

The Neural Feedback Response to Error As a Teaching Signal for the Motor Learning System

Scott T. Albert and Reza Shadmehr

Laboratory for Computational Motor Control, Department of Biomedical Engineering, Johns Hopkins School of Medicine, Baltimore, Maryland 21205

When we experience an error during a movement, we update our motor commands to partially correct for this error on the next trial. How does experience of error produce the improvement in the subsequent motor commands? During the course of an erroneous reaching movement, proprioceptive and visual sensory pathways not only sense the error, but also engage feedback mechanisms, resulting in corrective motor responses that continue until the hand arrives at its goal. One possibility is that this feedback response is co-opted by the learning system and used as a template to improve performance on the next attempt. Here we used electromyography (EMG) to compare neural correlates of learning and feedback to test the hypothesis that the feedback response to error acts as a template for learning. We designed a task in which mixtures of error-clamp and force-field perturbation trials were used to deconstruct EMG time courses into error-feedback and learning components. We observed that the error-feedback response was composed of excitation of some muscles, and inhibition of others, producing a complex activation/deactivation pattern during the reach. Despite this complexity, across muscles the learning response was consistently a scaled version of the error-feedback response, but shifted 125 ms earlier in time. Across people, individuals who produced a greater feedback response to error, also learned more from error. This suggests that the feedback response to error serves as a teaching signal for the brain. Individuals who learn faster have a better teacher in their feedback control system.

Key words: error-feedback response; motor control; motor learning

Significance Statement

Our sensory organs transduce errors in behavior. To improve performance, we must generate better motor commands. How does the nervous system transform an error in sensory coordinates into better motor commands in muscle coordinates? Here we show that when an error occurs during a movement, the reflexes transform the sensory representation of error into motor commands. To learn from error, the nervous system scales this feedback response and then shifts it earlier in time, adding it to the previously generated motor commands. This addition serves as an update to the motor commands, constituting the learning signal. Therefore, by providing a coordinate transformation, the feedback system generates a template for learning from error.

Introduction

When we hold an object in our hand, the mass of the object alters the dynamics of our arm, changing the relationship between the motor commands sent from the brain to the muscles of the arm, and the resulting motion of the hand (Shadmehr and Mussa-Ivaldi, 1994). If the object is unfamiliar to us, our movement will exhibit errors, producing a sensation in our proprioceptive and visual organs. That is, the brain expe-

riences errors in sensory coordinates. To improve performance, the brain must transform the sensory representation of error into better motor commands in muscle coordinates. How does the transformation from sensory coordinates of error to muscle coordinates of motor commands take place? That is, what signal serves as the teacher for the motor system?

Sensing error engages the proprioceptive and visual organs, but following a delay it also engages sensorimotor feedback pathways, producing reflexive and voluntary corrections that start as early as 50 ms into the reach, continuing until the hand arrives at its goal. These corrections represent a sensorimotor transformation that takes error in sensory coordinates and produces a feedback response in muscle coordinates. The feedback response is a sequence of motor commands that can, in principle, act as a template, providing the brain with an example of how to partially compensate for the unexpected dynamics (Kawato et al., 1987; Thoroughman and Shadmehr, 1999; Franklin et al., 2003, 2008; Milner and Franklin, 2005). However, testing this hypoth-

Received Jan. 14, 2016; revised March 8, 2016; accepted March 21, 2016.

Author contributions: S.A. and R.S. designed research; S.A. performed research; S.A. analyzed data; S.A. and R.S. wrote the paper.

This work was supported by grants from the National Institutes of Health (NS078311, NS095706) and the Office of Naval Research (N00014-15-1-2312).

The authors declare no competing financial interests.

Correspondence should be addressed to Scott Albert, Johns Hopkins University School of Medicine, 720 Rutland Ave., 416 Traylor Building, Baltimore, MD 21205. E-mail: salbert@jhmi.edu.

DOI:10.1523/JNEUROSCI.0159-16.2016

Copyright © 2016 the authors 0270-6474/16/364832-14\$15.00/0

esis is difficult because on any given movement, the motor commands are a mixture of what the brain correctly predicted about the dynamics of the object, and what the feedback pathways added in response to the unexpected motion of the arm. To determine the relationship between error and the learning that resulted from error, one must dissociate the motor commands that reflect a process of prediction, from the motor commands that reflect a process of within-movement feedback correction.

