



Supplementary Materials for

Human-in-the-Loop Optimization of Exoskeleton Assistance during Walking

Juanjuan Zhang, Pieter Fiers, Kirby A. Witte, Rachel W. Jackson,
Katherine L. Poggensee, Christopher G. Atkeson, Steven H. Collins

correspondence to: stevecollins@cmu.edu

This PDF file includes:

Materials and Methods
Figs. S1 to S11
Tables S1 to S4
Caption for Study Data
Caption for Optimization Code

Other Supplementary Materials for this manuscript include the following:

Study Data in the file archive [StudyData.zip](#)
Optimization Code in the file archive [OptimizationCode.zip](#)

Materials and Methods

Experimental Design

The purpose of this study was to test the effectiveness of a method for automatic identification of exoskeleton assistance patterns that minimize the metabolic energy used in human walking. We performed two sets of experiments: The main study was designed to test the effectiveness of the method across a range of individuals. A series of additional studies with single-subject design were used to test the effectiveness of the method across gait conditions, devices and objectives.

Eleven healthy individuals participated in the main study. Participants wore a torque-controlled, tethered exoskeleton on one ankle. The exoskeleton applied torque as a function of time when the foot was on the ground, defined by four control parameters that set the magnitude of peak torque, the timing of peak torque, the rise time and the fall time, constituting a control law. During the optimization phase of the experiment, a Covariance Matrix Adaptation Evolution Strategy (CMA-ES) was used to identify the control law that minimized the metabolic energy cost of walking for each participant. The metabolic rate corresponding to each control law was estimated by fitting a first-order model to two minutes of breath-by-breath respiratory data. CMA-ES settings, such as the number of control laws per generation (eight) and stopping criteria (typically four generations), were based on either the literature or the results of pilot testing. The mean of the final calculated generation was taken as the optimized control law. Separate validation trials with a double-reversal design were then conducted to compare optimized assistance to two baseline conditions: walking with the exoskeleton in a zero-torque mode and walking with a static controller from a prior experiment. In a subset of participants, additional tests were applied to check for convergence. The primary study outcomes were the differences in metabolic rate between the optimized condition and the zero-torque and static conditions. Two paired t-tests formed the corresponding primary statistical analysis. Secondary outcomes included the nature of the optimized assistance patterns and differences in metabolic rate between additional conditions and participant sub-groups, which were analyzed using paired or unpaired t-tests as appropriate.

Seven different conditions were tested in additional studies with single-subject design. Methods for these studies were similar to those of the main study, but included conditions with exoskeletons worn on both ankles, walking at slow, normal and fast speeds, walking uphill, walking while carrying a heavy load, and running. A final test optimized muscle activity rather than energy expenditure.

Participants in the Main Study

Eleven healthy adults ($N = 11$, 5 female and 6 male; age = 27.2 ± 4.2 [24-37] years; body mass = 69.6 ± 14.1 [50-93] kg; height = 1.75 ± 0.10 [1.60-1.88] m; mean \pm standard deviation [range]; table S1) participated in the study. One additional participant dropped out before completing the protocol, in part owing to hardware malfunctions. Sample size was chosen on the basis of data from previous studies. All participants provided written informed consent before participation, after the nature and possible consequences of the study were explained. The study protocol was approved and overseen by the Institutional Review Board of Carnegie Mellon University.

Exoskeleton Hardware in the Main Study

We used a tethered, one degree-of-freedom, torque-controlled ankle exoskeleton emulator (30) to apply ankle torques. Participants walked on a treadmill while wearing the exoskeleton on their right ankle. The primary reason for using this unilateral exoskeleton was to allow comparisons to a prior study with the same hardware (17). This choice also resulted in a simpler experimental protocol compared to tests with bilateral exoskeletons; fewer components needed maintenance, electromechanical failure during testing was less likely, and less time was required for users to put the exoskeleton on and take it off. We suspect that this choice made balance easier and reduced leg interference during swing, at a cost of less total assistance, but did not test these effects directly.

The ankle exoskeleton emulator is composed of an off-board control module and electric motor; a uni-directional Bowden cable transmission with a series leaf spring; and an exoskeleton frame that interfaces with the human foot and shank (cf. Fig. 2C; fig. S2).

The emulator used a high-speed control system (ACE1103, dSPACE) to sample sensors at 5000 Hz, filter sensor data at 200 Hz, and generate commands of desired motor velocity at 500 Hz. The motor unit included a low-inertia, 1.6 kW AC servo motor and a 5:1 planetary gear (BSM90N-175AD and GBSM90-MRP120-5, Baldor Electric). Motor input voltage was regulated by a driver running in velocity control mode (MFE460A010B, Baldor Electric). A digital optical encoder measured motor position (E8P, US Digital). The 100% rise time to peak motor velocity was 0.013 s.

A flexible, uni-directional Bowden cable was used as the transmission between the off-board motor and the exoskeleton end-effector worn on the leg. The cable was composed of a coiled-steel outer conduit (415310-00, Lexco Cable) and a 0.003 m diameter Vectran[®] inner rope, and was approximately 2 m in length. A custom leaf spring (fabricated from GC-67-UB fiberglass, Gordon Composites, Inc.) was attached at the end of the rope in series with the ankle joint to provide increased transmission compliance. This series-elastic transmission decoupled motor inertia from the exoskeleton and reduced interface impedance between the human and robotic portions of the system, thus improving torque tracking, comfort and safety. In the zero-torque mode, the cable was kept slack so as to transmit no forces to the user.

The exoskeleton end-effector applied forces on the front of the human shank below the knee, beneath the heel, and on the ground beneath the toe (cf. Fig. 2D; fig. S3). This generated the equivalent of an ankle plantarflexion torque, with torque magnitude proportional to transmission cable tension. Torque was measured using a load cell (LC201 Series; OMEGA Engineering) conditioned at 1000 Hz (CSG110, Futek). Joint angle was measured using a digital optical encoder (E8P, US Digital). The ankle exoskeleton has a peak allowable torque of 120 N m and weighs 0.83 kg.

Low-level exoskeleton torque was controlled using a combination of proportional control, damping injection, and iterative learning, with motor velocity commanded to the motor driver (31, 43). This low-level torque control approach does not rely on explicit models or integration, which makes it well-suited to the nonlinear, complex and time-varying dynamics of a human interacting with an exoskeleton during walking. Iteratively-learned feed-forward compensation mitigated steady-state errors by exploiting the cyclic nature of walking. With this low-level torque control approach, real-time torque tracking errors during walking can be as low as 1% of peak torque at steady state (31, 43).

Control Law Parameterization in the Main Study

We tested the optimization method on a control law that determined desired ankle exoskeleton torque as a function of the time since ground contact, normalized to stride period. Four control parameters defined peak torque, the time of peak torque, the rise time and the fall time for a curve comprising two cubic splines. Additional low-torque ramp-in and ramp-out patterns were included for the rest of the stance phase to improve torque tracking. Stride period was estimated online by low-pass filtering measured stride periods during walking.

The ankle torque curve had a hill-like shape that can be divided into four sections: a shallow, low-torque setup ramp; a rising s-like cubic spline linking the onset point to the peak; a falling arc-like cubic spline linking the peak to the removal point; and a shallow, low-torque settling ramp. The shape was primarily defined by four parameters: peak torque, τ_p , peak time, t_p , rise time, t_r , and fall time, t_f . Additional constant settings fully defined the shape. First, the desired torque at the beginning of the stride, defined as the instant of heel strike, was zero. Second, the desired torque at 65% of the stride period, which is typically slightly after toe-off, was zero. Third, at both torque onset and torque removal we enforced a small desired torque value, 2 N m, rather than zero. This avoided instability caused by torque measurement errors while tracking zero torque with a unidirectional cable transmission. Fourth, the slopes of the cubic splines at the onset and peak points were zero. Finally, the second derivative of the cubic spline at the removal point was zero. These constraints and the four parameters uniquely determined the torque at all instants during the stance period from heel strike to toe off. As soon as the foot came off the ground, a zero-torque ‘swing mode’ was applied by maintaining a small amount of slack in the cable.

We now define the shape of the stance-phase torque curve more precisely. We will use time-torque coordinates of the form [X%, Y N m], where X is time, normalized to percent stride, and Y is torque magnitude. During setup, desired torque was linearly interpolated between heel strike, [0%, 0 N m], and onset, [$(t_p - t_r)$ %, 2 N m]. Similarly, in the settling section, desired torque was linearly interpolated between torque removal, [$(t_p + t_f)$ %, 2 N m], and 65% of stride time, [65%, 0 N m]. If the user was still in the stance phase after 65% of the stride period, desired torque remained zero.

