Published in final edited form as:

Sci Robot. 2019 April 24; 4(29): . doi:10.1126/scirobotics.aav6079.

Ergodicity Reveals Assistance and Learning from Physical Human-Robot Interaction

Kathleen Fitzsimons¹, Ana Maria Acosta², Julius P.A. Dewald^{2,3,4}, Todd D. Murphey^{1,2}

¹Mechanical Engineering, Northwestern University, Evanston, IL 60208, USA

²Physical Therapy and Human Movement Science, Northwestern University, Chicago, IL 60611, USA

³Physical Medicine and Rehabilitation, Northwestern University, Chicago, IL 60611, USA

⁴Biomedical Engineering, Northwestern University, Evanston, IL 60208, USA

Abstract

This paper applies information theoretic principles to the investigation of physical human-robot interaction. Drawing from the study of human perception and neural encoding, information theoretic approaches offer a perspective that enables quantitatively interpreting the body as an information channel, and bodily motion as an information-carrying signal. We show that ergodicity, which can be interpreted as the degree to which a trajectory encodes information about a task, correctly predicts changes due to reduction of a person's existing deficit or the addition of algorithmic assistance. The measure also captures changes from training with robotic assistance. Other common measures for assessment failed to capture at least one of these effects. This information-based interpretation of motion can be applied broadly, in the evaluation and design of human-machine interactions, in learning by demonstration paradigms, or in human motion analysis.

Introduction

Hundreds of devices have been designed and built to facilitate forceful interactions between humans and an autonomous system for the purposes of training, safe collaboration, and physical task assistance [1, 2, 3]. The goal of these systems is to augment the capabilities of the human either by providing feedback that enhances the training of a person in a certain task or by eliminating an existing deficit, such as weakness or discoordination due to a neuromotor pathology. As a result, these robotic systems have unique requirements for sensing, actuation, and algorithmic design. In particular, the algorithmic component must be able to infer the quality of the measured behaviors or tasks performed by the joint human-robot system, implying the use of an appropriate metric as the basis for evaluating and

To whom correspondence should be addressed; t-murphey@u.northwestern.edu. .

Author Contributions: K.F. designed and performed experiments and wrote the paper. A.A. provided the data in Fig. 2. J.D. provided resources for the experiment and edited the paper. T.M. contributed to the experimental design, wrote, and edited the paper.

Data and materials availability: All (other) data needed to evaluate the conclusions in the paper are present in the paper or the Supplementary Materials.

modulating the human-robot interaction. However, it is unclear what metrics are appropriate for automating human-machine collaboration. Should one use measures from traditional robotic control such as trajectory error, biologically relevant measures (e.g., energy), or task-specific measures of motion quality? The choice of metric has implications beyond modulation of the interaction, including evaluation of the effectiveness of the physical interaction between the human and robot. One of the purposes of the interaction should be discernible improvements in performance from the assistance and learning from the assistance. In this paper, we show that ergodicity—a measure of the task information encoded by a movement—can predict the presence of robotic assistance and detect the training effect of assistance.

In the human body, sensory information and motor commands are transmitted by nerve fibers conducting action potentials from one synapse to the next. Information theory provides a means to measure the information contained in such signals and to characterize the communication channel [4]. Part of the difficulty of analyzing neural signals are the idiosyncratic sources of variability, but applying information theoretic principles to the nervous system has allowed us to understand and analyze neural coding and organization [5, 6, 7] as well as cognitive perception [8, 9, 10]. Here we provide evidence that the motions resulting from neuromotor signals can be understood as information-carrying signals themselves, and that information measures can also be used to analyze the movement and predict features of neuromotor control.

While signals analysis has been broadly applied to study the transfer of information from stimuli to cortex, principles from information theory are rarely used to model the output of motor commands. Instead theories of motor coordination are often developed on the basis of constrained optimization, often substituting the behavioral goal with the minimization of a measured quantity such as error or energy [11, 12]. These are useful metrics both because they allow us to reason about the underlying principles of neuromotor control and because many well-developed engineering techniques are based on minimizing these quantities [13, 14, 15, 16]. Therefore one can characterize human-like walking as an energy-minimizing trajectory [12] or reaching in the upper limb as a minimum jerk movement [17, 18]. When generic measures fail to capture important features of motion, they are supplemented with qualitative analysis (e.g., similarity to normal patterns of motion) [12] and task-specific measures of success such as work area [19], movement speed [20], or a combination of velocity threshold, aim, and maximum reach [21].

One of the reasons task-specific or outcome-based measures capture the qualitative description of the behavioral goal is that they are independent of the motion strategy, whereas energy or error is typically explicitly dependent on a specific desired trajectory. For instance, one might travel forward and then to the right to reach a target, or one could follow a diagonal path, resulting in the same level of task success using two disparate strategies. Even if the average of the two paths were used as a reference, the resulting desired trajectory may not convey the task goal, and the variance and other statistics of the set of motion strategies are not part of a typical control architecture. Stereotypical motion—such as reaching, self-feeding, and walking—have substantial variation between equally qualitatively successful trials, both within and between individuals. Because of the inherent

stochasticity in neuromotor commands and the resulting task executions, we use a distribution $\phi(x): \mathbb{R}^n \mapsto \mathbb{R}$ over a state space $\mathcal{X}$ to define a task goal. We assess a motion by asking *how much information about $\phi(x) \in \mathbb{R}, x \in \mathcal{X}$ is encoded in the movement $x(t): \mathbb{R} \mapsto \mathcal{X}$ where $x(0) = x_0$ and $x(t) \in \mathcal{X}$ for some time t .*