Here, we approached this problem by using an important tool: error-clamp trials (Scheidt et al., 2000). An error-clamp trial makes it possible to reliably guide the movement precisely along a reproducible trajectory. To measure the feedback response to error, we measured the motor commands sent to various muscles of the arm in an error-clamp trial, and then remeasured the commands when novel dynamics (a force field) introduced errors in the reaching movement. By comparing the time course of signals in the perturbation trial to the preceding error-clamp trial, we obtained a proxy for the neural feedback response to error. Following the perturbation trial, we again introduced an error-clamp trial. The change in the motor commands that occurred from the first error clamp to the second error clamp was a proxy for the learning that has occurred following the experience of error. We found that the learned motor commands were a scaled version of the feedback-generated commands, but shifted earlier in time. This suggested that the sensorimotor transformation that was provided by the feedback system, from sensory coordinates of error to muscle coordinates of action, acted as a teacher for the motor system, instructing it on how to improve its commands on the next movement.

Materials and Methods

We recruited $n = 57$ healthy, right-handed individuals to participate in our study (18–36 years of age, 31 females). The study was approved by the Johns Hopkins University School of Medicine Institutional Review Board and all subjects signed a consent form.

Experiment. Participants performed a center-out reaching task while holding the handle of a planar robotic arm. The forearm of each participant was supported by an arm rest that moved freely with the arm. The arm was obscured from view by a horizontal screen, upon which a projector displayed a cursor, serving as a proxy for hand position.

At the onset of each trial, the robot moved the hand to the start position, denoted by a circle 10 mm in diameter, whose location within the workspace remained fixed for the duration of the experiment. Once the hand entered the boundary of the starting position, a random intertrial-interval (ITI) elapsed, varying within the range of 300–700 ms. If the hand moved from the start position at any point during the ITI, the timer was reset. At the conclusion of the ITI, a target circle appeared 10 cm from the starting position, at an angle of 90° relative to the starting position. The target was also 10 mm in diameter and its appearance was accompanied by a short tone. The subject was instructed to move his or her hand to the target. The desired reach time was 500 ms, with a tolerance of ± 50 ms. Feedback regarding reach duration was provided after reach completion: the target turned red or blue if the movement duration was too short or too long, respectively. In addition, a tone accompanied the change in target color. For trials in which movement duration fell within the desired time interval, the target “exploded” in red and yellow concentric circles, a tone was played, and a point was added to a numerical score displayed at the bottom of the workspace. Subjects were instructed that the goal of the experiment was to score as many points as possible.

Our overall objective was to ask whether the feedback system that corrected for a perturbation during a movement produced a neural signal that became the teacher for the motor system, instructing it on how to predictively cancel the perturbation on the following trial. To test our hypothesis, we first measured the neuromotor activity in a given muscle [electromyography (EMG)] during an unperturbed movement [termed error-clamp trial 1 (EC1)]. On the next trial, we perturbed the reaching

movement via a force field. The difference in EMG between the perturbed trial and the preceding error-clamp trial was our proxy for the feedback-generated response to the perturbation. On the next trial, there was a 50% chance that the reach was in an error-clamp (EC2), and an equal chance that a consecutive perturbation occurred. If a second perturbation trial occurred, the following trial was always an error-clamp trial. The difference in activity between EC2 and EC1 was our proxy for learning, indicating the change in neuromotor activity due to experience of error in the preceding trial (or a pair of errors in the case of two consecutive perturbation trials).

The perturbations were standard velocity-dependent curl force fields that pushed the hand clockwise (CW) or counter-clockwise (CCW): $f = B\dot{x}$, where \dot{x} is the hand velocity vector, and $B = [0, -15; 15, 0] \text{ N} \cdot \text{s/m}$ or $B = [0, 15; -15, 0] \text{ N} \cdot \text{s/m}$. During an error-clamp trial, the hand path was confined to a straight trajectory between the start position and the target. To generate the error clamp, the robotic arm produced compensatory forces perpendicular to the hand trajectory in accordance with a stiff spring (spring coefficient, 6000 N/m; viscosity, 250 N · s/m).

The experiment (Fig. 1A) began with a block of 120 null field trials (data not shown in Fig. 1A). This was followed by two consecutive blocks (labeled Blocks 1 and 2) of 263 trials each (one block is shown in Fig. 1A). Blocks 1 and 2 began with 23 null field trials. Following this, one or two perturbation trials were sandwiched between pairs of error-clamp trials. Each type of perturbation (CW, CCW) and number of consecutive perturbations (one or two) was assayed 10 times, for a total of 40 triplet/quartet perturbations per block. The orientation and number of consecutive perturbations were pseudorandomly selected and counterbalanced so that subjects experienced an equal number of CW and CCW perturbations. Between each [error-clamp–perturbation–error-clamp] progression, either two or three null field trials were presented. The paradigm ensured that we could assess learning multiple times without accumulation of learning of either type of perturbation.