In the rising and falling cubic spline sections, desired torque was defined by two cubic splines, one from torque onset, [$(t_p - t_r)$ %, 2 N m], to peak, [t_p %, τ_p N m], and one from peak to removal [$(t_p + t_f)$ %, 2 N m]. Let the two splines be represented by

$$\begin{aligned}\tau_1(t) &= a_1 t^3 + b_1 t^2 + c_1 t + d_1 \\ \tau_2(t) &= a_2 t^3 + b_2 t^2 + c_2 t + d_2\end{aligned}$$

where t denotes the time elapsed since heel strike in the current stride, as a percentage of the stride period. a_i , b_i , c_i , and d_i ($i = 1$ or 2) are the eight parameters to be set. From the definitions above, we have

$$\begin{aligned}
\tau_1(t_p - t_r) &= 2 \\
\tau_1(t_p) &= \tau_p \\
\tau_2(t_p) &= \tau_p \\
\tau_2(t_p + t_f) &= 2 \\
\tau_1'(t_p) &= 0 \\
\tau_2'(t_p) &= 0 \\
\tau_1'(t_p - t_r) &= 0 \\
\tau_1''(t_p + t_f) &= 0
\end{aligned} \tag{1}$$

Solving this set of coupled equations gives us two unique cubic splines between torque onset and removal.

We used limits determined in pilot tests to set constraints on the four control parameters that defined the pattern of ankle torque. The maximum value of peak torque was 1 N m kg^{-1} , or one Newton meter for each kilogram of participant body mass. This reduced the likelihood of discomfort in cases with unusual timing. The lower bound on peak torque was 2 N m , which was the minimum value that could be tracked reliably. The rise time was constrained to between 10% and 40% of the stride period and the fall time was constrained to between 5% and 20% of the stride period. The lower limits on rise and fall times were chosen to avoid poor low-level torque tracking during rapid changes in desired torque. The upper limits were chosen to avoid discomfort associated with large torque impulses when peak torque was also high. Fall time limits were set to lower values than rise time limits based on the observation in pilot tests that faster drops in torque were tolerated better by the low-level controller while slower drops in desired torque were tolerated less well by the user. The time of peak torque was constrained to between 10% and 50-55% of stride, with the upper limit dependent on the magnitude of peak torque. When the magnitude of peak torque was 0.75 N m kg^{-1} or lower, the upper limit on time of peak torque was 55% of stride. As peak torques increased from 0.75 N m kg^{-1} to 1 N m kg^{-1} , the upper limit decreased from 55% to 50% of stride. This choice was made based on pilot testing, in which we observed that large, rapid drops in torque very late in the stance phase became uncomfortable for some participants. Given the approximate nature of these hand-selected parameter constraints, it is possible that a larger solution space could be achieved with further refinement. Sample code that takes parameter values, applies the above constraints and generates the resulting torque curves can be found in Optimization Code in the Supplementary Materials.

Online Optimization Strategy

We used a Covariance Matrix Adaptation Evolution Strategy (CMA-ES; 28) to identify the exoskeleton control laws that minimized the metabolic cost of walking for each participant. CMA-ES is stochastic, does not explicitly use derivative terms, and is often used for non-linear, non-convex optimization problems. The method is well-suited to human-in-the-loop optimization because it addresses noisy measurements, expensive objective function evaluations, nonlinear objective functions with unknown structures, and complex, subject-dependent human learning and adaptation processes. We did not

apply an exhaustive set of potential optimization strategies in simulations or pilot tests, but CMA-ES was by far the most effective of those strategies we did test. Here we provide an overview of the algorithm and rationale for its use in our method human-in-the-loop optimization.

CMA-ES first evaluates a group of candidate parameter sets which form the population of one generation. Candidate parameter sets are randomly selected according to a multivariate normal distribution of the parameters, characterized by a mean, a covariance matrix and a standard deviation. The mean of the distribution represents the current estimate of the optimal parameter values. After evaluating the current generation, the parameter sets are ranked in terms of performance and a rank-weighted average of the best sets becomes the mean of the next generation. The change in means is used to update the covariance matrix of the next generation, while additional multi-generational terms refine the covariance matrix and standard deviation to improve computational efficiency. This process is repeated from generation to generation to imitate natural selection. The mean of the final calculated generation serves as the best estimate of the optimal parameter values.

We chose CMA-ES based on simulations of the human-in-the-loop optimization problem and pilot-tests of candidate optimization strategies, which we derived from the literature. Optimizing multivariate exoskeleton and prosthesis assistance conditions from human metabolic rate measurements is challenging in multiple ways. First, metabolic rate is noisy, owing both to complicated human physiological and biochemical dynamics and to shortcomings in respiratory measurement hardware. CMA-ES is stochastic, which makes it less sensitive to noise than derivative-based methods such as gradient descent (44) and ‘hill-climbing’ methods such as Nelder-Mead (45).

A second challenge is that evaluation of candidate conditions is very expensive in terms of time and human effort. Measurement of metabolic rate requires on the order of minutes of respiratory data from a human interacting with the device, due to delays in the expression of energy used by muscles in expired gases (46). Often, multivariate optimization methods require a large number of function evaluations per step, and this number increases with the dimensionality of the control parameter space. For example, a quadratic approximation requires at least $1 + 2n$ evaluations assuming no interactions and $1 + 2n + 0.5(n^2 - n)$ evaluations including interactions, making it either $O(n)$ or $O(n^2)$, where n is the number of parameters being optimized. Gradient calculations and surrogate-based methods are similarly $O(n)$ or greater. In benchmarking problems, the CMA-ES method requires $4 + \text{floor}(3 \ln(n))$ evaluations per generation (47), making it potentially $O(\ln(n))$. In a forty-dimensional parameter space, for example, a full quadratic approximation would require at least 861 evaluations, while CMA-ES might use as few as 15 per generation, although this scaling might not hold in an adapting landscape and we did not test scaling in this study.

A third challenge is that the nature of the relationship between exoskeleton and prosthesis control parameters and human metabolic rate is not known in advance and may include complex nonlinearities and local minima. CMA-ES is stochastic and includes mechanisms to grow or shrink the standard deviation of the randomly-selected parameter values depending on the evolution of the mean over time. These features make CMA-ES robust against thresholds, discontinuities and local minima, as long as the initial values of the mean, covariance matrix, and standard deviation are well-chosen.

A fourth challenge is that humans exhibit complex, individualized learning and adaptation processes when using an exoskeleton or prosthesis. This puts gradient-based and quadratic approximation methods at a disadvantage, because calculating the gradient or Hessian requires substantial time, during which the human is changing. The calculated gradient or quadratic will therefore often be inaccurate, resulting in poor subsequent guesses at the optimal parameter values or requiring additional evaluations to recalculate the model of the space. CMA-ES is less sensitive to these problems because it uses only the rank order of the parameter sets, not the objective function values or their partial derivatives. This reduces the likelihood that, e.g., an early, poorly-adapted response to a potentially beneficial control law prevents movement of the mean in that direction. Human adaptation also presents challenges for methods that attempt to develop models of the space based on all available data, because data collected early in the adaptation process are likely to conflict with data from later in the process. The variant of CMA-ES used in this study updates the estimate of the optimal parameter values using only data from the current generation, reducing its sensitivity to erroneous, poorly-adapted responses. The covariance matrix update does have dependencies on the mean and covariance from prior generations, a feature that improves efficiency in time-invariant problems but might lead to reduced efficiency in human-in-the-loop optimization.

Online Optimization Parameters in the Main Study

In this study, we optimized four exoskeleton control parameters and therefore chose a population size of eight control laws per generation (using the previously described formula). This population size is intended to be robust and therefore applicable to a wide range of parameter spaces (28, 48). We aimed for four generations of optimization per participant, based on pilot tests suggesting that convergence was typically achieved in four generations or fewer. Four subjects experienced more or fewer generations of optimization (Optimization Trials; table S2). Parameter constraints were applied after the selection of control laws for each new generation. The optimized control law was defined by the final calculated (untested) mean parameter values. Note that this means participants never experienced the optimized control law during optimization.