There are a few natural ways to describe a task statistically rather than by specifying a goal state or a goal trajectory. If there is a particular goal state s , one can represent this as a Dirac delta function $\delta(x - s)$. Or, if the task definition is a consequence of measuring many instances of task execution, the collection of observations will form a distribution $\phi(x)$ in the domain $\mathcal{X}$. As more demonstrations of the target reaching task are added to the set of observations, the collective time spent at the goal state generates a higher peak at the state s , asymptotically approaching a delta function at the goal state. To quantify information content in a motion, one needs to measure a trajectory $x(t) \in \mathcal{X}$ that describes movement of the body. If $x(t)$ were itself a distribution across all of $\mathcal{X}$, one could use the Kullback-Leibler divergence D_{KL} [22] to measure how well $x(t)$ communicates information about $\phi(x)$.

However, $x(t)$ is a trajectory, taking on only one state at each time t , and as a consequence D_{KL} between $x(t)$ and $\phi(x)$ will be generally infinite.

To compare a trajectory $x(t)$ to a distribution $\phi(x)$ while avoiding the underlying problems with using D_{KL} , we use ergodicity—which relates the temporal behavior of a signal to a distribution. A trajectory $x(t)$ is *ergodic* with respect to a distribution $\phi(x)$ if, for every neighborhood $\mathcal{N} \subset \mathcal{X}$, the amount of time $x(t)$ spends in $\mathcal{N}$ is proportional to the measure of $\mathcal{N}$ provided by $\phi(x)$. On a long enough time horizon, measuring a perfectly ergodic $x(t)$ gives a complete description of $\phi(x)$. However, since a trajectory can only visit every neighborhood of $\mathcal{X}$ on an infinite time horizon, a finite time horizon $x(t)$ cannot be perfectly ergodic. Instead, we can ask that $x(t)$ be *maximally ergodic*, by introducing a metric on ergodicity, so that the *time-averaged* statistics of $x(t)$ best capture the statistics of $\phi(x)$ in a specified time horizon T , subject to system dynamics and constraints. Ergodicity can be measured by several metrics [23, 24]; here we use the spectral approach [25], which characterizes ergodicity by comparing spatial Fourier coefficients of $x(t)$ to coefficients of $\phi(x)$ —giving us the distance from ergodicity. In Fig. 1, we show two hypothetical cases of the trajectory of the center of mass during walking compared to an idealized reference trajectory based on typical gait patterns. The high quality execution is not temporally aligned with the reference and may represent faster or slower walking than the reference. Nonetheless, the time-averaged statistics of the trajectory match that of the reference distribution. The low quality execution provides an example of what one might obtain from an impaired individual with poor balance or motor coordination. The ergodic metric used here gives us the distance from ergodicity, such that trajectories which are highly ergodic, like the high quality execution in Fig. 1, have a lower ergodicity than those that are less ergodic.

We use this information measure to analyze two cases of assisted motion—where the lack of assistance may be interpreted as a deficit relative to the assisted condition. First, we looked

at data in supported reaching for a participant with an abnormal tendency to flex the elbow when lifting the arm at the shoulder, where we can see in Fig. 2 that the ergodic metric distinguishes between different levels of arm weight support—although error does not—even for a single participant.

Motivated by this individual result, we collected data from (healthy) human participants robotically assisted with a dynamic inversion and balance task, to see whether the presence of assistance and learning on the part of the participants can be detected. In both cases, the outcome is affirmative. Based on our task-specific measures, there is a clear distinction between the assisted and unassisted conditions. Analysis of the error measure does not show such a distinction based on assistance, but does detect a training effect based on the significantly lower RMS error of the training group in their second session compared to the control. The ergodic measure detects both the presence of assistance and the training effect. These results suggest that measures of the information encoded by a movement can be used to predict the presence of assistance and that such measures capture outcomes that would not otherwise be captured in task-specific performance measures.

Results

In order to evaluate training effect and presence of assistance, participants were tested in both assistance and no assistance modes. Each participant completed 2 sessions, approximately one week apart. Upon enrollment in the study, each participant was placed into 1 of 3 groups. If placed in the training group ($n=20$), the participant completed the first session with assistance and received no assistance in the second session. If a participant was placed in the non-training group ($n=20$), they performed the task without assistance in the first session and used the assistive interface in the second session. Finally, a control group ($n=13$) performed the task without assistance in both the first and second session. Participants were tasked with inverting and balancing a virtual cart-pendulum system as shown in Fig. 3 (often studied in nonlinear control [26, 27]). For the purposes of calculating ergodicity, a delta function $\delta(x - s)$ at the goal state s was used.

In this experiment, we implemented a form of assistance that can convert pure noise input into a successful task execution by comparing the noise input to that of an optimal controller [28]. The assistance acts as a filter similar to that described in [28] and [29], such that if user inputs agree with the optimal controller, user input is not modified by the interface. When user inputs do not agree, the robot physically rejects the input, providing feedback but not guidance. The user input—acceleration of the cart as measured at the robot end-effector—is either accepted or rejected at each instant based on whether or not the input vector is in the same direction as the input prescribed by an optimal controller. Note that the objective of the optimal controller is to minimize the error between the system trajectory and the goal state at the unstable equilibrium, $s = (\theta, \dot{\theta}) = (0, 0)$. The input vector calculated by the optimal controller is never implemented. It is only used as a filtering criterion. The input is completely rejected—replacing the user input with a zero-vector—when user inputs do not agree with the optimal controller. If the participant is a perfect actor, the assistance is completely transparent. Details of the assistance algorithm can be found in the Methods section.