Data recording and analysis. We recorded the position of the hand, velocity of the hand, force exerted by the hand on the robotic arm, and force applied via the torque motors at 200 Hz. The movement onset for each reach was determined via a velocity threshold of 35 mm/s. Trials in which the movement began < 200 ms after the target cue appeared were removed from the analysis (2.32% of trials). EMG was used to assess activity of four muscles of the upper arm and trunk, including the biceps, lateral head of the triceps, posterior deltoid, and pectoralis. We used EMG electrodes with a pre-amplifier at the recording head (Delsys), and sampled the resulting signal at 1000 Hz.

To determine an optimal position of recording for each muscle, the electrode position was varied until the largest dynamic range between resting state and contraction was detected. This region was marked for each muscle, the overlying skin for each targeted area was cleaned with isopropyl alcohol, and then Skin-Prep was applied to enhance adhesion of the electrode to the skin. Before application, the electrode was also cleaned with isopropyl alcohol, a double-sided adhesive skin interface was placed on the sensing apparatus, and an electrode preparation gel was applied to the electrode-sensing bars.

The EMG signal was bandpass filtered (10–250 Hz) using a fourth-order Butterworth filter and full-wave rectified. The filtered and rectified signal was smoothed by scaling the EMG amplitude at each time point by the root mean square of the signal in a 40 ms window centered at that time point. Following this preprocessing, we performed a within-subject, within-muscle normalization of each EMG trace by dividing the EMG amplitude at each time point by the average maximum EMG amplitude produced during the initial null trials of Blocks 1 and 2 (46 trials in total are included in this average). In other words, following this normalization, the units of EMG activity for each muscle of a given subject were represented with respect to the average maximum value recorded in that muscle during an unperturbed reaching movement of the same subject.

To compute the neural correlates of learning from error, we compared the EMG activity recorded in the error-clamp trial following the perturbation (EC2) to the error-clamp trial preceding the perturbation (EC1), for each triplet (or quartet) progression. This difference (EC2 minus EC1) represents the trial-to-trial change in the EMG following experience of an error. If the intervening trial was a single perturbation, we

termed this change as Learning 1. If the intervening trials were two perturbations, we termed this change as Learning 2.

To compute the neural feedback response to error, we first focused on triplet progressions (a perturbation trial between two error-clamp trials) and compared the activity measured in the perturbation trial (P1) with the activity measured in the preceding error-clamp trial (P1 minus EC1), and termed this difference Feedback 1 response. This difference represents how muscle activity was modified to counteract the perturbation during a perturbed movement, relative to an error-free reach. In quartet progressions (two perturbation trials between two error-clamp trials), we computed feedback responses in both the first and second perturbation trials. Importantly, for the feedback response to the second perturbation, we used EC2 from single-trial perturbations to estimate the feedforward command produced by the brain after single-trial learning, rather than EC1, which does not account for this learning.

Our hypothesis concerned the relationship between the time courses of learning and feedback responses. Temporal shifts relating learning and feedback were computed within subject via cross-correlation. In all cases, 700 ms temporal fragments of the learning traces were cross-correlated with 1100 ms fragments of the feedback response, beginning 200 ms before movement onset. The learning trace was padded with zeros at the end of the selected temporal fragment so that the learning and error fragments were of equal duration. The learning trace duration used for the cross-correlation was shortened relative to the error trace to reduce corruption of the cross-correlation from noise in the learning traces, which normally returned to baseline values at the conclusion of the reaching movement (i.e., 500 ms after movement onset). The optimal shift relating learning and error was found by identifying the time shift associated with the maximum of the cross-correlogram.

We asked at each moment of time into the reaching movement, how much the brain had learned from the feedback response. That is, we wished to answer whether there was greater learning from a specific part of the feedback signal (for example, its early part), or did the brain learn from the entire feedback signal. To answer this question, we first shifted each feedback response (independently for each subject and muscle pair) by the optimal shift determined via cross-correlation. Next we performed two separate analyses, one across subject and one within subject. In the former analysis, we looked at each muscle and field condition separately, and performed across-subject regressions of learning and feedback signals. Learning for a given muscle and field orientation was regressed onto the corresponding feedback response for that muscle and field orientation, independently at every time point. We identified time points for which these fits were statistically significant ($p < 0.05$) and possessed positive slope, signifying that learning and feedback were positively correlated at that point in time, across subjects. To determine the level of correlation between learning and feedback responses within each muscle, we linearly regressed the learning response onto the feedback response over the interval between -100 and 500 ms and consid-