We re-parameterized the optimization problem to improve the initial guess of the distribution defined by the covariance matrix, based on the results of pilot testing. Pilot tests suggested that the peak time and fall time, in units of percent stride, had smaller comfortable ranges than rise time, in units of percent stride, and peak torque, in units of Newton meters. We therefore mapped the original parameters $[\tau_p, t_p, t_r, t_f]$ into the space $[\varphi_1, \varphi_2, \varphi_3, \varphi_4] = [\tau_p, 2t_p, t_r, 2t_f]$ and optimized on φ .

We initialized the optimization with $[\varphi_1, \varphi_2, \varphi_3, \varphi_4] = [0.55m, 90, 25, 20]$, where m is participant body mass in kilograms. For a typical subject with mass of 70 kg, this resulted in an initial guess of $[\tau_p, t_p, t_r, t_f] = [38 \text{ N m}, 45\%, 25\%, 10\%]$. The initial covariance matrix was set as the size-four identity matrix and the standard deviation was initialized to 10. In pilot tests, we found that these values typically allowed subjects to complete the first generation with relative comfort, while covering a sufficient portion of the parameter space so as to span some of the eventual optimal parameter values. An example of the CMA-ES code that was used in our method can be found as Optimization Code in the Supplementary Materials.

Metabolic Rate Estimation

The metabolic rate corresponding to each exoskeleton control law was estimated by fitting a first-order dynamical model to two minutes of breath-by-breath metabolic measurements. We used an inverse dynamics approach similar to (22, 23), assuming a step change in actual metabolic rate with each change in control law. That is, we fit an exponential curve to measured metabolic rate and used the asymptote of this fit as our estimate of actual metabolic rate. The equations below detail the conversion from oxygen uptake and carbon dioxide production to metabolic rate, and the construction of the matrices used in the pseudoinverse operation for the least squares fit.

Breath-by-breath rates of O_2 consumption and CO_2 production were measured using a wireless, portable respirometry system (Oxycon Mobile; CareFusion). Data were communicated to the optimizer on the exoskeleton emulator system by first capturing them using screen automation software (Sikuli) from values displayed on the screen of the metabolics computer and then transmitting them via a serial channel.

Measured flow rates of O_2 and CO_2 were converted to a respiratory response term, related to metabolic rate, using a common equation (49). Respiratory output was calculated as $y(t) = 0.278 \cdot \dot{V}O_2(t) + 0.075 \cdot \dot{V}CO_2(t)$, where $y(t)$ is the respiratory response in Watts and $\dot{V}O_2(t)$ and $\dot{V}CO_2(t)$ are volumetric flow rates in mL min^{-1} . At steady state, the average of this respiratory response is equal to the estimated metabolic rate. However, the respiratory response does not immediately reflect changes in actual metabolic rate during transient periods (46), such as the period following a change in exoskeleton control laws. During such phases, the governing dynamics can be approximated as a first-order linear system (22).

We used an inverse-dynamics approach to estimate actual metabolic rate from noisy, transient respiratory response data. We treated actual metabolic rate, \dot{E} , as constant for each control law. A sequence of n breath-by-breath respiratory responses were measured, $y = [y(1) \ y(2) \ \cdots \ y(n)]$, at corresponding times, $t = [t(1) \ t(2) \ \cdots \ t(n)]$, for each control law. We then modeled y as related to \dot{E} in time according to

$$\dot{y}(t) = \frac{1}{\delta} (\dot{E} - y(t))$$

in which δ denotes a time constant. These continuous dynamics can be discretized as

$$y(i+1) = \frac{\delta - dt(i)}{\delta} y(i) + \frac{dt(i)}{\delta} \dot{E} \quad (2)$$

where $dt(i) = t(i+1) - t(i)$ is the time difference between samples. We then have

$$\begin{bmatrix} y(1) \\ y(2) \\ \vdots \\ y(n) \end{bmatrix} = A \begin{bmatrix} y_1 \\ \dot{E} \end{bmatrix}$$

in which y_1 is a constant related to the initial respiratory response and matrix $A \in \mathbb{R}^{n \times 2}$ is formed according to Eqn. 2 as follows. For $i=1$, we initialize the matrix as $A(1,1) = 1$ and $A(1,2) = 0$. For $i \in [2, 3, \dots, n]$, then we iteratively constructed the matrix as

$$A(i,1) = \frac{\delta - dt(i-1)}{\delta} \cdot A(i-1,1)$$

$$A(i,2) = \frac{\delta - dt(i-1)}{\delta} \cdot A(i-1,2) + \frac{dt(i-1)}{\delta}$$

Actual metabolic rate can then be estimated by

$$\begin{bmatrix} \hat{y}_1 \\ \hat{E} \end{bmatrix} = A^+ y$$

where A^+ is the pseudoinverse of A , and \hat{y}_1 and \hat{E} are estimates of the initial respiratory response and the actual steady-state metabolic rate. This estimate of metabolic rate was used as the cost function value associated with each control law during optimization. Sample metabolic rate estimation code can be found in Optimization Code in the Supplementary Materials.

There are two important parameters in this estimation of metabolic rate: the time constant of the first-order dynamics and the duration of the evaluation of each control law. We used a time constant of 42 s, an average value identified in a previous study (22), because it fit pilot data well. To identify an appropriate estimation period, we collected data in multiple six-minute walking bouts, each with a different, constant control law. We defined true metabolic rate as the estimate obtained using all six minutes of data. We then investigated the relationship between the duration of shorter virtual trial periods and the error in estimated metabolic rate compared to the true value. As expected, longer trials resulted in lower errors. We picked an evaluation period of two minutes because it seemed to strike a good balance between estimation errors (an average of 4% error) and trial duration (fig. S1). Shorter or longer evaluation periods would be expected to result in more or less noise per evaluation, but also more or fewer evaluations per hour, respectively, either of which could prove more efficient overall.

Optimization Trials in the Main Study

Participants walked on the treadmill at a constant speed of 1.25 m s⁻¹ while wearing the ankle exoskeleton on their right ankle in all tests. Participants first underwent a brief period of acclimation to the exoskeleton, in which peak torque was slowly increased from zero and they were advised to relax and not resist the device.

During the optimization phase of the study, participants experienced at least four generations of control laws, each consisting of eight two-minute periods. Participants were instructed to finish all four generations without stopping, beginning with two minutes of walking in the zero-torque mode to increase their initial respiratory response. As needed, the protocol was paused to provide subjects with rest or to adjust hardware. In these cases, the optimization was paused and resumed from the same point, following another warmup period. Eight participants experienced all four generations in the same data collection session. For the other three participants, optimization was paused before the end of the fourth generation due to hardware failure and was resumed on a subsequent day. Eight participants completed four additional generations of optimization as a test of whether the optimization had converged.

Four participants received more or less optimization to identify optimized assistance parameters. One participant (Subject 6) appeared to become trapped in a poor local

minimum due to unusual adaptation dynamics on the first day. We therefore reset the optimization and conducted an additional four generations of optimization on a second day. Validation trials were performed on both sets of optimized parameters on a third day. Another participant (Subject 10) appeared not to have converged after four generations and so was provided with an additional five generations of optimization. At the end of the ninth cumulative generation, it appeared that the participant was trapped in a poor local minimum. We therefore reset and conducted an additional four generations of optimization. Only the final optimized parameter values were included in validation. One participant (Subject 7) appeared to have converged at the end of three generations, and the mean from the fourth generation was used in validation trials. For one participant (Subject 8) we used the mean from the sixth generation in validation trials, rather than the fifth, due to experimenter error. Details of the experimental protocol for all participants are listed in table S2.

Validation Trials in the Main Study

In separate validation trials, we compared the metabolic cost of the optimized control law with two additional control conditions. In the ‘zero-torque’ condition, the unidirectional Bowden cable transmission of the exoskeleton was kept slack, thereby applying zero torque to the human ankle. This complete lack of torque was verified using load cell measurements. In the ‘static’ condition, we applied a pattern of torque that was found to reduce energy cost in a previous study (17). This condition is approximated by $[\tau_p, t_p, t_r, t_f] = [0.825m, 53, 17, 8]$, where m is participant body mass. One participant (Subject 2) was provided with the static condition as $[0.825m, 53, 25, 10]$ due to experimenter error. This happened to cause the static condition to be closer to the optimized condition (table S3), which would be expected to result in lower metabolic rate in the static condition for this subject. We also included a quiet standing condition, in which participants stood at rest on the treadmill while wearing the exoskeleton. Metabolic rate from the standing condition was subtracted from gross metabolic rate in walking conditions to obtain the energy cost of walking for each condition.

