Learning vs. minding: How subjective costs can mask motor learning

Chadwick M. Healy^{1,3*}, Max Berniker^{2&}, Alaa A. Ahmed^{1,3&}

¹ Departments of Integrative Physiology and Mechanical Engineering, University of Colorado, Boulder, Colorado, United States of America

² Advanced Product Development, Intuitive Surgical Inc. Sunnyvale, CA

³ Biomedical Engineering Program, University of Colorado, Boulder, Colorado, United States of America

*Corresponding Author:

Email: healycm@colorado.edu (CH)

&These authors contributed equally to this work.

Abstract

When learning new movements older adults tend to make larger kinematic errors than younger adults, interpreted as an age-related decline in motor-learning ability. However, this conclusion assumes that both older and younger adults use the same error-canceling strategy. Alternatively, it could be that older adults' higher errors can be explained by a difference in strategy, rather than a reduction in learning ability. Consider that error-cancelling strategies can incur higher effort costs. Older adults may be choosing to sacrifice error reduction in favor of a lower effort movement. We test this hypothesis using trajectories where subjects reached to targets in a force field. Utilizing the framework of optimal control theory, we infer subjective costs (i.e., strategies) and internal model accuracy (i.e., proportion of the novel dynamics learned) by fitting a model to each population's trajectory data. Our results demonstrate that the trajectories are not uniquely specified by a precise amount of learning, but rather through a combination of the amount learned and strategic differences represented by relative cost weights. Based on the model fits, younger adults learn between 60-85% of the novel dynamics, and older adults learn between 55-80%. Each model fit produces trajectories which match the experimentally observed data, where a lower proportion learned in the model is compensated for by increasing costs on kinematic errors relative to effort. This finding supports our hypothesis that older and younger adults could be learning to the same extent, but older adults have a higher relative cost on effort compared to younger adults. These results call into question the proposition that older adults learn less than younger adults and suggest that the metrics commonly used to probe motor learning paint an incomplete picture. Importantly, to accurately quantify the learning process the subjective costs of movements should be considered.

Author Summary

Here we show that how a person values effort versus error in their movements has an impact on their overall strategy for performing those movements and adapting to a novel environment. When error alone is considered as a measure of learning—a widely held assumption in the field of motor control—it appears that certain populations such as older adults are significantly worse at learning new motor tasks. However, using a novel framework, we are able to parse out differences in how much a population has learned, as well as how they subjectively value factors such as effort and error. In the case of older adults, we show that they could be learning as much as younger adults but exhibit larger errors because they care more about expending extra effort to reduce them.

Introduction

When people are introduced to a novel environment, they initially experience large errors relative to their intended performance. Through repeated practice, they systematically reduce these errors, learning how best to perform the task. Because the amount a person has learned is not directly assessable, we are reliant on indirect measures such as error that reflect the latent state of how much a person has learned. The traditional interpretation is that a person who performs with less error has learned more, while a person who performs with greater error has learned less. However, when comparing across subjects or populations, variation in strategies must be considered. In the field of motor control, greater error may not be a result of learning less, but rather, an optimal trade-off between movement errors and other similarly important factors.

The use of error reduction as an indicator for learning has led to useful insights into how individuals learn novel tasks. Error reduction is ubiquitous across a wide range of motor behaviors, including grip force adaptation [1], arm reaching [2-5], locomotor [6-8] and postural [9,10] adaptation tasks. While errors are characterized by both kinematic and kinetic metrics across these tasks, they are similar in that they reflect a deviation from desired or nominal behavior. The consensus across these paradigms is that we reduce our errors as we adapt our movements.

However, movement decisions and learning are not solely driven by error reduction. People optimize for other factors such as reward, effort, time, and risk. When people move faster, they move with greater error [11] and tend to have trajectories which minimize the variance of the end or target position [12]. Yet, people still choose to move faster in the presence of greater reward in both arm reaches [13] and saccades [14]. Reward is discounted by time, thus people and animals are willing to expend more effort to achieve an earlier reward [15-17]. In addition to error, studies have shown that effort is reduced through the learning process. [4,8], suggesting a simultaneous optimization of both these factors [18]. Studies have also shown learning rate is increased by increasing rewards [5], increasing consequence of an error [19], or by manipulating the level of uncertainty of state [20]. Traditionally interpreted as learning to a lower extent, increased residual error is influenced by the consequences of an error [21] or the level of effort [18]. These studies, in conglomeration, show that there is more at play in motor learning than just error reduction. When evaluating the extent of learning in motor learning studies, residual errors do not necessarily mean a reduced amount of learning has occurred, but rather, subjects could be optimizing for factors in addition to error.

Optimal feedback control theory offers a formalization of these trade-offs and has been used to describe human movements [22-25]. Using this framework, subjective costs are quantified and used to develop a control policy, offering insight beyond kinematics into the subjective strategies and decisions of each individual. Izawa and colleagues showed that in an adaptation task, higher error is acceptable, and indeed optimal, if the total movement cost includes not only error but effort as well [25]. Others have used an optimal control model to formalize trade-offs between reward and effort that could predict not only which movement would be made, but also how that movement would be performed [17]. Optimal control models offer insight that is deeper than a comparison of *how* the kinematics of two movements differ; by investigating the subjective costs used, these models offer an explanation for *why* the movements differ.