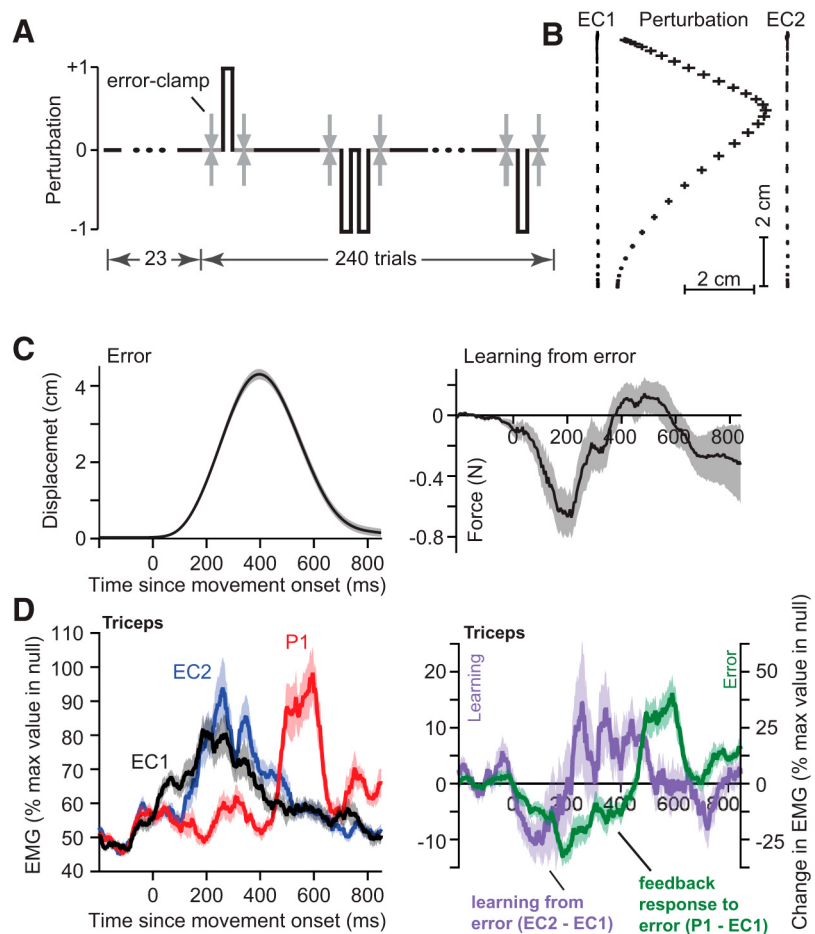


Figure 1. Experimental design and exemplar data from a single subject. **A**, Subjects performed a center-out reaching task to a single target. The experiment consisted of two blocks of 263 trials. Each block began with 23 trials, completed in a null field. During the next 240 trials, subjects encountered random CW and CCW velocity-dependent curl field perturbations. They encountered either a single perturbation or two perturbations in a row. Each perturbation or pair of perturbations was sandwiched by error-clamp trials. **B**, Reach trajectories in the error-clamp trials before and after the perturbation (EC1 and EC2). Error bars are ± 1 SEM in 25 ms intervals. **C**, Kinematic correlates of error and learning from a single CW perturbation. At left, the time course for the perpendicular displacement of the hand during CW perturbations is shown. At right, the learned force production from the single error is shown. The learned force is the net difference in the perpendicular forces produced in EC2 and EC1. Forces are normalized relative to the ideal force that would be produced given the subject's tangential velocity and the field strength. **D**, Triceps EMG activity during error-clamp and perturbation trials. At left, the triceps is active during the movement in EC1. As a result of the feedback response to the perturbation, the triceps activity is suppressed early in the CW field, and enhanced near movement termination (P1, red). Due to the experience of the error, the brain changes the triceps activity for the subsequent error-clamp trial (EC2, blue). At right, the error-feedback response is the trial-to-trial change in the triceps motor command due to the imposition of the CW field (green, P1 minus EC1). The learning response is the trial-to-trial change in the motor command from EC1 to EC2 (purple, EC2 minus EC1). Learning appears to be a time-shifted copy of the feedback response.

ered the R^2 value describing this regression for CCW and CW fields separately.

We next performed within-subject regressions of learning and feedback signals at each time point. For these regressions, we collapsed across muscles and field orientations. As we had recorded four muscles and two field orientations, each regression included a total of eight feedback–learning data points. We considered both the within-subject R^2 metric for this regression (which represents how much of the variation of the learning response is explained by the feedback response for that time point) as well as the slope of the regression (which represents the scaling factor relating feedback and learning). As a control condition, we quantified the baseline correlation between learning and feedback for each subject and each muscle from a dataset in which the feedback response was randomly shifted with respect to the learning response. We drew these random shifts from a uniform distribution between 0 and 400 ms, and shifted each of the eight feedback responses independently. The

