

## ROBOTICS

# Human-in-the-loop optimization of exoskeleton assistance during walking

Juanjuan Zhang,<sup>1,2</sup> Pieter Fiers,<sup>1</sup> Kirby A. Witte,<sup>1</sup> Rachel W. Jackson,<sup>1</sup> Katherine L. Poggensee,<sup>1</sup> Christopher G. Atkeson,<sup>3</sup> Steven H. Collins<sup>1,3\*</sup>

Exoskeletons and active prostheses promise to enhance human mobility, but few have succeeded. Optimizing device characteristics on the basis of measured human performance could lead to improved designs. We have developed a method for identifying the exoskeleton assistance that minimizes human energy cost during walking. Optimized torque patterns from an exoskeleton worn on one ankle reduced metabolic energy consumption by  $24.2 \pm 7.4\%$  compared to no torque. The approach was effective with exoskeletons worn on one or both ankles, during a variety of walking conditions, during running, and when optimizing muscle activity. Finding a good generic assistance pattern, customizing it to individual needs, and helping users learn to take advantage of the device all contributed to improved economy. Optimization methods with these features can substantially improve performance.

For more than a century, inventors and scientists have developed exoskeletons and active prostheses intended to improve human locomotor performance, particularly in terms of energy economy (1). Few approaches have been successful (2–8), however, and only modest enhancements have been achieved compared to the potential benefits expected on the basis of simulations (9–11). An overreliance on intuition and specialized hardware may be partially responsible for these shortcomings. Assistance strategies have typically been derived from mathematical models (12), biomechanics observations (13), and humanoid robots (14), but each simplifies important aspects of the human-robot system (15). Experiments have primarily been conducted using specialized prototypes that embed a single intuited functionality, with each prototype requiring years of development, limiting exploration to only a small set of potential assistance strategies. Compounding the challenge, physiological and neurological differences between individuals can cause divergent responses to the same device (16–18), and responses can change considerably during the course of adaptation (19, 20).

Methods for automatically discovering, customizing, and continuously adapting assistance could overcome these challenges, allowing exoskeletons and prostheses to achieve their potential. We have developed approaches in which device control is systematically varied during use so as to maximize human performance, which we call “human-in-the-loop optimization” (Fig. 1A). Such techniques can take inspiration from humans, who naturally

optimize their coordination patterns for energy cost and other aspects of locomotor performance (21). However, closing the loop on human performance is challenging. Objective functions based on measurements of human performance typically require lengthy evaluation periods and contain substantial noise. The best available estimate of metabolic energy cost, for example, requires about one minute of respiratory data per evaluation (22). The human portion of the system also has time-varying dynamics that make optimization difficult, including slow components of adaptation (19) and strong history dependence (20), reflecting complex neurocognitive factors (23). Control laws that are general enough to approximate globally optimal assistance strategies are likely to require multiple parameters per assisted joint (10), resulting in high-dimensional optimization problems. Initial efforts in this domain have demonstrated the ability to optimize a single gait or device parameter using line search (24) or gradient descent (25), but these methods are inefficient, being sensitive to drift and noise, and scale poorly, requiring many more evaluations for each new parameter to be optimized, particularly in the presence of parametric interactions (26). Many optimization methods that work well in simulation (27) are subject to these limitations; building an approximation of the system takes time, and the human changes during that time.

We have developed a sample-efficient method of identifying the exoskeleton control parameters that minimize the metabolic energy cost of human walking (Fig. 1B). During optimization, the user walks while the exoskeleton applies assistance. The exoskeleton periodically changes the pattern of assistance, defined by a control law, while metabolic rate is measured. Steady-state metabolic energy cost is estimated for each control law by fitting a first-order dynamical model to 2 min of transient metabolic data (fig. S1) (26).