Many motor learning studies implicitly assume that populations use identical strategies when learning new motor tasks, thus kinematic differences between populations are evidence for a difference in the ability to learn. Using error as a correlate for the amount learned, prior literature largely concludes that our ability to learn new motor tasks declines with age [26-31]. The underlying causes of a decline in ability to learn new motor tasks appear to be from multiple

factors, including motor variability [32,33], sensory deficits [34,35] and attention [36]. Yet the causes are inconsistent and are highly dependent on the task structure, complexity and familiarity [37]. Additionally, in some cases, older adults learn to the same extent as younger adults [38-41]. In other studies, we see older adults exhibit different movement strategies than younger adults, [41,42]. One study has shown that older adults put higher subjective value on effort than their younger counterparts [43]. The observed kinematic differences between older and younger adults [41], combined with their difference in strategies and subjective cost values, suggest that it is important to isolate and investigate these variables separately.

Here we ask whether observed age-related errors during motor learning in a velocity-dependent force field can be explained by a difference in subjective costs rather than a difference in learning. We focus specifically on a dynamics learning task, wherein subjects must reach in a velocity-dependent force field. In this task, learning involves a trade-off between effort and error reduction. We use previously published data [26] fitting an optimal feedback control model to younger and older adult trajectories. Using a simple model, we first demonstrate how differences in strategies, which we quantify through subjective costs, can give rise to changes in learned behavior usually interpreted as reduced learning. Next, we determine the range of strategies that can explain the observed behavior. Together, the results demonstrate a large overlap across the younger and older adults in the proportion learned if their subjective costs are different. In particular, the two groups appear to learn similar amounts, but older adults place a higher subjective cost on effort required to reduce kinematic errors.

Results

Modeling subjective cost trade-offs

Trade-offs between subjective costs and the amount of learning can produce equivalent control laws, and thus trajectories. This concept is best illustrated with a simple model. Let us consider a one-dimensional linear dynamical system:

$$\dot{x} = ax + bu \quad (1)$$

where x represents the state, u , the control, and a and b parameterize the dynamics. We can model learning as the process of estimating an internal model of the state dynamics where $\hat{a}$ is the estimate of the true value a , scaled by ε , the “proportion learned”.

$$\hat{a} = \varepsilon a \quad (2)$$

Given this internal model, we define the conventional linear quadratic cost, J , which penalizes state deviations from zero (kinematic error) and control (effort) with weights q and r , respectively.

$$J = J_q + J_r = \int_0^\infty qx^2 dt + \int_0^\infty ru^2 dt = \int_0^\infty (qx^2 + ru^2) dt \quad (3)$$

The solution to this cost is the well-known LQR solution, and the control, u , is expressed below:

$$u = -Kx = \frac{-1}{b} \left(\hat{a} + \sqrt{\hat{a} + \frac{b^2 q}{r}} \right) x = \frac{-1}{b} \left(\varepsilon a + \sqrt{\varepsilon a + \frac{b^2 q}{r}} \right) x \quad (4)$$

From this equation we see that similar control laws and resulting trajectories can be created through different combinations of the parameters, q , r , and ε . To illustrate, the controller obtained for one particular value of q , r , and ε can be identical to another controller with a reduced ε by increasing the penalty on control input, r , or decreasing the penalty on state, q , or a combination of the two. Accordingly, two different trajectories do not necessarily indicate a different internal

model of the dynamics; this could be a result of having differing costs. In other words, there is no unique mapping from the control law and trajectory to the proportion of the dynamics learned. These trade-offs are visualized in Fig 1. In the case of this simple model, q and r represent the subjective costs of a person or population and ϵ represents the proportion of the dynamics that person or population has learned. This relationship clearly demonstrates how changes in subjective costs across individuals can mask differences in how much they have learned.

Fig 1. Trade-offs between subjective costs and proportion learned can produce equivalent trajectories. Each line represents a unique solution produced by ratios of kinematic cost weight to effort cost weight of 2, 1 and 0.5 with perfect knowledge of the dynamics. As the extent of learning is reduced, moving left along the x-axis, we see that in order to produce an equivalent trajectory, represented by each line, the ratio of kinematic error costs to effort costs would need to increase. (A) shows the ratio of specific cost weights, q and r ; (B) shows the ratio of the total kinematic cost (J_q) combining q and x to the total effort cost combining r and u (J_r). Both (A) and (B) show that a trajectory with a lower proportion learned, but with a ratio of kinematic costs to effort costs is greater could be identical to a trajectory with a higher proportion learned but with a lower ratio of kinematics costs to effort costs. Conversely, if two groups used the same strategy (ratio of cost weights) and produced different trajectories, then they must have learned a different amount. For example, Point A and B are both equivalent trajectories because they fall along the same line, however, Point A has learned less than Point B. Thus, to produce the same trajectory as Point B, the ratio of kinematic costs to effort costs is higher for Point A. Next, Point B and Point C have learned to the same extent, but because Point B has a higher relative kinematic cost to effort cost than Point C, they produce different trajectories (i.e., they are not on the same line). These two comparisons show that kinematic observations alone do not correlate to a specific proportion learned.

Differences in reaching behavior between older and younger adults

Using the example illustrated above, we extend the model and apply it to an experimental dataset of younger and older adults performing a motor learning task [26]. In that study, younger and older adults made planar arm reaches while holding onto the handle of a robotic manipulandum (shown in Fig 2). After a baseline period of 200 reaching movements between a start position and target circle, the robot applied a velocity-dependent force to the hand, acting perpendicular to its velocity by the following equation:

$$\begin{bmatrix} f_x \\ f_y \end{bmatrix} = b \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix} \quad (5)$$

In Equation 5, b is the curl field gain. As is typical for this paradigm, the force perturbation led to large deviations (i.e., errors) from the baseline movements that were subsequently reduced over many reaches. Abrupt removal of the force field led to large deviations in the opposite direction. To measure the horizontal force exerted, subjects were exposed to a force channel trial every five trials throughout the experiment, where reaches appeared to travel in a straight-line path towards the target. While the trends in error onset and reduction were similar in both younger and older adults, there were distinct differences in how they chose to reach the target.