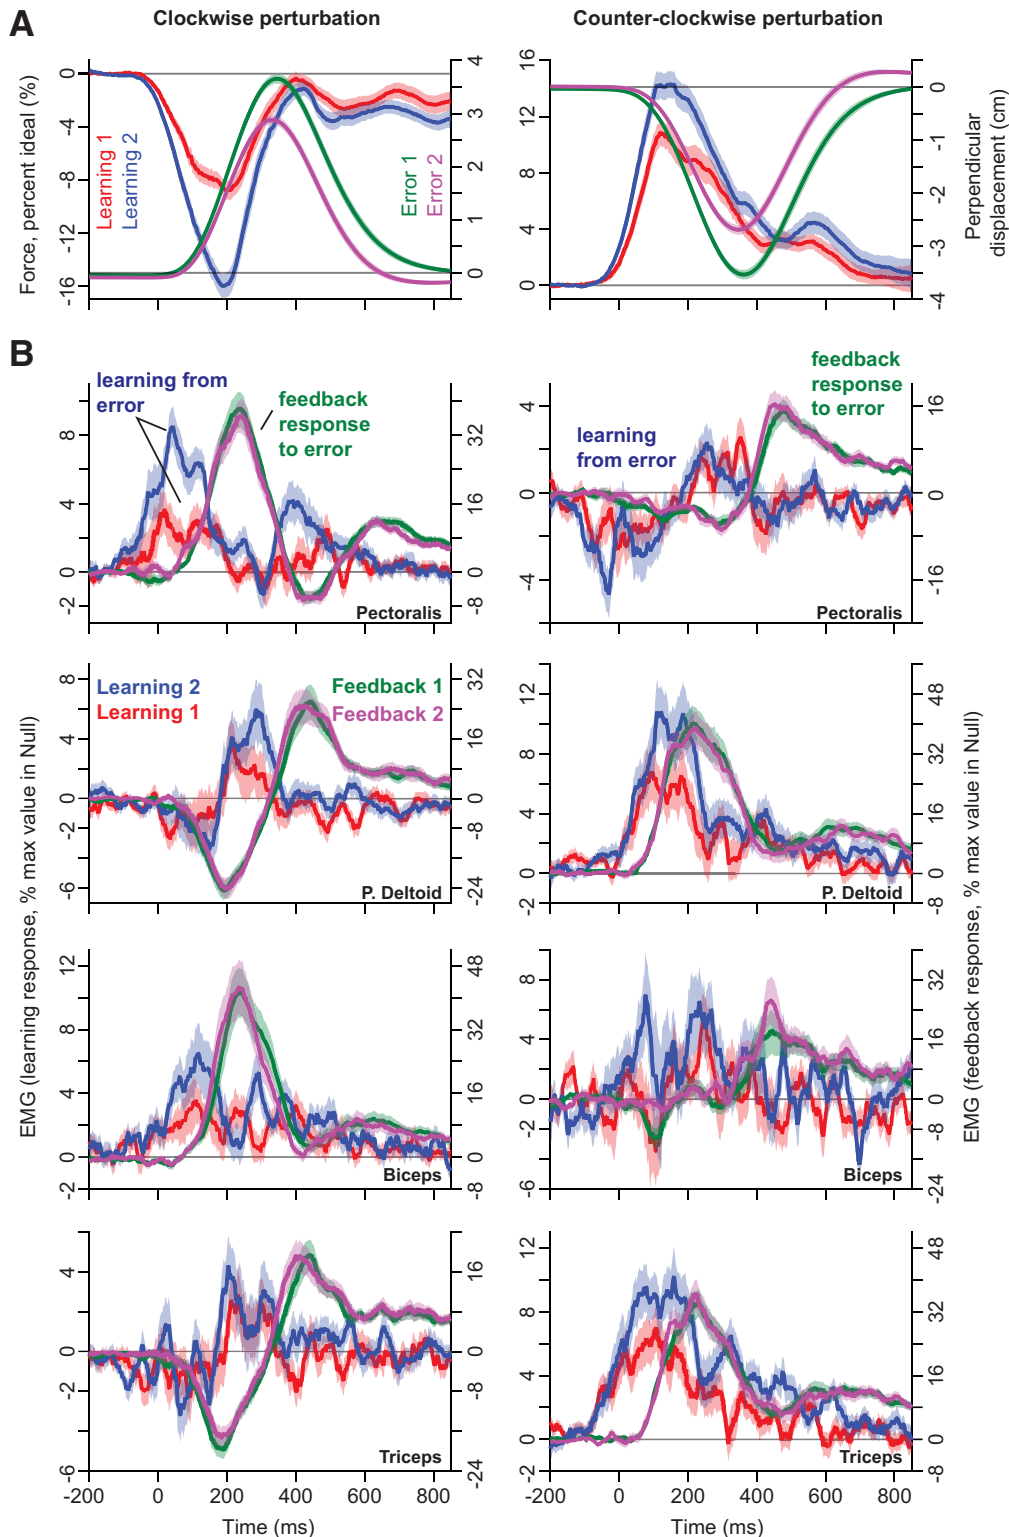


Figure 2. Learning and error-feedback responses across the population. The 0 ms time point denotes movement start. Error bars represent ± 1 SEM. **A**, Kinematic correlates of learning and error. During perturbation trials, subjects ($n = 57$) experienced large perpendicular displacements specific to the orientation of the force field (right axis). Displacement during the second perturbation (Error 2, purple) was smaller than that for the first perturbation (Error 1, green) due to partial learned compensation for the force field. As a result of the experience of error, subjects produced lateral forces against the error-clamp channel, in accordance with the field orientation (left axis). The net change in force production after the experience of two consecutive perturbations (Learning 2, blue) was larger in magnitude than single-trial learning (Learning 1, red), due to the accumulation of two single-trial learning events. **B**, EMG correlates of learning and error-feedback responses. Learning and error-feedback signals for four muscles of the upper arm and trunk are provided. Feedback responses (right axis) were in general much larger in magnitude than learning responses (left axis), as indicated by the 25% scaling factor relating the left and right axes for all muscles. The error-feedback response for the first perturbation (Feedback 1, green) is nearly identical to the error-feedback response for the second perturbation (Feedback 2, purple). Note that the reference error-free reach for Feedback 1 is the EC1 trial before the perturbation. For Feedback 2, it is the EC2 trial after single perturbations that accounts for single-trial learning. The learning signals were computed as the change in muscle activity (EC2 minus EC1) during the error-clamp reach before and after the perturbation(s). Learning from a single perturbation is shown in red and the accumulated learning from two consecutive perturbations is shown in blue. As would be expected, more is learned from two perturbations than one perturbation.

within-subject regression analysis described above was performed on the randomly shifted subject dataset and repeated 200 times, each time resampling shifts from the uniform distribution.

In terms of kinematic correlates of learning, we focused on the forces that subjects produced against the stiff spring that opposed lateral trajectory deviations during error-clamp trials. Using standard procedures, we compared this subject-produced force trace to the ideal force that would be required to compensate for the perturbation. In brief, the maximum tangential velocity attained during that trial was multiplied by the field magnitude in the preceding perturbation trial. Next, subject-produced force at each time point was normalized by this value and converted to a percentage.