Assistance adds task information

Several task-specific performance measures were recorded, including success rate, balance time, and time to success. The training and non-training group data was aggregated to evaluate the effect of assistance on the 40 participants in a counterbalanced fashion. Paired two-sample t-tests on these task-specific measures of the participants with and without assistance showed that the participants improved with the addition of assistance. The assisted trials had a higher success rate ($p < 0.001$, $t(39) = 12.314$), more cumulative time spent at the goal state ($p < 0.001$, $t(1199) = 26.519$), and reached the goal state more quickly ($p < 0.001$, $t(1199) = 17.202$). The trajectories generated in the assisted condition were also more ergodic with respect to the task distribution than those without assistance when the paired two sample t-test was performed ($p < 0.001$, $t(1199) = 11.261$). However, the root mean square (RMS) error of the trajectory from the goal state did not show a significant difference ($p = 0.094$, $t(1199) = 1.674$) between the assisted and unassisted conditions. This suggests that the assistance improved task-specific performance metrics and increased the task information encoded by the movement. A two-factor (assistance and block) analysis of variance also showed that the assistance had a significant effect in terms of the task-specific measures and the ergodic metric. However, the analysis of the RMS error revealed no significant effects. When we look at the spatial statistics (see Fig. 4) of the assisted trials versus the unassisted trials, we see that the assisted trials spend a larger proportion of time near the origin where the target distribution is centered.

Learning involves increasing task information

The effect of training was assessed by comparing the week 2 session of the trained group to the week 2 session of the control group. Although performance in the task-specific measures did not improve with training, the ergodicity and error of the trained group was significantly better than the control group. A two-sample t-test was performed on the task specific performance measures, finding no difference between trained group and untrained group in terms of their success rate ($p = 0.4280$), $t(31) = -0.8032$, time spent balanced ($p = 0.1687$, $t(988) = 1.378$), and time to success ($p = 0.1935$), $t(988) = 1.301$. The two-sample t-test of the RMS error showed a significant difference between the trained and control group ($p = 0.0499$, $t(988) = -1.963$), but the effect size was small ($d = 0.127$). The t-test of ergodicity ($p = 2.266 \times 10^{-4}$), $t(988) = -3.701$ also detected the difference, but with a larger effect size ($d = 0.237$). A two-factor (training group and block) analysis of variance also showed that block and the interaction between training group and block had a significant effect in terms of the RMS error and the ergodic metric. However, the task-specific measures revealed no significant effects. This indicates that the training effect is not captured by task-specific measures but is captured by error and ergodicity. While the training effect can be detected with error, the information measure detects it with a larger effect size.

Discussion

Our results suggest that the information encoded in bodily motion provides a language for describing changes due to assistance and learning in physical human-robot interaction. Using this framework, we found that an information measure was a better predictor of

changes in deficit compared to error, even when the robotic assistance was based on that error metric. Additionally we showed that when the error-based assistance was used to supplement training, task-specific measures failed to detect the effect of training, and the ergodic measure demonstrated a stronger statistical power to detect the performance changes due to training. When we consider the body as an information channel and bodily motion as an information-carrying signal, we can interpret these results in a way that captures phenomena that would not otherwise be captured by task-specific measures or standard measures such as error. This analysis could provide valuable insight into changes in performance over the course of training or therapy.

When we examined the effect of robotic assistance in reducing the deficit of a stroke participant, we found that simply reducing the deficit by arm weight support increased the information encoded in their reaching motions to multiple targets. Specifically, their motions were more ergodic with respect to a Dirac delta distribution at the target position. We found that task-oriented assistance using an error-based optimal control also reduced the information lost in task executions. These decreases in information loss indicate that the information channel (the human-robot pair) itself was improved by the addition of both forms of assistance.

Task-specific measures also reflected differences in task executions due to deficit. Clinically, task-specific measures are frequently used to assess deficit [30, 31, 32, 33], often motivated by the fact that more standard and generalizable measures, like error and energy, fail to predict the presence of deficit. As a consequence, deficit often must be assessed in very narrow experimental conditions where the measures are applicable. These task-specific measures have several negative consequences. First, they do not translate to other motions (e.g., an assessment strategy for reaching cannot be applied to interpretation of walking or even self-feeding). Moreover, and more importantly, task-specific measures do not admit the same level of principled interpretation as measures such as state error (which captures motor control accuracy) and energy (which captures metabolic efficiency). Measuring the task information encoded by a movement provides an information-theoretic framework for interpreting motion—in a principled manner like error or energy—by capturing the qualitative description of task while not implicitly prescribing a specific strategy for task completion.

