5-9-2022

General Variability Leads to Specific Adaptation Toward Energy Optimal Policies

Sabrina J. Abram
Simon Fraser University

Katherine L. Poggensee
Stanford University

Natalia Sánchez
Chapman University, sanchezaldana@chapman.edu

Surabhi N. Simha
Emory University

James M. Finley
University of Southern California

See next page for additional authors

Follow this and additional works at: https://digitalcommons.chapman.edu/pt_articles

Part of the Physical Therapy Commons

Recommended Citation
SJ. Abram, KL. Poggensee, N. Sánchez, SN. Simha, JM. Finley, SH. Collins, and JM Donelan. General variability leads to specific adaptation toward energy optimal policies. *Current Biology*. 2022;32(10):2222-2232.e5. https://doi.org/10.1016/j.cub.2022.04.015

General Variability Leads to Specific Adaptation Toward Energy Optimal Policies

Comments

Creative Commons License

Copyright

Authors

Sabrina J. Abram, Katherine L. Poggensee, Natalia Sánchez, Surabhi N. Simha, James M. Finley, Steven H. Collins, and J. Maxwell Donelan

General variability leads to specific adaptation toward optimal movement policies

Sabrina J. Abram[1], **Katherine L. Poggensee**[2], **Natalia Sánchez**[3], **Surabhi N. Simha**[4], **James M. Finley**[3,5,6], **Steven H. Collins**[2], **J. Maxwell Donelan**[7,8,9,*]

[1]School of Engineering Science, Simon Fraser University, Burnaby, BC V5A 1S6, Canada

[2]Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, USA

[3]Division of Biokinesiology and Physical Therapy, University of Southern California, Los Angeles, CA 90089, USA

[4]Department of Biomedical Engineering, Emory University, Atlanta, GA 30322, USA

[5]Department of Biomedical Engineering, University of Southern California, Los Angeles, CA 90089, USA

[6]Neuroscience Graduate Program, University of Southern California, Los Angeles, CA 90089, USA

[7]Department of Biomedical Physiology and Kinesiology, Simon Fraser University, Burnaby, BC V5A 1S6, Canada

[8]Twitter: @maxdonelan

[9]Lead contact

SUMMARY

Our nervous systems can learn optimal control policies in response to changes to our bodies, tasks, and movement contexts. For example, humans can learn to adapt their control policy in walking contexts where the energy-optimal policy is shifted along variables such as step frequency or step width. However, it is unclear how the nervous system determines which ways to adapt its control policy. Here, we asked how human participants explore through variations in their control policy to identify more optimal policies in new contexts. We created new contexts using exoskeletons that apply assistive torques to each ankle at each walking step. We analyzed four variables that spanned the levels of the whole movement, the joint, and the muscle: step frequency, ankle angle range, total soleus activity, and total medial gastrocnemius activity. We found that,

*Correspondence: mdonelan@sfu.ca.

across all of these analyzed variables, variability increased upon initial exposure to new contexts and then decreased with experience. This led to adaptive changes in the magnitude of specific variables, and these changes were correlated with reduced energetic cost. The timescales by which adaptive changes progressed and variability decreased were faster for some variables than others, suggesting a reduced search space within which the nervous system continues to optimize its policy. These collective findings support the principle that exploration through general variability leads to specific adaptation toward optimal movement policies.

In brief

Abram et al. find that the nervous system uses general variability to explore new ways of walking when wearing powered exoskeletons. It uses this exploration to identify specific variables that improve walking and to refine its search space of new gaits. It adapts along the variables that improve walking, converging on a new optimal gait.

INTRODUCTION

Humans are adept at learning optimal control policies. We use the term ''control policy'' to refer to the nervous system's mapping from states to the actions taken in those states.[1] The nervous system's actions are realized through motor commands, and its perceived states may range from the lengths of individual muscles to its estimate of the unevenness of the terrain. A control policy can be adapted to optimize an objective function, and when adaptation can no longer improve the objective, we refer to this as the ''optimal control policy.'' The objective function may consist of multiple terms, and the relative importance of these terms may depend on the task and context. For example, in the task of reaching, people can adapt their control policy to optimize an objective function consisting of error and effort.[2] In walking, studies have shown that the nervous system's objective function includes metabolic energetic cost among other terms and constraints such as stability or risk of falling.[3,4] When we measure the relationship between energetic cost and step frequency, we find that it is bowl shaped and that people prefer to walk with a step frequency that coincides with the minimum of this bowl.[5,6] In addition, when we reshape this relationship to shift the energy-optimal step frequency, we find that people adapt their step frequency to optimize energetic cost.[7] This same objective appears to influence the nervous system's control of other gait parameters such as step width, suggesting that the nervous system's continuous learning of optimal control policies for walking is a general phenomenon.[8,9]

Optimizing a control policy involves several steps. One step is to select a more optimal control policy from among candidate policies. Studies suggest that the nervous system greedily selects local solutions that improve on the underlying objective function.[10,11] Prior to selecting the next policy, the nervous system must first evaluate a set of candidate policies. One way to do so is to locally explore. Songbirds, for example, exhibit variability in vocal control, which appears to enable optimization of song performance.[12–14] Similar mechanisms underlie human vocal control.[15] That is, variability is not just an undesirable outcome of noise in the nervous system's control but also a means to explore and discover better outcomes.[16,17] However, variability can be costly—deviating from the previously optimal policy in contexts where the optimal policy has not shifted yields suboptimal

behavior. Even in contexts where the optimal policy has shifted, variability that is not in the same direction as the shift is suboptimal. It seems conceivable that the nervous system has some understanding of which aspects of its control policy to explore and adapt given prior experience. For example, in reaching arm movements, some people exhibit increases in baseline variability along aspects that are relevant to the task, and these selective increases in variability appear to enable faster learning.[18] However, it is useful to consider an alternative perspective where the nervous system has minimal prior experience that it can draw from and may need to vary many aspects of its control policy in order to learn which aspects give rise to better outcomes. The nervous system may need to determine which aspects of its control policy to adapt when learning to walk with an assistive device—a new context that introduces a control system external to that of the nervous system.

How the nervous system explores the space of control policies may influence how quickly it learns new optimal policies. Exploration faces a challenge—as the number of possible states and actions increases, the number of candidate control policies increases rapidly. This expansion of the space of control policies can impede learning as searching through and evaluating many policies takes time—a challenge known as the "curse of dimensionality."[1] The nervous system might overcome this challenge of the combinatorial complexity of the policy space by reducing the dimensionality of the policy space that it searches so it has fewer dimensions along which to explore for new optimal control policies.[19] How the nervous system does this is presently unknown.

Here, we determined how the nervous system explores through variations in its control policy to learn new optimal policies in new contexts. To accomplish this, we performed a post hoc analysis of data from our recent study where we gave participants experience with ankle exoskeleton assistance over 6 non-consecutive days.[20] This recent study found that people arrive at new energy-optimal policies when given sufficient experience with exoskeleton assistance.[20] On the other hand, the nervous system's mechanisms for converging on this new optimal policy are still unclear. Here, we tested three hypotheses about learning in new contexts: (1) the nervous system first explores by increasing general variability—where general variability refers to variability across many or all aspects of gait—to identify variables that improve its objective, (2) rather than decrease exploration across all variables in harmony, the nervous system selectively decreases variability along some variables more quickly than others in order to refine its control policy search space over time, and (3) the nervous system learns to adapt the magnitude of specific variables and exploit a new control policy that reduces energetic cost.

RESULTS

We created new contexts using ankle exoskeletons. We applied assistive torques to each ankle at each walking step by transmitting forces through a Bowden cable that was attached to an ankle lever on an ankle exoskeleton. High-powered off-board motors generated the forces that were transmitted through the cables, and high-frequency controllers commanded the motors to generate the desired torques. All participants experienced two main contexts while walking on an instrumented treadmill: without-assistance and with-assistance. In the without-assistance context, participants walked without the ankle exoskeletons (Figure 1A)

or while wearing the ankle exoskeletons but with slack cables and minimal applied torques (Figure 1B). In the with-assistance context, participants walked with the ankle exoskeletons while they generated torques that acted to extend each ankle during its stance phase. Figure 1C illustrates the general pattern of the applied torque, which was determined by a predefined control law that applied a constant magnitude of peak torque, while rise time, peak torque time, and fall time were constant percentages of the stride time.

To study the nervous system's learning mechanisms in new contexts, we gave participants experience with walking with exoskeleton assistance. On each day over 6 non-consecutive days, all participants walked without the ankle exoskeletons, with the ankle exoskeletons providing minimal applied torques, and with the ankle exoskeletons providing assistive torques. Participants completed these three conditions twice for 6 min each. The three conditions were first completed in random order, and then again in that same order, but reversed (e.g., CABBAC; Figures 1D and 1E). Participants completed additional exoskeleton assistance trials from the second day onward (Figure 1E), but we did not include these trials in our current analyses as they were designed to test the effect of different training protocols.[20] In these additional trials, some participants (n = 5) repeatedly experienced the predefined control law described above. Other participants (n = 5) experienced this predefined control law interspersed between human-in-the-loop optimization, where we used real-time measures of energetic cost to customize the control law parameters (Figure S1). Despite differences in these additional trials, all participants achieved similar reductions in energetic cost in response to the predefined control law that they repeatedly experienced on each day for 6 days (p = 0.62).[20] We therefore grouped participants (n = 10) and restricted our analyses to changes in response to the predefined control law that we define above as the with-assistance context. When accounting for the amount of experience with exoskeleton assistance, we include not only the time spent walking in the with-assistance trials that we analyzed but also the time spent walking in the additional trials because participants also experienced assistive ankle torques in human-in-the-loop optimization.

We studied exploration and adaptation in this new context by measuring changes at the levels of the whole movement, the joint, and the muscle. We focused our analysis on four variables: step frequency, ankle angle range during stance, total soleus activity, and total medial gastrocnemius activity (see STAR Methods). Step frequency can influence exoskeleton assistance through the timing of the assistive torque pattern because rise time, peak torque time, and fall time were all expressed as percentages of stride time in our control law parameterization. Ankle angle range may influence the power and work that the exoskeleton applies to the ankle by changing the angular displacement over which the assistive torque is applied. Lastly, the nervous system may learn to accept assistive torques at the ankle by lowering the contribution to the total ankle torque provided by the extensor muscles. The two primary ankle extensor muscles are soleus and gastrocnemius, and here, we analyzed both of their activities. We selected these variables *a priori*, based on preliminary evidence of how walking can take advantage of ankle exoskeleton assistance.[21–25] These variables reflect only some of the nervous system's control policy parameters—there are likely many other parameters that allow the nervous system to, for example, achieve a particular step frequency using many combinations of ankle angle range, total soleus activity, and total medial gastrocnemius activity.