After a prescribed number of control laws have been evaluated, forming one generation, a covariance matrix adaptation evolution strategy (CMA-ES) (28) is used to calculate the next generation of control laws to be tested. The mean of each new generation represents the best estimate of the optimal control parameter values, and the shape and size of the distribution are chosen to increase the likelihood of further improvement in subsequent generations. This optimization strategy is relatively tolerant of both measurement noise and human adaptation, because neither objective function values nor their derivatives are used directly, and each generation is evaluated independently. It scales well in benchmarking problems (28), with a suggested generation size that increases with the logarithm of the number of control parameters.

We tested our method by optimizing the pattern of assistive torque applied by an exoskeleton worn on one ankle during walking. We applied assistance at one ankle to allow comparisons to a prior study that used the same hardware (17). Ankle torque was determined by four parameters: peak torque, timing of peak torque, and rise and fall times (Fig. 2A) (26). This allowed for a wide range of possible torque patterns (Fig. 2B), including patterns approximating those previously found to be beneficial (3–5). This parameterization also implicitly allowed adjustment of features such as the timing and amount of positive joint work, which may be important to energy economy (29). Some torque patterns were not possible with this parameterization, such as those with multiple peaks. More complex patterns, defined by additional parameters, might allow better approximations of global optima at the cost of lengthier optimization periods. A generation size of eight was chosen on the basis of (28), and a target of four generations was set on the basis of pilot results (26). Torque was applied to one ankle using a versatile exoskeleton emulator system (30) (Fig. 2, C and D, and figs. S2 and S3) with precise low-level torque control (31). The emulator, inspired in part by other laboratory-based testbeds (32–34), allows a wide range of assistive behaviors to be applied in rapid succession, without the need to design or build new hardware (35).

We optimized assistance for 11 participants (subjects 1 to 11; table S1) as they walked on a treadmill at a normal speed ( $1.25 \text{ m s}^{-1}$ ). After optimization, we performed validation tests comparing optimized assistance with a fully passive “zero-torque” mode and with a “static” assistance condition. Static assistance approximated the best hand-tuned torque pattern for this device, which had previously resulted in a 6% reduction in energy cost compared to zero torque (17). The double-reversal validation test prevented confounding influences from measurement noise during optimization and trial order during validation (26). Participants were not exposed to any of the validation conditions during optimization, because optimized assistance was the weighted average of the best control laws from the final generation (26). The primary outcome was the energy cost of walking, defined as gross

<sup>1</sup>Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA. <sup>2</sup>School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore. <sup>3</sup>Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA.

\*Corresponding author. Email: stevecollins@cmu.edu

metabolic rate during walking minus the rate measured while standing still.

Optimized assistance substantially improved energy economy for all participants, confirming the effectiveness of the method. Optimized parameters were identified after four generations (64 min of walking) for all but two participants, who appeared to become trapped in local minima, requiring a reset of the algorithm and additional walking (128 and 208 min total; table S2). Optimized assistance reduced the metabolic cost of walking to  $2.16 \pm 0.38 \text{ W kg}^{-1}$ , down from  $2.84 \pm 0.40 \text{ W kg}^{-1}$  with zero torque (mean  $\pm$  standard deviation). Energy cost reductions ranged from


14.2 to 37.9% (fig. S4 and table S3), with an average reduction of  $24.2 \pm 7.4\%$  ( $t$  test,  $P = 1 \times 10^{-6}$ ,  $n = 11$ ; Fig. 3A). By the same measure, the largest average reduction provided by hand-tuned ankle exoskeletons has been 14.5%, with devices worn on both ankles (4), and the largest reduction for any exoskeleton has been 22.8%, with an exosuit acting at both hips and both ankles (8). Wearing the exoskeleton in zero-torque mode is about 10% more costly than walking in normal shoes without the exoskeleton (Fig. 3B) (17), suggesting about a 14% net improvement with optimized assistance compared to normal shoes. Extrapolating these results to an autonomous device is difficult,

with expected benefits from streamlining for a single control law and expected costs for carrying motors and batteries. However, a rough estimate (26) obtained by scaling an established autonomous exoskeleton (4) suggests improved performance with optimized assistance.


