



Delft University of Technology

Document Version

Final published version

Citation (APA)

Slade, P., Atkeson, C., Donelan, J. M., Houdijk, H., Ingraham, K. A., Kim, M., Kong, K., Poggensee, K. L., Collins, S. H., & More Authors (2024). On human-in-the-loop optimization of human-robot interaction. *Nature*, 633, 779-788.
<https://doi.org/10.1038/s41586-024-07697-2>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

On human-in-the-loop optimization of human–robot interaction

<https://doi.org/10.1038/s41586-024-07697-2>

Received: 20 November 2023

Accepted: 7 June 2024

Published online: 25 September 2024

 Check for updates

Patrick Slade^{1✉}, Christopher Atkeson², J. Maxwell Donelan³, Han Houdijk⁴, Kimberly A. Ingraham⁵, Myunghee Kim⁶, Kyoungchul Kong⁷, Katherine L. Poggensee^{8,9}, Robert Riener^{10,11}, Martin Steinert¹², Juanjuan Zhang¹³ & Steven H. Collins^{14✉}

From industrial exoskeletons to implantable medical devices, robots that interact closely with people are poised to improve every aspect of our lives. Yet designing these systems is very challenging; humans are incredibly complex and, in many cases, we respond to robotic devices in ways that cannot be modelled or predicted with sufficient accuracy. A new approach, human-in-the-loop optimization, can overcome these challenges by systematically and empirically identifying the device characteristics that result in the best objective performance for a specific user and application. This approach has enabled substantial improvements in human–robot performance in research settings and has the potential to speed development and enhance products. In this Perspective, we describe methods for applying human-in-the-loop optimization to new human–robot interaction problems, addressing each key decision in a variety of contexts. We also identify opportunities to develop new optimization techniques and answer underlying scientific questions. We anticipate that our readers will advance human-in-the-loop optimization and use it to design robotic devices that truly enhance the human experience.

The next phase of human technological development is expected to hinge on robotic systems that operate alongside people and interact with us in complex ways. These systems promise to profoundly improve many aspects of our lives, from mobility and health to productivity and leisure¹. To achieve these benefits, however, the robotic devices of the future will need to be far more effective at helping humans^{2–4}.

A critical challenge in delivering effective assistance to humans is our immense biological complexity. The human neuromuscular system involves dynamic interactions between thousands of motor units and trillions of synapses⁵. Humans exhibit individual diversity in our physiology, motor control and decision-making. We also change over time as we learn, heal and grow. Systems with this level of complexity are intrinsically difficult to understand, model and predict⁶. Learned models can capture important features of complex systems and individual users^{7,8}, but require more training data than are available at present. The robotic devices we develop must assist a human system that we do not fully comprehend.

Human-in-the-loop optimization can overcome these challenges to deliver effective assistance. In this process (Fig. 1), a user interacts with a robotic device while actual performance is measured. An algorithm systematically changes device characteristics and measures the resulting changes in human–robot performance, adjusting the device until the characteristics that maximize performance are identified. This

approach does not require an explicit model because the performance evaluations that guide the algorithm are made on the actual human–robot system. The optimizer continually tracks the individual user and co-adapts as they learn and change over time.

In some ways, human-in-the-loop optimization resembles a systematized and accelerated version of the traditional process for designing and tuning robots that interact closely with humans. Algorithms replace or augment low-level tinkering and intuition, more efficiently searching the space of potential designs and settings. Versatile hardware and software replace specialized systems, reducing the design cycle time from months or years to minutes or seconds. In these ways, human-in-the-loop optimization adds to existing design tools to decrease development time and risk.

Computer models of humans cannot yet replace experiments with humans in this process. Musculoskeletal simulations can operate quickly, easily and at low cost compared to experiments with humans. Unfortunately, these models are at present not accurate enough to predict human response to dynamic interactions with most robotic devices⁹. The main limitation is in simulating the nervous system, which is important to human behaviour but exceedingly difficult to model, even superficially. Simulated ‘digital twins’ may one day revolutionize device design, but empirical approaches are more effective at present.

¹Harvard John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA. ²The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA.

³WearTech Labs, Simon Fraser University, Burnaby, British Columbia, Canada. ⁴Department of Human Movement Sciences, University Medical Center Groningen, University of Groningen, Groningen, The Netherlands. ⁵Department of Electrical and Computer Engineering, University of Washington, Seattle, WA, USA. ⁶Mechanical and Industrial Engineering, University of Illinois Chicago, Chicago, IL, USA. ⁷Department of Mechanical Engineering, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea. ⁸Department of Rehabilitation Medicine, Erasmus Medical Center, Rotterdam, The Netherlands. ⁹Faculty of Mechanical Engineering, Delft University of Technology, Delft, The Netherlands. ¹⁰Sensory-Motor Systems Lab, ETH Zurich, Zurich, Switzerland. ¹¹Faculty of Medicine, University of Zurich, Zurich, Switzerland. ¹²Department of Mechanical and Industrial Engineering, Norwegian University of Science and Technology (NTNU), Trondheim, Norway. ¹³College of Artificial Intelligence, Institute of Robotics and Automatic Information System, Nankai University, Tianjin, China. ¹⁴Department of Mechanical Engineering, Stanford University, Stanford, CA, USA. ✉e-mail: slade@seas.harvard.edu; stevecollins@stanford.edu




Fig. 1 | Designing systems of human-in-the-loop optimization. During optimization, measurements of human–robot performance are used to iteratively improve characteristics of the robotic device, yielding better human–robot interactions. This Perspective is organized by the decisions that will define a new implementation of human-in-the-loop optimization, including those related to applications, hardware, parameterization, objectives, algorithms and workflows.

Machine learning could complement human-in-the-loop optimization when sufficient human–robot interaction data become available. With sufficiently large and diverse datasets, it may be possible to learn models that predict the best device for an individual user based on their personal characteristics and even make adjustments as they adapt over time. At present, however, not enough training data are available. Fortunately, human-in-the-loop optimization naturally generates diverse datasets suitable for machine learning.

Human-in-the-loop optimization has led to marked improvements in human–robot performance in research settings¹⁰. The most comprehensive studies have been conducted on exoskeletons^{11–15} and exosuits^{16–18} that improve the energy cost of human walking. The introduction of optimization in this domain led to fourfold improvements in performance for the same laboratory-based hardware¹¹ and twofold improvements in field benchmarks, across a range of conditions, in real-world settings¹⁹. The benefits come from a combination of discovery of generically effective device characteristics, individualizing assistance, and co-adaptation with the user¹². Similarly impressive results have been demonstrated in many other human–robot interaction problems, such as: identifying exoskeleton characteristics that maximize walking speed¹⁵ and user preference^{20,21}, minimize energy cost during running¹³ and squats²², minimize core muscle fatigue during lifting²³, minimize muscle activity during downhill walking²⁴ or bicep flexion²⁵ and best track joint kinematics²⁶ and impedance²⁷; identifying prosthetic-leg characteristics that minimize metabolic cost^{28,29} and best track joint kinematics³⁰; optimizing shoe characteristics^{31,32}; optimizing energy harvester performance³³; identifying energy-optimal step frequencies³⁴; optimizing kinematic synergies for upper-limb prostheses³⁵; and maximizing the realism of haptic rendering³⁶ and visual acuity of retinal prostheses³⁷.

Although human-in-the-loop optimization has primarily been used in gait research so far, it has the potential to improve robotic devices in a much wider range of applications and settings (Fig. 2). The initial application to lower-limb devices has been natural, given the small number of behaviours to be addressed and the high sensitivity to performance, but does not signify a limitation. The method has already begun to be applied effectively in upper-limb devices, medical implants and audiovisual cueing, and has excellent potential for surgical robots, cooperative warehouse robots, social robots and mobility systems. Similarly, although human-in-the-loop optimization was developed in academic institutions, it is equally applicable to private-sector and clinical settings. Designers could use it as a tool for product specification and customization, identifying effective designs faster and with

lower cost and risk than traditional intuition-driven processes. Clinicians and retailers could use it to deliver more effective and consistent experiences to patients and customers through prescription and customization processes or with software that automatically adapts as people use devices in everyday life.