Task specific measures are able to capture qualitative task success because they measure events at the goal state or task outcomes—making them independent of the strategy. In contrast, measurement of error is typically explicitly dependent on a reference trajectory, so error measures and error-based assistance prescribe a specific strategy for task completion. Measures on information are still independent of the strategy chosen, but can be expected to detect that a strategy is encoded in the movement even if the movement does not result in task success. A movement with a focused strategy should have higher task information than a movement without a focused strategy. In the case of the swing-up and inversion problem, the assistive algorithm forces participants to employ an error-reducing strategy, increasing the task information in the assisted movement as seen in Fig. 4. Participants attempting to maintain the error reducing strategy after training with the assistance explains the training effect that we see in the error and ergodic metrics (see Fig. 6). However, the error-based

assistance used in this experiment may not be the most effective strategy to improve task outcomes like balance time and success rate. Assistive algorithms that provide too much guidance—thereby reducing errors during training fail to result in improved training because both error [34] and kinematic variability[35] are critical to learning. Additionally, the feedback may lead users to learn the wrong task[36]. In this case, participants may have learned to reduce the cumulative error without learning the true task goal which is to reach the unstable equilibrium regardless of the RMS error incurred before or after reaching that state.

While this work demonstrates that the ergodic measure is useful for assessing human motion post-hoc, it could also be used as a powerful tool for modulating haptic and kinesthetic feedback during training. There is already some evidence that providing feedback based on relevant measures of movement quality is more effective than trajectory error-based assistance [37]. Using the ergodic metric, one can generate distribution-based representations of tasks from sets of imperfect demonstrations—providing an alternative means for robots to ‘learn’ from humans by applying information maximizing techniques [38, 39]. Rather than learning a policy or task objective, one could generate a reference distribution from recordings of human task executions and use the distribution as the task objective for a deterministic information-based model-predictive controller, allowing the human-robot pair to accomplish the task from different initial conditions and under various system constraints. Furthermore, providing assistance based on the ergodic measure allows us to build joint-human robot control policies that directly encode the natural variability of human motion, such that we do not need to restrict assistance to enforce a particular goal trajectory or make inferences about which movement parameters are relevant to the task.

These results suggest that the ergodic measure can augment error or energy measures in the study of biomechanical motion—providing fine-grained insight on the progression of robot-aided training and therapy. Specifically this study supports the idea that ergodicity is a principle of motion for interpreting and predicting animal motion with potential implications for the design of effective feedback and training strategies.

Materials and Methods

Quantifying Ergodicity.

One metric for quantifying ergodicity of the trajectory is the sum ε of the weighted square distance between Fourier coefficients of the distribution ϕ_k and the coefficients of the spatial Fourier transform of the trajectory c_k ,

$$\varepsilon = \sum_{k_1=0}^K \dots \sum_{k_n=0}^K \left| \Lambda_k c_k - \phi_k \right|^2, \quad (1)$$

where at each time the state is n -dimensional and there are $K + 1$ coefficients along each dimension [25]. The subscript k in eq. 1 is a multi-index over the coefficients of the multi-

dimensional Fourier transform. The coefficient $\Lambda_k = (1 + \|k\|^2)^{-s}$ where $s = \frac{n+1}{2}$ places larger weights on lower frequency information, creating a Sobolev norm [25]. Using Fourier basis functions of the form,

$$F_k(x) = \frac{1}{h_k} \prod_{i=1}^n \cos\left(\frac{k_i \pi}{L_i} x_i\right), \quad (2)$$

where h_k is a normalizing factor [25] and L_i is a measure of the length of the dimension, we can compute the coefficients of a spatial distribution or time-averaged trajectory using Eq. 3 and Eq. 4 respectively¹

$$\phi_k = \int_X \phi(x) F_k(x) dx \quad (3)$$

$$c_k = \frac{1}{T} \int_0^T F_k(x(t)) dt \quad (4)$$

The $x(t)$ defined here represents the set of ergodic states $x_e(t)$, or states which are relevant to the task, which may be a subset of the full set of dynamic states $x(t) \in \mathbb{R}^n$ (i.e., $n_e \leq n$). For example, for the cart-pendulum inversion task, the full dynamic state vector is $x(t) = [\theta(t), \dot{\theta}(t), x_{cart}(t), \dot{x}_{cart}(t)]$, but the relevant ergodic states for the inversion task are $x_e(t) = [\theta(t), \dot{\theta}(t)]$. For the comparisons made in this paper, RMS error was also calculated based on the ergodic states only.

New Arm Coordination Training Device.

The NACT3D device (Fig. 7) is a powerful haptic admittance-controlled robot that can be used to render virtual objects, forces, or perturbations in three degrees of freedom. This device is similar to that described in [40] and [41], and is used to quantify upper limb motor impairments and provide a means to modulate limb weight support during reaching. While in use, the participant is seated in a Biodex chair connected to the base of the NACT3D with their arm secured in a forearm-wrist-hand orthosis. The NACT3D is capable of exerting forces at this interaction point between the user and the robot in the x, y, and z directions only. The impedance control is updated at 1000Hz. The NACT3D can move its end effector within a workspace defined both by its design limits (a radius of approximately 0.6m around the participant's shoulder in the half plane in front of the participant's chest) and limits set by the investigators. The orthosis can rotate passively but no torque can be exerted by the

¹The choice of T

robot. At the point where the orthosis is mounted, a force-torque sensor measures the user input which is fed back to the admittance controller. The peak push-pull force that can be exerted by the device of the device at the end effector is approximately $4.7kN$. The force measured at the end effector is sent to a host computer for use in the assistance algorithm to compare the user input to the control policy and perform the filter update at a rate of $60Hz$.

Assistance Algorithm.