A general increase in variability upon initial exposure to new contexts

We quantified variability within each variable by high-pass filtering its signal to include timescales of 30 steps or less and then calculating the standard deviation of this filtered signal during the last 3 min of each 6-min trial (see STAR Methods). This method filtered out the relatively slow signal changes that we associate with adaptive changes, but not the relatively rapid changes that occur from step to step and over several steps that we associate with exploration. For most variables, we quantified each participant's without-assistance variability from the condition where they were not wearing the ankle exoskeletons. The exception was ankle angle range where we applied identical calculations but to the condition where participants were walking while wearing the ankle exoskeletons but with the devices applying minimal torques. The reason for this exception was that we needed the exoskeleton sensors to calculate ankle angle. We refer to the without-assistance variability averaged across the two trials on the first day as ''baseline variability.'' We use this baseline variability to normalize with-assistance variability.

Upon initial exposure to the exoskeleton assistance context, participants walked with increased variability across all variables that we analyzed. This was evident when comparing with-assistance variability averaged across the two first day trials to without-assistance baseline variability, with increases ranging from 21% to 279% (mean ± standard deviation, paired t test; step frequency: +56.2% ± 27.6%, p = 2.0 × 10^{-4}, Figure 2A; ankle angle range: +278.8% ± 85.0%, p = 6.8 × 10^{-7}, Figure 2B; total soleus activity: +21.1% ± 25.3%, p = 0.026, Figure 2C; total medial gastrocnemius activity: +46.0% ± 19.6%, p = 1.7 × 10^{-5}, Figure 2D; representative participant: Figure 2E).

A general decrease in variability with increased experience

As participants walked with exoskeleton assistance over multiple days, we determined how variability changed with this increased experience. We found that, as experience increased, participants walked with decreased variability across all variables that we analyzed. This was evident when comparing with-assistance variability measured on the last day to that measured on the first day, with decreases ranging from −18% to −39% (mean ± standard deviation, paired t test; step frequency: −38.5% ± 9.8%, p = 4.9 × 10^{-5}, Figure 2A; ankle angle range: −26.3% ± 22.8%, p = 6.5 × 10^{-3}, Figure 2B; total soleus activity: −18.3% ± 21.2%, p = 0.013, Figure 2C; total medial gastrocnemius activity: −19.7% ± 13.2%, p = 1.5 × 10^{-3}, Figure 2D). By the last day, participants' with-assistance variability was indistinguishable from their baseline variability for step frequency (−5.1% ± 17.3%, paired t test: p = 0.32) and total soleus activity (−3.3% ± 26.9%, paired t test: p = 0.39) but remained elevated for ankle angle range (+166.2% ± 52.5%, paired t test: p = 6.7 × 10^{-6}) and total medial gastrocnemius activity (+15.8% ± 15.8%, paired t test: p = 0.020). That is, the nervous system returned variability toward, and in some cases to, baseline variability with increased experience.

Adaptation occurs along specific variables, and these changes correlate with reduced energetic cost

We use the term ''adaptation'' to refer to changes in the magnitude of a variable that occur with experience. We quantified each variable's magnitude by averaging its signal

during the last 3 min of each 6-min trial. People appear to have an established policy for without-assistance contexts—we observed minimal changes in the magnitude of variables when comparing the first and last days of without-assistance walking (mean ± standard deviation, paired t test; step frequency: $-0.78\% \pm 2.6\%$, p = 0.38; ankle angle range: $5.6\% \pm 7.2\%$, p = 0.046; total soleus activity: $-2.7\% \pm 8.1\%$, p = 0.22; total medial gastrocnemius activity: $-4.0\% \pm 6.5\%$, p = 0.074). We refer to the without-assistance magnitude averaged over the two trials on the first day as the baseline value. We use this baseline value to normalize the measured with-assistance magnitudes. We estimated the energetic cost of each trial in the standard manner—using respiratory gas analysis of the last 3 min of each 6-min trial (see STAR Methods). We normalized energetic cost during with-assistance trials to each participant's energetic cost averaged over two without-assistance trials—when walking with the ankle exoskeletons applying minimal torques—on the first day. We used linear mixed-effects regression to estimate the slope of the relationship between a variable and energetic cost. This mixed-effects model used a single slope for each variable to estimate the relationship that is shared between participants while allowing for individual participant energetic cost intercepts (see STAR Methods).

Participants learned to adapt three of four variables and these three variables correlate with energetic cost. We found a significant correlation between energetic cost and step frequency (slope = 2.2, 95% CI [1.5, 2.8], p = 3.4×10^{-9}; Figure 3A), total soleus activity (slope = 0.69, 95% CI [0.44, 0.94], p = 2.9×10^{-7}; Figure 3C), and total medial gastrocnemius activity (slope = 0.36, 95% CI [0.18, 0.54], p = 1.7×10^{-4}; Figure 3D). The correlation between energetic cost and ankle angle range was not significant (slope = 0.069, 95% CI [-0.030, 0.17], p = 0.17; Figure 3B). As participants gained experience with walking with exoskeleton assistance, variables that correlated with energetic cost adapted in the direction that reduced cost. Comparing the first and last days of with-assistance trials, we found changes in step frequency ($-3.5\% \pm 3.4\%$, paired t test: p = 9.8×10^{-3}, Figure 3E), total soleus activity ($-10.4\% \pm 11.0\%$, paired t test: p = 0.023, Figure 3G), and total medial gastrocnemius activity ($-13.4\% \pm 12.4\%$, paired t test: p = 5.4×10^{-3}, Figure 3H), but not in ankle angle range ($-0.32\% \pm 11.5\%$, paired t test: p = 0.78, Figure 3F). Participants learned to exploit a new control policy that reduced energetic cost by $-25.0\% \pm 9.9\%$ (p = 3.0×10^{-5}, Figure 3I) when walking with ankle exoskeleton assistance on the last day compared with the first day. To be clear, adaptation along a specific variable that correlates with reductions in energetic cost suggests but does not prove that energetic cost drives this adaptation. We cannot rule out, for example, that the observed adaptation is due to optimization of different objectives that also correlate with changes in the observed variables.

Variability decreases quickly for quickly adapting variables

We modeled the timing of adaptive changes, as well as the timing of decreases in variability, as an exponential decrease from an initial value to a final steady-state value. We used nonlinear mixed-effects regression with a single time constant to estimate the time constant that is shared between participants while allowing for individual participant offsets (Equation 2 in STAR Methods). We then used bootstrapping to determine the dispersion of each time constant (see STAR Methods). We found that participants adapted along step

frequency with a time constant of 208.9 min (interquartile range [IQR] [150.0, 311.7], R-squared = 0.52; Figure 4A), total soleus activity with a time constant of 5.7 min (IQR [3.9, 8.6], R-squared = 0.47; Figure 4C), and total medial gastrocnemius activity with a time constant of 7.8 min (IQR [6.3, 10.2], R-squared = 0.77; Figure 4D). We found that variability decreased along step frequency with a time constant of 112.4 min (IQR [97.5, 129.9], R-squared = 0.70; Figure 4E), ankle angle range with a time constant of 87.4 min (IQR [70.5, 108.7], R-squared = 0.48; Figure 4F), total soleus activity with a time constant of 4.6 min (IQR [3.4, 6.3], R-squared = 0.58; Figure 4G), and total medial gastrocnemius activity with a time constant of 5.6 min (IQR [4.6, 6.9], R-squared = 0.63; Figure 4H).

Variables that adapted quickly also had rapid decreases in variability. We used bootstrapping to determine the dispersion of the time constant of adaptation for each variable and then tested for differences in time constants between variables (see STAR Methods). We found that total soleus activity and total medial gastrocnemius activity adapted with faster time constants than step frequency (ANOVA; total soleus activity versus step frequency: $p < 0.001$; total medial gastrocnemius activity versus step frequency: $p < 0.001$; total soleus activity versus total medial gastrocnemius activity: $p < 0.001$; Figure 5A). Next, we performed the same analysis but for time constants of variability. We found that variability in a similar way decreased faster for total soleus activity and total medial gastrocnemius activity than for step frequency (ANOVA; total soleus activity versus step frequency: $p < 0.001$; total medial gastrocnemius activity versus step frequency: $p < 0.001$; total soleus activity versus total medial gastrocnemius activity: $p < 0.001$; Figure 5B).

DISCUSSION

We provide insight into how the nervous system navigates a space of control policies to learn new optimal policies in new contexts. We created new contexts using ankle exoskeleton assistance and studied learning as energy optimization in human walking. We analyzed two processes—variability and adaptation—across four variables—step frequency, ankle angle range, total soleus activity, and total medial gastrocnemius activity. We found that, at the beginning of experience in new contexts, variability increased across all variables that we analyzed, and with increased experience, variability decreased across all variables. This appeared to lead to adaptive changes in the magnitude of specific variables, and these changes correlated with reduced energy cost. Adaptation progressed quickly for some variables but slowly for others, suggesting that the nervous system can independently control these variables. Variability progressed quickly for some variables but slowly for others, suggesting that the nervous system may optimize in a manner that reduces its control policy search space over time.

These findings generalize to other movement variables. We *a priori* selected four variables based on our understanding of how walking can take advantage of ankle exoskeleton assistance and without knowledge of how these variables changed over time or how these changes were associated with energetic cost. These four variables are only a subset of our measured dataset—which includes stride parameters, ground reaction forces, joint kinematics, and muscle activity—enabling us to test whether the conclusions we arrive at from our first analysis generalize to other variables. Toward this, we sampled four additional

variables: step width, peak ankle extension angle (which occurs during swing), total rectus femoris activity (a knee extensor and hip flexor muscle), and total biceps femoris activity (a knee flexor and hip extensor muscle). These four additional variables do not directly take advantage of ankle exoskeleton assistance but in the same way capture learning at the levels of the whole movement, the joint, and the muscle. Our additional analysis revealed that participants first increased and then decreased variability across three of these four variables: peak ankle extension angle, total rectus femoris activity, and total biceps femoris activity (Figure S2). Participants learned to adapt specific variables and these variables correlated with reductions in energetic cost (Figure S3). Some variables adapted faster than others, and this adaptation was accompanied by decreases in variability (Figure S4). We interpret these collective findings as supporting the principle that general variability leads to specific adaptation toward optimal control policies. Our data and code are open access, allowing others to test additional variables for the generalizability of these conclusions.