Optimized assistance patterns varied widely across participants (Fig. 3C), demonstrating the importance of customization. For example, the timing of optimized torque onset ranged from onset at 17% to onset at 37% of the stride period (Fig. 3D), or about half the testable range in this and prior studies (3). Optimized assistance did not fully replace human ankle torque, nor did it provide the maximum possible positive mechanical work, inconsistent with some predictions (4, 10). Optimized torque patterns (Fig. 3D) did share some qualitative features with each other, such as a peak torque that occurred at about 50% of stride, suggesting qualities that may be beneficial for most people and useful initial parameters for future optimizations. However, even subtle differences in torque can have large and unexpected effects on energy use, arising from complex interactions with the musculoskeletal and nervous systems (36, 37). Human-in-the-loop optimization accommodates this complexity.

Comparisons involving the static control law suggest three separate contributions to the success of this approach: discovering a good generic assistance pattern, customizing it to individual users, and facilitating motor learning. Static assistance (Fig. 3E) was similar to the average of the optimized control laws, and in our prior study (17), it delivered a 6% reduction in energy cost. This supports the idea that, once discovered, a good generic control law can be beneficial. Individually optimized assistance resulted in  $5.8 \pm 6.2\%$  lower metabolic rate than with the static control law ( $t$  test,  $P = 0.01$ ,  $n = 11$ ). Eight of 11 participants had a lower metabolic rate with optimized assistance, with the difference in rate ranging from a 3.3% increase to a 16.5% reduction. This confirms the benefits of customization. Static assistance resulted in a  $19.3 \pm 8.6\%$  reduction in energy cost compared to zero torque ( $t$  test,  $P = 4 \times 10^{-5}$ ,  $n = 11$ ), a larger improvement than in our prior study (17). The primary differences between these studies relate to the conditions during adaptation, which indicates an important role for facilitating motor learning. Participants had a similar duration of exposure to the exoskeleton in both studies, but in the prior study, participants were trained with a narrow range of eight static control laws, whereas during optimization, they experienced 32 diverse control laws. Some of these wide-ranging control laws, some of which participants noted were initially uncomfortable, may have forced them to explore new motor control strategies, which has been shown to be a necessary part of skill acquisition in some interventions (20).

This experiment was not designed to isolate the relative contributions of generic assistance, customization, or facilitating learning, but all were important to the large overall improvement in energy economy, and each could dominate



**Fig. 1. Human-in-the-loop optimization.** (A) Measurements of human performance are used to update device control so as to improve performance in the human portion of the system. (B) A method for minimizing the energy cost of human walking, in which various control laws are applied, metabolic (met.) rate is quickly estimated (est.) for each, costs are compared, and an evolution strategy is used to generate a new set of control laws to be tested, all during walking.  $p_1$  and  $p_2$  are hypothetical control parameters.




**Fig. 2. Control law and exoskeleton hardware.** (A) Parameterization of ankle torque. Each control law determined applied torque as a function of time, normalized to stride period, as a cubic spline defined by peak time, rise time, fall time, and peak torque. (B) Examples of possible torque patterns. (C) Exoskeleton emulator system. Off-board motor and control hardware actuated a tethered exoskeleton worn on one ankle while participants walked on a treadmill. (D) Ankle exoskeleton. Drive rope tension caused the device to push on the shank, heel, and toe contact, generating an ankle torque. High-quality images are shown in figs. S2 and S3.

in some contexts. With a new, simple device and a homogeneous population, the initial discovery of the best generic design would be critical. In a diverse patient population, benefits might come almost entirely from customization. With complex control architectures, facilitating motor learning could be decisive. In each case, human-in-the-loop optimization would provide a benefit, but for different reasons.