Validation trials were conducted in a double-reversal order. Each condition was tested twice, first in random order and then again in reverse order to mitigate the effects of adaptation or drift in respiratory measurements. Three validation tests (Subject 2 day 1, Subject 8 day 2 and Subject 9 day 2) were conducted in which the middle condition was only tested once. However, the singly-tested condition occurred in the middle of the sequence, meaning that compensation for adaptation and drift were still realized.

Validation trials were single-blind. Participants were not informed as to which exoskeleton condition they experienced, but it was necessary for one experimenter to be un-blinded to allow them to program the exoskeleton emulator appropriately.

During validation trials, each condition was presented for six minutes. Steady-state metabolic rate was determined as the average value from the final three minutes of walking, the most accurate and commonly-used method (22). The metabolic rate for each condition was then obtained by averaging steady-state metabolic rate from both presentations. Finally, metabolic rate was normalized to body mass to partially account for differences in participant size, and is presented in units of $W \text{ kg}^{-1}$.

For two participants, Subjects 2 and 8, static validation trials were conducted on a separate day from optimized validation trials. Day-to-day fluctuations in body chemistry and respirometry system calibration can result in differences in apparent oxygen consumption for

the same actual energy cost. To allow direct comparison of static and optimized conditions for these two participants, we normalized static metabolic rate to metabolic rate in the zero-torque condition, which is expected to be consistent across collection days; metabolic rate in the static condition was scaled by the ratio of metabolic rate in the zero-torque condition on the optimized validation day to that on the static validation day. Nearly identical results are obtained by either not accounting for day-to-day changes or by comparing the same-day percent changes. Comparisons between optimized and zero-torque conditions and between static and zero-torque conditions were unaffected, since zero-torque data were collected on the same day as optimized and static conditions for all participants.

The metabolic cost of walking, used to compare the effectiveness of control laws in validation trials, was defined as the total metabolic rate minus the rate for quiet standing. This is often referred to as net metabolic rate. The metabolic rate for quiet standing was $1.25 \pm 0.20 \text{ W kg}^{-1}$ in the main study.

Statistical Analysis in the Main Study

The two primary outcomes of this study are the change in metabolic rate with optimized assistance versus zero-torque, and the change in metabolic rate with optimized assistance versus static assistance. We designed the experiment to robustly identify these differences, and used two-sided, paired t-tests with $N = 11$ to obtain a measure of confidence (p-value) that the percent difference between these conditions was different from zero. We used a significance level of $\alpha = 0.05$ to determine statistical significance. We used a Jarque-Bera test to verify that metabolic rate data were normally distributed across participants in each of the compared conditions.

We also report metabolic cost data from a prior study (cf. Fig. 3B) to provide context for the primary results. In the previous study (17) the same hardware was used to apply the same zero-torque and static conditions on a different set of participants who did not undergo optimization. This provides a comparison point for the difference between zero-torque and static conditions in the present study, which suggests that facilitating human adaptation may have been an important factor in optimizing assistance. The prior study also included a condition not tested in the main study, in which no exoskeleton was worn. This comparison allows an inference to be made regarding the expected difference between optimized assistance and normal walking in the present study. Note that individual differences in metabolic cost between the two separate study participant groups led to differences in mean metabolic cost for the zero-torque condition, as well as large standard deviations for all conditions in both studies. This is a common feature of metabolic rate data, arising from physiological and neurological differences between individual humans. Differences in metabolic rate between conditions obtained for the same participant during the same collection period, however, are much more consistent, evident from participant-wise results in the present study (fig. S5; table S3). No statistical analyses were applied comparing the results of the two studies.

We made qualitative comparisons of optimized control law parameter values across participants. Optimized parameter values are presented in a box plot (cf. Fig. 3C), with whiskers indicating the range of optimized values and the median line showing typical optimized values. In the box plot, optimized parameter values are normalized to the allowable range set by the parameter constraints defined above. We also present the average measured torque with optimized assistance for each participant (cf. Fig. 3D). For each

participant, measured torque for each stride was normalized in time to percent stride and then averaged across all strides in one of the optimized assistance conditions from validation trials. The same process was used to calculate the average measured torque in the static and zero-torque conditions, with the result then averaged across all participants (cf. Fig. 3E).

We performed comparisons of a number of secondary outcomes, with preliminary observations intended to provide insights to inform future studies. Here, ‘change in metabolic rate’ refers to the difference in metabolic rate versus the zero-torque condition. For the subset of participants (N=8) who finished an additional four generations of optimization, the ‘convergence check’, we compared the change in metabolic rate with optimized assistance to the change in metabolic rate with convergence-check assistance using a paired t-test (fig. S9). We compared change in metabolic rate for participants who had prior experience with an exoskeleton (N = 5) to change in metabolic rate for participants with no prior experience (N = 6) using an unpaired t-test (fig. S10). We also looked for a correlation between exoskeleton mechanical work and human metabolic rate using linear regression, and report R^2 value as a measure of explained variance (fig. S11).

Estimate of Autonomous Device Mass

The net effect on performance resulting from use of an autonomous exoskeleton can only be determined with certainty by building and testing that device. Assistance from motors can reduce energy cost, but carrying the mass of the motors and batteries can increase energy cost (50) forming an initial penalty. Streamlining the device for one specialized purpose can decrease this initial penalty compared to a general-purpose emulator, leading to a greater net benefit. Removing tethers can also reduce resistance to natural leg motions, an important component of economical gait (51). The interactions between the initial penalty, the assistance pattern, and the design of the autonomous device can be complicated. However, we can obtain a rough estimate of the mass, initial penalty and net benefit by extrapolating from prior autonomous designs.

The most effective autonomous ankle exoskeleton to date (4) reduced the metabolic energy cost of walking by 10% compared to walking in normal boots. There was a 4.5% initial penalty for wearing the device on both ankles. The device weighed 1.06 kg per leg, including a 0.30 kg motor. Each exoskeleton produced a peak torque of 0.50 N m kg^{-1} and an average power of 0.15 W kg^{-1} per leg. A battery weighing 0.82 kg was carried in a backpack. In the present study, the optimized control law reduced metabolic energy cost by 24% on average compared to the zero-torque condition. Optimized assistance required a peak torque of about 0.76 N m kg^{-1} and an average net power of about 0.38 W kg^{-1} . We could obtain an estimate of the mass of a device capable of this mechanical output by extrapolating from (4); we could scale motor mass with average power, the mass of the remainder of the device with peak torque, and battery mass with total average power, noting that only one exoskeleton would require battery power. We could then obtain an estimate of the resulting initial penalty by either scaling it directly or using the relationships provided by (50), in either case accounting for the fact that the exoskeleton was worn on only one ankle. These approaches would suggest that scaling the hardware from (4) and applying optimized control could yield a substantially larger net improvement in metabolic cost (sample calculation in Study Data). It would be necessary to fabricate and test the scaled device to verify this idea. Nonetheless, the optimized control laws discovered in this experiment seem to have relevance to autonomous device design.

Additional Single-Subject Studies

We demonstrated the generality of the approach by conducting single-subject studies with different devices and objectives in several additional locomotion conditions. Tests included minimizing metabolic energy consumption during slow, normal, fast, uphill and loaded walking, minimizing metabolic energy consumption during running, and minimizing muscle activity during walking. The methods in these tests were identical to those used in the main experiment, described above, except for a few noted differences.

The six additional optimizations of metabolic rate were conducted using a different ankle exoskeleton, which was worn on both ankles rather than just one (fig. S4). These exoskeletons have structural differences from the device used in the main experiment, including a smaller overall envelope, a closer fit to the limb and less compliance in inversion-eversion and internal-external rotation (30). Using an exoskeleton on both legs presents strong differences from using an exoskeleton on only one leg, with differing implications for gait symmetry, inter-limb coordination, balance control, minimum step width, and maximum possible assistance, among other factors. In these experiments, both the left and right exoskeletons used the same control law.

One new participant (Subject 12; table S1) with no prior experience using an exoskeleton underwent optimization and validation under five new walking conditions: slow, normal, fast, uphill and loaded walking. The slow walking test was performed at 0.75 m s^{-1} . The maximum allowable peak timing was increased to 59% stride time in this condition, to allow for the longer double-support period of slow walking, and the optimization was run for eight generations, which we found was required for convergence. For the other walking conditions, the parameterization and number of generations were identical to the main study. The normal walking test was performed at 1.25 m s^{-1} . The fast walking test was performed at 1.75 m s^{-1} . The treadmill was inclined to a 10% grade with a speed of 1.25 m s^{-1} for the uphill walking test. The subject wore a weighted running vest (V-Force) with a mass of 19.4 kg (20% body mass) and walked at a speed of 1.25 m s^{-1} for the loaded walking condition. Validation trials were collected immediately following each optimization, with five minute periods of rest between trials. Validation conditions included walking without the exoskeleton, walking with the exoskeleton in zero-torque mode, and walking with optimized assistance. Each validation condition was performed twice and the order of the validation trials was block randomized (ABCCBA). The two normal walking trials were presented either first and last or third and fourth such that the exoskeleton was only donned and doffed one additional time.

