize both single and multiple control parameters using ankle exoskeletons (8, 9). In these cases, substantial metabolic reductions were achieved with the optimal parameter settings identified by either a one-dimensional (1D) gradient descent method using a pneumatically actuated ankle exoskeleton for a fixed 50 min (9) or a 4D Covariance Matrix Adaptation Evolution Strategy (CMA-ES) with an electromechanically actuated ankle exoskeleton for 83 ± 14 min (mean \pm SEM) (8). Although these achievements are impressive, there remain opportunities to explore different wearable assistive hardware, control parameterizations, applications to other joints, and alternative optimization methods that could improve sample efficiency.

We developed an experimental method to rapidly identify optimal control parameters in a 2D space that minimized the metabolic cost of walking (Fig. 1). This was achieved through the use of Bayesian optimization, an efficient global optimization strategy that is well suited to find the minima of objective functions that are noisy and expensive to evaluate (25–27). In a previous HIL study that optimized step frequency, we found that Bayesian optimization converged in half the time of a gradient descent method (28). For the current HIL scheme, a participant walked with hip extension assistance applied via a soft exosuit (Fig. 2A), a textile-based wearable device designed to apply forces across joints in parallel with human muscles (10, 29). The assistive profile was configured by multiple control parameters that were iteratively updated by the optimization and applied to the participant using a tethered actuation system with admittance force control (29). The optimization was initialized by

¹John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA. ²Wyss Institute for Biologically Inspired Engineering, Harvard University, Cambridge, MA 02115, USA.

*These authors contributed equally to this work.

†Corresponding author. Email: walsh@seas.harvard.edu (C.J.W.); scottk@seas.harvard.edu (S.K.)

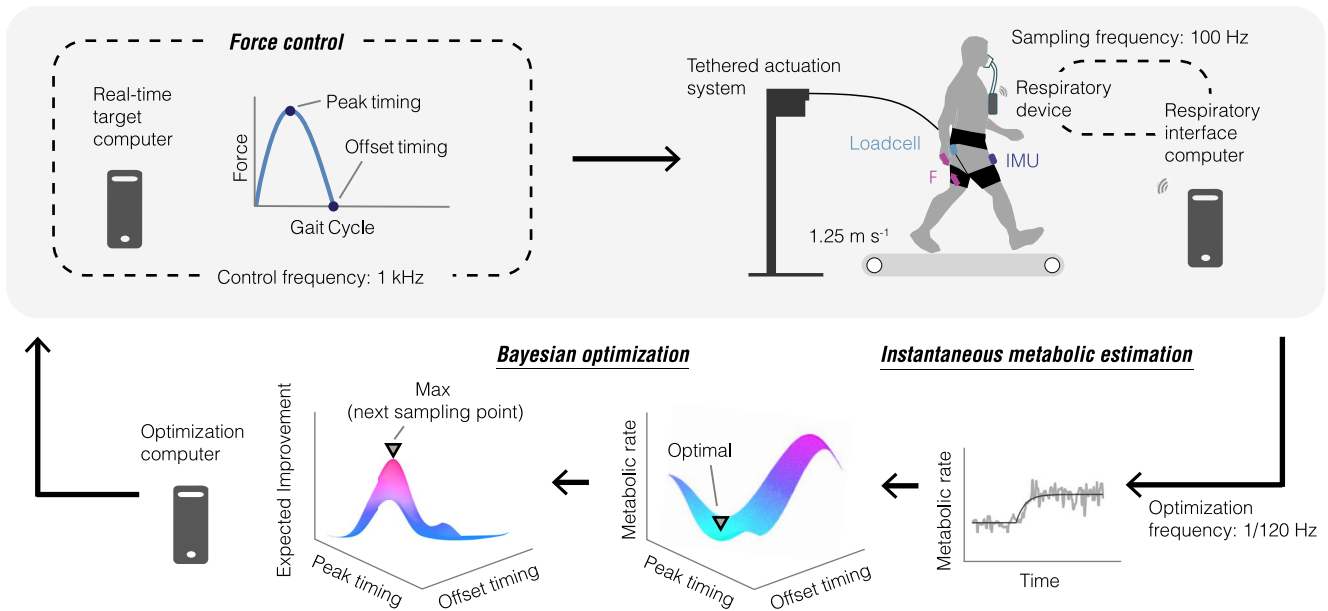


Fig. 1. Experimental setup for HIL Bayesian optimization. Bayesian optimization was used to adjust the control parameters of an assistive device to minimize the metabolic cost of walking. The metabolic rate was estimated from respiratory measurements and used to compute a posterior distribution of metabolic rate with respect to the free control parameters. The posterior was initially generated by evaluating six prefixed control parameters. Given the posterior at the current iteration, the control parameters with maximum EI were chosen and applied to the wearable device. This process was repeated until convergence. During this process, the configured force profiles were delivered through a soft hip exosuit with a tethered actuation system.