The assistance algorithm used in these experiments was Maxwell's Demon Algorithm (MDA). The MDA algorithm was proposed in [28] for noise-driven non-linear control based on the hypothesis that noisy inputs can be a rich source of control authority if filtered in a task-specific way. The MDA filter was implemented by combining a controller and a filter into a single computational unit that cancels noise samples not driving the system to the desired control direction. A modified version of the MDA algorithm that allows filtering of *user input* was implemented in [29]. Using this modified version, an interface for the NACT-3D was developed and implemented using Sequential Action Control (SAC) [42] as the nominal controller.

The filter, described in algorithm, works by evaluating the user input vector u_{user} and computing the value of a nominal controller u_c based on the current state of the system. Calculating the inner product between the user input and the nominal controller establishes whether or not the two vectors are in the same half plane (e.g. $\langle u_c, u_{user} \rangle > 0$). One can further specify that the user input vector must lie within a cone near the nominal control vector by specifying a maximum angle γ between u_{user} and u_c . If the user input lies in the same half space as u_c and within γ radians of u_c , then the filter does nothing. If the user input is not in the same half plane as u_c or not within γ radians of u_c , the input is rejected. In the case of the NACT3D, a force equal and opposite to the force of the user is exerted at the end-effector. This results in the interface being transparent when user inputs are accepted or velocity being held constant when inputs are rejected.

MDA approach for filtering user input—Initialize current time t_0 , sampling time t_s , final time t_f , input saturation u_{sat} and angle tolerance γ . $t_0 < t_f$ Get user input u_{user} Compute nominal controller value u_c Calculate inner product $\langle u_c, u_{user} \rangle$ Calculate angle ϕ between u_c and u_{user} $\langle u_c, u_{user} \rangle > 0$ and $|\phi| \leq \gamma$ $|u_{user}| < u_{sat}$ Use u_{user} as current input, $u_{curr} = u_{user}$ Apply saturated user input $u_{curr} = u_{sat}$ Completely “reject” u_{user} ($u_{curr} = 0$) Apply u_{curr} for $t \in [t_0, t_0 + t_s]$ $t_0 = t_0 + t_s$

Experimental Protocol and Analysis.

Fifty-three participants (17 males, 36 females) consented to participate in this study². At the beginning of each session, the system and task was demonstrated to the participant using a

²This study protocol was approved by the Northwestern University Institutional Review Board and all the participants signed an informed consent form.

video of a sample task completion. Participants were instructed to attempt to swing up the pendulum to the upward unstable equilibrium and balance there for as long as possible. Participants were instructed to continue to try to do this until the 30 seconds were over even if they succeeded at balancing near the equilibrium more than once. The full 30 seconds was used in calculation of all metrics, and the trajectory statistics were averaged over this time horizon. In each session, 30 trials were completed with short breaks upon request of the participant. In order to assess the effect of assistance, 40 participants were tested in two sessions (one week apart), one session with assistance and one session without assistance. The order of the sessions was randomized to account for any learning effects. Paired two-sample t-test were performed to evaluate differences between the session with assistance and sessions without assistance. Another group of 13 participants performed two unassisted sessions one week apart to establish a baseline for learning from unassisted practice. A two-sample t-test was evaluated to test the difference in means between the second session of the control group and the second session of the group of 20 participants who used assistance in their first session.

Statistical Results.—The results of the comparison of the unassisted and assisted trials using two sample t-tests are summarized in table 1. These tests pair samples from each participant according to the order in which they were performed in each session—accounting for the variance between participants. Additional analysis of the effect of assistance as participants progress through trials was performed by grouping individual trials into blocks of five trials, such that the effect of both the assistance and the trials could be assessed from independent and interaction effects.

The time spent at the balance position during each trial was analyzed with a 2 (assistance/no assistance) $\times$ 12 (blocks of 5 trials) mixed design ANOVA, which showed a significant main effect of the assistance mode

$F(1, 418) = 388.87$, $MSE = 1296.5$, $p < 2 \times 10^{-16}$, $Cohen's f = 0.76$. The main effect of the block was also significant $F(11, 418) = 2.196$, $MSE = 7.3$, $p = 0.0139$, $Cohen's f = 0.19$. The assistance and block interaction effect was not significant $F(10, 418) = 1.266$, $MSE = 4.2$, $p = 0.25$, $Cohen's f = 0.14$.

The time to success was also computed for each trial. This measure was analyzed with the same 2 $\times$ 12 mixed design ANOVA, and showed only a significant main effect of assistance mode $F(1, 418) = 224.922$, $MSE = 16629$, $p < 2 \times 10^{-16}$, $Cohen's f = 0.44$. The main effect of block was not significant $F(11, 418) = 0.809$, $MSE = 60$, $p = 0.63$, $Cohen's f = 0.09$. The assistance and block interaction effect also was not significant $F(10, 418) = 0.709$, $MSE = 52$, $p = 0.72$, $Cohen's f = 0.08$.

The mixed design ANOVA was also used to analyze the RMS error of the relevant states $(\theta, \dot{\theta})$ over each 30 second trial, and found no significant effects from any factor. The main effect of assistance was not significant $F(1, 418) = 1.367$, $MSE = 0.018$, $p = 0.24$, $Cohen's f = 0.05$. The main effect of block was also not significant $F(11, 418) = 1.399$, $MSE = 0.019$, $p = 0.17$, $Cohen's f = 0.18$. The interaction

of block and assistance was not significant either
 $F(10, 418) = 0.609$, $MSE = 0.008$, $p = 0.806$, $Cohen's f = 0.11$.