Running tests were performed on one participant who also completed the main experiment (Subject 2; table S1). For the running test, the torque trajectory had identical nodes, but the timing parameters were defined differently. Peak torque, τ_p , was defined as in the walking study. Peak time in the running study, t_{pr} , was defined as a percentage of the active period of the device, t_a , similar to the stance period, rather than as a percentage of stride. The active period, t_a , was the time between foot flat and 42% of the average stride time, which was measured online. Foot flat was defined as the moment after heel strike at which the ankle joint velocity crossed zero as sensed by the exoskeleton joint encoder. The onset and removal of torque were set by the time of onset, t_{on} , and time of removal, t_{off} , which were also defined relative to foot flat, rather than relative to the time of peak torque. Peak time, t_{pr} , was constrained to be between 20% and 75% of t_a . Onset time, t_{on} , was

constrained to be between 0% and $(t_{pr} - 0.20 \cdot t_a)$. Removal time was constrained to the range $(t_{off} + 0.20 \cdot t_a)$ to $(0.95 \cdot t_a)$. Peak torque was constrained to the range of 20 N m to 55 N m. We initialized the optimization with $[\tau_p, t_{on}, t_{pr}, t_{off}] = [35 \text{ N m}, 25, 55, 85]$. Running validation trials included two trials of optimized assistance and zero-torque and one trial with normal running shoes without the exoskeleton, with block randomization.

Optimization with muscle activity as the objective was performed on one new participant (Subject 13; table S1) with no prior experience using an exoskeleton. The exoskeleton, walking conditions, and optimization and validation methods were identical to the main experiment, with the exception that the objective was to minimize muscle activity rather than metabolic rate. Muscle activity was quantified as the root mean square of the processed electromyography signal taken from the lateral aspect of the right soleus (the largest calf muscle, which acts to plantarflex the ankle joint) during the stance phase. The exoskeleton therefore assisted the same joint from which muscle activity was measured. During optimization and validation, the signal was filtered with a high-pass second-order Butterworth filter with a cut-off frequency of 20 Hz, full-wave rectified, filtered with a low-pass second-order Butterworth filter with a cut-off frequency of 10 Hz, and then normalized to the peak of the average trajectory from normal walking without the exoskeleton. This processing approach was chosen to be consistent with electromyography techniques in other studies (19). Stance phase was determined using the heel switch and ankle encoder during optimization trials, and using ground reaction forces during validation trials. The root mean square (RMS) of soleus activity during stance was calculated for each stride independently, then averaged. In optimization trials, steady-state muscle activity was estimated as the average value from the final one minute of each condition. During validation, steady-state muscle activity was defined as the average from the final three minutes of each condition. Values in cf. Fig. 4 are normalized to the RMS value from normal walking without the exoskeleton.




Fig. S1. Metabolic rate estimation period versus estimate accuracy.

Estimating steady-state metabolic rate from a longer period of transient respiratory data for each control law leads to lower estimation error. We chose a duration of two minutes so as to strike a balance between accurate estimation for each control law, which should lead to better guesses at optimal control parameters and fewer evaluations before convergence, and shorter duration, which leads to less walking time per evaluation. Data shown are from all validation trials in the present study. The fitting procedure described in Metabolic Rate Estimation was applied to increasing amounts of breath-by-breath data and the corresponding estimates of actual metabolic rate were calculated. Error was defined as the difference between individual estimates and the value obtained using all six minutes of breath-by-breath data from the same condition, and then normalized to the same value. Mean error is black, median is blue. Maximum, mean and median error based on two minutes of data were 15.5%, 4.3% and 3.7%, respectively. In future algorithms, computational efficiency might be improved by allowing trials to have differing durations. For example the evaluation of bad control laws could be terminated early based on a rapid, high-variance estimate of the cost (52). This would be particularly easily implemented in optimization procedures like CMA-ES, in which only cost rank is used rather than the exact cost value.




Fig. S2. Photograph of the experimental setup.

The exoskeleton emulator system consists of: (1) a large off-board motor, actuating one end of a Bowden cable; (2) a tether, comprising a Bowden cable for transmitting mechanical force from the motor to the exoskeleton, and electrical wires for transmitting sensor information from the load cell, ankle joint encoder and heel switch to the control computer; (3) an ankle exoskeleton worn on the right leg in the main study (cf. Fig. 2D; fig. S3), or ankle exoskeletons worn on both legs in most of the single-subject studies (fig. S4); and (4) a real-time control computer that measured sensor information, computed desired ankle exoskeleton torque based on the control law being tested, and sent motor velocity commands to the motor drive to achieve the desired torque. Participants walked on (5) an instrumented split-belt treadmill with (6) emergency stop buttons for both the treadmill and the exoskeleton while wearing (7) a safety harness to prevent falls and (8) a wireless respirometry system, consisting of a mask and small backpack, to measure metabolic rate. The experimenter used (9) an interface computer to control the treadmill, set low-level control parameters, record data, start and stop the optimization, and set control laws during validation tests.



Fig. S3. Photograph of the ankle exoskeleton used in the main experiment.

The exoskeleton was worn on the right ankle. A load cell at the end of the ankle lever measured tension in the drive rope, which was used along with the length of the ankle lever arm to calculate exoskeleton torque. The ankle lever arm was made from fiberglass, forming a leaf spring that increased series elasticity between the motor and the ankle joint, which can improve torque tracking. A rotary encoder at the exoskeleton ankle joint was used to measure joint angle. A contact switch inside the heel of the shoe was used to sense the instant of heel strike at the beginning of the stride. The exoskeleton design is described in more detail in (30). A self-adhering sports wrap was used on the leg beneath the shank strap to prevent slipping and to improve comfort at the interface. The Bowden cable tether was routed to the side of the knee and hip joints, such that flexing of the cable corresponded to joint rotation, which minimizes interference with normal leg movements (35). The tether was routed away from the body to leave enough room for normal arm swinging, which also plays an important role in economical gait (53). The tether was supported by rubber bands attached to the handrails of the treadmill. The exoskeleton was integrated with a running shoe, and the matching unaltered shoe was worn on the other foot. Shoes were swapped out to accommodate participants of varying size, from 1.60 to 1.88 m height.

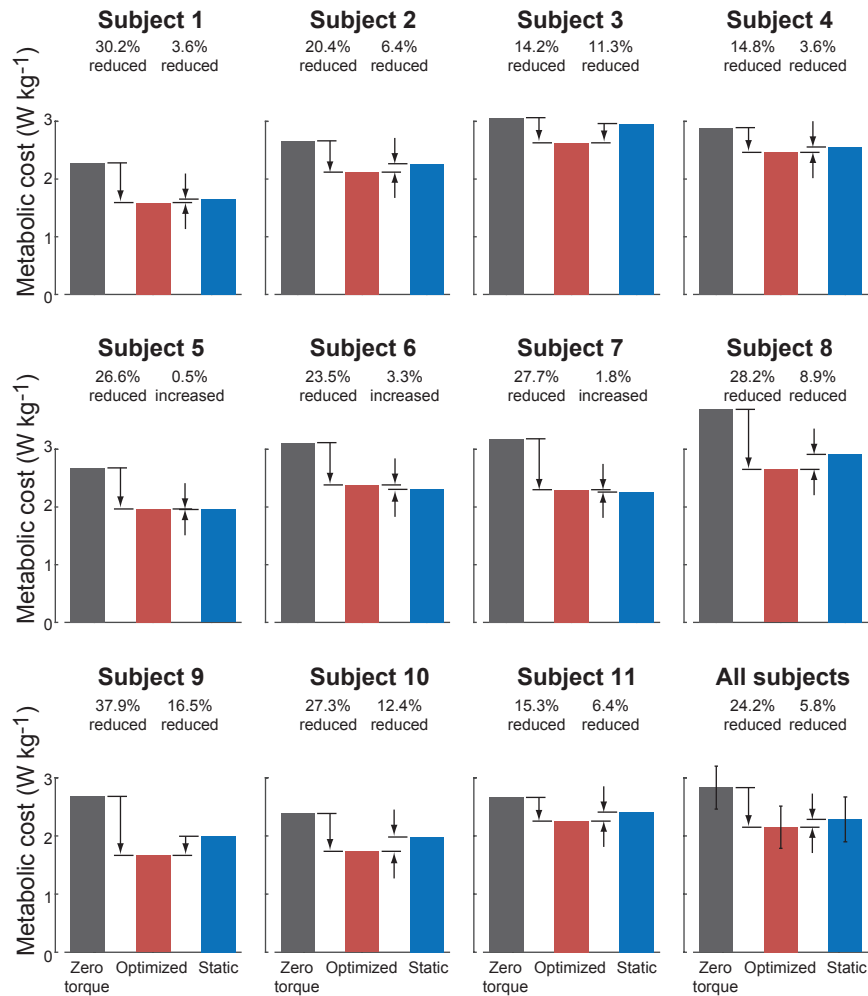


Fig. S4. Metabolic cost of walking for individual subjects in the main study.