We demonstrated the generality of this approach in single-subject studies with a different device and several additional locomotion conditions. One new participant (subject 12; table S1) wore exoskeletons (30) on both ankles (fig. S5). Optimized assistance reduced the metabolic cost of walking at a typical speed ( $1.25 \text{ m s}^{-1}$ ; 33% reduction versus zero torque, 25% versus normal shoes; Fig. 4B), walking at a faster speed ( $1.75 \text{ m s}^{-1}$ ; 39% reduction versus zero torque, 30% versus normal shoes; Fig. 4C), walking up-hill (10% incline; 26% reduction versus zero torque, 21% versus normal shoes; Fig. 4D), and loaded walking (vest weighing 20% of body mass; 15% reduction versus zero torque, 4% versus normal shoes; Fig. 4E). At a slow walking speed ( $0.75 \text{ m s}^{-1}$ ), the algorithm drove torque to its lower limit, resulting in a small reduction in energy cost compared to zero torque (Fig. 4A) and a 19% reduction compared to the initial control law (fig. S6). These results demonstrate the effectiveness of the method across different walking conditions, including cases where the best action is none, and confirm expected improvements compared to normal walking. We also applied the approach to running with exoskeletons on both ankles (subject 2; table S1) and found a 27% improvement in energy cost compared to the zero-torque mode and 13% energy savings compared to normal running shoes (Fig. 4F). This demonstrates the effectiveness of the method for different gaits. With another new participant (subject 13; table S1), we applied the approach to the minimization of calf muscle activity, rather than metabolic rate. Altered muscle activity can be a useful performance measure, for example, in rehabilitation. Soleus activity in the optimized condition was reduced by 41% compared to the zero-torque mode and 36% compared to walking in normal shoes (Fig. 4G). This demonstrates the ability to address multiple physiological objectives. Detailed protocols and results are provided in (26) and figs. S6 to S8. These tests hint at the potential for a new type of biomechanics study, in which human-in-the-loop optimization can be leveraged to compare the best possible outcomes for different devices or gait conditions or to test how various features of gait change with optimized performance.

We performed a test of convergence with a subset of participants in the main study ( $n = 8$ ) by continuing the optimization for an additional four generations, and we found only small changes in optimized parameters (table S3) and no further reduction in metabolic rate (fig. S9). Participants with prior experience using an exoskeleton ( $n = 5$ ) seemed to have greater improvements in energy economy (fig. S10), suggesting that longer-time-




**Fig. 3. Experimental results.** (A) Metabolic energy cost of walking for each condition, tested in validation trials. Optimized assistance resulted in the lowest metabolic rate and a large reduction compared to the zero-torque condition. Variability is primarily due to differences between participants. Bars are means, error bars are standard deviations, and asterisks denote statistical significance ( $P < 0.05$ ). (B) Results from a prior experiment using the same hardware (17), comparing the zero-torque condition with walking in normal shoes (no exoskeleton) or with static assistance. Static assistance provided a smaller benefit in the prior study. Bars are means, and error bars are standard deviations. (C) Optimized control law parameter values. Optimized values varied widely across participants. Values are normalized to their allowable range (26). Lines are medians, boxes cover the 25th to 75th percentiles, and whiskers show the range. (D) Optimized ankle exoskeleton torque pattern for each participant. Patterns varied widely and spanned a large portion of the allowable range. Lines are measured torque, normalized to stride time and body mass, averaged across strides. (E) Torque applied in the static and zero-torque conditions. The static pattern, based on (17), is similar to the optimized patterns but resulted in higher metabolic rate. Torque was negligible in the zero-torque mode. Lines are measured torque, normalized to stride time and body mass, averaged across strides and participants.

scale adaptation might be at play. Optimized exoskeleton control laws all included substantial net-positive mechanical work, but amounts varied strongly across participants and were lower with optimized assistance than with static assistance (fig. S11). Minimizing human energy consumption was therefore not equivalent to maximizing exoskeleton mechanical work, further illustrating how difficult it is to reason about how exoskeletons should work and what characteristics will be optimal.