These findings generalize to other movement contexts. We use the context of walking on a split-belt treadmill to again test how motor variability changes with motor learning. We reanalyzed a dataset where participants walked for 6 min in a familiar context with the treadmill belts moving at the same speed and then for 45 min in a new context with the treadmill belts moving at different speeds.[26] We found that participants increased variability in step length asymmetry during the initial 100 strides of split-belt walking compared with the baseline during tied-belt walking ($p = 2.0 \times 10^{-7}$; Figure 6). Participants then decreased variability in step length asymmetry during the final 100 strides of 45 min of split-belt walking compared with the initial 100 strides ($p = 1.2 \times 10^{-4}$; Figure 6). During this time, they also learned to adopt positive asymmetries and reduce energetic cost.[26] The process of adaptation in split-belt walking has been explained by more than one theory. One well established theory is that people adapt to minimize a sensory prediction error which is traditionally tracked using step length asymmetry—step length asymmetry gradually adapts when the treadmill perturbation is introduced and then shows after-effects when it is removed.[27,28] This observation is consistent with the hypothesis that sensory prediction error, or difference between predicted and observed outcomes, drives the nervous system's adaptation.[29] For example, prolonged exposure over multiple days to split-belt walking results in both recalibration of movement and of perception of the split-belt difference.[30] However, there is also evidence that adaptation is not driven by this process alone. Another theory is that people adapt to reduce energetic cost—foot placement gradually adapts to take advantage of positive work performed by the treadmill (due to one belt moving faster than the other belt) and reduce work performed by the legs.[31] This can explain the observation that participants adopt positive step length asymmetries.[26] Interestingly, these positive asymmetries were also observed during multiple day exposure, indicating that energetic cost reduction among other mechanisms might be at play during locomotor adaptation.[30] Importantly, we suspect the role of variability for adaptation is fairly independent of these candidate learning processes.[18] Here, we have demonstrated that changes in motor variability generalize not only to other movement variables but also to other movement contexts.

The nervous system appears to reduce the dimensionality of its search space over time. We infer the extent of the space in which the nervous system searches for new policies from the

extent of our measured gait, joint, and muscle variability. Variability first increases across nearly all biomechanical variables, suggesting a large search space. Variability then quickly decreases for some biomechanical variables as their adaptation progresses, suggesting a reduced space in which the nervous system optimizes fewer control policy parameters. The biomechanical variables we have chosen to measure do not map perfectly onto the nervous system's control policy parameters—it is unlikely that the nervous system's control policy has a parameter specifically for step frequency or total soleus activity. Consequently, if the nervous system were to stop exploring just one of its control parameters, we would likely measure this as a small decrease in variability across many biomechanical variables rather than a return to baseline variability in any single biomechanical variable, and because biomechanical variability is a consequence of many control parameters, that some biomechanical variables do return to baseline variability suggests a large reduction in the nervous system's control policy search space.

A reduced search space is superior to continuing with general exploration of learned policies. One reason for this is that continuing to explore variables that are already optimized results in needlessly increased costs. A second reason is that searching by general exploration may take a long time. That is, exploring more variables results in larger search spaces, and these search spaces fall victim to the curse of dimensionality where there are exponentially more candidate control policies to evaluate.[32] If the nervous system reduces the number of variables along which it explores, it has fewer combinations of states and actions to try and can therefore employ a more directed optimization for selecting more optimal control policies from among candidate policies.[10,11] In contrast to general exploration during learning of a new policy, learned policies often appear to have a low-dimensional structure. In motor coordination, for example, muscle activation patterns can be explained by a limited set of muscle synergies—or muscle activation patterns with consistent spatial and temporal characteristics.[33–36] Moreover, in neural systems such as the motor cortex of a monkey, relatively complex responses from individual neurons can be explained by relatively simple responses from a population of neurons.[37,38] Our findings show that, in new contexts, the nervous system can arrive at such a low-dimensional structure of control after first benefiting from general exploration. An interesting and open question is what elicits adaptation and why adaptation occurs quickly for some variables, and more slowly for others. There are several factors—which do not act in isolation—that can influence adaptation. Two factors may be the level of exploratory variability and the cost gradient—both can influence the range of cost savings that the nervous system experiences. Another factor may be how the biomechanical variables relate to the nervous system's control policy parameters—a variable that we measure may reflect one or many parameters that the nervous system optimizes, and the level of complexity may influence the timescale of adaptation we observe.

Some of our findings are consistent with findings in published literature. It is well established that energy optimization is a major objective in the nervous system's control of walking.[7,9,39] However, it is unclear how the nervous system navigates the many possible control policies to learn new optimal policies. We demonstrate that people explore through general variability to identify variables that improve the objective, and those that do not. People then refine the space in which they explore as they learn to adapt specific

variables and reduce energetic cost. We suspect that this understanding can be applied to the optimization of other cost functions too. For example, a multi-objective cost function might provide a more general view of optimization during non-steady-state walking by including terms such as energy, stability, or some weighted combination of these and other terms.[4,9] A recent study provided a theoretical basis for how people learn new optimal policies in new contexts such as walking with exoskeleton assistance and walking on a splitbelt treadmill— they found that prioritizing stability over short timescales and optimizing energy expenditure over long timescales, as well as using exploratory variability to estimate gradients, can explain learning in these contexts.[40] Here, we provide experimental evidence of this. The notion that the nervous system uses exploration to evaluate candidate control policies and that it selects policies that optimize energetic cost, is similar to what we find in our current study.

Not all variability is exploration. We must first differentiate exploration from other contributors to measured biomechanical variability, such as that arising from unintentional noise in the nervous system's control, variability in the forces that muscles produce, or unpredicted changes in the environment.[41] Müller and Sternad sought to do so in a throwing task, where participants could influence their performance of hitting a target by varying two parameters: angle and velocity at release.[42] They decomposed variability along these two parameters into components that differentiated task-relevant variability from stochastic noise. Here, we propose that baseline levels of variability in familiar contexts may mostly represent unintentional noise, either from the nervous system's control or some other source (i.e., stochastic noise in Müller and Sternad's terminology). This can then be used as a benchmark against which variability in new contexts is compared. We found that, upon first exposure to a new context, variability increased above baseline levels across all variables that we analyzed. With experience in the new context, variability gradually decreased and then plateaued across all variables that we analyzed, converging on baseline levels for some variables. We interpret variability as exploration when it is information seeking. That variability decreased as people learned to adapt the magnitude of specific variables— and that these two processes occurred over similar timescales for each variable—suggests that variability may reflect exploration. As in other studies that aim to determine how variability relates to intentional exploration, we recognize that we can never entirely rule out the alternative hypothesis that variability is unintentional. However, this alternative hypothesis seems less likely based on theory—a key concept in reinforcement learning is that exploration improves learning of new optimal policies—and evidence—higher levels of motor variability appear to enable faster learning of new optimal policies.[1,18] Plateauing of variability may indicate that the nervous system has settled on a new policy and has shifted from exploring candidate policies to exploiting the new preferred one. That variability remained elevated above baseline levels for some variables may reflect that the nervous system is still refining aspects of the control policy or, as we suspect, it may simply be additional variability introduced by imperfect torque control by the exoskeleton, which was not present when we established baseline levels.

A deeper understanding of the nervous system's mechanisms for learning can be used to both facilitate learning and customize training. We might facilitate learning by encouraging general exploration. That is, we might increase variability across many variables—through

strategies such as biofeedback—and then decrease variability along specific variables as people learn to adapt, reducing the nervous system's policy search space. We might also give experience with specific variables that affect energetic cost, indicating to the nervous system which variables are relevant to optimize.[43] To improve on the design of this study, future studies should seek to determine the energetic cost landscape of walking with ankle exoskeleton assistance by mapping the relationship between energetic cost and biomechanical variables. Future studies should also seek to develop methods for estimating energetic cost with increased time resolution to determine how variability in biomechanical variables relates to changes in energetic cost for energetic cost optimization. This can benefit those who seek to design wearable systems—such as orthoses, exoskeletons, and prosthetics—by facilitating learning, and then evaluating people's optimal responses to a range of designs. We also might customize training time by using baseline levels of variability as a benchmark to indicate when the nervous system initiates exploration, and when it shifts to exploiting a new policy. Determining the onset and termination of learning would be useful for coaching athletes who would benefit from knowing at which point they should transition to learning new skills in order to maximize their high capacity for training, and knowing when to terminate experience is especially important for rehabilitating those with mobility disorders who have a limited capacity for training.

STAR ★ METHODS

RESOURCE AVAILABILITY

Lead contact—Further information and requests for resources should be directed to and will be fulfilled by the lead contact, J. Maxwell Donelan (mdonelan@sfu.ca).

Materials availability—This study did not generate new unique reagents.

Data and code availability—The datasets generated during this study are available at: https://searchworks.stanford.edu/catalog?f%5Bcollection%5D%5B%5D=pp784wp5100.

The codes generated during this study are available at: https://github.com/SFULocomotionLab/GeneralVariabilitySpecificAdaptation.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

We included a total of ten participants (mean ± SD; age: 24 ± 2 years; body mass: 68.3 ± 11.1 kg; height: 1.7 ± 0.094 m; sex: 4 females, 6 males) in our study. To investigate the nervous system's learning mechanisms, we required that participants included in our study learned to reduce their energetic cost of walking with exoskeleton assistance. These ten participants consisted of two groups, which we randomly assigned participants to prior to the experiment. Each group of five participants completed similar protocols but with slight differences in additional trials that they completed from the second day onwards. Our previous study found that, despite these slight differences, these two groups learned to reduce their energetic cost of walking in response to a general pattern of assistive ankle torque on the last day compared to the first day (group 1: p = 0.016; group 2: p = 0.005). They also achieved similar reductions in energetic cost (p = 0.62).[20] We therefore grouped

participants (n = 10) and restricted our analyses to changes in response to this general pattern of assistive ankle torque. We excluded a third group of five participants as they did not meet our requirement of learning to reduce their energetic cost of walking on the last day compared to the first day (group 3: p = 0.73).[20] This was perhaps due to the nature of their additional trials, where they experienced many different patterns of assistive ankle torques. All participants were healthy and had no known gait or cardiopulmonary abnormalities. The Stanford Institutional Review Board approved the study protocol, and all participants gave their written, informed consent before participating in the study.