Some aspects of human-in-the-loop optimization are well understood and provide a powerful set of tools ready for deployment, whereas many aspects have hardly been explored and present exciting opportunities for research and development. In this Perspective, we provide guidance on established approaches to addressing new populations and tasks, hardware choices, parameter selection, objective definition, optimization algorithms, workflows and outcome evaluation. In each case, we also identify potent topics for future research. We emphasize examples from wearable and implantable robotic devices, but we expect that analogous approaches will be effective in any case with a tight coupling between human and machine dynamics. We expect the reader will find that human-in-the-loop optimization can help them to invent new classes of human–robot systems, advance the state of the art in their fields, improve the performance of products in real-world settings and deliver new benefits to patients and customers.

Potential applications

Human-in-the-loop optimization is likely to be effective whenever a human and a robotic device dynamically interact, performance can be quantified and device characteristics can be changed to affect performance. It is of greatest use when human complexity prevents approaches based on explicit models, for example, when the human nervous system has a decisive influence, when multiple device characteristics may be important and when performance standards are high. It requires a measurable outcome, or optimization ‘objective’, with a desired target, for example minimum oxygen consumption, maximum task execution speed or minimum deviation of a physiological signal from a target value. There must also be a device with changeable characteristics that affect performance, for example an exoskeleton with a controllable pattern of torque, an electrical stimulator with adjustable feedback parameters or an implanted device with settings that affect drug-release dynamics. Many combinations of users, tasks and devices meet these requirements.

We have only just scratched the surface of potential applications of human-in-the-loop optimization. There is a pressing need to use this technique to address mobility and manipulation among users with physical impairments resulting from ageing, osteoarthritis, neuropathy, stroke, limb amputation and cerebral palsy. Individualization may prove to be more important in these populations owing to their heterogeneity³⁸. Another critical application area is physical unloading for workers with strenuous jobs, for whom robotic devices promise to reduce the likelihood of both acute and chronic injuries^{39,40}. There are many opportunities to apply this technique in, for example, rehabilitation, biofeedback systems, shared-autonomy vehicles, cooperative robots, teleoperation, surgical training systems, sensory aids, implantable devices, caregiver robots, educational robots, social robots and even athletic equipment or other non-robotic devices. There is immense potential for impact through studying these methods in the contexts of product development, prescription and everyday use.

Implementing human-in-the-loop optimization for a new problem requires an understanding of the needs and limits of the user group, intrinsic dynamics of the task and features likely to be useful in new devices. At this stage of the process, intuition⁴¹, simple models⁴² and simulations^{43,44} can be valuable. For any new problem, some aspects of previous human-in-the-loop optimization approaches will need to be adjusted. For example, in a population in which fatigue is a concern, an optimization may be conducted over many shorter sessions. Such changes are to be expected and not a contraindication.



Fig. 2 | Potential applications for human-in-the-loop optimization.

a–f, Most suitable applications have yet to be explored. Opportunities include: advanced prosthetic limbs (**a**); implantable medical devices, such as insulin pumps (**b**); exoskeletons for augmentation or rehabilitation (**c**); teleoperated

robots in industrial or medical settings (**d**); social robots in educational or caretaking settings (**e**); and systems that share autonomy with humans, such as self-driving vehicles (**f**). Credit: SeventyFour Images/Alamy (**a**); Getty (**b**); Getty (**d**); BSIP SA/Alamy (**e**); Sundry Photography/Alamy (**f**).

Hardware considerations

Human-in-the-loop optimization can be used to improve an assistive robotic device during any stage of development (Table 1), including designing new devices, improving existing devices, prescribing custom devices and performing continual co-adaptation.

Designing new robotic devices

Early in the development of a new robotic device, versatile hardware complements human-in-the-loop optimization. Systems that allow exploration of a large range of device characteristics increase the likelihood of finding effective solutions. When tackling any new problem, the optimal device characteristics are unknown and may be unintuitive⁴⁵; we must acknowledge that our first design concepts are unlikely to benefit the user and we should expect to need to test many different approaches. Versatile testbeds such as emulators, lab-based systems designed to imitate many different portable devices^{46–48}, and reconfigurable portable devices^{49–51} are ideal at this stage of device exploration. Instrumenting these testbeds with low-cost, fieldable sensors can enable data-driven approaches at future stages.

Emulation and human-in-the-loop optimization are poised to become a powerful approach for commercial development of robotic devices. Once effective device specifications are identified in emulator experiments, they can be used to design portable devices with predictable benefits^{9,52}. To facilitate translation, researchers should develop better ways of addressing outcomes that are at present difficult to capture in an emulator. For example, the implied mass or cost of the emulated device could be included as a term in the optimization objective, leading to more viable product specifications. Other outcomes, such as patient adoption or ease of everyday use, could initially be partially explored through emulation, with separate user-centred

design activities complementing this process⁵³. Emulation paired with optimization does not yet address all aspects of productization but it excels at one of the most challenging aspects: delivering the intended functional benefit.

Improving existing devices

Human-in-the-loop optimization can improve the performance of existing robotic devices without making any changes to hardware^{19,23,26,33,54,55}. Large improvements are especially likely in cases in which the controller was based on intuition or bioinspiration. If the optimal characteristics are found to be at the limits of what is possible with the device, this information can help guide improvements to future hardware

Table 1 | Human-in-the-loop optimization can be used at any stage of development

Stage of development	Outcomes of human-in-the-loop optimization
Research and development	Device specifications
	Generically optimal controllers
	Human–robot interaction data
Product development	Revised device specifications
	Generically optimal controllers
	Human–robot interaction data
Device prescription or customization	Personalized device specifications
	Controllers personalized in the clinic or storefront
	Human–robot interaction data
Device deployment	Real-time personalized controllers
	Human–robot interaction data

versions. It is possible that many previous devices that ‘failed’ to deliver their intended benefits could yet succeed through the use of human-in-the-loop optimization^{56–58}.

Prescribing custom devices

Medical and consumer devices could be prescribed or customized using human-in-the-loop optimization by first identifying ideal characteristics using a versatile testbed and then assembling and programming a specialized device for the user. This approach has recently been trialled for prosthetic foot prescription⁵⁹ and may prove especially effective in customizing devices for heterogeneous groups⁶⁰. Although a computer-controlled fitting system may be most efficient, it is not required; manual adjustment of non-robotic hardware settings can achieve the same result, if performed systematically using optimization. For example, eyeglass lens prescriptions can be determined using a phoropter, while rocker shoe parameters that produce desired running mechanics can be identified using adjustable training shoes^{31,32}. Personalizing low-cost hardware in this way could be a powerful approach for solving design challenges in developing nations^{61–63}.

Parameterization

Parameterization is the process of defining the way in which the behaviour of a robotic device can be changed and the settings that specify those changes. Parameters may define control set points, control gains, virtual system dynamics or hardware settings to be adjusted manually. An ideal parameter set is both expressive, enabling exploration of many human–robot interactions and increasing the likelihood of finding effective ones, and low dimensional, reducing the size of the optimization problem and thereby the time required to identify the optimal values. So far, human-in-the-loop optimization has addressed between two and 22 parameters. In defining an architecture to optimize, designers draw from many sources, including biology, simple models, computational simulations, analogous systems and existing designs (Fig. 3).