Optimized assistance reduced metabolic rate compared to the zero-torque condition for all subjects, with an average reduction of $24.2 \pm 7.4\%$ (t-test: $P = 1 \cdot 10^{-6}$; $N = 11$) and a range of 14.2-37.9%. Optimized assistance led to lower metabolic rate compared to the static condition for eight of eleven subjects, with an average reduction in metabolic rate of $5.8 \pm 6.2\%$ (t-test: $P = 0.01$), ranging from a 3.3% increase to a 16.5% reduction. Cases in which metabolic rate was lower with static assistance may be explained by similarity between static and optimized control laws for these participants combined with noise in the measurement of metabolic rate. Bars are means of net metabolic rate from each presentation of conditions in double-reversal validation trials.



Fig. S5. Photograph of the ankle exoskeletons used in single-subject studies.

The exoskeletons used in most single-subject studies had strain gages on the ankle levers to measure applied torque. These exoskeletons had no series springs and relied on compliance in the Bowden cable tether, shank straps and heel rope to provide series elasticity. Rotary encoders at the exoskeleton ankle joints were used to measure joint angles. Contact switches inside the heels of the shoes were used to sense the instant of heel strike at the beginning of each stance period. The exoskeleton design is described in more detail in (30). Self-adhering sports wraps were used on the legs beneath the shank straps to prevent slipping and improve comfort at the interface. In the running study, additional sports wraps were used over the shank strap to increase resistance to downward migration. The Bowden cable tethers were routed to the side of the knee and hip joints, such that flexing of the cable corresponded to joint rotation, which minimizes interference with normal leg movements (35). The tether was routed away from the body to leave enough room for normal arm swinging, which also plays an important role in economical gait (53). Tethers were supported by rubber bands connected to the handrails of the treadmill.




Fig. S6. Optimized exoskeleton torque patterns for single-subject walking studies.

(A) Optimized bilateral ankle exoskeleton torque patterns for slow, normal, and fast walking. Slow walking (0.75 m s^{-1}) resulted in optimized control parameters with almost zero torque: $[\tau_p, t_p, t_r, t_f] = [4 \text{ N m}, 51\%, 30\%, 8\%]$. During validation, we included the initial control law to allow direct comparison with the optimized control law; energy cost was 22% higher with the initialization parameters, confirming that driving the torque to zero improved metabolic energy cost at this speed. Normal walking (1.25 m s^{-1}) resulted in optimized control parameters $[\tau_p, t_p, t_r, t_f] = [52 \text{ N m}, 55\%, 26\%, 10\%]$. Fast walking (1.75 m s^{-1}) resulted in optimized parameters of $[\tau_p, t_p, t_r, t_f] = [66 \text{ N m}, 52\%, 30\%, 11\%]$. (B) Optimized bilateral ankle exoskeleton torque patterns for uphill and loaded walking. Walking uphill on an inclined treadmill (10% grade) resulted in optimized control parameters with later onset: $[\tau_p, t_p, t_r, t_f] = [57 \text{ N m}, 55\%, 13\%, 10\%]$. Loaded walking (20% body weight carried in a running vest) resulted in optimized control parameters of $[\tau_p, t_p, t_r, t_f] = [52 \text{ N m}, 53\%, 33\%, 12\%]$. Solid lines are measured torque of the right exoskeleton, normalized to stride time and averaged across strides. Measured torque of the left exoskeleton was nearly identical. Optimized torque patterns varied widely across walking conditions and spanned a large portion of the allowable space.




Fig. S7. Optimized exoskeleton torque pattern for the single-subject running study. Optimized control parameters were $[\tau_p, t_{on}, t_{pr}, t_{off}] = [55 \text{ N m}, 7\%, 75\%, 95\%]$. The solid line is measured torque of the right exoskeleton, normalized to stride time and averaged across strides. Measured torque of the left exoskeleton was nearly identical. This pattern of torque had a much earlier onset than the optimized patterns for walking.

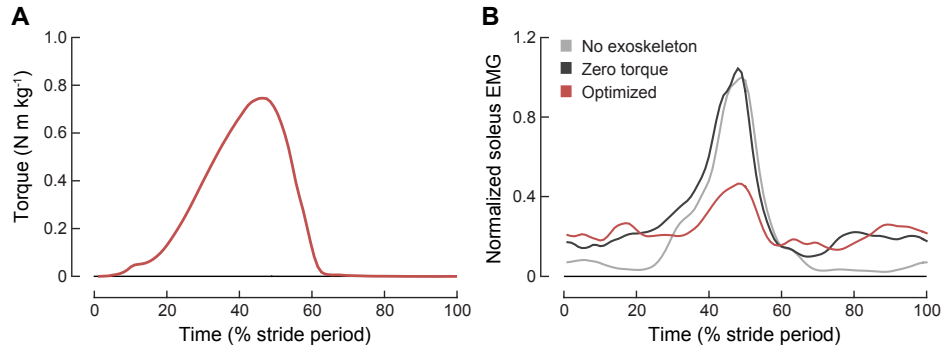


Fig. S8. Results from the single-subject study on minimizing muscle activity.

(A) Optimized ankle exoskeleton torque pattern when the objective function was muscle activity. The optimized parameters were $[\tau_p, t_p, t_r, t_f] = [58 \text{ N m}, 44\%, 35\%, 17\%]$, which encodes a substantially earlier onset than when optimizing on metabolic rate at this speed. The solid line is measured torque, normalized to stride time and averaged across strides.

(B) Processed electromyography (EMG) patterns for normal walking without the exoskeleton, walking in the zero-torque condition, and walking with optimized assistance. Data are from the soleus on the same leg as the exoskeleton. Solid lines are processed electromyography patterns of the lateral aspect of the soleus, normalized to stride time and averaged across strides. Data are normalized to the peak of the average pattern for normal walking without the exoskeleton. The magnitude of the electromyography pattern was substantially reduced during the stance period when optimized exoskeleton assistance was applied.




Fig. S9. Check of convergence to optimized metabolic rate.

We performed a test of convergence with a subset of participants ($N = 8$) in the main study by continuing the optimization for an additional four generations. Additional optimization did not lead to further reductions in metabolic cost; energy cost with the convergence-check control law was slightly higher than with the optimized control law (1.9% difference compared to zero-torque; paired t-test: $P = 0.03$) and optimized parameters changed little (table S3). Bars are means, error bars are standard deviations.

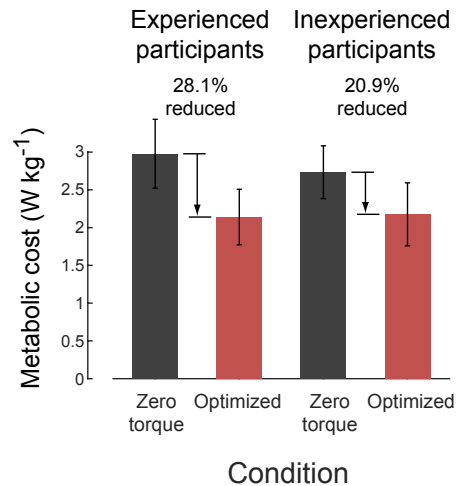


Fig. S10. Results for experienced versus inexperienced participants.

Participants with prior experience using an exoskeleton (N = 5) seemed to obtain greater improvements in energy economy compared to inexperienced participants (N = 6). Prior exoskeleton experience seemed to result in lower optimized metabolic rate, with an average reduction compared to the zero-torque condition that was 7.3% larger, although the comparison was not statistically significant (unpaired t-test: P = 0.1). This suggests that longer-term adaptation, neural reorganization or even growth might play a role in maximizing the benefits of exoskeleton assistance. Bars are means, error bars are standard deviations.