METHOD DETAILS

Exoskeleton hardware—Participants walked on an instrumented treadmill while wearing a bilateral, tethered ankle exoskeleton emulator that applied assistive torques to each ankle at each walking step. This system is described in more detail in our previous work.[20,44] In brief, we used an off-board controller to command the desired force to an off-board electric motor via a motor driver.[45] This force was then transmitted through a Bowden cable to the end of the ankle lever on the exoskeleton, which in turn applied an ankle plantarflexion torque to the participant. This ankle plantarflexion torque was achieved by the tension in the cable, combined with the moment arm of the ankle lever on the exoskeleton, producing forces on the body where it interfaces with the exoskeleton—the heel, the shank, and the toe.[44]

Exoskeleton controller—Our control system produced desired torque patterns as a function of the user's stride time. As detailed in [20], we used a high-speed controller running at 1000 Hz to achieve the predefined torque pattern (Speedgoat, Liebefeld, Switzerland). This controller sampled from sensors, calculated time since heel strike as well as desired force at that time, and then commanded that force to the motor. We used contact switches on each heel to measure heel strikes and calculate stride time as the difference in time between consecutive heel strikes with the same leg (Pololu, NV, USA; McMaster-Carr, IL, USA). We used strain gauges to measure the tension in each cable and calculate torques by multiplying tension with the moment arm of the ankle lever on the exoskeleton (Omega Engineering, CT, USA). The combination of real-time measured stride times and torques allowed for accurate torque tracking that was a function of time since heel strike, normalized to average stride time. We used average stride time as a filter for large, perhaps inaccurate, changes in stride time that may result in undesirable torques. We calculated average stride time as:

$$t_{avg} = t_{avg}(1 - \mu) + t_{stride}\mu \qquad \text{(Equation 1)}$$

where $\mu = 0.9$. Our controller was also designed to include many features that prioritized the user's comfort.[20] For example, at the beginning of all walking periods with exoskeleton assistance, the peak torque was slowly increased from zero to that of the predefined torque pattern to ensure that participants were not perturbed by the assistance. During each swing phase, the controller initiated a 'swing mode' where the cable tracked the ankle angle with added slack. And during each stance phase, the torque was slowly increased from zero after heel strike and decreased to zero before toe off as prescribed by the parameterization of the control law.

Exoskeleton control law—We parameterized the control law to achieve a range of customized torque profiles, as well as one predefined torque profile. Because our original study was designed to test the relative benefits of training with a predefined control law and a customized control law,[20] some participants received training with repeated exposure to the predefined control law, whereas other participants received training with human-in-the-loop optimization of the control law and only periodic exposure to the predefined control law. All control laws were defined by four parameters: magnitude of peak torque (Nm), timing of peak torque (% stride), rise time (% stride), and fall time (% stride). We defined the magnitude of peak torque to be a function of each participant's body mass. We determined the predefined control law in a pilot experiment prior to the main experiment.[20] In this pilot experiment, ten participants—which were not the participants in our main experiment—completed one day of habituation with the bilateral ankle exoskeletons followed by one day of human-in-the-loop optimization of the control law. The predefined control law for the main experiment was defined as the average optimized parameters from this pilot study: magnitude of peak torque was 0.54M where M is participant's mass in kilograms, timing of peak torque was 52.9% of stride, rise time was 26.2% of stride, and fall time was 9.8% of stride.

Measurements—We measured ground reaction forces, joint kinematics, and muscle activity to quantify variables that people may optimize. First, we measured ground reaction forces and moments while participants walked on an instrumented split-belt treadmill (Bertec, Columbus, OH, USA). We used ground reaction forces and moments to calculate the center of pressure and identify foot contact events as the rapid fore-aft translation in center of pressure during double support. Second, we measured ankle angle in all trials with the exoskeleton using a rotary magnetic encoder mounted on the ankle joint of the exoskeleton (Renishaw, Gloucestershire, UK). We zeroed the encoder during standing on each day, and calculated ankle flexion and extension as the angle from neutral position in the sagittal plane. Third, we collected electromyography (EMG) data from medial gastrocnemius, lateral gastrocnemius, soleus, tibialis anterior, rectus femoris, vastus medialis, biceps femoris, and semitendinosus using surface electrodes on each muscle for both legs, and during all walking trials (Delsys, Boston, MA, USA). For the muscles that we considered in our analyses, we inspected EMG data to exclude channels with poor signal quality. Lastly, we used a respiratory gas analysis system to measure rates of oxygen consumption and carbon dioxide production (Cosmed Quark CPET, Rome, Italy). We calculated gross metabolic power using the standard Brockway equation, and then subtracted each participant's resting metabolic power measured on the same day to obtain net metabolic power.[46] We calculated resting metabolic power as the average metabolic power during the final three minutes of a five-minute standing resting period at the beginning of each day of testing. We collected EMG at 1000 Hz and all other measures at 500 Hz.

Behavioral task—The protocol consisted of a testing session on each day for a total of 6 days. During all testing sessions, participants walked on an instrumented treadmill at a constant speed of 1.25 m/s. For conditions that involved exoskeleton assistance, we instructed participants to ''walk comfortably'' and to ''let the device do the work for you''.

We gave participants experience with walking with exoskeleton assistance over multiple days. On each day, participants experienced at least 3 conditions: walking without the ankle exoskeletons (Figure 1A), walking with the ankle exoskeletons providing minimal applied torques via slack cables (Figure 1B), and walking with the ankle exoskeletons with the predefined control law generating assistive ankle torques at each walking step (Figure 1C). Some participants had an additional walking condition from the second day onwards, where they experienced the final control law from human-in-the-loop optimization during additional trials on that same day. Participants completed these conditions twice for six minutes each. The conditions were first completed in random order, and then again in that same order, but reversed (e.g. CABBAC; Figures 1D and 1E). When comparing between conditions or days, we averaged a given variable over the final three minutes of each six-minute trial, and then across the two repeated six-minute trials on each day.

Participants completed additional trials which altered the training that each group received. We did not include these trials in our analyses as they were designed primarily to test the effect of different training protocols in our previous study.[20] They were also designed to increase participants' experience with the predefined control law. Participants completed these trials from the second day onward, prior to the main six-minute trials (Figure 1E). In brief, one group (n = 5) repeatedly experienced the predefined control law at each walking step over 72 minutes (Figure S1). A second group (n = 5) experienced two minutes of the predefined control law followed by 16 minutes of human-in-the-loop optimization, and repeated this four times resulting in 72 minutes of training (Figure S1). During periods of human-in-the-loop optimization, participants experienced a series of eight control laws for two minutes each. We selected these control laws based on our estimate of the optimal control law, which was determined by an algorithm that ranked previously experienced control laws by their respective energetic cost measurements. The process of human-in-the-loop optimization is described in more detail in our previous work.[20,44]

QUANTIFICATION AND STATISTICAL ANALYSIS

We wrote custom MATLAB scripts to process and analyze the data, as well as perform statistical comparisons and generate figures included in this manuscript.

Step frequency calculation—We quantified the variability and magnitude of step frequency. This variable can influence exoskeleton assistance through the timing of the assistive torque pattern because rise time, peak torque time, and fall time were all expressed as percentages of stride time in our control law parameterization. We calculated step frequency by identifying foot contact events and then taking the inverse time difference between consecutive steps. We determined the variability within step frequency by applying a third-order, high-pass, bidirectional digital Butterworth filter with a cut-off frequency of 0.033 steps^{-1} (period of 30 steps), and then calculated the standard deviation of this filtered signal during the last three minutes of each six-minute trial. We used MATLAB's filtfilt command to perform zero-phase digital filtering. We determined the cut-off frequency to be 0.033 steps^{-1} based on previous studies of variability in human walking,[47] as well as visual inspection of the power spectrum. We calculated the magnitude of step frequency as the average of this signal during the last three minutes of each six-minute trial.

Ankle angle calculation—We quantified the variability and magnitude of ankle angle range during stance. This variable can influence the power and work that the exoskeleton applies to the ankle by changing the angle over which the torque is applied. We calculated ankle angle range during stance by first time-locking ankle angle to heel strike events, and then calculating the difference between the maximum and minimum ankle angles during stance. We determined the variability within ankle angle range by applying a high-pass filter with a cut-off frequency of 0.033 steps^{-1} (third-order Butterworth), and then calculated the standard deviation of this filtered signal during the last three minutes of each six-minute trial. We calculated the magnitude of ankle angle range as the average of this signal during the last three minutes of each six-minute trial.

Ankle extensor muscle activity calculation—We quantified the variability and magnitude of total ankle extensor muscle activity. The nervous system may learn to accept assistive torques at the ankle by lowering the contribution to the total ankle torque provided by the extensor muscles. We selected two variables at the level of the muscle as there are two primary ankle extensor muscles—soleus and gastrocnemius. For each muscle on each leg, we applied a high-pass filter with a 20 Hz cut-off (third-order Butterworth), rectified the signal, and then applied a low-pass filter with a 6 Hz cut-off (third-order Butterworth).[48] We next time-locked the signal to heel strike events, divided each stance phase into 100 evenly spaced segments, and normalized each muscle's activity for each participant to their same day average peak activation while walking without the ankle exoskeletons. We calculated total soleus activity and total medial gastrocnemius activity by integrating each muscle's activity during stance at each walking step. Similar to our previous analyses, we quantified the variability of total soleus activity and total medial gastrocnemius activity by applying a high-pass filter with a cut-off frequency of 0.033 steps^{-1} (third-order Butterworth), and then calculated the standard deviation of each filtered signal during the last three minutes of each six-minute trial. We calculated the magnitude of each muscle's total activity as the average of its signal during the last three minutes of each six-minute trial.

Additional variables—We analyzed four additional variables: step width, peak ankle extension angle, total rectus femoris activity, and total biceps femoris activity. We calculated step width by identifying foot contact events and then taking the difference between the lateral centers of pressures—which we determined by dividing the left and right lateral moments by their vertical forces—for consecutive steps. We calculated peak ankle extension angle by first time-locking ankle angle to heel strike events, and then calculating the peak angle during the stride. Similar to our previous analysis of muscle activity, we calculated total rectus femoris activity and total biceps femoris activity by first applying a high-pass filter, rectifying the signals, and then applying a low-pass filter. We next time-locked the signals to heel strike events and normalized each muscle's activity to its average peak activation while walking without the ankle exoskeletons on the same day. We calculated total muscle activity by integrating each muscle's activity during stance at each walking step. Lastly, we quantified the variability and magnitude of each additional variable in the same way as our original four variables. That is, we calculated the variability of each variable by applying a high-pass filter with a cut-off frequency of 0.033 steps^{-1} (third-order Butterworth), and then calculated the standard deviation of this filtered signal during the

last three minutes of each six-minute trial. We calculated the magnitude of each variable by averaging its signal during the last three minutes of each six-minute trial. We analyzed these additional variables and reported these additional results in Figures S2–S4.