Finally, to determine how the relationship between feedback and learning might vary with temporal variation in the magnitude of the feedback response, we considered the fact that in some trials, a subject might produce a strong feedback response to the perturbation, whereas in other trials the same subject might produce a weak response. Did the variability in the feedback response correspond to the variability in the learning response? To answer this question, we separated the data for each subject and each muscle into two classes that corresponded to high and low feedback gains. We will refer to these classes as the “large” and “small” feedback responses, respectively. To construct these two labels, we considered each muscle and each subject separately and focused on the Feedback 1 EMG traces. For agonist muscle feedback responses, we computed the mean Feedback 1 response over the interval between 150 and 450 ms after the start of movement. This interval was selected because it best captured differences in the gain of early agonist activity. For antagonist muscle feedback responses, we selected a longer averaging window defined by the range 0–600 ms. This wider interval was selected to include both early inhibition in antagonist responses during the perturbation as well as excitation that occurred near movement termination. A perturbation trial was labeled as high feedback response if its Feedback 1 EMG trace exceeded the median Feedback 1 response observed in that muscle and that subject over the appropriate time interval (similarly for the low feedback response label). We computed the mean feedback responses for these two labeled datasets, and then, for each labeled feedback trace, we computed the learning trace that immediately followed. In addition, we considered kinematic correlates of these responses, corresponding to the maximum perpendicular displacement during the perturbation (for feedback) and maximum error-clamp force production (for learning). We used *t* tests to determine whether there existed a difference between these kinematic parameters, and expressed their difference as ratios (high-feedback trials/low-feedback trials).

Results

We asked whether the feedback response that corrected for a perturbation during a reach produced a signal that acted as a teacher for the motor system, instructing it on how to predictively cancel the perturbation on the following trial. Our experiment employed triplet or quartet progressions of error-clamp, perturbation, error-clamp trials, as illustrated in Figure 1A. The average hand paths for the error-clamp trials and perturbation trials are provided for a typical subject in Figure 1B. The time course of the perturbation-induced displacement perpendicular to the direction of motion is shown in Figure 1C (at left). Following experience of this error, the nervous system altered the motor commands that it produced on the very next trial. To