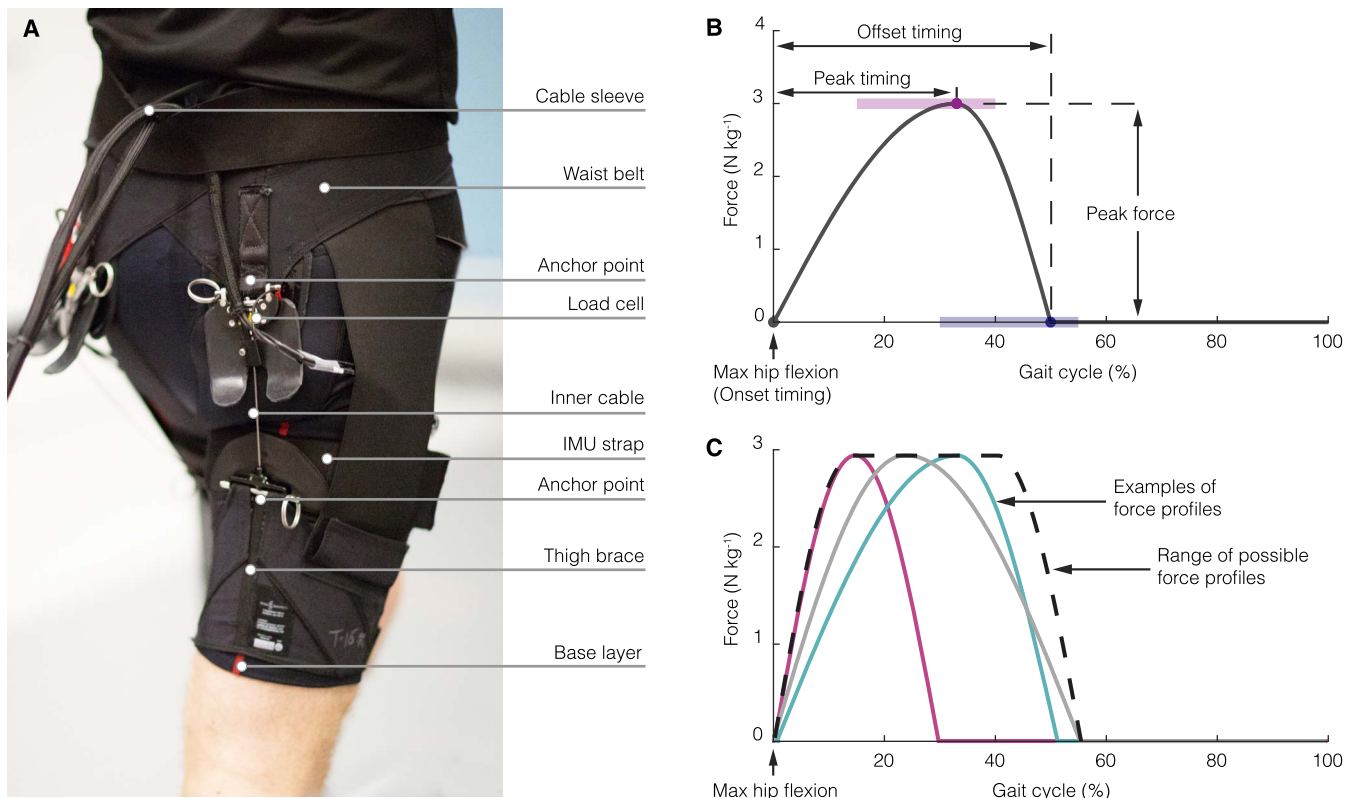


Fig. 2. Soft exosuit and assistive hip force profile. (A) The hip soft exosuit. A hip extension moment was generated by pulling the inner cable to create a tension between two anchor points. (B) Parameterization of hip force profile. The hip force profile was chosen to be a combination of two parameterized sinusoidal curves joined at the peak. Peak force was set to 30% of body weight, and onset timing was fixed to the time of maximum hip flexion. Peak and offset timing were actively adjusted by the optimization to determine the shape of the force profile as a function of gait percentage. Shaded purple and blue bars represent the range of peak and offset timing, respectively. (C) Examples of feasible hip force profiles.

obtaining the metabolic cost of a prescribed number of assisted conditions with respect to pseudo-randomly selected control parameters from an evenly distributed parameter space. On the basis of this information, the optimization iteratively estimated the participant's metabolic cost distribution using a Gaussian process (fig. S1) (30) and selected control parameters for the next iteration by maximizing the expected improvement (EI) (25, 26). At each iteration, the metabolic cost was estimated by fitting a first-order dynamic model to 2 min of transient metabolic data (31). After a set number of iterations, the control parameters corresponding to the minimum value of the metabolic landscape (the mean of the metabolic cost distribution) represented the optimal values.

The hip assistive profile was a combination of two halves of sinusoidal curves joined at their peaks. This profile was defined by two fixed parameters (peak force and onset timing) and two free parameters (peak timing and offset timing; Fig. 2B) that were adjusted by the optimization method. We fixed peak force to 30% body weight to ensure comfort during the long walking test while still maintaining assistance high enough to achieve sufficient metabolic reduction. Previous work demonstrated that higher assistance magnitude resulted in larger metabolic benefits for both hip (10, 11) and ankle (6), and the forces we evaluated here were approximately within the range of previous evaluations. We also fixed the onset timing at the maximum hip flexion event based on previous hip studies, which showed the largest metabolic reductions with onset timing set close to maximum hip flexion (7, 10). Further, we found less intraparticipant variability of metabolic cost and lower signal-to-noise ratio when varying onset timing in our pilot testing (table S1). For the purposes of this study, we defined the start of the gait cycle using the maximum hip flexion event. Peak and offset timing were bounded within 15 to 40% and 30 to 55%, respectively, in this newly defined gait cycle. It is worth noting that the maximum hip flexion event was on average at 86.2% of the conventional gait cycle, defined with heel strike as 0% (table S2). The offset timing was constrained to occur at least 15% later than the peak timing. The range and constraint of peak and offset timing (Fig. 2C) were chosen by slightly extending the average range of the biological hip extension moment (32) while considering limitations on the ramp-up speed of assistance that our soft exosuit was capable of achieving. This configuration was able to shape our profiles similar to the hip assistive profiles used in previous assistive device studies (4, 7, 10, 11).

We conducted a single-day experiment on eight participants (table S3), optimizing the assistance timings as they walked on a treadmill at 1.25 m s^{-1} . For the optimization, 6 iterations (six pairs of peak and offset timing) were evaluated for the optimization initialization and 14 iterations were followed to adjust the tuning. These numbers of iterations were chosen based on our simulation results (fig. S2). After optimization, we performed a validation test confirming the optimal condition found during the optimization process, then compared both optimal condition and validation test with a no-suit condition. The primary analysis included (i) the net metabolic cost of walking, defined as the gross metabolic rate during walking minus the rate measured during quiet standing; (ii) the convergence time across participants; (iii) participant-specific optimal timings; (iv) participant-specific optimal assistive profiles; (v) participant-specific metabolic landscape, defined as the mean of the metabolic cost distribution with respect to the peak and offset timing along with participant-specific probability of improvement landscape, interpreted as the likelihood of exceeding the largest metabolic reduction.

RESULTS

Metabolic rate

Participant-specific optimal assistance substantially improved energy economy for all participants by reducing the net metabolic cost of walking to $2.26 \pm 0.13 \text{ W kg}^{-1}$ and $2.27 \pm 0.18 \text{ W kg}^{-1}$ for the optimal and validation conditions, respectively, from $2.75 \pm 0.18 \text{ W kg}^{-1}$ for the no-suit condition (mean \pm SEM). Net metabolic reduction of the validation condition ranged from 6.7 to 33.9%, with an average reduction of $17.4 \pm 3.2\%$ (mean \pm SEM; paired t test, $P = 0.003$; Fig. 3A and table S4).