Sensitivity analysis—Our choice of high-pass filter cut-off frequency makes an assumption about the timescale of changes in variability. We performed a sensitivity analysis of this high-pass filter cut-off frequency to determine the effect of this assumption. In our main analysis, we used a high-pass filter cut-off frequency of 0.033 steps^{-1} (period of 30 steps). Here, we performed the same analysis but with high-pass filter cut-off frequencies of 0.1 steps^{-1} (period of 10 steps) and 0.02 steps^{-1} (period of 50 steps). This sensitivity analysis revealed that changes in variability over timescales ranging from 10–50 steps do not impact our findings of general variability. That is, across all variables that we analyzed, we observed increases in with-assistance variability on the first day compared to without-assistance baseline, and then decreases in with-assistance variability on the last day compared to the first day. We summarized these results in Table S1.

Statistics—We compared variability in the with-assistance condition to variability in the without-assistance condition for each variable. We refer to variability in the without-assistance condition averaged across the two six-minute trials on the first day as ''baseline variability''. For most variables, we quantified each participant's baseline from the condition where they walked without the ankle exoskeletons. For ankle angle range, we instead used the condition where participants walked with the ankle exoskeletons but with the devices applying minimal torques as we needed the exoskeleton sensors to calculate ankle angle. We calculated with-assistance variability on the first day by averaging variability across the two six-minute trials where participants walked with the ankle exoskeletons with the predefined control law generating assistive ankle torques at each walking step. For each variable, we used a one-tailed paired Student's t-test to determine whether with-assistance variability on the first day was higher than baseline variability. In all statistical analyses, we used a significance level of 0.05.

We determined how participants modified with-assistance variability along each variable as they gained experience walking with exoskeleton assistance. We calculated with-assistance variability on the last day by averaging variability across the two six-minute trials. For each variable, we used a one-tailed paired Student's t-test to determine whether with-assistance variability on the last day was lower than the first day. We used a two-tailed paired Student's t-test to determine if with-assistance variability on the last day had converged on baseline variability. Lastly, we normalized each participant's with-assistance variability in each six-minute trial to their baseline variability.

We determined how participants adapted the magnitude of each variable as they gained experience walking with exoskeleton assistance. We calculated the magnitude of each variable on the first day and the last day by averaging magnitude across the two six-minute with-assistance trials on each day. We used a two-tailed paired Student's t-test to determine whether participants adapted the magnitude of each variable on the last day compared to the first day. We normalized the magnitude of each variable in each six-minute with-assistance trial to each participant's baseline value. We calculated the baseline value by averaging the

magnitude of each variable across the two six-minute without-assistance trials on the first day.

We estimated the slope of the relationship between a variable and energetic cost. We used each variable's normalized values during all with-assistance trials and their respective metabolic costs. For each six-minute trial, we calculated steady-state metabolic cost by averaging net metabolic power during the final 3 minutes. We normalized metabolic cost during all with-assistance trials to each participant's metabolic cost during without-assistance trials on the first day, which we calculated by averaging metabolic cost over the two six-minute trials of walking with ankle exoskeletons applying minimal torques. We used a linear mixed-effects regression model to estimate participants' shared relationship between a given variable and energetic cost, while allowing for individual differences in their energetic cost intercepts. When plotting this linear model with individual participant data, we subtracted each participant's random effects term (offset) from their data to better illustrate the fixed effects term (slope) that was the focus of this analysis. For each variable, we used a Student's t-test to test if the slope of the linear model was different from zero.

We modeled the timing of adaptive changes, as well as the timing of decreases in variability, for each variable as an exponential decrease from an initial value to a final steady-state value. We used nonlinear mixed-effects regression of the form:

$$Y(t) = a \times e^{-t/\tau} + b \qquad \text{(Equation 2)}$$

where t is the experience calculated as the total amount of walking time with exoskeleton assistance and $Y(t)$ is the model output, which is the magnitude or variability of a given variable. We determined the total amount of walking time with exoskeleton assistance as the time spent walking in the with-assistance trials that we analyzed, as well as the time spent walking in the additional trials where participants experienced assistive ankle torques in human-in-the-loop optimization. We estimated the time constant (τ), amplitude (a), and offset (b) model parameters using nonlinear optimization. We used a mixed-effects model to estimate a single time constant (τ) that is shared between participants while allowing for individual participant offsets (b). When plotting this exponential model with individual participant data, we subtracted each participant's random effects term (offset) from their data to better illustrate the fixed effects term (time constant) that was the focus of this analysis.

We compared time constants between variables. First, we used bootstrapping to estimate the dispersion of each time constant.[49,50] We used the model output from Equation 2 to calculate each participant's residuals as the difference between their data points and the model output at these time points. We sampled from each participant's residuals with replacement, and then added each participant's residuals to the model output at these time points to simulate 10 new participants. We fit the exponential model to 10 new participants to simulate a new experiment and estimate a new time constant, and then repeated this process 10000 times for the time constants of adaptation, as well as for the time constants of variability, for each variable. For each time constant, we report the median and interquartile range (IQR), calculated as the difference between 75th and 25th percentiles. We did not observe adaptation in ankle angle range during stance and therefore excluded this variable

from this analysis. Next, we tested for differences in time constants of adaptation between step frequency, total soleus activity, and total medial gastrocnemius activity. We used a Kruskal–Wallis one-way ANOVA to test for differences in time constants of adaptation between variables—the bootstrapped time constants did not follow a normal distribution (Anderson–Darling test; $p = 5.0 \times 10^{-4}$)—and then performed a multiple comparison test (Dunn–Šidák correction) of the time constants. We repeated the same analysis but for time constants of variability. We report p-values for total soleus activity versus total medial gastrocnemius activity, total soleus activity versus step frequency, and total medial gastrocnemius activity versus step frequency.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

REFERENCES

1. Sutton RS, and Barto AG (2018). Reinforcement Learning: an Introduction (MIT Press).

2. Block HJ, and Bastian AJ (2010). Sensory reweighting in targeted reaching: effects of conscious effort, error history, and target salience. J. Neurophysiol 103, 206–217. [PubMed: 19846617]

3. Rogers MW, and Mille M-L (2003). Lateral stability and falls in older people. Exerc. Sport Sci. Rev 31, 182–187. [PubMed: 14571957]

4. Bauby CE, and Kuo AD (2000). Active control of lateral balance in human walking. J. Biomech 33, 1433–1440. [PubMed: 10940402]

5. Minetti AE, Capelli C, Zamparo P, Di Prampero PE, and Saibene F (1995). Effects of stride frequency on mechanical power and energy expenditure of walking. Med. Sci. Sports Exer 27, 1194–1202.

6. Umberger BR, and Martin PE (2007). Mechanical power and efficiency of level walking with different stride rates. J. Exp. Biol 210, 3255–3265. [PubMed: 17766303]

7. Selinger JC, O'Connor SM, Wong JD, and Donelan JM (2015). Humans can continuously optimize energetic cost during walking. Curr. Biol 25, 2452–2456. [PubMed: 26365256]

8. Donelan JM, Kram R, and Kuo AD (2001). Mechanical and metabolic determinants of the preferred step width in human walking. Proc. Biol. Sci 268, 1985–1992. [PubMed: 11571044]

9. Abram SJ, Selinger JC, and Donelan JM (2019). Energy optimization is a major objective in the real-time control of step width in human walking. J. Biomech 91, 85–91. [PubMed: 31151794]

10. Emken JL, Benitez R, Sideris A, Bobrow JE, and Reinkensmeyer DJ (2007). Motor adaptation as a greedy optimization of error and effort. J. Neurophysiol 97, 3997–4006. [PubMed: 17392418]

11. Ganesh G, Haruno M, Kawato M, and Burdet E (2010). Motor memory and local minimization of error and effort, not global optimization, determine motor behavior. J. Neurophysiol 104, 382–390. [PubMed: 20484533]

12. Tumer EC, and Brainard MS (2007). Performance variability enables adaptive plasticity of 'crystallized'adult birdsong. Nature 450, 1240–1244. [PubMed: 18097411]

13. Sober SJ, and Brainard MS (2012). Vocal learning is constrained by the statistics of sensorimotor experience. Proc. Natl. Acad. Sci. USA 109, 21099–21103. [PubMed: 23213223]

14. Kuebrich BD, and Sober SJ (2015). Variations on a theme: songbirds, variability, and sensorimotor error correction. Neuroscience 296, 48–54. [PubMed: 25305664]

15. Tachibana RO, Xu M, Hashimoto R-I, Homae F, and Okanoya K (2020). Spontaneous variability predicts adaptive motor response in vocal pitch control. Preprint at bioRxiv 10.1101/2020.06.06.138263.