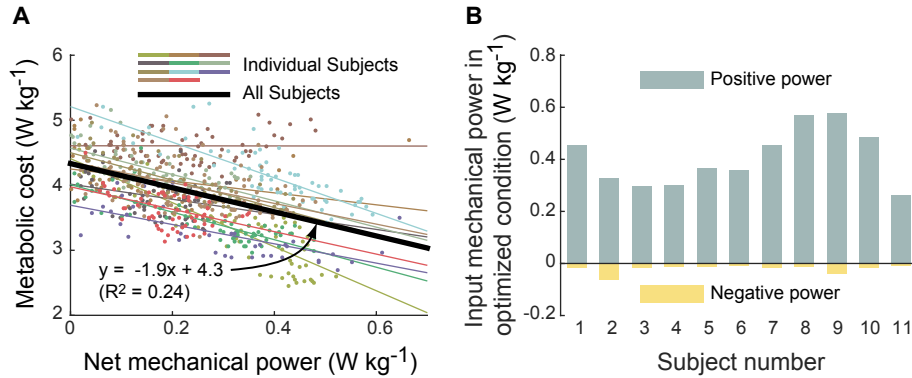


Fig. S11. Exoskeleton mechanical power.

Metabolic rate was negatively correlated with net exoskeleton power during optimization and all optimized control laws encoded substantial net exoskeleton power in the main study, suggesting that positive net power may be a common component of optimal ankle exoskeleton assistance for these gait conditions and participants. However, net power was not maximized during optimization, and static assistance resulted in higher net mechanical power than optimized assistance for all participants ($0.45 \pm 0.09 \text{ W kg}^{-1}$ vs. $0.38 \pm 0.11 \text{ W kg}^{-1}$, or 23% higher; t-test: $P = 0.006$; $N = 11$) but also resulted in higher metabolic energy cost. Optimizing assistance was therefore not equivalent to maximizing net exoskeleton power, even for an average controller. Net power with optimized control also varied greatly across participants, ranging from 0.25 to 0.54 W kg^{-1} . Both of these findings are inconsistent with the theory that the net benefit of an assistive device equals the balance between the net mechanical work provided and the cost of carrying added mass, sometimes called an ‘augmentation factor’ (4). **(A)** Scatter plot of metabolic rate versus net exoskeleton power, defined as the integral of instantaneous ankle exoskeleton power over the stride divided by stride period, normalized to participant body mass, for all control laws tested during optimization trials in the main study. We observed a correlation between net exoskeleton power and metabolic rate for all but one participant. On average, this relationship explained about 24% of the variance in metabolic rate. Note that this is not a strong test for a relationship between metabolic rate and mechanical power because the power values arise from a separate process that also affects metabolic rate. **(B)** Positive and negative mechanical power, defined as positive or negative work per stride divided by stride period, normalized to body mass, measured during the optimized control law condition in validation trials. Optimized assistance encoded substantial positive power for all participants, but the magnitudes varied widely. Some optimized control laws encoded negative mechanical power, but most were associated with little dissipation. Net mechanical power values for all conditions are available in Study Data.

Table S1. Participant characteristics.

Relevant characteristics of all study participants.

Participant	Gender	Body mass (kg)	Height (m)	Age (yrs)	Exoskeleton experience
Main Study					
Subject 1	F	57	1.63	28	NO
Subject 2	M	80	1.84	37	YES
Subject 3	F	50	1.60	27	NO
Subject 4	M	83	1.88	22	NO
Subject 5	M	86	1.80	28	YES
Subject 6	M	70	1.73	26	NO
Subject 7	F	68	1.78	26	YES
Subject 8	F	59	1.68	24	YES
Subject 9	F	55	1.65	32	YES
Subject 10	M	65	1.80	24	NO
Subject 11	M	93	1.88	25	NO
Mean ± SD	6 M, 5 F	69.6 ± 14.1	1.75 ± 0.10	27.2 ± 4.2	5 Y, 6 N
Single-Subject Studies					
Subject 12	M	97	1.93	22	NO
Subject 13	M	77	1.75	25	NO

Table S2. Experimental protocol for the main study.

Protocol details for each participant. In validation order sequences, ‘O’ refers to the optimized condition, ‘Z’ refers to the zero-torque condition, ‘S’ refers to the static condition, ‘C’ refers to the convergence-check condition, ‘P’ refers to the pre-reset condition just before reset (Subject 6 only), and D refers to the alternate-day validation check (Subject 8 only). In single-subject studies, participants completed each gait condition in one day, with four generations of optimization (eight for slow walking) followed by the double-reversal random-order validation tests on the same day, including optimized, zero-torque and no exoskeleton (and initial parameters for slow walking). Please see [Materials and Methods](#) for details.

Participant	Day	Protocol description
Subject 1	Day 1:	Optimization (2 generations)
	Day 2:	Optimization (2 remaining generations) Convergence check (3 additional generations)
	Day 3:	Convergence check (1 additional generation) Validation (order: ZSCOOC SZ)
Subject 2	Day 1:	Optimization (4 generations) Validation (order: ZOZ)
	Day 2:	Convergence check (4 additional generations) Convergence validation (order: CZSSZC)
Subject 3	Day 1:	Optimization (4 generations)
	Day 2:	Convergence check (4 additional generations)
	Day 3:	Validation (order: ZSCOOC SZ)
Subject 4	Day 1:	Optimization (4 generations) Convergence check (1 additional generations)
	Day 2:	Convergence check (3 additional generations) Validation (order: ZSCOOC SZ)
Subject 5	Day 1:	Optimization (4 generations)
	Day 2:	Convergence check (4 additional generations)
	Day 3:	Validation (order: ZSCOOC SZ)
Subject 6	Day 1:	Optimization (4 generations)
	Day 2:	Reset optimization (4 generations)
	Day 3:	Validation (order: ZSOPPOSZ)
Subject 7	Day 1:	Optimization (3 generations) Convergence check (1 additional generation)
	Day 2:	Convergence check (4 additional generations)
	Day 3:	Validation (order: ZCSOOSCZ)

Subject 8	Day 1:	Optimization (2 generations)
	Day 2:	Optimization (3 remaining generations) Validation (order: ZOZ)
	Day 3:	Convergence check (3 additional generations)
	Day 4:	Convergence validation (order: ZSCDDCSZ)
Subject 9	Day 1:	Optimization (4 generations) Convergence check (1 additional generation)
	Day 2:	Validation (order: SOZOS)
Subject 10	Day 1:	Optimization (4 generations)
	Day 2:	Continued optimization (5 addl. generations) Reset optimization (4 generations)
	Day 3:	Validation (order: ZSOOSZ)
Subject 11	Day 1:	Optimization (4 generations)
	Day 2:	Convergence check (4 additional generations)
	Day 3:	Validation (order: ZSCOOC SZ)

Table S3. Validation trial settings and outcomes for the main study.

Control law settings used in validation trials, with the corresponding measured metabolic rate. Metabolic rate is the average of both double-reversal presentations. All data are available in the Study Data file archive.