The evolutionary strategy that we used was more effective than other methods that we tried, but it seems likely that improved techniques could be developed. In early pilot tests of related methods, model-based optimization techniques were ineffective because of sensitivity to noise and adaptation dynamics. We have illustrated this problem by generating quadratic approximations of data for each participant in the main experiment, which explain little variance and make unreasonable estimates of optimal parameter values (table S4). New candidate algorithms should tolerate high measurement noise, facili-

tate human adaptation, and require very few evaluations before converging. It remains to be seen how the duration of the optimization period will scale with additional parameters for this and other algorithms.

We expect human-in-the-loop optimization to improve the effectiveness of assistive devices in many contexts. Successful optimization of a four-dimensional control law in a relatively short time suggests that optimizing exoskeletons and prostheses of greater complexity may be possible. Optimizing a similar number of parameters in a feedback control structure, such as a neuromuscular model (38), or switching between optimized modes (39) could enhance performance under changing locomotor conditions. Successfully reducing both metabolic rate and muscle activity suggests that alternate objective functions with similar properties could be optimized—for example, related to speed (40), endurance (41), balance, or overall satisfaction. Versatile emulator systems could be used to identify optimal device characteristics during a prescription process, and customized mobile devices could then be fabricated. In daily life,



**Fig. 4. Single-subject studies under alternate conditions.** (A) Slow walking ( $0.75 \text{ m s}^{-1}$ ). (B) Normal walking ( $1.25 \text{ m s}^{-1}$ ). (C) Fast walking ( $1.75 \text{ m s}^{-1}$ ). (D) Uphill walking (10% grade). (E) Loaded walking (load equal to 20% of body mass). (F) Running ( $2.68 \text{ m s}^{-1}$ ). All speeds, grades, loads, and gaits were with bilateral ankle exoskeletons. (G) Optimizing soleus muscle activity, rather than metabolic rate, during normal-speed walking. In each case, one participant was tested ( $n = 1$ ). The method identified assistance patterns that substantially improved the target outcome in all circumstances. Optimized torque patterns can be found in figs. S6 to S8.

a proxy measure such as heart rate or muscle activity (42) could be used for optimization, providing noisier but more abundant performance data. These approaches have scope to improve mobility for people with a wide range of distinct physiological needs, from individuals with chronic stroke to athletes.