Interesting problems in parameterization

Systematic methods of parameterization would represent a notable breakthrough. Human-in-the-loop optimization is at present systematic in exploring a given parameterization, but if parameters are not well chosen the optimizer is limited to ineffective solutions. One potential approach is nested optimization, in which different parameterizations would be generated and assessed, slowly converging to a parameterization that strikes a balance between expressiveness and dimensionality appropriate to the application. Alternatively, past optimization data could be analysed to uncover relationships between parameters, combine parameters, eliminate redundant parameters and introduce new parameters, drawing on statistical methods widely used in machine learning. Automated parameterization could occur at the beginning of the research process, with a large initial investment leading to better understanding and long-term outcomes. Automated control parameterization may also be possible during continuous use of devices in the real world, in which a large volume of data is available and decision-making algorithms could systematically balance between exploration and exploitation of parameterizations.

Parameterization affects the dynamics of human–robot interaction, particularly in terms of stability and adaptation, and basic scientific study of these relationships would be valuable for many domains. For example, it would be useful to determine the relative difficulty of learning how to use an assistive device under position, velocity or torque control, or with time-based, spring-like, damper-like or mass-like dynamics, or with various types and amounts of delay. Scientific experiments comparing different control architectures would be valuable and could be explicitly explored through an optimization framework.

Different types of control architecture should be explored to enable robotic devices to generalize to different interaction tasks, especially

under non-steady conditions. So far, human-in-the-loop optimization has primarily been performed on controllers that set device torque as a function of time, but this is not an intrinsic feature of the method. Other ways of relating states to actions, such as phase-based⁶⁴ or virtual-model control⁶⁵, are also defined by parameters that could be optimized and may allow a single controller to be effective across many tasks and transitions. New optimization techniques could be developed to address more general control types and could optimize performance for a large set of tasks simultaneously. Alternatively, tasks can be discretized and addressed separately^{19,66–68} and methods developed to smoothly interpolate between tasks and contexts.

Optimization objectives

The objective measure that is optimized should be the outcome of actual importance, measured as directly as possible, during the course of the optimization. This could be a measure of functional task performance, a physiological signal, a subjective outcome, a device characteristic with implications for product success, a combination of these, or any other outcome of interest (Table 2). The objective measure must be available to the optimizer quickly enough to allow it to choose new parameter values to test, whether the evaluation requires seconds or weeks.


A common misstep is to optimize a term that is easily measured but is later determined not to correlate well with the true objective. Poorly defined objective functions often lead optimizers to exploit unexpected system dynamics to give undesirable solutions. For example, drone racing controllers that directly optimize the high-level objective of time to complete a course achieve better performance than controllers that use explicit intermediate representations, such as tracking a trajectory⁶⁹. Similarly, in assistive devices, positive mechanical power may be correlated with user energy cost, but maximizing one does not minimize the other^{12,32}. Optimizing a substitute for the outcome of interest, without first validating that it is a reliable surrogate measure, undermines the primary benefit of human-in-the-loop optimization: addressing the full complexity of the real human–robot system. Choosing a surrogate measure based on a priori models or intuition reintroduces the types of error that human-in-the-loop optimization aims to overcome.

In cases in which direct measurement of the objective is not possible, a validated, data-driven surrogate can be used. This scenario is likely to arise when performing optimization outside a controlled setting, in which some measurements become impractical. In such cases, a large set of data, including both candidate surrogates and ground-truth measurements, can first be collected in the laboratory. The minimal set of sensors needed for accurate estimation can then be determined^{70–72} and data-driven models suitable for real-world optimization developed¹⁹. For many optimization algorithms, the data-driven model need not estimate the value of the outcome but instead only whether the outcome was better in one condition or another, which may result in greater accuracy⁷³. Capturing the relative differences between two conditions can simplify the relationship a data-driven model must learn to effectively distinguish which conditions are most beneficial. For example, exoskeleton kinematic data can be used to estimate the difference in metabolic cost between two assistance conditions more accurately than using kinematic data to estimate the absolute metabolic cost associated with each condition and then comparing them¹⁹. To verify that the data-driven measure generalizes sufficiently, optimization should be performed using both the surrogate and ground-truth measures, and the optimal ground-truth outcomes compared, before deployment.


Open questions on optimization objectives

Many aspects of human–robot interaction are not yet well captured by objective measurements, whereas others are difficult to measure


a Virtual model control




b Model-based control




c Cyclic control



d Collaborative control



e Trajectory control



f Shared autonomy




Fig. 3 | Example parameterizations. Any parameter that affects human–robot interaction is a suitable target for human-in-the-loop optimization. Example parameters are shown in blue. **a**, Parameters that define virtual systems can be optimized, for example the virtual stiffness and damping of a knee prosthesis. **b**, Parameters that define system dynamics can be optimized, for example gains and delays involved in mapping microphone recordings to auditory nerve stimulation in a cochlear implant. **c**, Parameters that define cyclic control can be optimized, for example the magnitude and timing of peak ankle exoskeleton

torque during walking. **d**, Parameters that define interaction dynamics can be optimized, for example task-dependent forces and torques applied to an object during cooperative manipulation. **e**, Parameters that define trajectories can be optimized, for example the motion of a caretaker robot arm during assisted feeding. **f**, Parameters affecting shared autonomy can be optimized, for example terms in the model used to predict future user actions in a self-driving car or piece of construction equipment.

online or outside the laboratory. For example, there is at present no consensus measure of balance ability that could be used to optimize wearable devices intended to reduce fall risk^{74,75}. Other potentially useful objective measures that are difficult to directly measure include joint loading, injury risk, fatigue and pain, each of which has complex, individual-specific dynamics that can change over time. Fundamental scientific research is needed to develop new metrics, followed by

engineering to develop estimation approaches that rely on low-cost, wearable sensors for real-world use.

A topic ripe for theoretical consideration is how long to spend on each measurement. Minimizing the total optimization time is desirable, especially for patient populations in which exhaustion is a limiting factor. Gathering more data per measurement requires more time but may improve measurement accuracy and lead to more beneficial

Table 2 | Examples of human–robot performance optimization objectives

Physiological measurements	Clinical outcomes	Work performance
Kinematics	Clinical test scores	Task completion time
Environmental forces	Activity capacity	Task accuracy
Musculotendon forces	Mental acuity	Production cost
Joint contact forces	Daily step count	Environmental impact
Muscle activity	Balance confidence	Human task complexity
Metabolic rate	Walking speed	Injury risk
Neural activity	Visual acuity	Human cognitive load
Fatigue	Range of motion	Rated perceived effort
Eye gaze	Muscle strength	Social engagement
Blood pressure	Organ health	User emotional valence
Blood glucose level	Pain	Realism of haptic rendering
Cortisol levels	Satisfaction	Student learning rate

optimization updates. On the other hand, providing updates to the optimizer more rapidly with less accurate measurements might improve performance more quickly, especially early in the optimization process. Finding the appropriate trade-off in the optimization rate is an important research question.

The objectives of the human and the optimization process must be sufficiently well aligned to obtain a beneficial solution. If the person and the optimizer have goals that differ in particular ways, and if they respond to each other on similar timescales, the optimizer can converge to undesirable outcomes, oscillate or even diverge⁷⁶. This phenomenon has been observed in experiments with treadmills that act to quickly enforce unusual relationships between walking speed and step frequency, leading to strange, but consistent, human gait choices^{77,78}. Further game theoretic analysis of human-in-the-loop optimization could help us to avoid these conditions and possibly improve outcomes⁷⁹.

A better understanding of what users truly want from their robotic devices would lead to better selection of optimization objectives and more positive societal impact. In any human–robot interaction, there are many different, possibly competing, objectives, the balance of which will change on the basis of internal and external factors. For example, when walking, a person may care about stability, economy, speed, pain or aesthetics^{80,81}, and the importance of each of these terms will change when, for example, physically fatigued or in a risky environment⁸². As another example, for exoskeleton users with paraplegia, the outcomes of speed, energy cost and crutch load can compete with each other⁸³. Basic scientific study of the revealed objectives of the nervous system^{84,85} will help map out the terms of interest for human–robot interaction. Comparisons of parameters optimized for different objectives could reveal competing objectives or shared solutions³¹.