Fig 2. Experimental setup and conditions. (A) Subjects controlled a cursor on a computer screen by making reaches while grasping a robotic manipulandum, reaching anteriorly from one circle to a single target circle. (B) Subjects experienced 900 trials of varying dynamics. *Baseline* and *washout* phases have no perturbations, while *learning* imparted a curl field force specified in Equation 5 and represented by the gray arrows pushing them perpendicular to the direction of their reach. *Channel trials* were interspersed throughout the experiment, where subjects moved along a straight-line path to the target within a force channel. These trials are used to estimate the extent to which subjects were compensating for the dynamics by measuring the force that subjects pushed into the force channel wall, represented in gray. Details of how the channel trial is implemented are summarized in the Methods.

For our analysis, we quantify performance using three common metrics of learning: maximum perpendicular error, maximum perpendicular force, and a trajectory-derived adaptation

index. Maximum perpendicular error is measured in non-channel trials and is calculated as horizontal deviation from a straight-line trajectory from the start position to the target. Maximum perpendicular force is measured in channel trials and is a coarse reflection of learning as it is a measure of the anticipatory force the subject is generating to counter the force field. The adaptation index normalizes the anticipatory force by the velocity of the movement, purportedly correlating to a subject's estimate of the curl field gain. However, this metric is also a reflection of desired error cancellation and will be influenced by subjective strategies. We focus on performance at four phases of the experiment: the last five trials of the baseline period (late baseline), the first and last five trials in the learning period (early learning and late learning), and the first five trials after the perturbation was removed (early washout).

Younger adults exhibited smaller perpendicular position errors than older adults in both late baseline (mean±s.e., -0.84 ± 0.15 cm versus -0.94 ± 0.10 cm) and late learning (-0.92 ± 0.29 cm versus -1.17 ± 0.26 cm). They also exhibited greater maximum perpendicular force at late learning (12.39 ± 1.63 N versus 7.13 ± 1.48 N), and a higher adaptation index at late learning (0.857 ± 0.053 versus 0.651 ± 0.060). Notably, maximum perpendicular force and the adaptation index are significantly different between older and younger adults ($P = 0.031$ and $P = 0.017$, respectively, using unpaired t-tests), consistent with the conventional interpretation that older adults learn less than younger adults [26]. However, this conclusion does not consider the potential strategic differences between older and younger adults that may also cause these observed differences.

Model fits to reach trajectories

To describe these trajectories, we model the limb as a point mass that moves in a two-dimensional plane. Similar to the simplified model described above, we assume an internal model of the curl force field is parameterized by the gain. The model's foundation is adapted from Izawa et al. 2008, where this internal model of the state dynamics (what we call "proportion learned") was used to calculate the control law [25]. However, our model assumes that the state is deterministic and perfectly observed and includes higher derivative terms analogous to muscle activation filters [24]. The model uses a linear dynamical system, which includes hand position, velocity, force, rate of force, and target position as state variables. The model cost function penalizes hand position error from target, hand force, rate of change of hand force, second derivative of hand force (control input), terminal position error, terminal velocity, terminal force, and terminal rate of change of force. We fit trajectories from the model to the experimentally observed trajectory data, where a single controller was used to describe all four phases of the experiment (late baseline, early learning, late learning, and early washout), and only the model's value for proportion learned was allowed to vary between phases. Critically, we assume that within a population, the strategy remains the same throughout the course of the experiment where each phase uses the same cost parameters.

Using this method, we found model fits that accurately described the experimental data for both younger and older adults. The results from the best fit model for each age group are summarized in Fig 3, which shows the model fits' spatial trajectories (Fig 3A) and each of the learning metrics for each phase (Fig 3B). For both subject groups, the model-generated trajectories in both late learning and early washout had learning metrics fall within the 95% confidence intervals of the experimentally observed trajectories.

Fig 3. Model fits for younger and older adults. (A) Model-produced trajectories for late baseline, early learning, late learning, and early washout are compared to data from older and younger adults. Data from the experiment are shown in black with the 95% confidence interval shaded in gray. Model fits are shown in green for younger adults and blue for older adults. (B) Maximum perpendicular error, adaptation index, and maximum perpendicular force are compared for each phase. Experimental data represented in gray bars with error bars representing 95% confidence intervals, and model fits are represented in green and blue for younger and older, respectively.

Some learning metrics did not fall within the 95% confidence interval for the late baseline and early learning phases. The maximum perpendicular error and maximum perpendicular force for both older and younger adults fell outside this range. These phases, however, are exemplary in capturing the limitations of our model. Some of the natural curvature seen in the trajectories, a symptom of biomechanical constraints and the dynamics of the robotic manipulandum, is difficult to capture with a point mass. These differences are small, as shown by the resulting spatial plots of the trajectories and are acceptable because our focus is on later phases of the experiment (late learning and early washout), where the model captures subject behavior more reliably. Taken together, these results demonstrate that a model that assumes subjective costs do not change over the course of the experiment can reliably capture subject behavior.

Model-derived range of learning

The model-based proportion learned, as previously defined, provides a latent metric for the internal model of the dynamics. When fitting all four phases, we find solutions that qualitatively match the learning process. The late baseline phase found a proportion learned of close to zero for both younger and older adults. Younger adults had proportion learned values 12.1%, 85.9% and 87.0% for early learning, late learning, early washout respectively, while older adults learned -4.03%, 68.1% and 78.0%. This matches expectations, where early learning should be close to zero, and late learning and early washout should be roughly equivalent.