Convergence time

The optimization converged for all participants during the optimization process (fig. S3). The convergence time was on average 21.4 ± 1.0 min (mean \pm SEM), ranging from 18 to 24 min.

Optimal timing

Participant-specific optimal peak and offset timings spread over about half of the feasible region of the control parameters (Fig. 3B). Most of the optimal timings were on the boundaries of the parameter ranges, with three participants having their optima at the latest peak and offset timing.

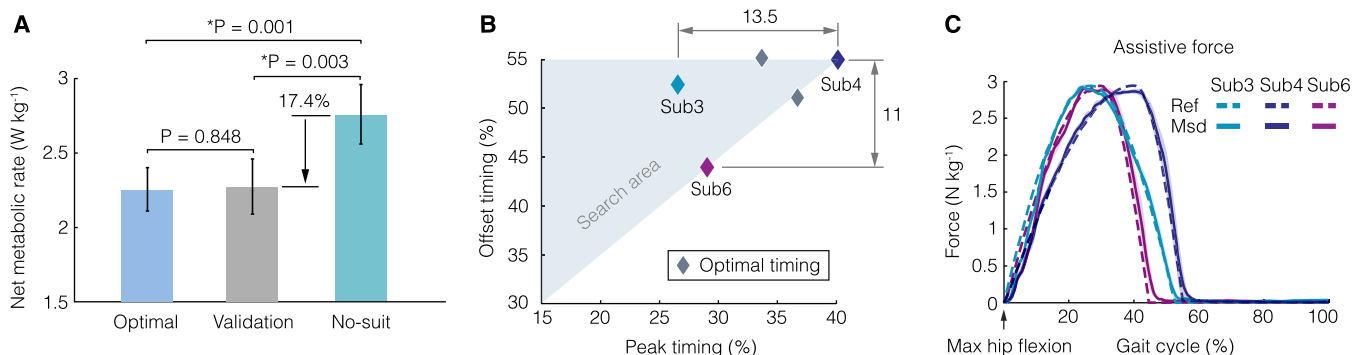


Fig. 3. Experimental results. (A) The net metabolic rate for each condition. Optimal: Minimum mean value of the posterior distribution (metabolic landscape). Validation: Metabolic rate of 5-min walking with optimized assistance. No-suit: Metabolic rate of 5-min walking with a regular pair of pants. Bars are means, error bars are SEMs, and asterisks denote statistical significance. (B) Feasible parameter region and optimal timing values for all participants. Optimal timings were varied across participants, and three participants shared the same optimal timings at the latest peak and offset timing. (C) Optimal assistive force profiles for participants 3, 4, and 6. Dashed and solid lines are reference and measured forces normalized by body mass, averaged across 10 strides during the last minute of the validation condition. The maximum hip flexion event was used to initialize the gait cycle in this study.

Optimal assistive force profile

For the validation condition, the averaged delivered peak force was 215.6 ± 10.1 N (2.84 ± 0.02 N kg⁻¹, mean \pm SEM). The average root mean square error of the optimal assistive force tracking of the validation condition was 4.1%. For a clear representation, only three representative optimal force tracking samples with the most different optimal timings from the validation condition are shown in Fig. 4C, whereas all optimal force profiles are shown in fig. S4.

Metabolic landscape and probability of improvement landscape

The representative participant-specific metabolic landscapes (Fig. 4, A to C) further illustrated the interparticipant variability with respect to timings. The participants' metabolic landscapes, represented as Gaussian process posteriors, showed substantial visual differences. To quantitatively summarize the differences between the participants' metabolic landscapes, we computed the probability that each participant's optimal parameters would reduce the metabolic cost of other participants according to each participant's posterior landscape (Fig. 4, D to F). This analysis suggested that, in general, one participant's optimal peak and offset timing were likely to be suboptimal for another.

DISCUSSION

With the optimized hip extension assistance obtained from HIL Bayesian optimization, the average net metabolic reduction was 17.4% compared with walking without the device. Using a similar hip exosuit to assist loaded walking, our group previously showed an average reduction of 8.5% compared with an unpowered condition (10). Another study with a tethered hip exoskeleton using pneumatic actuators demonstrated

average metabolic reductions of 10.3 and 9.7% when assisting either hip extension or hip flexion, respectively, compared with an unpowered condition (7). Last, a recent study assisting both hip extension and flexion simultaneously with an autonomous electromechanical hip exoskeleton reported an average metabolic reduction of 21.1% when compared with walking without an exoskeleton (4). The result of this study suggests that substantial metabolic reductions can be achieved by solely assisting hip extension with optimized assistance and indicates the potential improvement of assisting both flexion and extension with hip assistive devices.

The average convergence time of our HIL Bayesian optimization was 21.4 min. Short convergence time could be important in some cases to mitigate widely observed inaccuracies stemming from cardiopulmonary drift and participant fatigue (22). This result also suggests that HIL Bayesian optimization could be applied to wearable devices designed for strenuous activities or clinical populations with limited physical strength—both cases where participant endurance is a limiting factor (23).

The variability shown in the optimized assistance profiles demonstrates the importance of individualization. The participant-specific metabolic landscapes and the probability of improvement generated by the Bayesian optimization further illustrate the interparticipant variability with respect to timings.

The nonparametric HIL Bayesian optimization was more effective than the model-based naïve grid search. We illustrated this problem by generating quadratic approximations with the first 10 iterations of the data from the optimization process, which was the average amount of data needed for the convergence of Bayesian optimization (table S5). The comparison showed that the model-based naïve grid search made unreasonable estimates of optimal parameter value in the high measurement noise environment.