Human-in-the-loop inverse optimization frameworks could be developed to infer the optimization criterion of the human user. Human responses to different device characteristics could be used to identify the importance that an individual places on each term in their own optimization criterion in various contexts. Such approaches could draw on techniques developed for humanoid robots⁸⁶ and could potentially automate the process of defining the optimization objective.

In some human–robot interactions, several objectives can be addressed by simply asking users what they prefer. Preference-based optimization can be effective when the user is consciously aware of the effects of the device, as in lower-limb exoskeletons^{20,21,36,87}, and the process can both engage and empower users⁸⁸. This approach presents several attractive opportunities for future research. People are far more

consistent when comparing two recent conditions than when assigning absolute scores⁸⁸. Algorithms that both take pairwise comparisons and operate in large parameter spaces⁸⁹ could allow optimization of a wider range of devices or account for several objectives at once⁹⁰. Users must pay attention to the behaviour of the robotic device and provide feedback, a potential drawback for systems designed to be ‘invisible’ to the user. Methods of seamlessly addressing preference feedback in real-world settings would be valuable. In some cases, user preference may not be consistent with other goals, such as long-term health objectives or features of task performance, which might be addressed by constraining the search space or incorporating further terms into the optimization objective. We expect that a preference for using a robotic device will prove a good indicator of eventual adoption.

Optimization algorithms

An algorithm for human-in-the-loop optimization must operate in real time, be robust to noise, adapt to changes in the person and use data efficiently (Fig. 4). By ‘real time’, we mean that the algorithm applies parameters, measures outcomes and then decides which parameters to test next, all as the device is being used. The algorithm must be robust to noise because the complex dynamics of the human body often lead to high-frequency changes in physiological signals that do not reflect important changes in the objective. Algorithms must be adaptive because people continually change as they learn and grow, which can lead to changes in the relationship between parameter values and the objective. They must be sample efficient because each objective evaluation is expensive, in terms of human time and other costs, compared with simulation-based optimizations.

So far, four categories of algorithms have been used for human-in-the-loop optimization: evolutionary algorithms, surrogate optimization, gradient-based methods and reinforcement learning. Human-in-the-loop optimization has most often used evolutionary algorithms. Covariance matrix adaptation evolution strategy (CMA-ES) is a sample-based, local optimizer for continuous domains that has been used frequently. CMA-ES is tolerant of both measurement noise and human adaptation because it makes decisions primarily based on the ranking of the most recently tested conditions. It is sample efficient, scaling with the number of parameters as $\log(N)$ (ref. 91). CMA-ES requires little algorithmic parameter tuning, uses little computation and performs well on nonlinear and non-convex problems⁹². CMA-ES should be initialized to cover a relatively large range of parameters to avoid becoming trapped in local minima in complex objective spaces. CMA-ES has been effective when applied to several assistive devices^{12,14,17,22,46}, patient populations^{28,93} and performance objectives^{11,15,22,23,33,54}, with up to 22 parameters¹⁴. Another evolutionary algorithm, particle swarm optimization, has also been used⁹⁴ and merits more investigation.

Surrogate optimization methods include variations of Bayesian optimization, a global optimization approach that yields estimates of both the optimal parameters and a model of the objective space. With properly tuned hyperparameters, it may be robust to noise and human adaptation⁹⁵. It is sample efficient but has computational complexity that scales as N^3 , which makes it challenging to apply to larger problems⁹⁶ (approximately 20 or more parameters). Bayesian optimization has been effective in several human-in-the-loop optimization experiments^{16,27–29,36,87,97,98}. It is unclear when CMA-ES or Bayesian optimization will be better for a new human-in-the-loop optimization application⁹⁵.

Gradient-based methods have been used to perform human-in-the-loop optimization of a single parameter in two studies^{34,35}. Gradient-based approaches are challenging to use for sample-based optimization because measurement noise is amplified by the finite differencing of samples used to approximate gradients. Estimating the gradient over a larger interval can reduce the effects of measurement noise but this may require more data than is practical to collect.




Fig. 4 | Challenges for human-in-the-loop optimization algorithms.

a, The algorithm must tolerate noisy measurements, because the human will generally contribute substantial noise to measurements of the optimization objective. **b**, The algorithm must tolerate time-varying adaptations in the

human system, because the optimal solution changes as the human adapts on several timescales. **c**, The algorithm must be sample efficient, especially for high-dimensional parameterizations, because human–robot interaction time is limited.

Gradient-based approaches also scale poorly with the number of parameters.

Reinforcement learning, although not strictly an optimization algorithm, uses a similar structure: developing a policy that determines state-dependent actions to maximize a reward function. Reinforcement learning methods have been implemented in two studies optimizing human–robot interactions^{25,30} but more research is needed to explore the potential benefits of this state–action–reward formulation.

Optimizing the optimizer

There are many opportunities for domain experts to apply and compare existing algorithms, develop new algorithms and tuning approaches, and build infrastructure to speed up these processes. Most optimization algorithms⁹² have not yet been tested on human-in-the-loop optimization. Direct experimental comparisons of existing algorithms could quantify trade-offs, in terms of sample efficiency, robustness to human adaptation and incorporation of previous data to inform initial optimization parameters, to improve algorithm selection for new problems. It is likely that new algorithms designed specifically for human-in-the-loop optimization will achieve better performance than any existing algorithms, which have primarily been developed for dissimilar domains. Methods for automated tuning of optimization algorithm parameters, or optimizing the optimizer, could remove another source of human bias and increase the likelihood of success. To allow preliminary assessment and comparison of algorithms in simulation, it would be valuable to have simple, high-level models of the human nervous system, capturing features such as complexity, hidden states, noise and adaptation.

Experiments and workflows

A human-in-the-loop optimization experiment usually includes an acclimation phase, an optimization phase, stopping when convergence criteria are met, and a separate evaluation of the optimal parameters identified. The acclimation and optimization periods should allow sufficient time for the human to adapt to the robotic device. There are many trade-offs to consider and substantial pilot testing is usually required to adjust the problem definition and workflow to ensure that optimal parameters are reachable, the user is well adapted and the optimizer converges.

In experiments evaluating the benefits of optimal human–robot interactions, the optimal set of device parameters must be evaluated separately to avoid positively biasing results. The parameters believed to be optimal should be identified first and then tested in a separate validation to determine their effects. A common misstep is to retrospectively select the best outcomes observed during optimization. This will overstate the benefits, because measurements of human

outcomes are noisy and selecting the best outcomes after noise has taken effect can only operate in a beneficial direction. In mathematical terms, $\text{mean}(\max(\text{randn}(m, n))) > 0$, whereas the true mean is zero. This form of post hoc selection bias must be avoided.

The amount of time provided for human adaptation can have a substantial impact on the efficacy of the solution. When exposed to a new robotic device, humans first respond quickly based on their existing neural control and then slowly adjust their control and, in some cases, body structure (Fig. 4b). Some human–robot interaction problems may require little to no time for adaptation, whereas others may take weeks or months. For example, people exhibit distinct fast (seconds) and slow (hours) responses to large changes in walking conditions⁷⁷, with indications of muscular remodelling over longer (days) exposures¹². The optimizer may need more time to converge following human adaptation, although the optimizer will have a ‘warm start’⁹⁹. The benefits of robotic interaction can be expected to substantially improve with exposure, with diminishing returns.

Iterative refinement of the problem definition and workflow are typically needed to achieve a good balance between the duration, robustness and efficacy of the optimization. For example, pilot tests may suggest that some parameters are unnecessary, always converging to the same value, or that some portions of the parameter space never contain optima. Adjusting the initialization, parameterization or parameter constraints accordingly may speed convergence but could also exclude optimal parameter values or diminish benefits for some new participants. After refinement, the workflow can be translated to new contexts relatively easily, for example from the laboratory to real-world conditions. When devices are deployed at scale in the real world, this refinement could be crowdsourced; information from existing users could benefit future users with over-the-cloud updates.