Although a single best solution for modeling the data was found, we sought to determine the sensitivity of these model fits. Specifically, we asked whether different amounts of learning could predict similar learning metrics as previously analyzed (maximum perpendicular error, maximum perpendicular force, and adaptation index). To investigate this, rather than leaving the proportion learned as a free parameter, we held it constant and allowed the subjective cost weights to freely vary. We then fit trajectory data from late learning and early washout for both subject groups using that single, fixed amount of learning and varying subjective cost weights. We repeated this analysis for a range of proportion learned, from 0.4 to 1.1 in 0.05 increments. A model fit was deemed acceptable if the model-generated trajectories for both phases had learning metrics that fell within 95% confidence intervals of the real trajectory data.

As visualized in Fig 4, we found solutions for younger adults with a model-based proportion learned ranging from 60% to 85% of the dynamics, and for older adults, 55% to 80% of the dynamics: an overlap in proportion learned from 60% to 80%. If there was no overlap between the proportion learned of older and younger adults, then we could confidently state that the two populations had learned different amounts. However, because the resulting ranges overlapped, the differences in behavior between younger and older adults could be due to a difference in subjective costs, rather than a difference in ability to learn.

Fig 4. Adaptation metrics for model fits across different amounts of learning. Learning metrics from the data are shown as dark gray bars with black 95% confidence intervals. The light gray region highlights the acceptable range for the model to be statistically similar to the data for that phase and metric. The yellow area highlights the region that has statistically similar model fits across all three learning metrics, for both phases within a subject group (younger: green, older: blue). Each model fit is represented as a point, where acceptable model fits are highlighted in bold and unacceptable model fits are a lighter shade. Fits for all three learning metrics (A) maximum perpendicular error, (B) adaptation index, and (C) maximum perpendicular error, and (D)

each fit's resulting trajectories are shown with the actual trajectories shown in black with 95% confidence ellipses in gray.

The differences in subjective costs

Model fits suggest that older and younger adults may have learned the same amount, but differences in their subjective costs resulted in different reaching behaviors. We speculated whether the interaction of these subjective costs involved a consistent trade-off between effort and error, that could help define a general strategy for each population. To investigate this interaction, we used the same method to produce numerous additional model fits for both younger and older adults across their overlapping range of proportion learned (60% to 80%). We then analyzed the ratio of their kinematic costs (position and velocity terms) relative to their effort costs (force, derivative of force, and control input) to investigate how these and proportion learned interact. If the predictions matched those presented in Fig 1, we would expect that greater relative costs on kinematic error could mask a deficit in learning. Similarly, we would observe that older adults, who exhibit higher error, would have consistently higher costs on effort relative to kinematic errors than younger adults.

First, we see that for equivalent trajectories within each subject group, as the proportion learned increases, the ratio of kinematic costs to effort costs decreases (Fig 5). Both younger and older adults have significantly negative slopes (younger: $P = 2.1 \times 10^{-7}$; older: $P = 8.3 \times 10^{-5}$). This matches the predictions laid out by the simple model described in Fig 1, validating that the more complex model exhibits this same predicted behavior. Additionally, the ratio of kinematic to effort costs, across this range of learning is significantly greater for younger adults than older adults (younger: $[-1.49, -0.06]$, older: $[-3.44, -2.50]$). This further supports our explanation of the observed differences between younger and older adults: older adults could have learned to the same extent as younger adults; however, older adults placed a greater premium on reducing effort compared to reducing kinematic errors than younger adults.

Fig 5. Older adults have a lower cost on kinematics relative to effort than younger adults.

The ratio of kinematic costs to effort costs are shown for younger (green) and older (blue) adults. Each dot represents a model fit that is statistically indistinguishable from the trajectory data for that subject group for the prescribed proportion learned. The cost ratio for the median solution per subject group for a given proportion of learning is visualized by the solid lines. A higher ratio of kinematic costs to effort costs (larger value in the y-axis) can be interpreted as having an increased sensitivity to effort.

Discussion

Our analysis suggests that larger kinematic errors do not necessarily imply less learning, and that we must consider subjective strategies when assessing learning. Using our model, we find that older and younger adults reaching trajectories can be explained with similar amounts of learning, despite their large kinematic differences. These differences in motor behavior may be attributed to older adults caring more about effort relative to kinematic error than younger adults. Our model assumes that objective effort is similar in younger and older adults, which is consistent with measured metabolic power in Huang and Ahmed 2014 [26]; thus, our results suggest that it is the subjective weighting of effort that differs. An alternative explanation is that the objective effort costs are higher in older adults compared to younger adults. Locomotion studies have shown older adults incur greater metabolic cost, an objective measure of effort, compared to younger adults. [44,45] While metabolic cost has been measured in both younger and older adults performing motor learning tasks, we await a direct comparison of the metabolic cost between groups. Our findings, in their current form, cannot distinguish between a higher subjective effort cost versus a higher objective effort cost in older adults.

Additionally, our results question the validity of adaptation index used in Huang and Ahmed 2014 [26] as a measure of learning in curl field experiments. We show that the same trajectory can be produced across different proportions learned, and because the adaptation index is calculated using trajectory data, the adaptation index will not correlate to the model-based, latent state of proportion learned. This suggests that adaptation index may be a poor indicator of learning between groups, as it inherently assumes a single strategy where kinematic error-canceling is more highly weighted than effort.