Participant	Control law parameters [τ_p, t_p, t_r, t_f] [N m, % Stride, % Stride, % Stride]		Metabolic rate (W kg ⁻¹)
Subject 1	Zero torque:	N.A.	2.28
	Static:	[46.3, 53.0, 17.0, 8.0]	1.65
	Optimized:	[43.5, 50.1, 19.1, 13.2]	1.59
	Convergence check:	[46.6, 51.3, 18.3, 11.8]	1.59
Subject 2	Zero torque:	N.A.	2.66
	Optimized:	[67.5, 50.1, 29.6, 11.8]	2.12
	Alt-day zero torque	N.A.	2.71
	Alt-day static:	[65.3, 53.0, 25.0, 10.0]	2.31
	Alt-day conv. check:	[66.2, 50.8, 25.4, 9.7]	2.19
Subject 3	Zero torque:	N.A.	3.06
	Static:	[40.8, 53.0, 16.8, 8.0]	2.96
	Optimized:	[31.9, 52.8, 21.2, 6.23]	2.63
	Convergence check:	[26.6, 51.3, 13.6, 10.5]	2.77
Subject 4	Zero torque:	N.A.	2.89
	Static:	[67.4, 53.0, 17.2, 8.0]	2.56
	Optimized:	[52.0, 52.6, 27.4, 11.0]	2.46
	Convergence check:	[49.8, 51.1, 27.6, 10.2]	2.48
Subject 5	Zero torque:	N.A.	2.68
	Static:	[70.5, 53.0, 17.2, 8.0]	1.96
	Optimized:	[72.8, 49.9, 24.5, 10.3]	1.97
	Convergence check:	[71.7, 50.0, 34.0, 12.7]	2.01
Subject 6	Zero torque:	N.A.	3.12
	Static:	[57.1, 53.0, 17.2, 8.0]	2.31
	Pre-reset:	[32.5, 54.9, 22.1, 9.42]	2.41
	Optimized:	[56.1, 49.2, 11.7, 12.8]	2.38
Subject 7	Zero torque:	N.A.	3.18
	Static:	[55.5, 53.0, 17.2, 8.0]	2.26
	Optimized:	[56.0, 48.4, 25.3, 14.8]	2.30
	Convergence check:	[52.7, 51.2, 25.1, 13.7]	2.31
Subject 8	Zero torque:	N.A.	3.69
	Optimized:	[49.9, 50.0, 32.8, 14.7]	2.65

	Alt-day zero torque:	N.A.	3.54
	Alt-day static:	[48.1, 53.0, 17.0, 8.0]	2.79
	Alt-day optimized:	[48.2, 50.6, 33.6, 13.8]	2.88
	Alt-day conv. check:	[53.6, 49.5, 33.5, 15.4]	2.94
Subject 9	Zero torque:	N.A.	2.68
	Static:	[44.9, 53.0, 17.0, 8.0]	1.99
	Optimized:	[47.2, 48.5, 23.2, 15.9]	1.67
Subject 10	Zero torque:	N.A.	2.39
	Static:	[53.0, 53.0, 17.0, 8.0]	1.98
	Optimized:	[45.3, 51.1, 22.8, 13.1]	1.74
Subject 11	Zero torque:	N.A.	2.66
	Static:	[75.9, 53.0, 17.4, 8.0]	2.41
	Optimized:	[67.8, 51.2, 17.7, 11.8]	2.25
	Convergence check:	[58.5, 51.6, 13.7, 11.0]	2.39

Table S4. Quadratic approximations of the energy cost landscape.

For each participant in the main study, we used metabolic rate estimates from all control laws presented during the optimization phase of the experiment to conduct a model-based offline optimization. We assumed a four-variable quadratic function without interactions, $\dot{E} = c_1\tau_p^2 + c_2\tau_p + c_3t_p^2 + c_4t_p + c_5t_r^2 + c_6t_r + c_7t_f^2 + c_8t_f + c_9$, and solved for the coefficients, $[c_1 \ c_2 \ \dots \ c_9]$, that resulted in the least squared error. We then identified the quadratic model-based estimate of optimal parameter values, limited to the range of allowable parameter values. This constraint was frequently encountered, because many of the quadratic approximations had negative coefficients on quadratic terms. We then calculated the percent error of the estimated optimal values, defined as the difference between the quadratic model-based estimate of the optimal parameters and the experimentally optimized parameter values, divided by the range of allowable values. Average errors were [24%, 16%, 21%, 37%].

Participant	Quadratic approximation	R ²	Model optima	Error (% of range)
Subject 1	$-0.0012\tau_p^2 + 0.0576\tau_p$ $+0.0008t_p^2 - 0.1224t_p$ $+0.0035t_r^2 - 0.1427t_r$ $+0.0099t_f^2 - 0.2313t_f$ $+9.8160$	0.7	[56.5, 55.0, 20.5, 11.6]	[23, 11, 5, 11]
Subject 2	$0.0005\tau_p^2 - 0.0587\tau_p$ $+0.0023t_p^2 - 0.2708t_p$ $+0.0012t_r^2 - 0.0516t_r$ $+0.0035t_f^2 - 0.0996t_f$ $+14.4496$	0.5	[58.5, 55.0, 21.0, 14.2]	[11, 11, 29, 16]
Subject 3	$0.0006\tau_p^2 - 0.0330\tau_p$ $+0.0055t_p^2 - 0.5585t_p$ $+0.0001t_r^2 - 0.0027t_r$ $+0.0002t_f^2 + 0.0118t_f$ $+18.9850$	0.2	[27.5, 50.5, 11.0, 5.0]	[9, 5, 34, 8]

Subject 4	$0.0001\tau_p^2 - 0.0116\tau_p$ $-0.0002t_p^2 - 0.0168t_p$ $+0.0018t_r^2 - 0.1058t_r$ $+0.0013t_f^2 - 0.0513t_f$ $+7.4808$	0.5	[68.5, 55.0, 29.5, 19.4]	[20, 5, 7, 56]
Subject 5	$-0.0004\tau_p^2 + 0.0419\tau_p$ $-0.0011t_p^2 + 0.0359t_p$ $-0.0000t_r^2 + 0.0057t_r$ $+0.0016t_f^2 - 0.0731t_f$ $+3.5803$	0.7	[2.0, 55.0, 10.0, 20.0]	[82, 11, 48, 65]
Subject 6	$-0.0004\tau_p^2 + 0.0298\tau_p$ $+0.0057t_p^2 - 0.6132t_p$ $-0.0003t_r^2 + 0.0099t_r$ $+0.0016t_f^2 - 0.0771t_f$ $+20.6465$	0.7	[2.0, 54.0, 40.0, 20.0]	[44, 2, 60, 71]
Subject 7	$0.0001\tau_p^2 - 0.0169\tau_p$ $+0.0032t_p^2 - 0.3511t_p$ $+0.0007t_r^2 - 0.0299t_r$ $-0.0004t_f^2 - 0.0228t_f$ $+14.6816$	0.8	[66.0, 55.0, 21.5, 20.0]	[15, 15, 13, 35]
Subject 8	$0.0001\tau_p^2 - 0.0373\tau_p$ $+0.0033t_p^2 - 0.3434t_p$ $+0.0004t_r^2 - 0.0293t_r$ $-0.0010t_f^2 + 0.0241t_f$ $+14.7401$	0.8	[59.0, 52.5, 36, 20.0]	[18, 4, 8, 41]

Subject 9	$0.0006\tau_p^2 - 0.0485\tau_p$ $+0.0062t_p^2 - 0.6346t_p$ $+0.0013t_r^2 - 0.0610t_r$ $0.0035t_f^2 - 0.1297t_f$ $+21.9363$	0.7	[41, 51.5, 23.5, 18.2]	[11, 7, 1, 15]
Subject 10	$-0.0003\tau_p^2 + 0.0115\tau_p$ $-0.0027t_p^2 + 0.1986t_p$ $+0.0002t_r^2 - 0.0092t_r$ $-0.0014t_f^2 - 0.0142t_f$ $+1.3776$	0.5	[65.0, 10.0, 28.5, 20.0]	[30, 91, 19, 46]
Subject 11	$0.0011\tau_p^2 - 0.1334\tau_p$ $+0.0048t_p^2 - 0.5456t_p$ $+0.0006t_r^2 - 0.0187t_r$ $+0.0029t_f^2 - 0.1004t_f$ $+24.0035$	0.7	[63.5, 55.0, 14.5, 17.6]	[5, 8, 11, 39]

Study Data (separate file)

Complete study data are provided in the file archive StudyData.zip. Data from both the optimization and validation phases of both the main experiment and single-subject studies are found in the Matlab archives MainStudy.mat and SingleSubjectStudies.mat. Optimization data include subject, generation number, generation mean, trial day, control law parameter values, net mechanical power, root mean square muscle activity, and estimated steady-state metabolic rate, as appropriate, for each control law tested. Validation data include subject, condition name, control law parameter values, net mechanical power, average measured exoskeleton torque trajectory, muscle activity trajectory, root mean square muscle activity, and estimated steady-state metabolic rate, as appropriate, for each condition tested during validation. The readme.txt file provides a more detailed description of the file structure and examples of how to access and process data. Matlab code in the sample_processing.m file demonstrates how to access and process study data, including recreating all outcomes and figures from the main study and single-subject studies reported in the main text and supplementary materials.

Optimization Code (separate file)

Sample versions of the code that comprise the optimization method are provided in the file archive OptimizationCode.zip. Included are sample scripts for applying parametric constraints and generating torque curves, estimating steady-state metabolic rate from breath by breath data from each control law, and applying the CMA-ES algorithm to determine the next generation of control laws to test. Code is provided in Matlab format, with instructions and descriptions as comments within each file.