16. Sternad D (2018). It's not (only) the mean that matters: variability, noise and exploration in skill learning. Curr. Opin. Behav. Sci 20, 183–195. [PubMed: 30035207]

17. Dhawale AK, Smith MA, and Ölveczky BP (2017). The role of variability in motor learning. Annu. Rev. Neurosci 40, 479–498. [PubMed: 28489490]

18. Wu HG, Miyamoto YR, Gonzalez Castro LN, Ölveczky BP, and Smith MA (2014). Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. Nat. Neurosci 17, 312–321. [PubMed: 24413700]

19. Niv Y, Daniel R, Geana A, Gershman SJ, Leong YC, Radulescu A, and Wilson RC (2015). Reinforcement learning in multidimensional environments relies on attention mechanisms. J. Neurosci 35, 8145–8157. [PubMed: 26019331]

20. Poggensee KL, and Collins SH (2021). How adaptation, training, and customization contribute to benefits from exoskeleton assistance. Sci. Robot 6, eabf1078. [PubMed: 34586837]

21. Gordon KE, Kinnaird CR, and Ferris DP (2013). Locomotor adaptation to a soleus EMG-controlled antagonistic exoskeleton. J. Neurophysiol 109, 1804–1814. [PubMed: 23307949]

22. Gordon KE, and Ferris DP (2007). Learning to walk with a robotic ankle exoskeleton. J. Biomech 40, 2636–2644. [PubMed: 17275829]

23. Kao P-C, Lewis CL, and Ferris DP (2010). Invariant ankle moment patterns when walking with and without a robotic ankle exoskeleton. J. Biomech 43, 203–209. [PubMed: 19878952]

24. Sawicki GS, and Ferris DP (2008). Mechanics and energetics of level walking with powered ankle exoskeletons. J. Exp. Biol 211, 1402–1413. [PubMed: 18424674]

25. Sawicki GS, and Ferris DP (2009). Powered ankle exoskeletons reveal the metabolic cost of plantar flexor mechanical work during walking with longer steps at constant step frequency. J. Exp. Biol 212, 21–31. [PubMed: 19088207]

26. Sánchez N, Simha SN, Donelan JM, and Finley JM (2021). Using asymmetry to your advantage: learning to acquire and accept external assistance during prolonged split-belt walking. J. Neurophysiol 125, 344–357. [PubMed: 33296612]

27. Reisman DS, Block HJ, and Bastian AJ (2005). Interlimb coordination during locomotion: what can be adapted and stored? J. Neurophysiol 94, 2403–2415. [PubMed: 15958603]

28. Roemmich RT, Long AW, and Bastian AJ (2016). Seeing the errors you feel enhances locomotor performance but not learning. Curr. Biol 26, 2707–2716. [PubMed: 27666970]

29. Shadmehr R, Smith MA, and Krakauer JW (2010). Error correction, sensory prediction, and adaptation in motor control. Annu. Rev. Neurosci 33, 89–108. [PubMed: 20367317]

30. Leech KA, Day KA, Roemmich RT, and Bastian AJ (2018). Movement and perception recalibrate differently across multiple days of locomotor learning. J. Neurophysiol 120, 2130–2137. [PubMed: 30183471]

31. Sánchez N, Simha SN, Donelan JM, and Finley JM (2019). Taking advantage of external mechanical work to reduce metabolic cost: the mechanics and energetics of split-belt treadmill walking. J. Physiol 597, 4053–4068. [PubMed: 31192458]

32. Bellman R (1957). Dynamic Programming (Princeton University Press).

33. Ting LH, and Macpherson JM (2005). A limited set of muscle synergies for force control during a postural task. J. Neurophysiol 93, 609–613. [PubMed: 15342720]

34. Ting LH (2007). Dimensional reduction in sensorimotor systems: a framework for understanding muscle coordination of posture. Prog. Brain Res 165, 299–321. [PubMed: 17925254]

35. Safavynia SA, and Ting LH (2013). Sensorimotor feedback based on task-relevant error robustly predicts temporal recruitment and multidirectional tuning of muscle synergies. J. Neurophysiol 109, 31–45. [PubMed: 23100133]

36. Steele KM, Jackson RW, Shuman BR, and Collins SH (2017). Muscle recruitment and coordination with an ankle exoskeleton. J. Biomech 59, 50–58. [PubMed: 28623037]

37. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, and Shenoy KV (2012). Neural population dynamics during reaching. Nature 487, 51–56. [PubMed: 22722855]

38. Cunningham JP, and Yu BM (2014). Dimensionality reduction for large-scale neural recordings. Nat. Neurosci 17, 1500–1509. [PubMed: 25151264]

39. Simha SN, Wong JD, Selinger JC, and Donelan JM (2019). A mechatronic system for studying energy optimization During walking. IEEE Trans. Neural Syst. Rehabil. Eng 27, 1416–1425. [PubMed: 31107655]

40. Seethapathi N, Clark B, and Srinivasan M (2021). Exploration-based learning of a step to step controller predicts locomotor adaptation. Preprint at bioRxiv 10.1101/2021.03.18.435986.

41. Todorov E, and Jordan MI (2002). Optimal feedback control as a theory of motor coordination. Nat. Neurosci 5, 1226–1235. [PubMed: 12404008]

42. Müller H, and Sternad D (2004). Decomposition of variability in the execution of goal-oriented tasks: three components of skill improvement. J. Exp. Psychol. Hum. Percept. Perform 30, 212–233. [PubMed: 14769078]

43. Selinger JC, Wong JD, Simha SN, and Donelan JM (2019). How humans initiate energy optimization and converge on their optimal gaits. J. Exp. Biol 222, jeb198234. [PubMed: 31488623]

44. Witte KA, and Collins SH (2020). Chapter 13. Design of lower-limb exoskeletons and emulator systems. In Wearable Robotics, Rosen J, and Ferguson PW, eds. (Academic Press), pp. 251–274.

45. Zhang J, Cheah CC, and Collins SH (2015). Experimental comparison of torque control methods on an ankle exoskeleton during human walking IEEE International Conference on Robotics and Automation (ICRA), 5584–5589.

46. Brockway JM (1987). Derivation of formulae used to calculate energy expenditure in man. Hum. Nutr. Clin. Nutr 41, 463–471. [PubMed: 3429265]

47. Collins SH, and Kuo AD (2013). Two independent contributions to step variability during over-ground human walking. PLoS One 8, e73597. [PubMed: 24015308]

48. Jackson RW, and Collins SH (2015). An experimental comparison of the relative benefits of work and torque assistance in ankle exoskeletons. J. Appl. Physiol 119 (1985), 541–557. [PubMed: 26159764]

49. Duhamel A, Bourriez JL, Devos P, Krystkowiak P, Destée A, Derambure P, and Defebvre L (2004). Statistical tools for clinical gait analysis. Gait Posture 20, 204–212. [PubMed: 15336292]

50. Pataky TC, Vanrenterghem J, and Robinson MA (2015). Zero- vs. one-dimensional, parametric vs. non-parametric, and confidence interval vs. hypothesis testing procedures in one-dimensional biomechanical trajectory analysis. J. Biomech 48, 1277–1285. [PubMed: 25817475]

Highlights

- In new contexts, the nervous system explores through increases in gait variability

- With experience, the nervous system selectively adapts specific aspects of gait

- Simultaneously, the nervous system reduces exploration along these aspects

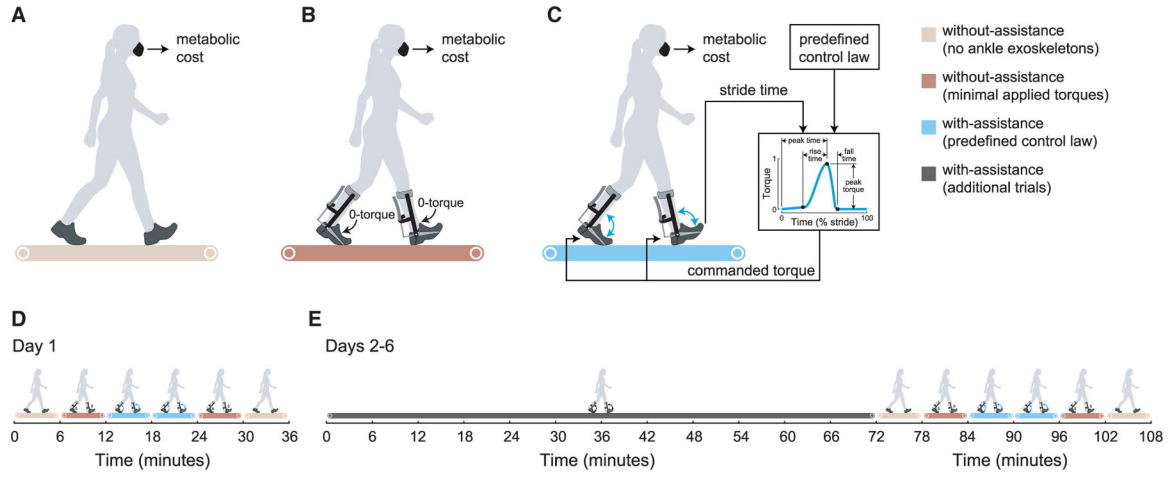- The nervous system's adaptive changes reduce the energetic cost of walking

Figure 1. Experimental design and protocol

All participants experienced two main contexts: without-assistance and with-assistance.

(A and B) In the without-assistance context, participants (A) walked without the ankle exoskeletons or (B) walked while wearing the ankle exoskeletons with minimal applied torques via slack cables.

(C) In the with-assistance context, participants walked while wearing the ankle exoskeletons which used a predefined control law to generate assistive ankle torques at each walking step.

(D) On day 1, all participants completed these three conditions twice.