The ergodic metric was computed over each 30 second trials using the relevant states $(\theta, \dot{\theta})$ and was analyzed using the 2 x 12 mixed design ANOVA. The only significant main effect was assistance mode $F(1, 418) = 62.51$, $MSE = 6.90$, $p = 2.38 \times 10^{-14}$, $Cohen's f = 0.35$.

Block was not a significant main effect

$F(11, 418) = 1.31$, $MSE = 0.144$, $p = 0.218$, $Cohen's f = 0.17$, and the interaction of assistance and block was not significant $F(10, 418) = 0.691$, $MSE = 0.076$, $p = 0.73$, $Cohen's f = 0.12$.

These ANOVAs show that the task-specific measures and the ergodic metric detected the effect of assistance with a moderate effect size, while error did not distinguish between the assisted and unassisted conditions over the course of each session.

The comparison of the trained and control group using two sample t-tests is summarized in table 2. Unlike the assisted and unassisted trials, the groups in these tests are independent and therefore, the samples are not paired. The progress of the two groups over the second session (Fig. 6) was analyzed further by performing mixed design ANOVAs on the training group (between participants) and block (within participants).

The balance time of the control group and the trained group in the second session was analyzed with a 2 (training groups) x 6 (blocks) mixed design ANOVA, which showed no significant main effects or interactions effects. The main effect of training group was not significant $F(1, 31) = 1.202$, $MSE = 1.25$, $p = 0.28$, $Cohen's f = 0.08$. The main effect of block also was not significant $F(5, 155) = 2.018$, $MSE = 0.44$, $p = 0.079$, $Cohen's f = 0.11$, nor was the interaction of training and block significant

$F(5, 155) = 1.05$, $MSE = 0.23$, $p = 0.39$, $Cohen's f = 0.08$.

The mixed design 2 x 6 ANOVA design was also applied to the time to success, and the main effect of training group was not significant

$F(1, 31) = 0.334$, $MSE = 103.4$, $p = 0.567$, $Cohen's f = 0.05$. The main effect of block was not significant either $F(5, 155) = 1.34$, $MSE = 66.32$, $p = 0.25$, $Cohen's f = 0.09$. The interaction effect of block and training group also was not significant

$F(5, 155) = 1.34$, $MSE = 66.50$, $p = 0.25$, $Cohen's f = 0.09$.

The same mixed design ANOVA was used to analyze the RMS error in each trial. The main effect of block was significant $F(5, 155) = 4.336$, $MSE = 0.011$, $p = 0.001$, $Cohen's f = 0.19$, but the main effect of training was not significant

$F(1, 31) = 0.76$, $MSE = 0.035$, $p = 0.39$, $Cohen's f = 0.15$. The interaction effect of training group and block also was not significant

$F(5, 155) = 1.61$, $MSE = 0.004$, $p = 0.16$, $Cohen's f = 0.12$.

The analysis of the ergodic metric using the mixed design ANOVA revealed a significant main effect of block $F(5, 155) = 2.88$, $MSE = 0.08$, $p = 0.0163$, $Cohen's f = 0.15$, and a significant interaction effect of block and training group

$F(5, 155) = 2.33$, $MSE = 0.06$, $p = 0.045$, $Cohen's f = 0.14$. The main effect of training was not significant $F(1, 31) = 1.056$, $MSE = 0.49$, $p = 0.312$, $Cohen's f = 0.17$.

These ANOVAs demonstrate that the task-specific measures were not sensitive to either the improvement made by the participants throughout the second session or the benefit of the feedback provided to the trained group in the first session. The error measure indicates that users performed better over the course of the second session. In Fig. 6, the control group performed worse at the beginning of the second session than it did at the end of the first session, and their performance increased in terms of error over the course of the session. The trained group also improves moderately during the second session. The ANOVA of the ergodic metric is also able to detect the significant improvement during the second session by the control group as well as the interaction effect of group and training. This interaction is a result of the trained group performing better at the beginning of the second session and maintaining that performance, while the control group eventually reached the same level of performance.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Hendrik A. Dewald for collecting the data in Fig. 2.

Funding: Supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program and by National Science Foundation under grant 1637764. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the DoD, NDSEG program, or of the National Science Foundation.

References

- [1]. Gupta A, O'Malley MK, Patoglu V, Burgar C, *Int. Journal of Robotics Research* 27, 233 (2008).
- [2]. De Santis A, Siciliano B, De Luca A, Bicchi A, *Mechanism and Machine Theory* 43, 253 (2008).
- [3]. Mörtl A, et al., *Int. Journal of Robotics Research* 31, 1656 (2012).
- [4]. Shannon CE, *The Bell System Technical Journal* 27, 379 (1948).
- [5]. Panzeri S, Schultz SR, Treves A, Rolls ET, *Proc. Royal Society of London B: Biological Sciences* 266, 1001 (1999).
- [6]. Averbeck BB, Latham PE, Pouget A, *Nature Reviews Neuroscience* 7, 358 (2006). [PubMed: 16760916]
- [7]. Quiroga RQ, Panzeri S, *Nature Reviews Neuroscience* 10, 173 (2009). [PubMed: 19229240]
- [8]. Zimmermann M, *Human physiology* (Springer, 1989), pp. 166–173.
- [9]. Yanagisawa H, *Int. Design Eng. Tech. Conf. and Comp. and Info. in Eng. Conf.*.
- [10]. Sims CR, *Science* 360, 652 (2018). [PubMed: 29748284]
- [11]. Todorov E, *Nature Neuroscience* 7, 907 (2004). [PubMed: 15332089]
- [12]. Collins S, Ruina A, Tedrake R, Wisse M, *Science* 307, 1082 (2005). [PubMed: 15718465]
- [13]. Fisher ME, Huang FC, Wright ZA, Patton JL, *IEEE Int. Conf. on Engineering in Medicine and Biology* (2014), pp. 6953–6956.
- [14]. Marchal-Crespo L, Reinkensmeyer DJ, *J. NeuroEngineering and Rehabilitation* 6, 20 (2009).
- [15]. Reinkensmeyer DJ, Dietz V, *Neurorehabilitation Technology* (Springer, 2016), pp. xvii–xxii.