The time and effort required to develop and implement a human-in-the-loop optimization protocol may seem a deterrent but will often reduce the overall time to achieve a given level of human–robot performance. For example, one study of a complex lower-limb exoskeleton used about 15 two-hour pilot testing sessions to develop a protocol that required three two-hour optimization sessions per subject¹⁰⁰. Although these times are daunting, the study also demonstrated the largest improvement in energy cost with an exoskeleton so far, whereas previous attempts with similar devices had yielded no benefit⁴¹. Because of the challenges in achieving effective human–robot interaction, rapid but unsystematic development approaches are likely to yield suboptimal results or fail entirely, adding years to the overall development effort. Fortunately, the time required for developing and implementing human-in-the-loop optimization protocols also decreases with experience. For example, a recent exoskeleton study¹⁹ used information from earlier studies¹² to require minimal pilot testing and only one hour of optimization per participant while achieving equivalent performance.

Designing better experiments and workflows

Convergence criteria represent a substantial open problem in human-in-the-loop optimization. So far, most optimizations are designed to end after a preset duration, in some cases based on pilot testing. However, this approach does not address variability in human adaptation times¹⁰¹ or in optimizer convergence times, which can be inconsistent in the presence of noise⁹⁵. The strict convergence criteria often used in computer simulations tend to require a large number of evaluations, making them impractical in human experiments. Improved stopping criteria would be both meaningful, for example related to objectives or parameters, and implementable in practice. These could save experimenter time, increase consistency in estimated benefits and improve exploration–exploitation trade-offs¹⁰² in continuous optimization problems.

Our field would benefit from basic scientific studies on the effects of several environmental factors on adaptation to human–robot interaction. For example, we do not yet know whether entertainment, distraction, boredom or perceived utility of the experimental task affect the speed or quality of adaptation in laboratory or real-world settings¹⁰³. It is not known whether multiday training sessions are more effective, by allowing consolidation of motor skills into a more permanent representation, or whether long gaps between training sessions affect the retention of learned skills. Indicators of motor learning with robotic devices can differ between real-world and laboratory conditions¹⁰⁴, which could lead to differences in optimal training approaches. We also do not know whether small changes in robotic device characteristics between conditions or sessions, for example owing to a change in prosthesis control parameters or the realignment of a prosthetic socket, impede transfer or facilitate more robust human motor programmes, similar to the effects of ‘domain randomization’ techniques in machine learning.

There is a substantial need for theories of training people to use robotic devices. Although much research has addressed the training of robots^{105,106}, for example using demonstration, intervention and evaluative feedback to augment reinforcement learning¹⁰⁷, little work has considered the training or coaching of their human users. Most human–robot interaction studies rely on unstructured exploration by users with incidental instruction from experimenters. Yet human adaptation can be the dominant factor in time to convergence¹², differences in training can affect adaptation time¹⁰⁴ and quality¹⁰⁸, and users can sometimes require particular prompts to trigger adaptation¹⁰⁹. Fundamental motor control studies are needed to improve our understanding of the mechanisms driving adaptation and to allow us to reliably estimate level of adaptation¹¹⁰, for example, using signals from the limbs, muscles or brain^{111,112}. New methods are needed to rapidly and reliably train users, for example, using automated curriculum generation^{113,114}, skill inference¹¹⁵, biofeedback targeting patterns learned from expert users^{116,117} or even coaching that is optimized online in parallel with the device. Training users more effectively may result in more consistent responses during optimization, allowing for simpler, faster optimization algorithms.

Another opportunity is continuous human-in-the-loop optimization during everyday life. Methods for opportunistically testing parameters and measuring outcomes as people engage in various activities throughout the day will allow robotic devices to continually co-adapt with people. Longitudinal optimization studies, over weeks or months, could investigate the effects of device interactions on outcomes such as muscular strength, activity level, fitness and other long-term health outcomes. Continuous device use would naturally provide more exposure to infrequent one-shot tasks, such as sit-to-stand transitions, which would allow optimization without condensed, fatigue-inducing laboratory experiments. It may also be possible to leverage data from many users simultaneously, parallelizing the process to speed the identification of good generic interventions in domains with outcome measures

that require unusually lengthy evaluation, such as novice skill training and rehabilitation.

The future of personalized human–robot interaction

Some form of human–robot interaction could provide benefits to nearly anyone in their daily lives (Fig. 2). Robotic components provide an excellent complement to human physiology; electric motors can be more controllable and efficient than muscles¹¹⁸, sensors can be faster and more accurate than human perception^{119,120} and there are many effective ways of interfacing humans with robots⁵³. A critical challenge is determining what a robotic device should do to assist the person using it. Although traditional development approaches can overcome this challenge^{83,121–123}, human-in-the-loop optimization lets us overcome it faster, more easily and at lower cost, helping unlock the potential of existing technologies for assisting people.

Which human–robot technologies will first appear in our everyday lives? The next notable leaps may come in wearable robotic devices that improve mobility for clinical populations, enhance rehabilitation, increase physical activity or boost productivity and safety in industrial settings; portable medical devices that provide continuous health monitoring and personalized non-physical interventions, co-adapting with users over long timescales in their daily lives; robots that share autonomy in challenging tasks, such as remote surgery or piloting; or cooperative robots that assist people in their homes as caregivers, provide personalized education, or facilitate social interaction or development. No matter what problem we use robots to improve next, human-in-the-loop optimization can make the human experience better.

- Demir, K. A., Döven, G. & Sezen, B. Industry 5.0 and human-robot co-working. *Procedia Comput. Sci.* **158**, 688–695 (2019).
- Farina, D. et al. Toward higher-performance bionic limbs for wider clinical use. *Nat. Biomed. Eng.* **7**, 473–485 (2023).
- Sawicki, G. S., Beck, O. N., Kang, I. & Young, A. J. The exoskeleton expansion: improving walking and running economy. *J. Neuroeng. Rehabil.* **17**, 25 (2020).
This review presents a timeline of lower-limb exoskeleton development and performance enhancements.
- Crea, S. et al. Occupational exoskeletons: a roadmap toward large-scale adoption. Methodology and challenges of bringing exoskeletons to workplaces. *Wearable Technol.* **2**, e11 (2021).
- Uchida, T. K. & Delp, S. L. *Biomechanics of Movement: The Science of Sports, Robotics, and Rehabilitation* (MIT Press, 2021).
- Ghez, C. & Krakauer, J. in *Principles of Neural Science* 4th edn (eds Kandel, E. R., Schwartz, J. H. & Jessell, T. M.) 653–673 (McGraw-Hill, 2000).
- Halilaj, E. et al. Machine learning in human movement biomechanics: best practices, common pitfalls, and new opportunities. *J. Biomech.* **81**, 1–11 (2018).
- Alili, A. et al. A novel framework to facilitate user preferred tuning for a robotic knee prosthesis. *IEEE Trans. Neural Syst. Rehabil. Eng.* **31**, 895–903 (2023).
- Franks, P. W. et al. in *Proc. 2020 8th IEEE RAS/EMBS International Conference for Biomedical Robotics and Biomechatronics (BioRob)* 700–707 (IEEE, 2020).
This study demonstrates the shortcomings of simulation-based optimization of human-robot interactions.
- Diaz, M. A. et al. Human-in-the-loop optimization of wearable robotic devices to improve human–robot interaction: a systematic review. *IEEE Trans. Cybern.* **53**, 7483–7496 (2022).
- Zhang, J. et al. Human-in-the-loop optimization of exoskeleton assistance during walking. *Science* **356**, 1280–1284 (2017).
This study highlights the effectiveness of human-in-the-loop optimization for increasing the benefits of an exoskeleton.
- Poggensee, K. L. & Collins, S. H. How adaptation, training, and customization contribute to benefits from exoskeleton assistance. *Sci. Robot.* **6**, eabf1078 (2021).
This study highlights the importance of human adaptation in achieving effective human–robot interaction.
- Witte, K. A., Fiers, P., Sheets-Singer, A. L. & Collins, S. H. Improving the energy economy of human running with powered and unpowered ankle exoskeleton assistance. *Sci. Robot.* **5**, eaay9108 (2020).
- Bryan, G. M. et al. Optimized hip–knee–ankle exoskeleton assistance reduces the metabolic cost of walking with worn loads. *J. Neuroeng. Rehabil.* **18**, 161 (2021).
- Song, S. & Collins, S. H. Optimizing exoskeleton assistance for faster self-selected walking. *IEEE Trans. Neural Syst. Rehabil. Eng.* **29**, 786–795 (2021).
- Ding, Y., Kim, M., Kuindersma, S. & Walsh, C. J. Human-in-the-loop optimization of hip assistance with a soft exosuit during walking. *Sci. Robot.* **3**, eaar5438 (2018).
This study illustrates the use of Bayesian optimization for human-in-the-loop optimization.
- Kim, J. et al. Reducing the energy cost of walking with low assistance levels through optimized hip flexion assistance from a soft exosuit. *Sci. Rep.* **12**, 11004 (2018).