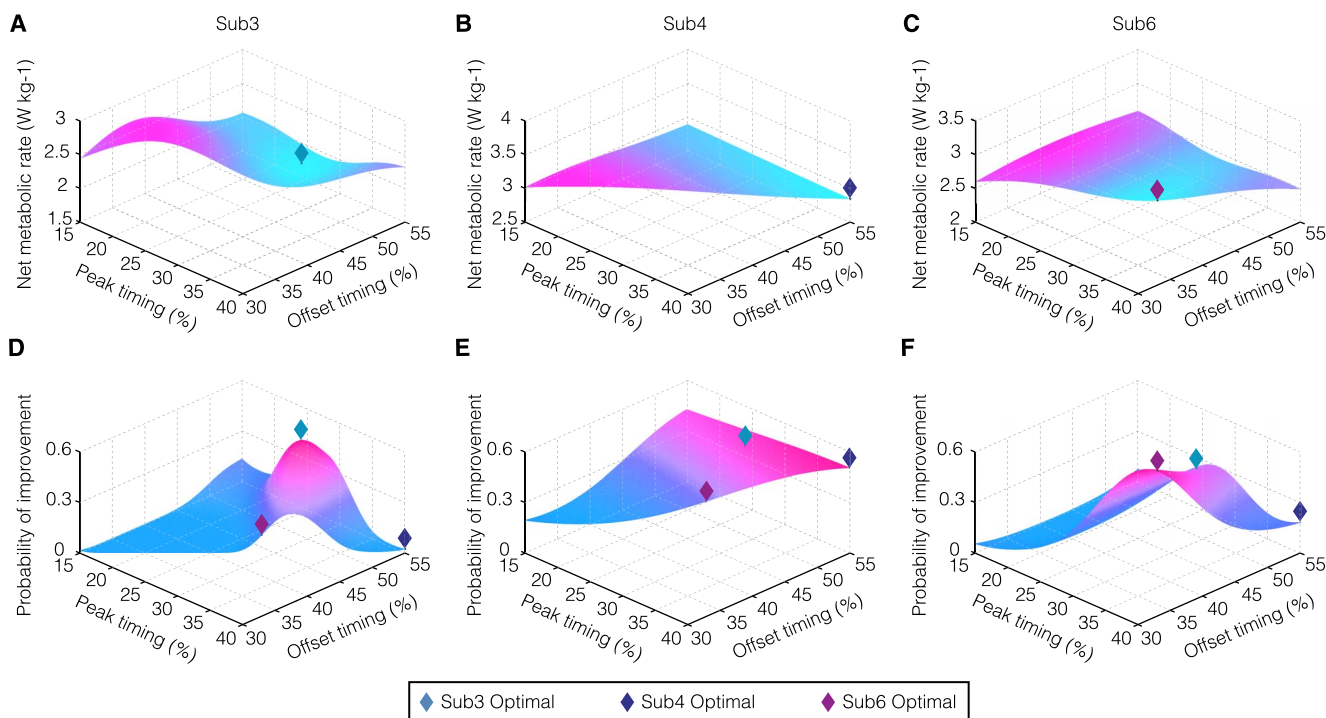


Fig. 4. Participant-specific metabolic landscape and probability of improvement landscape. (A to C) Metabolic landscapes (the mean of the metabolic cost posterior distribution with respect to peak and offset timing) for participants 3, 4, and 6. Diamonds indicate the locations of participant-specific optimal timings. (D to F) Probability of improvement landscapes (capturing the probability of reducing metabolic cost beyond the identified optimal) for participants 3, 4, and 6.

The optimized assistance did not maximize the duration of force to maximize the positive mechanical power for the hip joint. This may be partially because assistance with a late offset timing may hinder hip flexion (32). However, most of the optimal timings for participants were on the boundaries of the parameter ranges, which may suggest that, with a larger parameter search area, further reductions in metabolic cost could be obtained. Currently, the selected parameter range was constrained by the limited ramp-up speed of assistance with our exosuit and a cautious approach to ensure that the assistance profile did not greatly exceed limits of the average range of biological hip extension moment (32). To reduce the likelihood that the optimal values are caught on the boundaries, future studies could expand the feasible parameter range by improving the exosuit stiffness to increase the ramp-up speed of the assistance and having participant-specific search areas based on training performance.

Another limitation of the current optimization is the lengthy sampling time for each measurement, which could prevent straightforward extension to higher dimensional parameterizations. It may be beneficial to add additional flexibility to the optimization not only to choose the exploration points but also to adjust the length of sampling time (33). In addition to adaptive sampling time, it may be useful to use musculo-skeletal models to provide an initial estimation of the metabolic landscape, which could reduce the number of samples required to find low-cost parameters. In addition, the smoothness and regularity assumptions imposed by the Gaussian process kernel function may not be valid for all metabolic landscapes and wearable devices, but in our experiments, these landscapes were well approximated using a squared exponential kernel and a single global noise parameter. Last, because Bayesian optimization uses all available data to compute the posterior metabolic distribution and acquisition function, additional methods such as “data forgetting” would have to be used to deal with human adaption effects (34).

HIL optimization holds promise to improve the performance of wearable robotic devices for a wide range of tasks. The presented method shows a substantial metabolic reduction and suggests the possibility of optimizing wearable devices using low-dimensional control parameterization. The short convergence time would enable researchers to apply this method to individualize control parameters during strenuous tasks or for people with limited physical strength or endurance. Using a noisy respiratory signal as the objective function of the optimization indicates that this method can be applied to other alternate physiological or biological signals, such as using kinematic symmetry to optimize wearable devices for poststroke patients or using balance-related measurements to optimize prostheses. The participant-specific metabolic landscapes and probability of improvement landscapes demonstrate the significant variability between participants and suggest that participant-specific optimal timing provides the highest probability of achieving the largest metabolic reduction, further highlighting the benefit of individualization.

MATERIALS AND METHODS

Experimental design

This was a single-day protocol without training sessions. To minimize the effects of adaptation, we recruited eight participants who had previous experience walking with the exosuit at least two times before. Participants walked without load on an instrumented treadmill (Bertec) at 1.25 m s^{-1} wearing a respiratory measurement device (COSMED; fig. S5). These conditions were chosen partially to lessen fatigue effects of the relatively long walking protocol, and the constant walking speed

allowed the comparisons between studies (7, 10, 11). Each participant went through five conditions (fig. S6): (i) a 5-min quiet standing condition, (ii) a 5-min no-suit condition, (iii) a 40-min optimization condition intersected by two 3-min warm-up periods and 5-min rest periods, (iv) a 5-min validation condition with the optimal timing, and (v) a 5-min no-suit condition. Both warm-up periods were assisted walking with the same assistive profiles used in the follow-up iteration of the optimization condition. During the no-suit condition, participants walked with a regular pair of pants (mass, 715 g), which was chosen to assess the metabolic benefits from walking with active assistance to walking with normal clothes, similar to configuration in our previous hip assistance study (10). Resting breaks were given between all conditions besides the break during the optimization condition. Considering the relatively long walking time (61 min), two no-suit conditions were designed at both the beginning and the end of the protocol as a visual check of the possible fatigue reported by the participants.