Of note, our finding may be unique to the force-field adaptation paradigm. Using visuomotor rotations may eliminate any effort-centric strategic differences. However, visuomotor rotation experiments probe different mechanisms in the motor learning domain. Force-field adaptation tasks use proprioceptive feedback to make online and trial-to-trial corrections, while visuomotor rotations use visual information as a feedback signal. Because force-field adaptation tasks probe both effort and error simultaneously, the paradigm may better emulate motor tasks encountered in the real world.

Notably, our model was purposefully simple, so differences between subject groups are more easily interpretable. However, as with all modeling studies, there is a trade-off between biological realism and model complexity. It could be the case that a different, or more realistic model would result in different findings. A higher fidelity model, reducing the influence of underlying assumptions, could be accomplished through a few different means.

The first improvement lies in the dynamical model. For instance, the use of a non-linear model of the arm could offer improvements over a point mass. As already discussed, this could improve the fits in the late baseline phase. This type of model could exhibit some natural deviation from the centerline and reduce costs penalizing non-straight reaches. As a second improvement, we could include off-diagonal terms in our cost functions to capture interaction terms between state variables. This would improve the quality of our fits, but in turn, make the model prone to over-fitting and make the subjective costs less interpretable. Finally, developing a model that incorporates motor noise, sensory uncertainty or delay may offer an alternative explanation to the differences in behavior. [46] This could be a driving force in their movement strategy that could be reflected in learning rate or trajectory differences, and potentially mask differences in learning or subjective costs. Overall, these improvements to the model could help tease out the specific differences between subject groups, and ultimately determine whether populations are learning less or compensating in a different manner.

While it is important to extract certain behaviors from observed trajectories such as arm reaches, using metrics such as maximal values cannot account for the temporal, stochastic, and highly dynamic factors surrounding human motor control. Strong conclusions about how two populations learn differently should be extremely thorough and consider multiple metrics. The powerful framework of optimal control enables us to compare temporal data to temporal data, extract valuable information from these models, and estimate the hidden value of how much a person has learned. Deducing whether a difference in behavior is due to a difference in learning, a difference in subjective costs, or a combination of the two is still unanswered; however, we offer a framework which can probe these differences.

Our results show that subjective movement strategies can mask the latent variable of how much a person has learned. Additionally, we have shown that both older and younger adults adapt their reaches to a curl field, but whether they definitively learn to different extents remains unclear. If younger and older adults learn to the same extent, our model offers a plausible explanation that behavior differences between older and younger adults are caused by older adults caring more

about effort relative to kinematic errors. We show that using learning metrics alone gives insufficient insight into the adaptation process. In future studies investigating how much a person or population has learned, it is imperative to consider their implicit strategies.

Materials and Methods

Experimental Setup

This experiment used data from Huang and Ahmed 2014, which investigated differences in learning between older and younger adults [26]. We will briefly review the experiment here and refer the reader to the original publication for greater detail. Eleven older adults (mean±s.d., age 73.8±5.6 years) and 15 younger adults (23.8±4.7 years) made targeted reaching movements while grasping the handle of a robotic arm. The experimental setup is visualized in Fig 2A. Subjects were seated with their right forearm cradled and the computer screen set at eye level. Reaches were made in the anterior and posterior directions and were restricted to movement times of 300 – 600 ms.

Fig 2B outlines the various dynamics subjects experienced throughout the course of the experiment. Each subject made 900 reaches with the robotic arm, 450 anteriorly, 450 posteriorly. For the middle 500 reaches, subjects were exposed to a velocity-dependent force (curl) field. The forces imparted by the robotic arm per Equation 5, where the forces, f_x and f_y , imparted on the hand were proportional and perpendicular to the velocity of the hand, v_x and v_y , scaled by the curl field gain, b . In this experiment, b is -20 N-s/m. The progression of trials was broken into three blocks: 200 trials with no forces (baseline), 500 trials with the curl field on (learning), and a final 200 trials with no forces (washout).

Throughout the 900 trials, subjects were exposed to one force channel trial every five trials, pseudo-randomly dispersed. The channel trial forced subjects to reach in a straight line, while simultaneously measuring the forces exerted on the robotic arm. The robot arm enforced the channel trial using a horizontal force relative to the horizontal position and velocity, summarized in the equation below:

$$f_x = -2000p_x - 50v_x \quad (6)$$

As subjects adapted their reaches to the curl field, they began to anticipate the perturbing forces. On channel trials, subjects' anticipatory compensation resulted in exerting force against the channel, opposite the direction of the curl field perturbation. The measured force trace, often analyzed in conjunction with the velocity trace, provides insight into how well subjects have learned the novel dynamics. The exact process of estimating this value is detailed in the *Learning Metrics* section below.

Data Preparation

Trajectory Data

To analyze the performance of our model, we considered only the outward reaches from the final five trials in the late baseline phase and late learning phase, and the outward reaches from the first five trials of the early learning and early washout phase. All position, velocity and force data were collected at 200 Hz. In order to be included in the analysis, the subject must have reached the target between 250 ms (50 samples) and 1500 ms (300 samples) after the target appeared. Movement onset was defined as when the anterior velocity was ≥ 0.03 m/s towards the target, and movement termination was defined as when the cursor was within the target area and anterior velocity was ≤ 0.03 m/s. If the trial never reached the target area, and/or did not sufficiently slow

down, it was still included, but subjected to the reach time criteria mentioned above. The trials that met these criteria were averaged across each subject with each resampled to the mean trial length of that subject's trajectories. To obtain group averages for the younger and older subject groups, the mean trial length across subjects was calculated, each subject's mean trajectory was normalized to that trial length, then averaged across subjects. Channel trials were considered separately from the non-channel trials but used the same criteria and processing method. Due to the pseudo-randomly dispersed trials in the original experiment and the criteria set for an acceptable reach, there was typically one or zero channel trials for a subject in the five trials considered for each phase. However, it was essential to only consider the first five reaches, especially within early washout, because adaptation and de-adaptation occur very quickly. We chose these stricter inclusion criteria to more accurately represent model assumptions, thus performance metric values presented here are slightly different than the numbers in the original manuscript; however, they do not alter the conclusions made in the previous manuscript.