#### REFERENCES AND NOTES

- W. Cornwall, *Science* **350**, 270–273 (2015).
- H. M. Herr, A. M. Grabowski, *Proc. R. Soc. B* **279**, 457–464 (2012).
- P. Malcolm, W. Derave, S. Galle, D. De Clercq, *PLOS ONE* **8**, e56137 (2013).
- L. M. Mooney, E. J. Rouse, H. M. Herr, *J. Neuroeng. Rehabil.* **11**, 151 (2014).
- S. H. Collins, M. B. Wiggin, G. S. Sawicki, *Nature* **522**, 212–215 (2015).
- J. R. Koller, D. A. Jacobs, D. P. Ferris, C. D. Remy, *J. Neuroeng. Rehabil.* **12**, 97 (2015).
- K. Seo, J. Lee, Y. Lee, T. Ha, Y. Shim, in *2016 IEEE International Conference on Robotics and Automation* (Institute of Electrical and Electronics Engineers, 2016), pp. 4628–4635.
- B. T. Quinnivan *et al.*, *Sci. Robot.* **2**, eaah4416 (2017).
- A. J. van den Bogert, *Biomed. Eng. Online* **2**, 17 (2003).
- M. L. Handford, M. Srinivasan, *Sci. Rep.* **6**, 19983 (2016).
- T. K. Uchida *et al.*, *PLOS ONE* **11**, e0163417 (2016).
- S. H. Collins, A. D. Kuo, *PLOS ONE* **5**, e9307 (2010).
- W. van Dijk, H. van der Kooij, *Rob. Autom. Mag.* **21**, 56–61 (2014).
- H. Kazerooni, R. Steger, L. Huang, *Int. J. Robot. Res.* **25**, 561–573 (2006).
- A. J. Ijspeert, *Science* **346**, 196–203 (2014).
- K. E. Zelik *et al.*, *IEEE Trans. Neural Syst. Rehabil. Eng.* **19**, 411–419 (2011).
- R. W. Jackson, S. H. Collins, *J. Appl. Physiol.* **119**, 541–557 (2015).
- R. E. Quesada, J. M. Caputo, S. H. Collins, *J. Biomech.* **49**, 3452–3459 (2016).
- K. E. Gordon, D. P. Ferris, *J. Biomech.* **40**, 2636–2644 (2007).
- J. C. Selinger, S. M. O'Connor, J. D. Wong, J. M. Donelan, *Curr. Biol.* **25**, 2452–2456 (2015).
- R. M. Alexander, *Physiol. Rev.* **69**, 1199–1227 (1989).
- J. C. Selinger, J. M. Donelan, *J. Appl. Physiol.* **117**, 1406–1415 (2014).
- T. R. Makin, F. de Vignemont, A. A. Faisal, *Nat. Biomed. Eng.* **1**, 0014 (2017).
- W. Felt, J. C. Selinger, J. M. Donelan, C. D. Remy, *PLOS ONE* **10**, e0135342 (2015).
- J. R. Koller, D. H. Gates, D. P. Ferris, C. D. Remy, “Body-in-the-loop optimization of assistive robotic devices: A validation study,” paper presented at Robotics: Science and Systems XII, Ann Arbor, MI, 18 to 22 June 2016; available at [www.roboticsproceedings.org/rss12/p07.pdf](http://www.roboticsproceedings.org/rss12/p07.pdf).
- Materials and methods are available as supplementary materials.
- M. Srinivasan, A. Ruina, *Nature* **439**, 72–75 (2006).
- N. Hansen, in *Towards a New Evolutionary Computation*, J. A. Lozano, P. Larrañaga, I. Inza, E. Bengoetxea, Eds. (Springer, 2006), pp. 75–102.
- A. D. Kuo, J. M. Donelan, A. Ruina, *Exerc. Sport Sci. Rev.* **33**, 88–97 (2005).
- K. A. Witte, J. Zhang, R. W. Jackson, S. H. Collins, in *2015 IEEE International Conference on Robotics and Automation* (Institute of Electrical and Electronics Engineers, 2015), pp. 1223–1228.
- J. Zhang, C. C. Cheah, S. H. Collins, in *Bioinspired Legged Locomotion: Models, Concepts, Control and Applications*, M. Sharbafi, A. Seyfarth, Eds. (Elsevier, 2017), chap. 5.
- D. P. Ferris, G. S. Sawicki, A. Domingo, *Top. Spinal Cord Inj. Rehabil.* **11**, 34–49 (2005).
- D. J. Reinkensmeyer *et al.*, *J. Rehabil. Res. Dev.* **43**, 657–670 (2006).
- J. F. Veneman *et al.*, *IEEE Trans. Neural Syst. Rehabil. Eng.* **15**, 379–386 (2007).
- J. M. Caputo, S. H. Collins, *J. Biomech. Eng.* **136**, 035002 (2014).
- D. J. Farris, B. D. Robertson, G. S. Sawicki, *J. Appl. Physiol.* **115**, 579–585 (2013).
- R. W. Jackson, C. L. Dembia, S. L. Delp, S. H. Collins, *J. Exp. Biol.* **150**, 11 (2017).
- N. Thattai, H. Geyer, *IEEE Trans. Biomed. Eng.* **63**, 904–913 (2016).
- H. Huang *et al.*, *IEEE Trans. Biomed. Eng.* **58**, 2867–2875 (2011).
- W. Hoogkamer, S. Kipp, B. A. Spiering, R. Kram, *Med. Sci. Sports Exerc.* **48**, 2175–2180 (2016).
- S. Galle, P. Malcolm, W. Derave, D. De Clercq, *Eur. J. Appl. Physiol.* **114**, 2341–2351 (2014).
- M. Hamaya, T. Matsubara, T. Noda, T. Teramae, J. Morimoto, *Pattern Recognit. Lett.* **10.1016/j.patrec.2017.04.007** (2017).