Participants

Eight healthy male adults ($n = 8$; age, 30.3 ± 7.1 years; mass, 76.5 ± 8.9 kg; height, 1.77 ± 0.05 m; mean \pm SD; table S3) participated in this study. Sample size was chosen based on the data from previous studies (10, 11). The study was approved by the Harvard Longwood Medical Area Institutional Review Board, and all methods were carried out in accordance with the approved study protocol. All participants provided written informed consent before their participation and after the nature and the possible consequences of the studies were explained.

Soft exosuit

The soft exosuit used in this study was designed to solely assist hip extension. The textile components of the hip exosuit consisted of a spandex base layer (mass, 181 g), a waist belt (mass, 275 g; fig. S7), two thigh braces (mass, 2×69 g; fig. S8), and two elastic straps (mass, 2×46 g) for mounting inertial measurement units (IMUs; mass, 2×13 g). Bowden cables and sensor wires including expandable braided cable sleeves for each leg (mass, 2×328 g) were tied together at the waist and connected to the actuation platform. The participant supported about half of the weight of the Bowden cable assembly. All textile components (size medium) and half of the weight of the Bowden cable assembly had a total mass of 0.859 kg. The stiffness evaluation of the soft exosuit used in this study is shown in (29).

Actuation platform

A tethered actuation system with two modular actuators was used to generate assistive forces. Each actuator consisted of one customized frameless brushless motor (Allied Motion), a customized spiroid gear set (ITW Heartland), a 90-mm-diameter pulley, and other supportive structures (29). Bowden cable was used to transmit the force from the actuator to the hip joint. On the actuator side, the Bowden cable sheath connected to the frame of the pulley cover and the inner cable attached to the pulley. On the exosuit side, the Bowden cable sheath connected to the anchor point on the bottom of the waist belt and the inner cable connected to the anchor point on the top of the thigh piece. When the actuator retracts, the distance between the two anchor points is shortened, generating a force to assist hip extension.

Sensing and control

Two IMUs (VN-100 Rugged IMU, VectorNav Technologies) attached to the front of each thigh detected the maximum thigh flexion angle to

segment the stride (10, 35). Stride time was measured as the time between two consecutive maximum hip flexion events (35). By using the average stride time from the previous two steps, the reference force profile was scaled for each stride. The actual force signal was measured by two load cells (LSB200, FUTEK Advanced Sensor Technology) placed in series with the Bowden cables on each leg. Combined with the actuator position signals measured by the encoders (AS5134, Ams) mounted on the back of the customized brushless motors, an admittance controller with feedforward models was implemented to track the force profiles with different peak and offset timings. The detailed controller design, frequency response, and force tracking evaluation with different ramping speeds are presented in (29).

Instantaneous metabolic estimation

The metabolic rate was estimated by fitting a first-order dynamic model to 2 min of transient metabolic data (21). The mathematical representation in the frequency domain takes the form

$$Z(s) = H(s)R(s) \quad (1)$$

where $Z(s)$ is the measured metabolic cost, $R(s)$ is the instantaneous metabolic cost f^{inst} in frequency domain, and $H(s)$ is the first-order dynamic model $H(s) = 1/(\tau s + 1)$ with a time constant $\tau = 42$ s (31). In the discrete-time domain, Eq. 1 can be written as

$$z(i+1) = \frac{(\tau - dt(i))}{\tau} z(i) + \frac{dt(i)}{\tau} f^{\text{inst}} \quad (2)$$

where i is the number index of the measured breath and $dt(i)$ is the time duration between the i th and $(i+1)$ th breath. After measuring z and dt for 2 min, we obtained f^{inst} by first calculating the change of the instantaneous metabolic rate from the last condition and then minimizing the error between the model estimation and measurements using least squares (21).

Bayesian optimization

Bayesian optimization is an efficient global optimization method that is particularly well suited to optimizing unknown objective functions that are expensive to evaluate (25–27, 36). It takes advantage of the information provided by the time history by computing a posterior distribution of cost as a function of the optimization variables and then using acquisition functions computed on this posterior to select the next points to evaluate. A prior belief over the objective function distribution is defined using mean and covariance functions. The posterior distribution of the objective function is iteratively computed in closed form when new data become available. Using this model, the algorithm balances exploitation with uncertainty reduction to guide exploration (37).

In our study, we initialized the optimization by evaluating instantaneous metabolic cost f^{inst} for six iterations with different pairs of prefixed peak and offset timing, which were pseudo-randomly selected from evenly spaced timing intervals (fig. S9). This initialization is a common practice to avoid biased sampling that could lead to premature convergence (25). After initial evaluation, the optimization calculated the metabolic landscape, $f(\mathbf{x})$, using Gaussian processes (30, 34), where the parameter $\mathbf{x} = [x_p, x_o]$ consisted of peak and offset timing. Given the calculated landscape, the next sampling timing was selected by maximizing EI, which naturally balances exploration and exploitation (25, 26). With the metabolic rate of the newly sampled timing added to the data set, the metabolic landscape was refined again for selecting the next sampling timing.

This process was repeated for 14 iterations. In total, there were 20 iterations in the optimization process including 6 iterations of initialization, and fig. S10 shows one sample optimization process described above on iterations 6, 7, and 20.

The metabolic landscape, $f(\mathbf{x})$, was modeled using a Gaussian process. The prior of the Gaussian process is represented by mean, $\mu(\mathbf{x})$, and covariance, $k(\mathbf{x}, \mathbf{x}')$, functions. As is standard practice, we used zero mean and the anisotropic squared exponential kernel for the covariance function (25),

$$k(\mathbf{x}, \mathbf{x}') = \sigma^2 \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}')^T M (\mathbf{x} - \mathbf{x}')\right) \quad (3)$$

where σ^2 is the metabolic rate (signal) variance and M is a diagonal matrix consisting of the length scale parameters of peak and offset timing, l_1 and l_2 . Intuitively, the signal variance captures the overall magnitude of the cost function variation, and the length scales capture the sensitivity of the metabolic rate with respect to changes in peak and offset timing. Assuming that metabolic cost has an additive, independent, and identically distributed noise, the samples can be expressed as