Learning Metrics

From the time-normalized average trajectories, we calculated commonly used learning metrics: maximum perpendicular error, maximum perpendicular force, and adaptation index. Maximum perpendicular error is calculated as the largest absolute perpendicular deviation from a theoretical straight line that connects the start position to the target position. Often, the net perpendicular deviation is considered, where the average perpendicular deviation in the training phase, or baseline phase, is subtracted from the perpendicular error in the novel environment. Using the net deviation accounts for any natural curvature in reaches due to the biomechanical constraints of the arm and allows for a better within-subject analyses. In this experiment, however, we were more concerned about the comparison between subject groups, so total perpendicular error was a more appropriate metric.

Because this experiment was a force-adaptation task, it was also useful to consider how the anticipatory force changed over the course of the experiment. Using data collected from channel trials, we calculated the maximum force that each subject pushed against the channel, perpendicular to the direction of movement. A positive value indicates a force in the right-hand direction against the perturbing force, while a negative value indicates a force in the left-hand direction with the perturbing force. Accordingly, this value can be compared to an expected maximum perpendicular force, calculated from the maximum velocity and curl field gain.

A more comprehensive metric for adaptation, assuming the strategy is to reach as straight as possible, is to compare the entire force trace to the entire velocity trace in a channel trial. In a curl field trial, the horizontal force is applied to the hand proportional to the vertical velocity through the scalar value, b , as per Equation 4. If a trajectory accurately anticipates and compensates for these forces, the force trace measured in a channel will be approximately equal to b times the velocity profile. If the dynamics are underestimated, the force profile will be equal to the velocity profile scaled by some value less than b . Thus, using the measured horizontal force trace and dividing it by the vertical velocity trace, the scalar value that best approximates this linear relationship is an estimate for the amount of adaptation that has occurred, which we call the adaptation index.

Arm Reach Model

We employed a discrete-time, two-dimensional, finite-horizon, linear quadratic optimal control model using a symmetrical point mass to describe the arm reaching trajectories. The model included an internal estimate of the state dynamics to calculate the control law (referred to as the

proportion learned, above), where all states were deterministically observable and there was no system variability or uncertainty. The system dynamics are governed by the following equation

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t \quad (7)$$

where $\mathbf{A}$ and $\mathbf{B}$ are the dynamics of the system, and $\mathbf{x}$ is the state vector defined as follows:

$$\mathbf{x} = [p_x \ p_y \ v_x \ v_y \ f_x \ f_y \ \dot{f}_x \ \dot{f}_y \ T_x \ T_y]^T \quad (8)$$

The variables p_x and p_y represent hand position, v_x and v_y , hand velocities, f_x and f_y are hand forces, $\dot{f}_x$ and $\dot{f}_y$ are the rate of change of hand forces, and T_x and T_y are target position. The motor commands, u_x and u_y are the second derivative of force. The matrices $\mathbf{A}$ and $\mathbf{B}$ encapsulate the dynamics of the curl field and channel trials, as outlined in Equations 5 and 6. Additionally, a separate, but similar matrix, $\hat{\mathbf{A}}$, represents the internal model of the dynamics, which includes a subject's estimate of the curl field gain instead of the true value.

The cost function used to calculate the control law is defined as:

$$J = \sum_{t=1}^{T-1} (\mathbf{x}_t^T \mathbf{Q} \mathbf{x}_t + \mathbf{u}_t^T \mathbf{R} \mathbf{u}_t) + \mathbf{x}_T^T \mathbf{\Phi} \mathbf{x}_T \quad (9)$$

where the matrices $\mathbf{Q}$, $\mathbf{R}$, and $\mathbf{\Phi}$ are symmetric matrices, which penalize state tracking, control input and terminal state, respectively. The subjective costs for a movement are contained within these matrices and are used to formulate a movement plan. The control sequence, determined by the control law, is calculated with these cost matrices and the estimated dynamics, $\hat{\mathbf{A}}$ and $\mathbf{B}$. Trajectories were simulated by running the control sequence forward in time through the true dynamics (either the with or without a curl field or within a force channel).

Trajectory Matching

When fitting the model to the experimental results, we approached the trajectory matching problem using optimization techniques. We sought to minimize an objective function that loosely described the data through varying the cost weights in $\mathbf{Q}$, $\mathbf{R}$, and $\mathbf{\Phi}$ used to calculate the optimal control model and trajectory. We used MATLAB's constrained minimization function, *fmincon*, designed for nonlinear optimization problems. The minimizing solution was obtained by comparing results from multiple restarts with randomized initial parameter values. Of the multiple restarts, the solution that resulted in the smallest value of the objective function was chosen.

Objective Function

Our initial analysis considered four phases of the experiment: late baseline, early learning, late learning and early washout, and the fit sensitivity and cost ratio analyses considered only late learning and early washout. Within each phase, both channel trials and non-channel trials were considered in the objective function. Within the objective function, the weighted sum of the z-scores of the data's end point, maximum perpendicular position (error), adaptation index, and maximum perpendicular force are penalized.