18. Haufe, F., Wolf, P. & Riener, R. Human-in-the-loop optimization of a multi-joint wearable robot for movement assistance. *Proc. Autom. Med. Eng.* **1**, 023 (2020).
19. Slade, P., Kochenderfer, M. J., Delp, S. L. & Collins, S. H. Personalizing exoskeleton assistance while walking in the real world. *Nature* **610**, 277–282 (2022).
This study demonstrates a data-driven method for human-in-the-loop optimization and provides an example of optimization under naturalistic conditions.
20. Ingraham, K. A., Remy, C. D. & Rouse, E. J. The role of user preference in the customized control of robotic exoskeletons. *Sci. Robot.* **7**, eabj3487 (2022).
21. Lee, U. H. et al. User preference optimization for control of ankle exoskeletons using sample efficient active learning. *Sci. Robot.* **8**, eadg3705 (2023).
22. Kantharaju, P. et al. Reducing squat physical effort using personalized assistance from an ankle exoskeleton. *IEEE Trans. Neural Syst. Rehabil. Eng.* **30**, 1786–1795 (2022).
23. Pang, M. et al. Stiffness optimization based on muscle fatigue and muscle synergy for passive waist assistive exoskeleton. *Robot. Intell. Autom.* **43**, 209–224 (2023).
24. Koginov, G. et al. Human-in-the-loop personalization of a bi-articular wearable robot's assistance for downhill walking. *IEEE Trans. Med. Robot. Bionics* **6**, 328–339 (2023).
25. Hamaya, M., Matsubara, T., Noda, T., Teramae, T. & Morimoto, J. Learning task-parameterized assistive strategies for exoskeleton robots by multi-task reinforcement learning. In *IEEE International Conference on Robotics and Automation (ICRA)* 5907–5912 (IEEE, 2017).
26. Liu, R. et al. Adaptive symmetry reference trajectory generation in shared autonomy for active knee orthosis. *IEEE Robot. Autom. Lett.* **8**, 3118–3125 (2023).
27. Li, Z., Li, Q., Huang, P., Xia, H. & Li, G. Human-in-the-loop adaptive control of a soft exo-suit with actuator dynamics and ankle impedance adaptation. *IEEE Trans. Cybern.* **53**, 7920–7932 (2023).
28. Kantharaju, P. et al. Framework for personalizing wearable devices using real-time physiological measures. *IEEE Access* **11**, 81389–81400 (2023).
29. Wen, T. C., Jacobson, M., Zhou, X., Chung, H. J. & Kim, M. in *Proc. 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* 3431–3436 (IEEE, 2020).
30. Wen, Y., Si, J., Brandt, A., Gao, X. & Huang, H. H. Online reinforcement learning control for the personalization of a robotic knee prosthesis. *IEEE Trans. Cybern.* **50**, 2346–2356 (2019).
31. Tankink, T., Carloni, R. & Hijmans, J. M. & Houdijk, H. Human-in-the-loop optimization of rocker shoes via different cost functions during walking. *J. Biomech.* **166**, 112028 (2024).
This study provides an example of human-in-the-loop optimization of a non-robotic device.
32. Tankink, T., Houdijk, H. & Hijmans, J. M. Human-in-the-loop optimized rocker profile of running shoes to enhance ankle work and running economy. *Eur. J. Sport Sci.* **24**, 164–173 (2024).
33. Huang, G., Lin, S. & Xie, L. Human-in-the-loop optimization of knee-joint biomechanical energy harvester to maximize power generation with minimal user effort. *Energy Convers. Manage.* **283**, 116913 (2023).
34. Felt, W., Selinger, J. C., Donelan, J. M. & Remy, C. D. “Body-in-the-loop”: optimizing device parameters using measures of instantaneous energetic cost. *PLoS One* **10**, e0135342 (2015).
This study provides an example of an early, gradient-based approach to human-in-the-loop optimization.
35. Garcia-Rosas, R., Tan, Y., Oetomo, D., Manzie, C. & Choong, P. Personalized online adaptation of kinematic synergies for human-prosthesis interfaces. *IEEE Trans. Cybern.* **51**, 1070–1084 (2019).
36. Catkin, B. & Patoglu, V. Preference-based human-in-the-loop optimization for perceived realism of haptic rendering. *IEEE Trans. Haptics* **16**, 470–476 (2023).
37. Fauvel, T. & Chalk, M. Human-in-the-loop optimization of visual prosthetic stimulation. *J. Neural Eng.* **19**, 036038 (2022).
This study provides an example of user preference as an optimization objective, in this case applied to a retinal prosthesis.
38. Sánchez, N. et al. Multi-site identification and generalization of clusters of walking behaviors in individuals with chronic stroke and neurotypical controls. *Neurorehabil. Neural Repair* **37**, 810–822 (2023).
39. Lamers, E. P., Yang, A. J. & Zelik, K. E. Feasibility of a biomechanically-assistive garment to reduce low back loading during leaning and lifting. *IEEE Trans. Biomed. Eng.* **65**, 1674–1680 (2017).
40. Nuesslein, C. et al. Comparing metabolic cost and muscle activation for knee and back exoskeletons in lifting. *IEEE Trans. Med. Robot. Bionics* **6**, 224–234 (2023).
41. Kazerooni, H., Racine, J.-L., Huang, L. & Steger, R. in *Proc. 2005 IEEE International Conference on Robotics and Automation* 4353–4360 (IEEE, 2005).
This study describes an early exoskeleton that did not improve user performance despite extensive investment, illustrating the risks of a traditional development approach.
42. Garcia, M., Chatterjee, A., Ruina, A. & Coleman, M. The simplest walking model: stability, complexity, and scaling. *J. Biomech. Eng.* **120**, 281–288 (1998).
43. Dembia, C. L., Silder, A., Uchida, T. K., Hicks, J. L. & Delp, S. L. Simulating ideal assistive devices to reduce the metabolic cost of walking with heavy loads. *PLoS One* **12**, e0180320 (2017).
44. Sivi, C. et al. Offline assistance optimization of a soft exosuit for augmenting ankle power of stroke survivors during walking. *IEEE Robot. Autom. Lett.* **5**, 828–835 (2020).
45. Jackson, R. W. & Collins, S. H. An experimental comparison of the relative benefits of work and torque assistance in ankle exoskeletons. *J. Appl. Physiol.* **119**, 541–557 (2015).
46. Caputo, J. M. & Collins, S. H. A universal ankle-foot prosthesis emulator for human locomotion experiments. *J. Biomech. Eng.* **136**, 035002 (2014).
47. Witte, K. A., Zhang, J., Jackson, R. W. & Collins, S. H. in *Proc. 2015 IEEE International Conference on Robotics and Automation (ICRA)* 1223–1228 (IEEE, 2015).
48. Anderson, A. et al. A robotic emulator for the systematic exploration of transtibial biarticular prosthesis designs. Preprint at <https://doi.org/10.36227/techrxiv.24417310.v1> (2023).
49. Portnova, A. A., Mukherjee, G., Peters, K. M., Yamane, A. & Steele, K. M. Design of a 3D-printed, open-source wrist-driven orthosis for individuals with spinal cord injury. *PLoS One* **13**, e0193106 (2018).
50. Severin, A. C. et al. Case report: adjusting seat and backrest angle improves performance in an elite paralympic rower. *Front. Sports Act. Living* **3**, 625656 (2021).