(E) On days 2–6, all participants completed additional trials that were followed by the original three conditions twice.
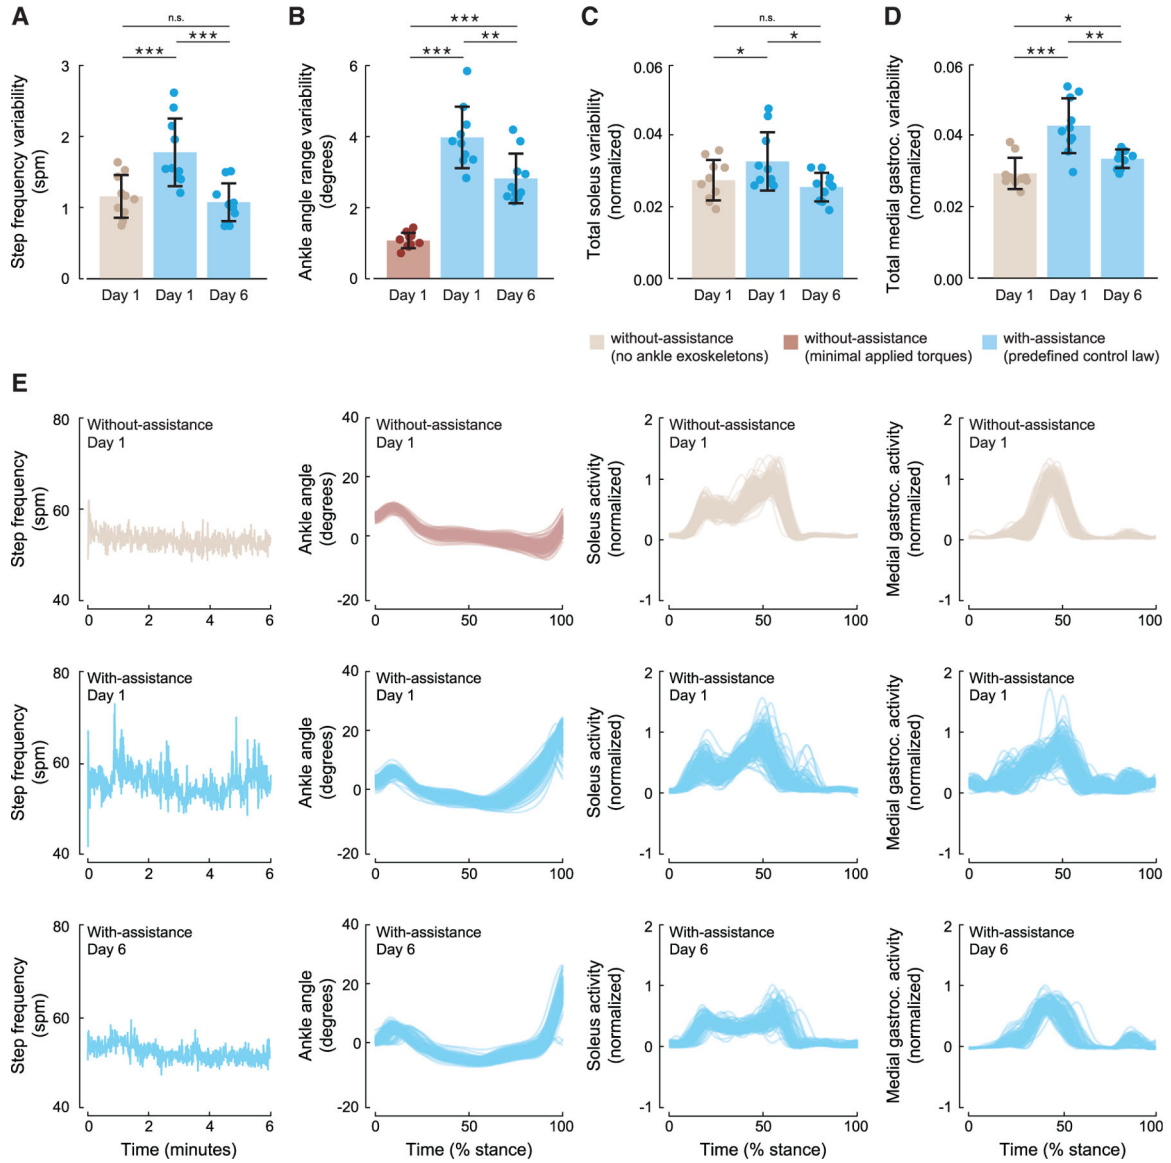
See also Figure S1.

Figure 2. Changes in variability as participants gain experience with exoskeleton assistance

(A–D) (A) Step frequency variability, (B) ankle angle range variability, (C) total soleus variability, and (D) total medial gastrocnemius variability. Variability in muscle activity is expressed as a fraction of the peak activation (see STAR Methods). Walking without the ankle exoskeletons is shown in beige, walking with the ankle exoskeletons with minimal applied torques is shown in red, and walking with the ankle exoskeletons with the predefined control law generating assistive torques at each walking step is shown in blue. Circular markers are participant averages, which we calculated by averaging each variable's variability across two 6-min trials on a given day. Bar height is average across participants, error bars represent one standard deviation, and asterisks indicate statistically significant differences between conditions or days using the notation: *** for p < 0.001, ** for p < 0.01, * for p < 0.05, and n.s. for not significant.

(E) Measured step frequency in strides per minute (spm) as a function of walking time in a given 6-min trial, and measured ankle angle, soleus activity, and medial gastrocnemius activity for the left leg as a function of stance phase at each walking step in a given 6-min trial. Data are from one representative participant.
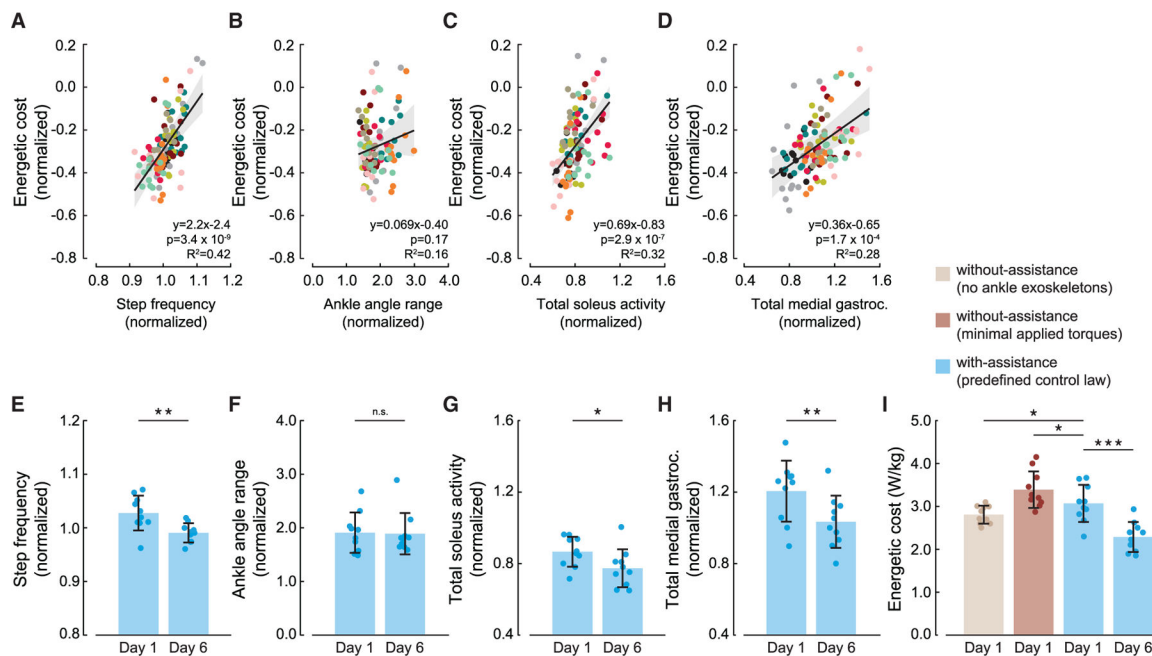
See also Figure S2.

Figure 3. Changes in magnitude of variables that reduce energetic cost

(A–D) Correlation between energetic cost and (A) step frequency, (B) ankle angle range during stance, (C) total soleus activity, and (D) total medial gastrocnemius activity during with-assistance trials across all days. We normalized energetic cost during all with-assistance trials to each participant's energetic cost during without-assistance trials—when walking with ankle exoskeletons with minimal applied torques—on the first day. We normalized each variable's with-assistance magnitude to each participant's without-assistance baseline value. The solid black lines are linear mixed-effects models with 95% confidence intervals (gray shading). Individual participants are represented by distinct colored circles. Each participant experienced the same number of with-assistance trials—2 per day for 6 days. To better illustrate a variable's correlation with energetic cost as determined by the linear mixed-effects model, we subtracted each participant's random effects intercept term from their energetic cost data, which changes their cost data between variables (see STAR Methods).

(E–H) Magnitude of (E) step frequency, (F) ankle angle range, (G) total soleus activity, and (H) total medial gastrocnemius activity during with-assistance trials on the first day and on the last day.

(I) Energetic cost during without-assistance trials on the first day, as well as during with-assistance trials on the first day and the last day. Without-assistance trials include walking without the ankle exoskeletons (beige) and walking with the ankle exoskeletons with minimal applied torques (red). With-assistance trials are walking with the ankle exoskeletons generating assistive torques at each step (blue). Bar height is average across participants, error bars represent one standard deviation, and asterisks indicate statistically significant differences between days using the notation: *** for $p < 0.001$, ** for $p < 0.01$, * for $p < 0.05$, and n.s. for not significant.
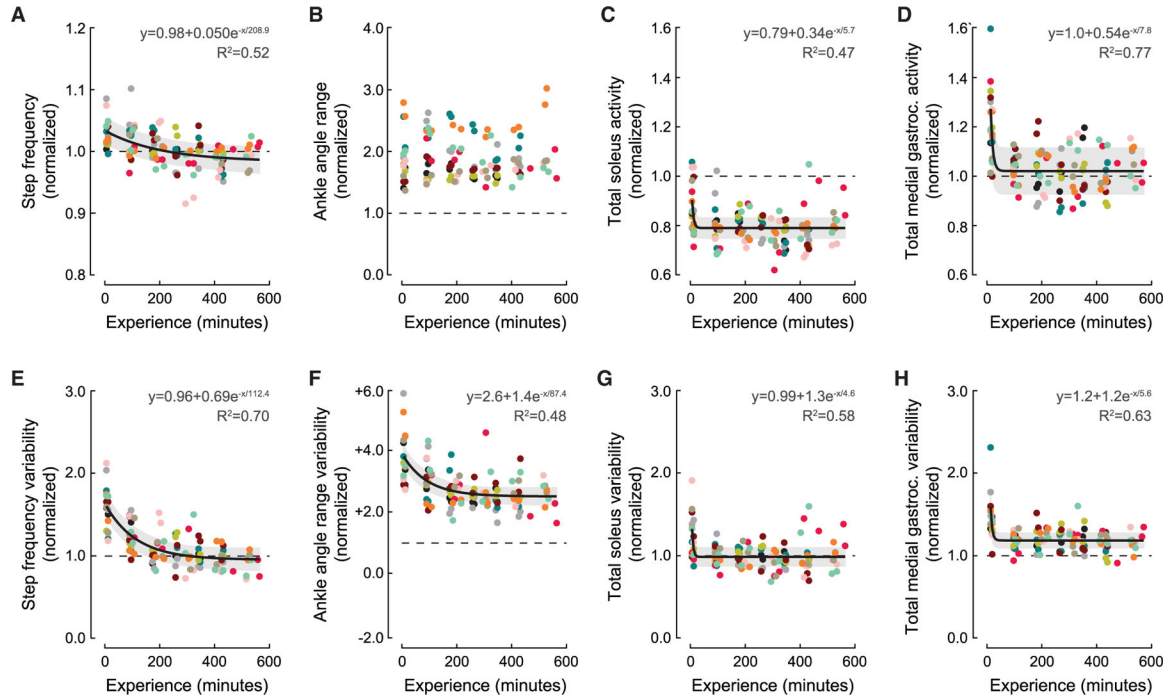
See also Figure S3.