$$f^{\text{inst}}(\mathbf{x}) = f(\mathbf{x}) + \varepsilon, \varepsilon \sim N(0, \sigma_{\text{noise}}^2) \quad (4)$$

where σ_{noise}^2 is the noise variance. Given the Gaussian process prior and data set \mathbf{D} , the posterior metabolic cost distribution f_*^{inst} was calculated for a parameter \mathbf{x}_* as $f_*^{\text{inst}}(\mathbf{x}_*) \equiv f_*^{\text{inst}} \sim N(E[f_*^{\text{inst}}], s_*^2)$. The mean and variance are calculated as

$$E[f_*^{\text{inst}}] = \mathbf{k}_*^T (K + \sigma_{\text{noise}}^2 I)^{-1} \mathbf{y} \quad (5)$$

$$s_*^2 = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T (K + \sigma_{\text{noise}}^2 I)^{-1} \mathbf{k}_* \quad (6)$$

where $\mathbf{k}_* = [k(\mathbf{x}_1, \mathbf{x}_*), \dots, k(\mathbf{x}_n, \mathbf{x}_*)]^T$ and K is the positive definite kernel matrix, $[K]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$.

We optimized hyperparameters ($\theta = [\sigma \ l_1 \ l_2 \ \sigma_{\text{noise}}]$) at each iteration by maximizing log marginal likelihood of the data collected ($\mathbf{D} = \{\mathbf{X}, \mathbf{y}\}$, $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^T \in R^{N \times 2}$, $\mathbf{y} = [f_1^{\text{inst}}, \dots, f_n^{\text{inst}}]^T \in R^N$) using Matlab's `fmincon` function with 10 random initializations to avoid poor local minima.

The peak and offset timing, x_p , x_o , were selected by maximizing the expected reduction in the metabolic cost over the best timing previously assessed $\max(f_{\text{best}} - f_*^{\text{inst}}, 0)$ (25). EI, which balanced between predictive minimum points and high uncertainty (25, 27), took the following form

$$\text{EI}[\mathbf{x}_*] = (f_{\text{best}} - E[f_*^{\text{inst}}])\phi(u_*) + s_*\varphi(u_*) \quad (7)$$

where $f_{\text{best}} = \min_{i=1, \dots, N} E[f^{\text{inst}}(\mathbf{x}_i)]$, $u_* = (f_{\text{best}} - E[f_*^{\text{inst}}])/s_*$, and $\phi(\cdot)$ and $\varphi(\cdot)$ were the cumulative distribution function and probability density function of the normal distribution, respectively. The EI was set to zero when s_* was zero. At each iteration, the next sampling timing was selected by maximizing EI using Matlab's `fmincon` while enforcing the constraint that the offset timing be at least 15% later than the peak timing, $x_o - x_p \geq 15\%$. We again used 10 random restarts to avoid poor local minima. We note that, as the dimensionality increases, the number

of random restart points required to reliably maximize EI would likely need to increase.

Metabolic measurement and analysis

Respiratory data were collected throughout the protocol. Metabolic rates from the quiet standing, first no-suit, validation, and second no-suit conditions were calculated from the last 2 min of carbon dioxide and oxygen rates using a modified Brockway equation (38). For the optimization process, the instantaneous metabolic estimations for each 2-min measurement period were also collected. Net metabolic rate and net metabolic landscape were obtained by subtracting the quiet standing metabolic rate, then normalizing by each participant's body mass. The metabolic reduction of the validation condition was obtained by subtracting the net metabolic rate of the validation condition from the net metabolic rate of the second no-suit condition and then dividing the result by the net metabolic rate of the second no-suit condition. The second no-suit condition was chosen for the comparison of metabolic reduction because it is the closest no-suit condition to the validation condition. The metabolic reduction of the optimal condition was obtained with the same calculation by replacing the net metabolic rate of the validation condition with the minimum value from the net metabolic landscape generated by the optimization. One participant's data were not included in the metabolic analysis because of fatigue reported by the participant during the protocol, where the net metabolic rate of the second no-suit condition increased by 32.4% compared with the first no-suit condition.

Convergence time analysis

The convergence time for each participant was calculated in a post hoc analysis (fig. S3 and table S4). We defined the convergence of the optimization with the following two conditions: (i) Two consecutive iteration-to-iteration changes of maximum metabolic reduction in percentage from the metabolic landscape fell below our preset convergence threshold ($t_m = 4\%$), and (ii) two consecutive iteration-to-iteration changes of hyperparameters from the Gaussian process fell below our preset convergence threshold ($t_h = 3$). The convergence threshold for the changes of metabolic reduction t_m was chosen based on the previous study (8), which has shown an average error of 4% on this instantaneous metabolic estimation. The convergence threshold t_h was obtained from a separate simulation study. For this simulation, a generative model of metabolic landscape with added noise was first created. The noise was generated by Matlab's `awgn` function with a signal-to-noise ratio of 8.8 obtained from our pilot test (table S1). With this model, we ran Bayesian optimization for 50 iterations and calculated the iteration-to-iteration changes in the maximum metabolic reduction and the hyperparameters. The maximum changes for the metabolic reduction were set to 2% while evaluating the changes of all hyperparameters. We repeated the simulation 100 times and found that the metabolic reduction threshold was met when the threshold for all hyperparameters was set to 3.

Ground reaction force

Ground reaction forces (GRFs) were collected via the instrumented split-belt treadmill (Beltec) and synced with the actuation platform using the motion capture system (Qualisys AB). All the GRF force data were filtered with a zero-lag fourth-order low-pass Butterworth filtered with a 5- to 15-Hz optimal cutoff frequency that was selected using a custom residual analysis algorithm (32). A customized Matlab script was created using GRFs to segment the percentage of the gait cycle defined by the maximum hip flexion based on the detected heel strikes.

Statistics

Means and SEM of the net metabolic rate were calculated for each condition. According to the Jarque-Bera test (significance level $\alpha = 0.05$; Matlab), the collected data followed the normal distribution ($P > 0.3$). Therefore, we conducted a mixed-model, two-factor analysis of variance (ANOVA; random effect, participant; fixed effect, test condition) to test the effect across different conditions including optimal, validation, and no-suit conditions (significance level $\alpha = 0.05$; Matlab). For the outcome of the ANOVA test, it showed a significant difference of the net metabolic rate between conditions. We used paired t tests for the comparison between the conditions to identify which conditions exacted a significant change in the net metabolic rate (39).