51. Sanz-Pena, I., Jeong, H. & Kim, M. Personalized wearable ankle robot using modular additive manufacturing design. *IEEE Robot. Autom. Lett.* **8**, 4935–4942 (2023).
52. Sloop, L. H. et al. Effects of a soft robotic exosuit on the quality and speed of overground walking depends on walking ability after stroke. *J. Neuroeng. Rehabil.* **20**, 113 (2023).
53. Walsh, C. Human-in-the-loop development of soft wearable robots. *Nat. Rev. Mater.* **3**, 78–80 (2018).
54. Xu, L. et al. Reducing the muscle activity of walking using a portable hip exoskeleton based on human-in-the-loop optimization. *Front. Bioeng. Biotechnol.* **11**, 1006326 (2023).
55. Kong, H. M. A Personalized Quasi-passive Ankle Exoskeleton Using Human-in-the-loop Optimization Approaches Doctoral dissertation, KTH Royal Institute of Technology (2023).
56. Hybart, R., Villancio-Wolter, K. S. & Ferris, D. P. Metabolic cost of walking with electromechanical ankle exoskeletons under proportional myoelectric control on a treadmill and outdoors. *PeerJ* **11**, e15775 (2023).
57. Kinsey, H., Upton, E. & Young, A. Towards meaningful community ambulation in individuals post stroke through use of a smart hip exoskeleton: a preliminary investigation. *Assist. Technol.* **36**, 198–208 (2023).
58. Fang, Y., Orekhov, G. & Lerner, Z. Improving the energy cost of incline walking and stair ascent with ankle exoskeleton assistance in cerebral palsy. *IEEE Trans. Biomed. Eng.* **69**, 2143–2152 (2021).
59. Caputo, J. M. et al. Robotic emulation of candidate prosthetic foot designs may enable efficient, evidence-based, and individualized prescriptions. *J. Prosthet. Orthot.* **34**, 202–212 (2022).
60. Welker, C. G., Voloshina, A. S., Chiu, V. L. & Collins, S. H. Shortcomings of human-in-the-loop optimization of an ankle-foot prosthesis emulator: a case series. *R. Soc. Open Sci.* **8**, 202020 (2021).
61. Arelekatti, V. N. M. & Winter, A. G. V. in *Proc. 2015 IEEE International Conference on Rehabilitation Robotics (ICORR)* 350–356 (IEEE, 2015).
62. Mattson, C. A. & Winter, A. G. Why the developing world needs mechanical design. *J. Mech. Des.* **138**, 070301 (2016).
63. Eikevåg, S. W., Erichsen, J. F. & Steinert, M. in *Proc. The Engineering of Sport 14* 1–2 (International Sports Engineering Association, 2022).
64. Quintero, D., Villarreal, D. J., Lambert, D. J., Kapp, S. & Gregg, R. D. Continuous-phase control of a powered knee–ankle prosthesis: amputee experiments across speeds and inclines. *IEEE Trans. Robot.* **34**, 686–701 (2018).
65. Geyer, H. & Herr, H. A muscle-reflex model that encodes principles of legged mechanics produces human walking dynamics and muscle activities. *IEEE Trans. Neural Syst. Rehabil. Eng.* **18**, 263–273 (2010).
66. Varol, H. A., Sup, F. & Goldfarb, M. Multiclass real-time intent recognition of a powered lower limb prosthesis. *IEEE Trans. Biomed. Eng.* **57**, 542–551 (2009).
67. Simon, A. M. et al. Configuring a powered knee and ankle prosthesis for transfemoral amputees within five specific ambulation modes. *PLoS One* **9**, e99387 (2014).
68. Tran, M., Gabert, L., Cempini, M. & Lenzi, T. A lightweight, efficient fully powered knee prosthesis with actively variable transmission. *IEEE Robot. Autom. Lett.* **4**, 1186–1193 (2019).
69. Song, Y., Romero, A., Müller, M., Koltun, V. & Scaramuzza, D. Reaching the limit in autonomous racing: optimal control versus reinforcement learning. *Sci. Robot.* **8**, eadg1462 (2023).
70. Slade, P., Kochenderfer, M. J., Delp, S. L. & Collins, S. H. Sensing leg movement enhances wearable monitoring of energy expenditure. *Nat. Commun.* **12**, 4312 (2021).
71. Revi, D. A., Alvarez, A. M., Walsh, C. J., De Rossi, S. M. & Awad, L. N. Indirect measurement of anterior-posterior ground reaction forces using a minimal set of wearable inertial sensors: from healthy to hemiparetic walking. *J. Neuroeng. Rehabil.* **17**, 82 (2020).
72. Ramadurai, S., Jeong, H. & Kim, M. Predicting the metabolic cost of exoskeleton-assisted squatting using foot pressure features and machine learning. *Front. Robot. AI* **10**, 1166248 (2023).
73. Flach, P. & Matsubara, E. in *Dagstuhl Seminar Proceedings* Vol. 71611–10 (Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2008).
74. Wang, W., Raitor, M., Collins, S., Liu, C. K. & Kennedy, M. in *Proc. 2023 IEEE International Conference on Robotics and Automation (ICRA)* 10483–10489 (IEEE, 2023).
75. Eveld, M. E., King, S. T., Vailati, L. G., Zelik, K. E. & Goldfarb, M. On the basis for stumble recovery strategy selection in healthy adults. *J. Biomech. Eng.* **143**, 071003 (2021).
76. Chasnov, B. J., Ratliff, L. J. & Burden, S. A. Human adaptation to adaptive machines converges to game-theoretic equilibria. Preprint at <https://arxiv.org/abs/2305.01124> (2023).
77. Snaterse, M., Ton, R., Kuo, A. D. & Donelan, J. M. Distinct fast and slow processes contribute to the selection of preferred step frequency during human walking. *J. Appl. Physiol.* **110**, 1682–1690 (2011).
78. Finley, J. M., Bastian, A. J. & Gottschall, J. S. Learning to be economical: the energy cost of walking tracks motor adaptation. *J. Physiol.* **591**, 1081–1095 (2013).
79. Nikolaidis, S., Nath, S., Procaccia, A. D. & Srinivasa, S. in *Proc. 2017 ACM/IEEE International Conference on Human-Robot Interaction* 323–331 (IEEE, 2017).
80. Medrano, R. L., Thomas, G. C., Margolin, D. & Rouse, E. J. The economic value of augmentative exoskeletons and their assistance. *Commun. Eng.* **2**, 43 (2023).
81. Brown, G. L., Seethapathi, N. & Srinivasan, M. A unified energy-optimality criterion predicts human navigation paths and speeds. *Proc. Natl Acad. Sci.* **118**, e2020327118 (2021).
82. Umker, T., Lamothe, C. J., Houdijk, H., van der Woude, L. H. & Beek, P. J. Postural threat during walking: effects on energy cost and accompanying gait changes. *J. Neuroeng. Rehabil.* **11**, 71 (2014).
83. Park, K. W., Choi, J. & Kong, K. Iterative learning of human behavior for adaptive gait pattern adjustment of a powered exoskeleton. *IEEE Trans. Robot.* **38**, 1395–1409 (2022).
This study illustrates the potential for human-robot interaction to improve mobility for individuals with severe impairments.
84. Antos, S. A., Kording, K. P. & Gordon, K. E. Energy expenditure does not solely explain step length–width choices during walking. *J. Exp. Biol.* **225**, jeb243104 (2022).
85. McDonald, K. A., Cusumano, J. P., Hieronymi, A. & Rubenson, J. Humans trade off whole-body energy cost to avoid overburdening muscles while walking. *Proc. R. Soc. B* **289**, 20221189 (2022).