Figure 4. Timescales of changes in variability and changes in magnitude

(A–D) Changes in magnitude of (A) step frequency, (B) ankle angle range during stance, (C) total soleus activity, and (D) total medial gastrocnemius activity as experience increased. (E–H) Changes in variability of (E) step frequency, (F) ankle angle range during stance, (G) total soleus activity, and (H) total medial gastrocnemius activity as experience increased. For each variable, we normalized with-assistance magnitude and variability to each participant's without-assistance baseline levels (dashed horizontal lines). The solid black lines are exponential model fits with 95% confidence intervals (gray shading). We did not fit an exponential model to (B) as we did not observe changes in magnitude of ankle angle range during stance. Individual participants—represented by distinct colored circles —experienced the same number of with-assistance trials—2 per day for 6 days—but at different experience times depending on the design of their additional trials. To better illustrate the exponential fit, we subtracted each participant's random effects offset from the plotted data. The elapsed experience time includes not only the time spent walking in the with-assistance trials that we analyzed here but also the time spent walking in the additional exoskeleton training trials.
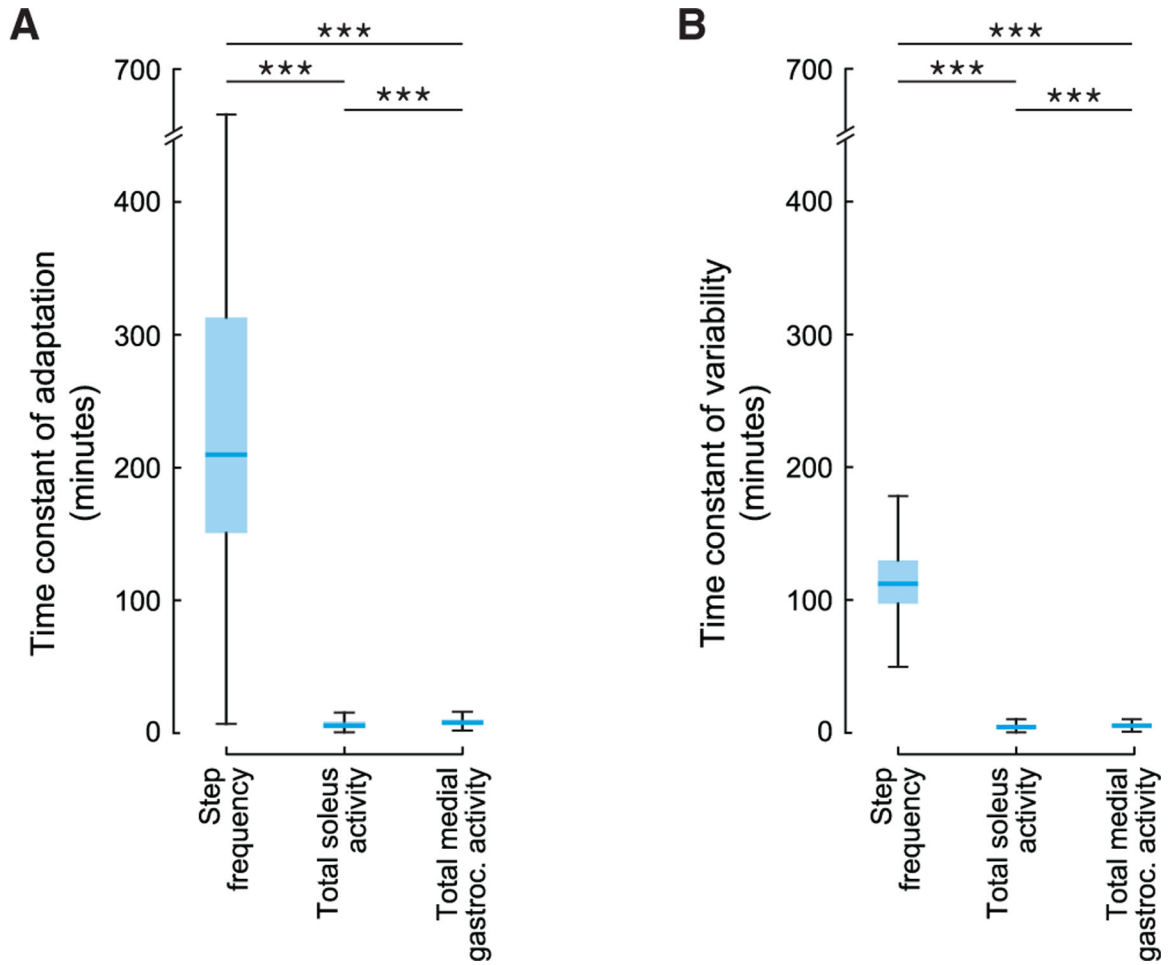
Figure 5. Differences in timescales between variables

(A) Time constants of adaptation and (B) time constants of variability for step frequency, total soleus activity, and total medial gastrocnemius activity during with-assistance trials (blue). We did not observe adaptation in ankle angle range during stance and therefore excluded this variable from this analysis. The central mark indicates the median, the bottom edge of the box indicates the lower quartile (25th percentile), and the top edge of the box indicates the upper quartile (75th percentile). Error bars extend to the most extreme data points not considered outliers, which we define as more than 1.5 times the interquartile range away from the edges of the box. Asterisks indicate statistically significant differences between variables using the notation: *** for $p < 0.001$.
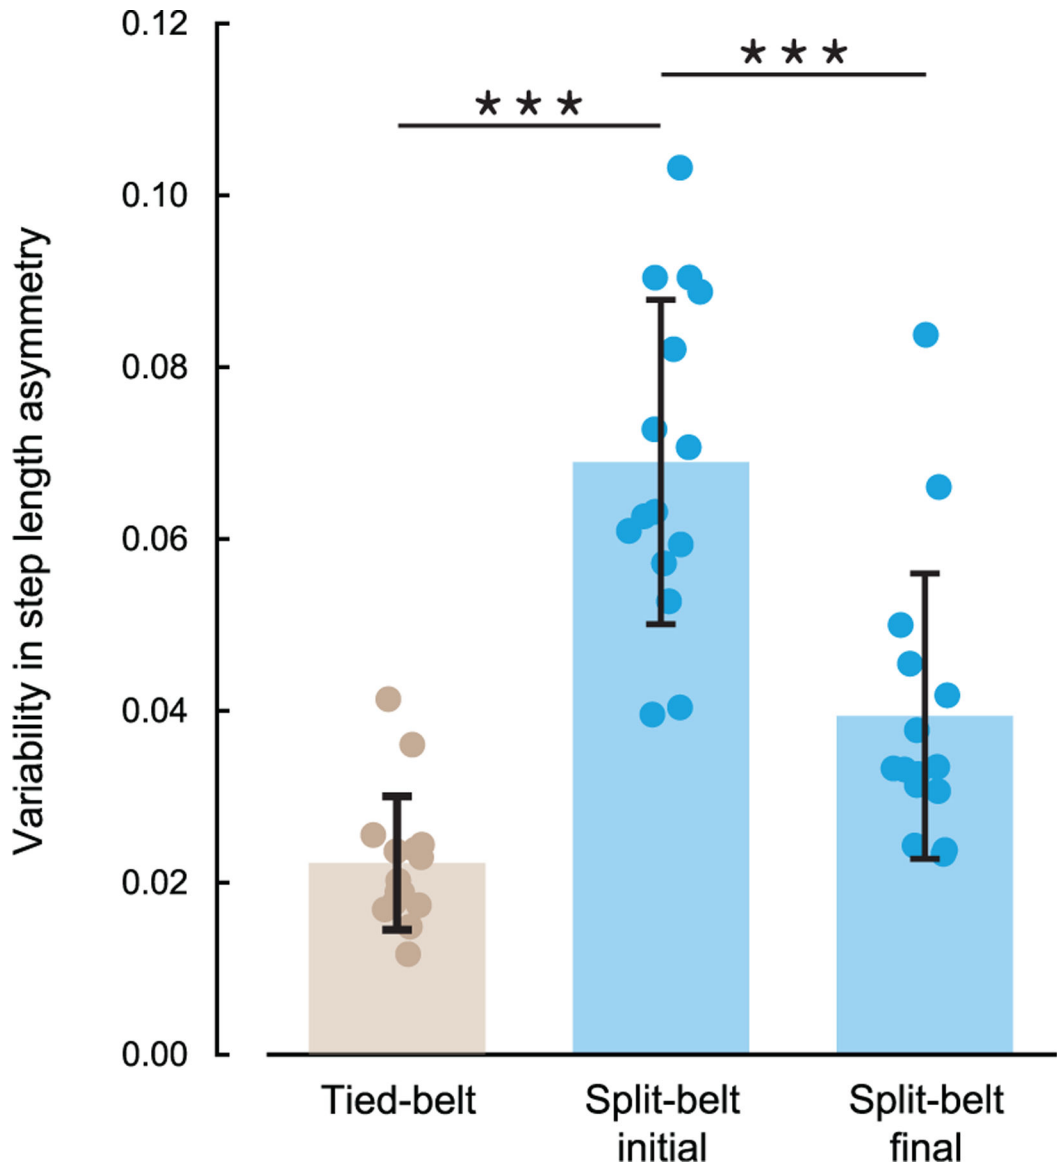
See also Figure S4.

Figure 6. Changes in variability in response to the new context of split-belt walking
Variability in step length asymmetry during the first 100 strides of tied-belt walking (beige), during the first 100 strides of split-belt (blue), and during the last 100 strides of split-belt (blue). Bar height is average across participants (n = 15), error bars represent one standard deviation, and asterisks indicate statistically significant differences using the notation: *** for p < 0.001.

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited data		
Device, metabolic, and EMG data for participant 1	This paper	https://purl.stanford.edu/st957pf8319
Device, metabolic, and EMG data for participant 2	This paper	https://purl.stanford.edu/mw935fz1170
Device, metabolic, and EMG data for participant 3	This paper	https://purl.stanford.edu/yr312kt5378
Device, metabolic, and EMG data for participant 4	This paper	https://purl.stanford.edu/hq152jn6095
Device, metabolic, and EMG data for participant 5	This paper	https://purl.stanford.edu/mh986bj2257
Device, metabolic, and EMG data for participant 6	This paper	https://purl.stanford.edu/jj710vy7867
Device, metabolic, and EMG data for participant 7	This paper	https://purl.stanford.edu/ww452xb7000
Device, metabolic, and EMG data for participant 8	This paper	https://purl.stanford.edu/mm626wf3265
Device, metabolic, and EMG data for participant 9	This paper	https://purl.stanford.edu/zr858qp8088
Device, metabolic, and EMG data for participant 10	This paper	https://purl.stanford.edu/hs191pw6736
Software and algorithms		
MATLAB 2019a	Mathworks	https://www.mathworks.com/