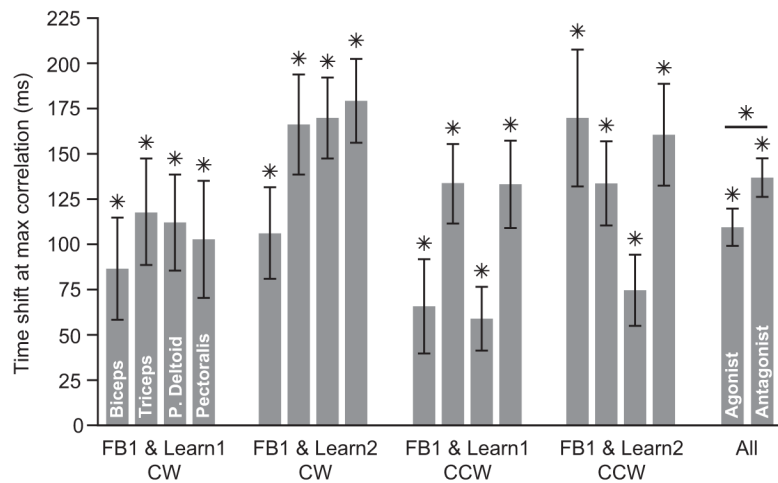


Figure 3. Temporal shifts relating learning and feedback responses. Within-subject cross-correlations were used to determine the time shifts for which the feedback responses and learning responses were maximally correlated. EMG time courses of learning from a single perturbation and from two consecutive perturbations were cross-correlated with EMG time courses of feedback responses to the first perturbation. Positive values indicate that feedback responses lagged behind the learning response. Shifts were averaged across subjects and error bars represent ± 1 SEM. Each group of four bars represents a particular learning–feedback condition. Each bar in these groups represents a different muscle (left to right: biceps, triceps, posterior deltoid, and pectoralis). From left to right, the field orientations for each four-bar group are as follows: CW, CW, CCW, and CCW. From left to right, the number of perturbations experienced for each four-bar group were as follows: one, two, one, and two. The right-most bars represent data collapsed across all perturbations and separated for muscles that were agonist (responded early to restore the trajectory) or antagonist.

visualize this change, we compared the forces produced in the error-clamp trial preceding the perturbation to the forces produced in the error-clamp trial following the perturbation. The change in the motor commands produced a force pattern that was opposite in direction to that of the displacement (Fig. 1C, right).

By analyzing the temporal patterns of muscle activity in the error-clamp and perturbation conditions, we obtained neural correlates of feedback response to error, as well as trial-to-trial learning. Example traces of EMG activity in the triceps for a typical participant are shown in the left column of Figure 1D. In the error-clamp trial that preceded the perturbation (EC1), the triceps gradually increased its activity, peaking midmovement at ~ 200 ms. In the perturbation (Fig. 1D, red trace, left column, P1), the triceps activity was inhibited relative to EC1 for the majority of the reach, but then demonstrated a sharp excitation as the movement was terminated.

To compute the feedback response to error (Fig. 1D, green curve, right column), we subtracted the EMG time course in the error-clamp trial (EC1) from the EMG time course in the perturbation trial (P1 minus EC1). In this participant, the CW displacement produced a feedback response that included an early inhibition of triceps (Fig. 1D, green curve, right column), followed by a late excitation of the same muscle. To compute the learning response, we compared the trial-to-trial change in the EMG signal in EC1 and EC2, the error-clamp trials before and after the perturbation trial (Fig. 1D, blue trace, left column). This difference (EC2 minus EC1) represents the trial-to-trial change in the motor command as a result of experiencing a single trial of error. We observed that the learning response (Fig. 1D, purple curve, right column) appeared to be a scaled version of the feedback response, but shifted earlier in time.

Group-averaged kinematic and EMG traces for CW and CCW perturbations are shown in Figure 2. The kinematic and force data are shown in Figure 2A, where we have plotted the error