86. Mombaur, K., Truong, A. & Laumond, J. P. From human to humanoid locomotion—an inverse optimal control approach. *Auton. Robots* **28**, 369–383 (2010).
87. Tucker, M. et al. in *Proc. 2020 IEEE International Conference on Robotics and Automation (ICRA)* 2351–2357 (IEEE, 2020).
88. Ingraham, K. A., Tucker, M., Ame, A. D., Rouse, E. J. & Shepherd, M. K. Leveraging user preference in the design and evaluation of lower-limb exoskeletons and prostheses. *Curr. Opin. Biomed. Eng.* **28**, 100487 (2023).
89. Brunner, C., Fischer, A., Luig, K. & Thies, T. Pairwise support vector machines and their application to large scale problems. *J. Mach. Learn. Res.* **13**, 2279–2292 (2012).
90. Astudillo, R. et al. in *Proc. ICML 2023 Workshop The Many Facets of Preference-Based Learning* (ICML, 2023).
91. Hansen, N. in *Towards a New Evolutionary Computation. Studies in Fuzziness and Soft Computing*, Vol. 192 (eds Lozano, J. A., Larrañaga, P., Inza, I. & Bengoetxea, E.) 75–102 (Springer, 2006).
92. Kochenderfer, M. J. & Wheeler, T. A. *Algorithms for Optimization* (MIT Press, 2019).
93. Lakmazaheri, A. et al. Optimizing exoskeleton assistance to improve walking speed and energy economy for older adults. *J. Neuroeng. Rehabil.* **21**, 1 (2024).
94. Han, H. et al. Selection of muscle-activity-based cost function in human-in-the-loop optimization of multi-gait ankle exoskeleton assistance. *IEEE Trans. Neural Syst. Rehabil. Eng.* **29**, 944–952 (2021).
95. Kutulakos, Z. & Slade, P. Simulating human-in-the-loop optimization of exoskeleton assistance to compare optimization algorithm performance. Preprint at *bioRxiv* <https://doi.org/10.1101/2024.04.05.587982> (2024).
96. Antonova, R., Rai, A. & Atkeson, C. G. in *Proc. 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)* 22–28 (IEEE, 2016).
97. Kim, M. et al. Human-in-the-loop Bayesian optimization of wearable device parameters. *PLoS One* **12**, e0184054 (2017).
98. Kim, M. et al. in *Proc. 2019 International Conference on Robotics and Automation (ICRA)* 9173–9179 (IEEE, 2019).
99. Denning, P. J. Working sets past and present. *IEEE Trans. Softw. Eng.* **1**, 64–84 (1980).
100. Franks, P. W. et al. Comparing optimized exoskeleton assistance of the hip, knee, and ankle in single and multi-joint configurations. *Wearable Technol.* **2**, e16 (2021).
101. Vasudevan, E. V., Torres-Oviedo, G., Morton, S. M., Yang, J. F. & Bastian, A. J. Younger is not always better: development of locomotor adaptation from childhood to adulthood. *J. Neurosci.* **31**, 3055–3065 (2011).
102. Macready, W. G. & Wolpert, D. H. Bandit problems and the exploration/exploitation tradeoff. *IEEE Trans. Evol. Comput.* **2**, 2–22 (1998).
103. McAllister, M. J., Blair, R. L., Donelan, J. M. & Selinger, J. C. Energy optimization during walking involves implicit processing. *J. Exp. Biol.* **224**, jeb242655 (2021).
104. Hybart, R. & Ferris, D. Gait variability of outdoor vs treadmill walking with bilateral robotic ankle exoskeletons under proportional myoelectric control. *PLoS One* **18**, e0294241 (2023).
105. Waldherr, S., Romero, R. & Thrun, S. A gesture based interface for human-robot interaction. *Auton. Robots* **9**, 151–173 (2000).
106. Landi, C. T., Ferraguti, F., Fantuzzi, C. & Secchi, C. in *Proc. 2018 IEEE International Conference on Robotics and Automation (ICRA)* 3279–3284 (IEEE, 2018).
107. Xiao, X. et al. APPL: adaptive planner parameter learning. *Robot. Auton. Syst.* **154**, 104132 (2022).
108. Kristoffersen, M. B., Franzke, A. W., van der Sluis, C. K., Murgia, A. & Bongers, R. M. The effect of feedback during training sessions on learning pattern-recognition-based prosthesis control. *IEEE Trans. Neural Syst. Rehabil. Eng.* **27**, 2087–2096 (2019).
109. Wong, J. D., Selinger, J. C. & Donelan, J. C. Is natural variability in gait sufficient to initiate spontaneous energy optimization in human walking? *J. Neurophysiol.* **121**, 1848–1855 (2019).
110. Abram, S. J. et al. General variability leads to specific adaptation toward optimal movement policies. *Curr. Biol.* **32**, 2222–2232 (2022).
111. Song, S., Haynes, C. A. & Bradford, J. C. Human cortical, muscular, and kinematic gait adaptation with novel use of an ankle exoskeleton. Preprint at *Research Square* <https://doi.org/10.21203/rs.3.rs-2675191/v1> (2023).
112. Jacobsen, N. A. & Ferris, D. P. Electrocutaneous activity correlated with locomotor adaptation during split-belt treadmill walking. *J. Physiol.* **601**, 3921–3944 (2023).
113. Mu, T., Goel, K. & Brunskill, E. in *Proc. 31st Conference on Neural Information Processing Systems (NIPS 2017)* (Curran Associates, 2017).
114. Ghonasgi, K. et al. in *Proc. 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* 771–776 (IEEE, 2021).
115. Byeon, S., Choi, J., Zhang, Y. & Hwang, I. Stochastic-skill-level-based shared control for human training in urban air mobility scenario. *ACM Trans. Hum.-Robot Interact.* (in the press).
116. Srivastava, M., Biyik, E., Mirchandani, S., Goodman, N. & Sadigh, D. Assistive teaching of motor control tasks to humans. *Adv. Neural Inf. Process. Syst.* **35**, 28517–28529 (2022).
117. Kim, M. et al. Visual guidance can help with the use of a robotic exoskeleton during human walking. *Sci. Rep.* **12**, 3881 (2022).
118. Madden, J. D. Mobile robots: motor challenges and materials solutions. *Science* **318**, 1094–1097 (2007).
119. Burden, S. A., Libby, T., Jayaram, K., Sponberg, S. & Donelan, J. Why animals can outrun robots. *Sci. Robot.* **9**, eadi9754 (2024).
120. Riener, R., Rabezzana, L. & Zimmermann, Y. D. Do robots outperform humans in human-centered domains? *Front. Robot. AI* **10**, 1223946 (2023).
121. Collins, S. H., Wiggin, M. B. & Sawicki, G. S. Reducing the energy cost of human walking using an unpowered exoskeleton. *Nature* **522**, 212–215 (2015).
122. Lee, H. J. et al. A wearable hip assist robot can improve gait function and cardiopulmonary metabolic efficiency in elderly adults. *IEEE Trans. Neural Syst. Rehabil. Eng.* **25**, 1549–1557 (2017).
123. Mooney, L. M., Rouse, E. J. & Herr, H. M. Autonomous exoskeleton reduces metabolic cost of human walking during load carriage. *J. Neuroeng. Rehabil.* **11**, 80 (2014).

Acknowledgements We thank C. Walsh for editorial suggestions.

Competing interests S.H.C. and P.S. are inventors on patents that cover the human-in-the-loop optimization concept. S.H.C. is an inventor on patents for the emulator systems discussed in this paper.

Additional information

Correspondence and requests for materials should be addressed to Patrick Slade or Steven H. Collins.

Peer review information *Nature* thanks the anonymous reviewers for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© Springer Nature Limited 2024