- [16]. Patton JL, Stoykov ME, Kovic M, Mussa-Ivaldi FA, Experimental Brain Research 168, 368 (2006). [PubMed: 16249912]
- [17]. Hogan N, J. Neuroscience 4, 2745 (1984).
- [18]. Flash T, Hogan N, J. Neuroscience 5, 1688 (1985).
- [19]. Ellis MD, Sukal-Moulton T, Dewald JPA, Neurorehabilitation and Neural Repair 23, 862 (2009). PMID: . [PubMed: 19454622]
- [20]. Kahn L, Rymer W, Reinkensmeyer D, IEEE Int. Conf. Engineering in Medicine and Biology (2004), vol. 1, pp. 2722–2725.
- [21]. Krebs HI, et al., Autonomous Robots 15, 7 (2003).
- [22]. Kullback S, Leibler RA, The Annals of Mathematical Statistics 22, 79 (1951).
- [23]. Scott SE, Redd TC, Kuznetsov L, Mezi I, Jones CK, Physica D: Nonlinear Phenomena 238, 1668 (2009).
- [24]. Scott SE, Journal on Applied Dynamical Systems 12, 1948 (2013).
- [25]. Mathew G, Mezi I, Physica D: Nonlinear Phenomena 240, 432 (2011).
- [26]. Isidori A, IEEE Trans. on Automatic Control 45, 1817 (2000).
- [27]. Liberzon D, Switching in Systems and Control (Springer, 2012).
- [28]. Tzorakoleftherakis E, Murphey TD, IEEE Conf. on Decision and Control (CDC) (2015), pp. 4368–4374.
- [29]. Fitzsimons K, Tzorakoleftherakis E, Murphey TD, American Control Conference (ACC) (2016), pp. 4397–4402.
- [30]. Huang VS, Krakauer JW, J. NeuroEngineering and Rehabilitation 6, 5 (2009).
- [31]. Hsieh YW, et al., Stroke 40, 1386 (2009). [PubMed: 19228851]
- [32]. Lin KC, Chen YA, Chen CL, Wu CY, Chang YF, Neurorehabilitation and Neural Repair 24, 42 (2010). [PubMed: 19729583]
- [33]. Kitago T, et al., Neurorehabilitation and Neural Repair 27, 99 (2013). [PubMed: 22798152]
- [34]. Thoroughman KA, Shadmehr R, Nature 407, 742 (2000). [PubMed: 11048720]
- [35]. Koenig AC, Riener R, Neurorehabilitation Technology (Springer, 2016), pp. 161–181.
- [36]. Li Y, Patoglu V, O'Malley MK, Transactions on Applied Perception 6, 3 (2009).
- [37]. Jantscher WH, et al., IEEE Haptics Symposium (IEEE, 2018), pp. 241–246.
- [38]. Vergassola M, Villermaux E, Shraiman BI, Nature 445, 406 (2007). [PubMed: 17251974]
- [39]. Miller LM, Silverman Y, MacIver MA, Murphey TD, IEEE Trans. Robot. 32, 36 (2016).
- [40]. Stienen AH, McPherson JG, Schouten AC, Dewald JP, IEEE Int. Conf. on Rehabilitation Robotics (2011), pp. 1–6.
- [41]. Ellis MD, Lan Y, Yao J, Dewald JP, J. NeuroEngineering and Rehabilitation 13, 95 (2016).
- [42]. Ansari AR, Murphey TD, IEEE Trans. Robot. 32, 1196 (2016).
- [43]. R Core Team R: A Language and Environment for Statistical Computing, R Foundation for Statistical Computing, Vienna, Austria (2016).

Fig. 1: Illustration of Motion Signals and Statistics using the Center of Mass in Walking. For the task of walking on a line, we can distinguish between two hypothetical cases—a high quality execution (A) and a low quality execution (B) by tracking the vertical and medio-lateral displacement of the person’s center of mass. These displacements can be characterized as motion signals (C) with a reference or desired trajectory that is based on typical gait patterns. As a trajectory, the high quality execution does not exactly track the reference trajectory in time, but when we look at the Fourier reconstruction of the trajectory statistics (D), we can see that the high quality execution is very similar to the reference distribution. In contrast, the low quality execution has spatial statistics that are very different from the reference distribution.

Fig. 2: Target-Reaching Trials of a Stroke Participant.

A stroke patient was asked to reach to one of three targets (EE=Elbow Extension, SF = Shoulder Flexion, RF=Reach Forward) in different areas of their workspace. The ergonomic Measure (left) provides clear distinction between the level of full arm support and partial or no arm support in the case of both EE and RF (as indicated by the circled data). The error measure (right) provides little distinction between the fully supported case and the partially supported. Each marker represents a trial from the same individual.