#### ACKNOWLEDGMENTS

We thank V. Chiu and J. Caputo for thoughtful conversations on optimization approaches; T. Nguyen, A. Fatschel, P. Franks, and C. Morales for assistance with data collection and photography; H. Geyer for use of data collection equipment; Gordon Composites for donating materials; and M. Donelan, M. Lansberg, L. Lau, D. Lentink, K. Sreenath, M. Srinivasan, G. R. Tan, G. Torres-Oviedo, the reviewers, and others for editorial suggestions. This material is based on work supported by the National Science Foundation under grant IIS-1355716 to S.H.C. and Graduate Research Fellowships under grant DGE-1252522 to K.A.W., R.W.J., and K.L.P., and a grant from Nike to S.H.C. The data reported in this paper are provided as data S1, and code samples are provided as data S2. S.H.C., J.Z., and C.G.A. conceived the method; S.H.C. and J.Z. designed the main study; J.Z. conducted experiments and analyzed data for the main study; P.F., K.A.W., R.W.J., K.L.P., and S.H.C. conducted the single-subject studies, with P.F. leading the running study and K.L.P. leading the muscle activity study; S.H.C. and J.Z. drafted the manuscript; and all authors edited the manuscript. J.Z. and S.H.C. are inventors on patent applications 62/495,690 and 62/392,270 submitted by Carnegie Mellon University that respectively cover the online optimization approach and exoskeleton torque control approach used in this study. K.A.W., R.W.J., and S.H.C. are inventors on patent application 62/392,263 submitted by Carnegie Mellon University that covers the exoskeleton emulator used in this study.

#### SUPPLEMENTARY MATERIALS

[www.sciencemag.org/content/356/6344/1280/suppl/DC1](http://www.sciencemag.org/content/356/6344/1280/suppl/DC1)  
Materials and Methods  
Figs. S1 to S11  
Tables S1 to S4  
References (43–53)  
Data S1 and S2

30 November 2016; resubmitted 13 February 2017  
Accepted 3 May 2017  
10.1126/science.aal5054

## 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 and Steven H. Collins

*Science* **356** (6344), 1280-1284.  
DOI: 10.1126/science.aal5054

### Optimum human input

Exoskeletons can be used to augment human abilities—for example, to lift very heavy loads or to provide greater endurance. For each user, though, a device will need to be adjusted for optimum effect, which can be time-consuming. Zhang *et al.* show that the human can be included in the optimization process, with real-time adaptation of an ankle exoskeleton (see the Perspective by Malcolm *et al.*). By using indirect calorimetry to measure metabolic rates, the authors were able to adjust the torque provided by the device while users were walking, running, and carrying a load.

*Science*, this issue p. 1280; see also p. 1230

#### ARTICLE TOOLS

<http://science.sciencemag.org/content/356/6344/1280>

#### SUPPLEMENTARY MATERIALS

<http://science.sciencemag.org/content/suppl/2017/06/21/356.6344.1280.DC1>

#### RELATED CONTENT

<http://science.sciencemag.org/content/sci/356/6344/1230.full>  
<http://stm.sciencemag.org/content/scitransmed/5/210/210rv2.full>  
<http://stm.sciencemag.org/content/scitransmed/5/210/210ps15.full>

#### REFERENCES

This article cites 41 articles, 5 of which you can access for free  
<http://science.sciencemag.org/content/356/6344/1280#BIBL>

#### PERMISSIONS

<http://www.sciencemag.org/help/reprints-and-permissions>

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