In order to find the trajectory that minimizes the objective function, the weights on each cost parameter were varied. Cost parameters that penalized the horizontal hand position, vertical hand position, hand force, rate of change of hand force, second derivative of hand force (control input), terminal position, terminal velocity, terminal force, and terminal rate of change of force were allowed to vary. Each of these parameters were along the diagonal of cost matrices $\mathbf{Q}$, $\mathbf{R}$, and $\mathbf{\Phi}$. Constraining these matrices to be diagonal limited the quality of fits thus our solutions represent a lower-bound on the quality of fits achievable. Additionally, the value of the internal model's proportion learned for each phase was allowed to vary in finding the best fit. However, in the second analyses, we find how sensitive the best fit trajectory was, the internal model of the dynamics was fixed at ranges from 40 to 110% of the curl field gain in the late learning phase,

while the other parameters were still allowed to vary. Ultimately, this produced a range of learning values and subjective costs that accurately describe the data.

Model Validation and Analysis

Simulated reaches whose learning metrics fell within the 95% confidence intervals of the experimentally obtained metrics were considered statistically indistinguishable from the data. Only solutions that had all adaptation metrics that fell within this range for both late learning and early washout were considered acceptable solutions. After these solutions were found, they were checked by comparing their spatial plots to that of the data.

First, we compared the resulting ranges of acceptable values of proportion learned. If there was no overlap, then we could conclude the two populations had learned a different amount. If there was overlap, then investigating how the specific cost parameters differ between subject groups could offer an explanation for the differences in reaching behavior.

To more easily interpret the motor strategy differences, individual costs were categorized and combined into two types: kinematic or effort costs. Kinematic costs included costs on position (perpendicular error, distance to target, end position). Effort costs included force and the derivatives of force states ($f_x, f_y, \dot{f}_x, \dot{f}_y, \mathbf{u}$). Costs for each state were normalized by the sum of the squared states for to account for the difference in number of samples between subject groups. Each category of costs was summed together, then to account for a potential uniform scaling of cost weights, normalized kinematic cost is divided by the normalized effort cost to create a metric that encapsulated each group's subjective value of kinematic versus effort costs, referred to as the cost ratio. For each subject group, the log transform of the set of all cost weight ratios for all proportions learned were sampled with replacement 10,000 times to provide the estimated 95% confidence intervals.

Acknowledgements

A special thanks to Helen Huang for sharing her dataset.

References

1. Flanagan JR, Wing AM. The role of internal models in motion planning and control: evidence from grip force adjustments during movements of hand-held loads. *The Journal of Neuroscience*. 1997 Feb 15;17(4):1519-28.
2. Shadmehr R, Mussa-Ivaldi FA. Adaptive representation of dynamics during learning of a motor task. *The Journal of Neuroscience*. 1994 May 1;14(5):3208-24.
3. Krakauer JW, Ghilardi MF, Ghez C. Independent learning of internal models for kinematic and dynamic control of reaching. *Nature neuroscience*. 1999 Nov 1;2(11):1026-31.
4. Huang HJ, Kram R, Ahmed AA. Reduction of metabolic cost during motor learning of arm reaching dynamics. *Journal of Neuroscience*. 2012 Feb 8;32(6):2182-90.
5. Nikooyan AA, Ahmed AA. Reward feedback accelerates motor learning. *Journal of neurophysiology*. 2015 Jan 15;113(2):633-46.
6. Reisman DS, Block HJ, Bastian AJ. Interlimb coordination during locomotion: what can be adapted and stored? *Journal of neurophysiology*. 2005 Oct;94(4):2403-15.
7. Morton SM, Bastian AJ. Cerebellar contributions to locomotor adaptations during splitbelt treadmill walking. *Journal of Neuroscience*. 2006 Sep 6;26(36):9107-16.
8. Finley JM, Bastian AJ, Gottschall JS. Learning to be economical: the energy cost of walking tracks motor adaptation. *The Journal of physiology*. 2013 Feb;591(4):1081-95.
9. Pienciak-Siewert A, Horan DP, Ahmed AA. Trial-to-trial adaptation in control of arm reaching and standing posture. *Journal of neurophysiology*. 2016 Dec 1;116(6):2936-49.
10. Ahmed AA, Wolpert DM. Transfer of dynamic learning across postures. *Journal of neurophysiology*. 2009 Nov;102(5):2816-24.
11. Fitts PM. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of experimental psychology*. 1954 Jun;47(6):381.
12. Harris CM, Wolpert DM. Signal-dependent noise determines motor planning. *Nature*. 1998 Aug;394(6695):780-4.
13. Summerside EM, Shadmehr R, Ahmed AA. Vigor of reaching movements: reward discounts the cost of effort.
14. Xu-Wilson M, Zee DS, Shadmehr R. The intrinsic value of visual information affects saccade velocities. *Experimental Brain Research*. 2009 Jul;196(4):475-81.
15. Shadmehr R, Huang HJ, Ahmed AA. A representation of effort in decision-making and motor control. *Current biology*. 2016 Jul 25;26(14):1929-34.
16. Manohar SG, Chong TT, Apps MA, Batla A, Stamelou M, Jarman PR, Bhatia KP, Husain M. Reward pays the cost of noise reduction in motor and cognitive control. *Current Biology*. 2015 Jun 29;25(13):1707-16.
17. Rigoux L, Guigon E. A model of reward- and effort-based optimal decision making and motor control. *PLoS Comput Biol*. 2012 Oct 4;8(10):e1002716.
18. Emken JL, Benitez R, Sideris A, Bobrow JE, Reinkensmeyer DJ. Motor adaptation as a greedy optimization of error and effort. *Journal of neurophysiology*. 2007 Jun;97(6):3997-4006.
19. Wei K, Körding K. Uncertainty of feedback and state estimation determines the speed of motor adaptation. *Frontiers in computational neuroscience*. 2010 May 11;4:11.
20. Pienciak-Siewert A, Horan DP, Ahmed AA. Role of muscle coactivation in adaptation of standing posture during arm reaching. *Journal of neurophysiology*. 2020 Feb 1;123(2):529-47.