Fig. 3: Experimental System.

Participants controlled the cart position x_c directly and indirectly controlled the angle θ and angular velocity $\dot{\theta}$ of the cart-pendulum system (left). The goal state used to calculate the RMS error was $(\theta, \dot{\theta}) = (0, 0)$ and the distribution used as the task definition for the information measure was a Dirac delta function at $(\theta, \dot{\theta}) = (0, 0)$ (right).

Fig. 4: Assistance adds information.

The histogram of unassisted trajectories (left) has its highest density at $\theta = \pm \pi$ which is the farthest point from the goal state. The rest of the distribution is diffuse over the state space. Although the histogram of the assisted trajectories (right) also has a high density at $\theta = \pm \pi$, the distribution is not as diffuse as that of the unassisted trajectories. There are bands of high density spreading outward from the goal state $(\theta, \dot{\theta}) = (0, 0)$. The spatial statistics of the assisted trajectories are more similar to the reference distribution in Fig. 3, because there is a high density at and around the goal state. This suggests that assistance increased the task information encoded in the movement. This outcome is captured by measuring the ergodicity of the trajectories in each group with respect to the reference distribution. The mean ergodicity of the unassisted trajectories is 0.739, and the mean ergodicity with assistance is 0.631. This lower number indicates that less information is lost in the assisted motion than the unassisted motion.

Fig. 5: Learning increases information.

The histogram of week 2 control trajectories (left) has its highest density at $\theta = \pm \pi$ which is the farthest angle from the goal state at $(\theta, \dot{\theta}) = (0, 0)$. The control trajectories also spend time near the goal state, but to a lesser extent. The histogram of trained trajectories (right) also has high density near $\theta = \pm \pi$, but there are large bands of high density in the region $-1.5 \leq \theta \leq 1.5$ and $-4 \leq \dot{\theta} \leq 4$. These bands make the statistics of the trained group closer to the spatial statistics of the reference distribution in Fig. 3. We quantify how well these statistics match that of the reference by calculating the ergodicity. The trained trajectories are on average more ergodic ($\mu = 0.705$) than the controls ($\mu = 0.751$). In other words, the trained motions communicate information about the task goal more effectively than the control motions.

Fig. 6: Comparison of the control and trained group performance progress over training.

The statistical comparisons of the trained and control groups excluded the data from the first session (gray) to avoid including the effects of the assistance algorithm itself. For the task specific-measures (top row), there was no difference between the two groups, and block had not significant effect on performance. For the error and ergodic metric, block has a significant effect, especially in the control group. Under both measures, the control group performance was worse at the beginning of the second session (the first two blocks in white), but by the end of the session performed as well as the trained group. Ergodicity enables one to see the difference between the treated and untreated group, and both error and ergodicity allow one to see learning as a function of block.

Fig. 7:

The New Arm Coordination Training 3D (NACT-3D) device provides haptic feedback in three dimensions to simulate a specified inertial model via admittance control. A force-torque sensor at the end-effector provides input to the admittance control loop. During this experiment, high stiffness virtual springs were used to restrict user motion to a straight line corresponding to the path of the cart in the virtual display (bottom left). The display provided real-time visual state feedback about the cart-pendulum system the user was attempting to invert.

Table 1:
Paired two-sample t-tests comparing unassisted and assisted trials.

Hypothesis testing was performed in *R* [43] by subtracting the unassisted condition from the assisted condition, showing improvement due to assistance in all five measures. Both the task-specific measures and ergodicity—measured as the distance from ergodicity—capture the effect of the assistance. Note that the degree of freedom (df) for success rate is 39 since there is only one rate per participant.

	No Assistance		Assistance				
	<i>n</i> =20		<i>n</i> =20				
Measure	μ	SD	μ	SD	<i>t</i>	df	<i>p</i>
Success Rate	0.3475	0.163	0.792	0.182	12.314	39	5.16210^{-15}
Balance Time	0.191	0.411	1.661	1.913	26.519	1199	2.541×10^{-122}
Time to Success	25.333	7.648	20.068	7.824	-17.202	1199	1.92610^{-59}
Error	0.632	0.062	0.626	0.102	-1.674	1199	9.433×10^{-2}
Ergodicity	0.739	0.191	0.631	0.283	-11.261	1199	4.954×10^{-28}

Table 2:
Two-sample t-tests of week 2 control trials and week 2 trained trials.

Hypothesis testing was performed in *R* [43] by comparing the means of the control group to the means of the trained group. Error and ergodicity—measured as the distance from ergodicity—were the only measures that revealed a significant improvement in the mean between the trained group and the control group. Note that the degree of freedom (df) for success rate is 31 since there is only one rate per participant.

	Control		Post-Training				
	<i>n</i> =13		<i>n</i> 20				
Measure	μ	SD	μ	SD	t	df	<i>p</i>
Success Rate	0.438	0.380	0.380	0.200	−0.8032	31	0.4280
Balance Time	0.295	0.243	0.243	0.502	1.378	988	0.1687
Time to Success	24.319	24.981	24.981	7.794	1.301	988	0.1935
Error	0.629	0.621	0.621	0.058	−1.963	988	0.0499
Ergodicity	0.751	0.705	0.705	0.177	−3.701	988	2.266×10 ^{−4}