SUPPLEMENTARY MATERIALS

robotics.sciencemag.org/cgi/content/full/3/15/ear5438/DC1

Fig. S1. Illustration of 1D Gaussian process.

Fig. S2. Simulation results on the number of iterations needed for the optimization.

Fig. S3. Convergence analysis.

Fig. S4. Optimized hip extension force profiles for all participants.

Fig. S5. Experimental setup.

Fig. S6. Experimental protocol.

Fig. S7. Structure of the waist belt component.

Fig. S8. Structure of the thigh brace.

Fig. S9. Pseudo-randomly sampled timings for the initialization of Bayesian optimization.

Fig. S10. Optimization process.

Table S1. Signal-to-noise ratio and variations of metabolic cost of pilot tests.

Table S2. Onset timing.

Table S3. Participant characteristics.

Table S4. Metabolic rates, optimal timing, and convergence timing for each participant.

Table S5. Quadratic approximation of metabolic landscape.

REFERENCES AND NOTES

1. S. H. Collins, M. B. Wiggin, G. S. Sawicki, Reducing the energy cost of human walking using an unpowered exoskeleton. *Nature* **522**, 212–215 (2015).
2. F. A. Panizzolo, I. Galiana, A. T. Asbeck, C. Siviyy, K. Schmidt, K. G. Holt, C. J. Walsh, A biologically-inspired multi-joint soft exosuit that can reduce the energy cost of loaded walking. *J. Neuroeng. Rehabil.* **13**, 43 (2016).
3. L. M. Mooney, E. J. Rouse, H. M. Herr, Autonomous exoskeleton reduces metabolic cost of human walking during load carriage. *J. Neuroeng. Rehabil.* **11**, 80 (2014).
4. J. Lee, K. Seo, B. Lim, J. Jang, K. Kim, H. Choi, Effects of assistance timing on metabolic cost, assistance power, and gait parameters for a hip-type exoskeleton. *IEEE Int. Conf. Rehabil. Robot.* **2017**, 498–504 (2017).
5. P. Malcolm, W. Derave, S. Galle, D. De Clercq, A simple exoskeleton that assists plantarflexion can reduce the metabolic cost of human walking. *PLOS ONE* **8**, e56137 (2013).
6. B. T. Quinlivan, S. Lee, P. Malcolm, D. M. Rossi, M. Grimmer, C. Siviyy, N. Karavas, D. Wagner, A. Asbeck, I. Galiana, C. J. Walsh, Assistance magnitude vs. metabolic cost reductions for a tethered multiarticular soft exosuit. *Sci. Robot.* **2**, eaah4416 (2017).
7. A. J. Young, J. Foss, H. Gannon, D. P. Ferris, Influence of power delivery timing on the energetics and biomechanics of humans wearing a hip exoskeleton. *Front. Bioeng. Biotechnol.* **5**, 4 (2017).
8. J. Zhang, P. Fiers, K. A. Witte, R. W. Jackson, K. L. Poggensee, C. G. Atkeson, S. H. Collins, Human-in-the-loop optimization of exoskeleton assistance during walking. *Science* **356**, 1280–1284 (2017).
9. J. R. Koller, D. H. Gates, D. P. Ferris, C. D. Remy, 'Body-in-the-Loop' optimization of assistive robotic devices: A validation study. *Robot. Sci. Syst.* (2016).
10. Y. Ding, F. A. Panizzolo, C. Siviyy, P. Malcolm, I. Galiana, K. G. Holt, C. J. Walsh, Effect of timing of hip extension assistance during loaded walking with a soft exosuit. *J. Neuroeng. Rehabil.* **13**, 87 (2016).
11. Y. Ding, I. Galiana, A. T. Asbeck, S. M. M. De Rossi, J. Bae, T. R. T. Santos, V. L. de Araujo, S. Lee, K. G. Holt, C. Walsh, Biomechanical and physiological evaluation of multi-joint assistance with soft exosuits. *IEEE Trans. Neural Syst. Rehabil. Eng.* **25**, 119–130 (2017).
12. R. W. Jackson, S. H. Collins, An experimental comparison of the relative benefits of work and torque assistance in ankle exoskeletons. *J. Appl. Physiol.* **119**, 541–557 (2015).