21. Trent III MC, Ahmed AA. Learning from the value of your mistakes: evidence for a risk-sensitive process in movement adaptation. *Frontiers in computational neuroscience*. 2013 Aug 23;7:118.
22. Todorov E. Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural computation*. 2005 May 1;17(5):1084-108.
23. Franklin DW, Wolpert DM. Computational mechanisms of sensorimotor control. *Neuron*. 2011 Nov 3;72(3):425-42.
24. Nagengast AJ, Braun DA, Wolpert DM. Optimal control predicts human performance on objects with internal degrees of freedom. *PLoS computational biology*. 2009 Jun 26;5(6):e1000419.
25. Izawa J, Rane T, Donchin O, Shadmehr R. Motor adaptation as a process of reoptimization. *The Journal of Neuroscience*. 2008 Mar 12;28(11):2883-91.
26. Huang HJ, Ahmed AA. Older adults learn less, but still reduce metabolic cost, during motor adaptation. *Journal of neurophysiology*. 2014 Jan 1;111(1):135-44.
27. Buch ER, Young S, Contreras-Vidal JL. Visuomotor adaptation in normal aging. *Learning & Memory*. 2003 Jan 1;10(1):55-63.
28. Etnier JL, Landers DM. Motor performance and motor learning as a function of age and fitness. *Research Quarterly for Exercise and Sport*. 1998 Jun 1;69(2):136-46.
29. McNay EC, Willingham DB. Deficit in learning of a motor skill requiring strategy, but not of perceptuomotor recalibration, with aging. *Learning & Memory*. 1998 Jan 1;4(5):411-20.
30. Seidler-Dobrin RD, Stelmach GE. Persistence in visual feedback control by the elderly. *Experimental Brain Research*. 1998 Apr 1;119(4):467-74.
31. Vandevoorde K, Orban de Xivry J-J. Internal model recalibration does not deteriorate with age while motor adaptation does. *Neurobiol Aging*. 2019;80:138–153. pmid:31170534
32. Bruijn SM, Van Impe A, Duysens J, Swinnen SP. Split-belt walking: adaptation differences between young and older adults. *Journal of neurophysiology*. 2012 Aug 15;108(4):1149-57.
33. Cooke JD, Brown SH, Cunningham DA. Kinematics of arm movements in elderly humans. *Neurobiology of aging*. 1989 Apr 30;10(2):159-65.
34. Christou EA. Aging and variability of voluntary contractions. *Exercise and sport sciences reviews*. 2011 Apr;39(2):77.
35. Lord SR, Dayhew J. Visual risk factors for falls in older people. *Journal of the American Geriatrics Society*. 2001 May 1;49(5):508-15.
36. Goble DJ, Coxon JP, Wenderoth N, Van Impe A, Swinnen SP. Proprioceptive sensibility in the elderly: degeneration, functional consequences and plastic-adaptive processes. *Neuroscience & Biobehavioral Reviews*. 2009 Mar 31;33(3):271-8.
37. Voelcker-Rehage C. Motor-skill learning in older adults—a review of studies on age-related differences. *European Review of Aging and Physical Activity*. 2008 Jan 24;5(1):5.
38. Wishart LR, Lee TD, Cunningham SJ, Murdoch JE. Age-related differences and the role of augmented visual feedback in learning a bimanual coordination pattern. *Acta psychologica*. 2002 Jun 30;110(2):247-63.

39. Trewartha KM, Garcia A, Wolpert DM, Flanagan JR. Fast but fleeting: adaptive motor learning processes associated with aging and cognitive decline. *Journal of neuroscience*. 2014 Oct 1;34(40):13411-21.
40. Reuter EM, Pearcey GE, Carroll TJ. Greater neural responses to trajectory errors are associated with superior force field adaptation in older adults. *Experimental gerontology*. 2018 Sep 1;110:105-17.
41. Cesqui B, Macrì G, Dario P, Micera S. Characterization of age-related modifications of upper limb motor control strategies in a new dynamic environment. *J Neuroeng Rehabil*. 2008;5(1):31.
42. Morgan M, Phillips JG, Bradshaw JL, Mattingley JB, Iansek R, Bradshaw JA. Age-related motor slowness: simply strategic? *Journal of gerontology*. 1994 May 1;49(3):M133-9.
43. Sidney KH, Shephard RJ. Perception of exertion in the elderly, effects of aging, mode of exercise and physical training. *Perceptual and Motor Skills*. 1977 Jun 1;44(3):999-1010.
44. Larish DD, Martin PE, Mungiole M. Characteristic patterns of gait in the healthy old. *Ann N Y Acad Sci*. 1988;515(1 Central Deter):18-32.
45. Martin PE, Rothstein DE, Larish DD. Effects of age and physical activity status on the speed-aerobic demand relationship of walking. *J Appl Physiol*. 1992;73(1):200-6.
46. Sherback M, Valero-Cuevas FJ, D'Andrea R. Slower visuomotor corrections with unchanged latency are consistent with optimal adaptation to increased endogenous noise in the elderly. *PLoS computational biology*. 2010 Mar 12;6(3):e1000708.