13. P. Malcolm, S. Lee, S. Crea, C. Sivi, F. Saucedo, I. Galiana, F. A. Panizzolo, K. G. Holt, C. J. Walsh, Varying negative work assistance at the ankle with a soft exosuit during loaded walking. *J. Neuroeng. Rehabil.* **14**, 62 (2017).
14. T. K. Uchida, A. Seth, S. Pouya, C. L. Dembia, J. L. Hicks, S. L. Delp, Simulating ideal assistive devices to reduce the metabolic cost of running. *PLOS ONE* **11**, e0163417 (2016).
15. M. Kim, S. H. Collins, Once-per-step control of ankle push-off work improves balance in a three-dimensional simulation of bipedal walking. *IEEE Trans. Robot.* **33**, 406–418 (2017).
16. W. van Dijk, H. van der Kooij, E. Hekman, A passive exoskeleton with artificial tendons: Design and experimental evaluation. *IEEE Int. Conf. Rehabil. Robot.* **2011**, 5975470 (2011).
17. C. L. Dembia, A. Silder, T. K. Uchida, J. L. Hicks, S. L. Delp, Simulating ideal assistive devices to reduce the metabolic cost of walking with heavy loads. *PLOS ONE* **12**, e0180320 (2017).
18. G. Lee, J. Kim, F. A. Panizzolo, Y. M. Zhou, L. M. Baker, I. Galiana, P. Malcolm, C. J. Walsh, Reducing the metabolic cost of running with a tethered soft exosuit. *Sci. Robot.* **2**, eaan6708 (2017).
19. R. E. Quesada, J. M. Caputo, S. H. Collins, Increasing ankle push-off work with a powered prosthesis does not necessarily reduce metabolic rate for transtibial amputees. *J. Biomech.* **49**, 3452–3459 (2016).
20. P. Malcolm, D. M. Rossi, C. Sivi, S. Lee, B. T. Quinlivan, M. Grimmer, C. J. Walsh, Continuous sweep versus discrete step protocols for studying effects of wearable robot assistance magnitude. *J. Neuroeng. Rehabil.* **14**, 72 (2017).
21. W. Felt, J. C. Selinger, J. M. Donelan, C. D. Remy, “Body-in-the-loop”: Optimizing device parameters using measures of instantaneous energetic cost. *PLOS ONE* **10**, e0135342 (2015).
22. G. H. Markovitz, J. W. Sayre, T. W. Storer, C. B. Cooper, On issues of confidence in determining the time constant for oxygen uptake kinetics. *Br. J. Sports Med.* **38**, 553–560 (2004).
23. L. N. Awad, J. Bae, K. O’Donnell, S. M. M. De Rossi, K. Hendron, L. H. Sloat, P. Kudzia, S. Allen, K. G. Holt, T. D. Ellis, C. J. Walsh, A soft robotic exosuit improves walking in patients after stroke. *Sci. Transl. Med.* **9**, eaai9084 (2017).
24. J. C. Selinger, S. M. O’Connor, J. D. Wong, J. M. Donelan, Humans can continuously optimize energetic cost during walking. *Curr. Biol.* **25**, 2452–2456 (2015).
25. E. Brochu, V. M. Cora, N. de Freitas, A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. <http://arxiv.org/abs/1012.2599> (2010).
26. H. J. Kushner, A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *J. Basic Eng.* **86**, 97–106 (1964).
27. J. Mockus, V. Tiesis, A. Zilinskas, The application of Bayesian methods for seeking the extremum, in *Towards Global Optimisation 2*, L. C. W. Dixon, G. P. Szegö, Eds. (Elsevier, 1978), vol. 2, pp. 117–129.
28. M. Kim, Y. Ding, P. Malcolm, J. Speeckaert, C. J. Sivi, C. J. Walsh, S. Kuindersma, Human-in-the-loop Bayesian optimization of wearable device parameters. *PLOS ONE* **12**, e0184054 (2017).
29. G. Lee, Y. Ding, I. Galiana Bujanda, N. Karavas, Y. M. Zhou, C. J. Walsh, Improved assistive profile tracking of soft exosuits for walking and jogging with off-board actuation. *Proc. IEEE/RSSJ Int. Conf. Intell. Robot. Syst.*, 1699–1706 (2017).
30. C. E. Rasmussen, C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, 2006).
31. J. C. Selinger, J. M. Donelan, Estimating instantaneous energetic cost during non-steady-state gait. *J. Appl. Physiol.* **117**, 1406–1415 (2014).
32. D. A. Winter, *Biomechanics and Motor Control of Human Movement* (Wiley, 2005), chap. 7.
33. N. Balakrishnan, *Methods and Applications of Statistics in Clinical Trials: Planning, Analysis, and Inferential Methods* (Wiley, 2014), chap. 24.
34. F. Pérez-Cruz, S. Van Vaerenbergh, J. J. Murillo-Fuentes, M. Lázaro-Gredilla, I. Santamaria, Gaussian processes for nonlinear signal processing: An overview of recent advances. *IEEE Signal Process. Mag.* **30**, 40–50 (2013).
35. Y. Ding, I. Galiana, C. Sivi, F. A. Panizzolo, C. Walsh, IMU-based iterative control for hip extension assistance with a soft exosuit. *Proc. IEEE Int. Conf. Robot. Autom.*, 3501–3508 (2016).
36. R. Lorenz, R. P. Monti, I. R. Violante, C. Anagnostopoulos, A. A. Faisal, G. Montana, R. Leech, The automatic neuroscientist: A framework for optimizing experimental design with closed-loop real-time fMRI. *Neuroimage* **129**, 320–334 (2016).
37. B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, N. de Freitas, Taking the human out of the loop: A review of Bayesian optimization. *Proc. IEEE* **104**, 148–175 (2016).
38. J. M. Brockway, Derivation of formulae used to calculate energy expenditure in man. *Hum. Nutr. Clin. Nutr.* **41**, 463–471 (1987).
39. S. A. Glantz, *Primer of Biostatistics* (McGraw-Hill, 2005).

Acknowledgments: We thank H. Su, A. Long, G. Lee, J. R. Forster, K. Pawel, D. Ryan, A. E. Erdheim, and S. K. Sullivan for their assistance with the hardware and experiments. We thank C. Sivi, B. Quinlivan, S. Lee, J. Bae, I. Galiana, P. Malcolm, and X. Zhao for their feedback on the paper. We thank our study participants who gave their time for this research. **Funding:** This research is based on the work supported by the Defense Advanced Research Projects Agency, Warrior Web Program (contract no. W911NF-14-C-0051), the Wyss Institute, and the John A. Paulson School of Engineering and Applied Science at Harvard University. **Author contributions:** Y.D., M.K., S.K., and C.J.W. designed the experiment. Y.D. designed force controller and integrated the HIL system. M.K. implemented the Bayesian optimization. Y.D. and M.K. performed the experiments. Y.D., M.K., S.K., and C.J.W. analyzed and interpreted the data. Y.D., M.K., S.K., and C.J.W. prepared the manuscript. All authors provided critical feedback on the manuscript. **Competing interests:** Patents describing the exosuit components documented in this article have been filed with the U.S. Patent Office. Y.D., M.K., S.K., and C.J.W. are inventors of at least one of the following patent/patent applications: U.S. 9,351,900, U.S. 14/660,704, U.S. 15/097,744, U.S. 14/893,934, PCT/US2014/068462, PCT/US2015/051107, and PCT/US2017/042286, filed by Harvard University. Harvard University has entered into a licensing and collaboration agreement with ReWalk Robotics. C.J.W. is a paid consultant to ReWalk Robotics. The other authors declare that they have no competing interests. **Data and materials availability:** Contact C.J.W. for source code and other materials.

Submitted 30 November 2017

Accepted 5 February 2018

Published 28 February 2018

10.1126/scirobotics.aar5438

Citation: Y. Ding, M. Kim, S. Kuindersma, C. J. Walsh, Human-in-the-loop optimization of hip assistance with a soft exosuit during walking. *Sci. Robot.* **3**, eaar5438 (2018).

Human-in-the-loop optimization of hip assistance with a soft exosuit during walking

Ye Ding, Myunghee Kim, Scott Kuindersma, and Conor J. Walsh

Sci. Robot. **3** (15), eaar5438. DOI: 10.1126/scirobotics.aar5438

View the article online

<https://www.science.org/doi/10.1126/scirobotics.aar5438>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.

Copyright © 2018 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works