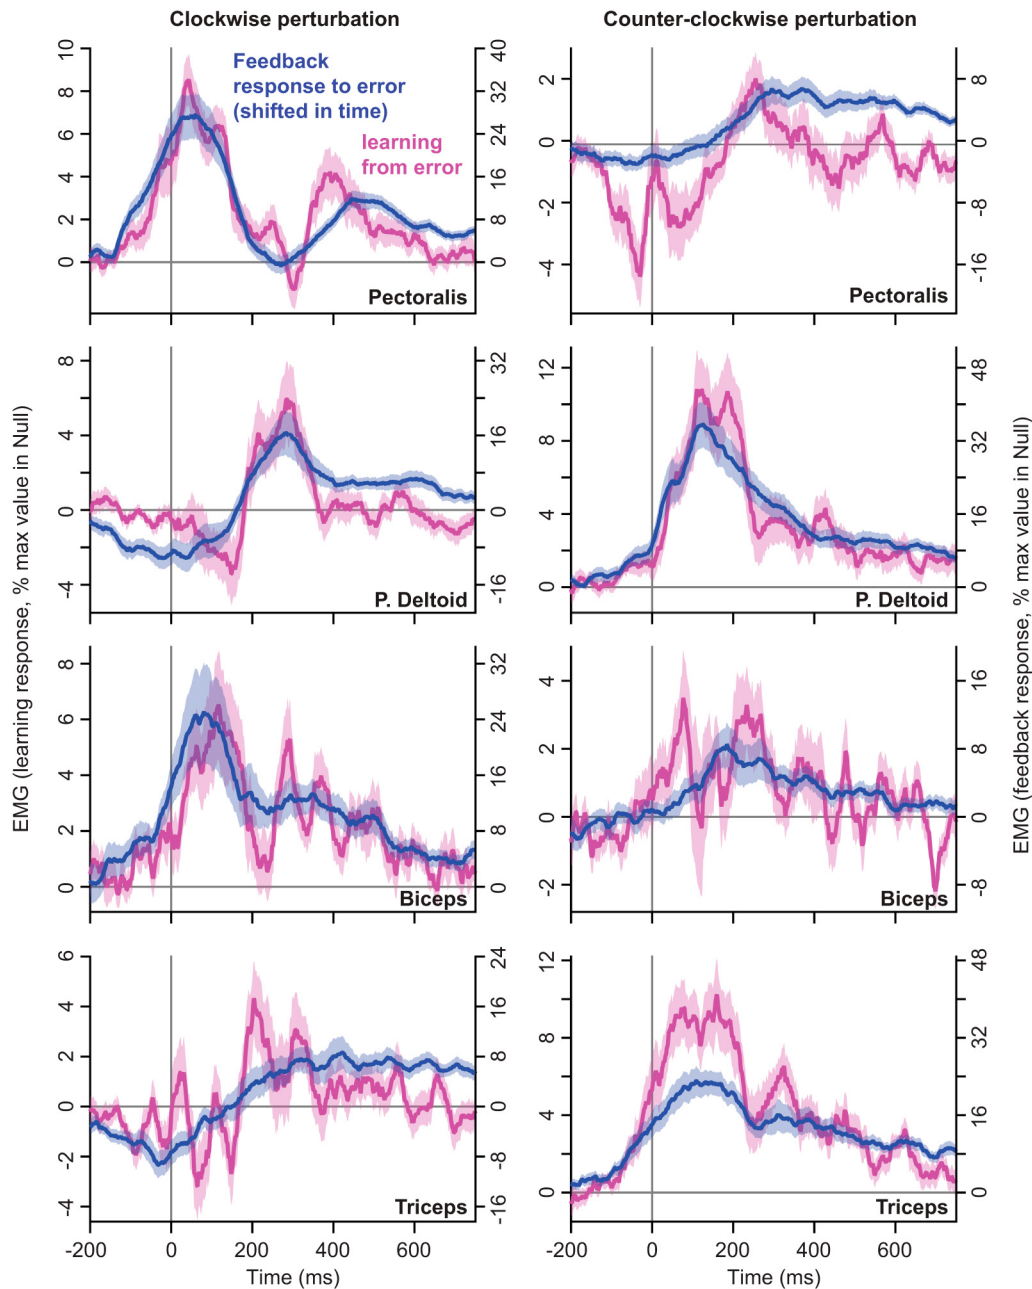


Figure 4. Learning resembles time-shifted copies of the error-feedback response. The accumulated EMG learning responses for two consecutive errors (purple, left axis) are compared with shifted feedback responses to the first perturbation (blue, right axis). The feedback responses were shifted independently for each muscle–field orientation–subject trio, to maximally align them with their corresponding learning responses, according to cross-correlation analysis of the learning–feedback time courses. Clear correspondence between the learning time course and the feedback time course is illuminated in this shifted feedback space. Here 0 ms represents movement onset for the nonshifted learning error-clamp trials. The behavior for CW perturbations and CCW perturbations is shown in the left and right columns, respectively. Each figure displays a different muscle. Shaded error bars represent ± 1 SEM at each recorded time point.

induced by the first perturbation (Error 1), and the resulting trial-to-trial change in force produced in the subsequent error-clamp trial (Learning 1). In trials in which a second perturbation followed the first, the errors were smaller (Fig. 2A; Error 2 vs Error 1, peak displacement, $p < 10^{-23}$ for both fields). Similarly, in trials in which a second perturbation followed the first, learning following two perturbations was larger (Fig. 2A; Learning 2 vs Learning 1, peak force, $p < 10^{-8}$ for CW field, $p < 10^{-4}$ for CCW field).

Figure 2B illustrates the EMG measures of error-feedback and learning responses. Perhaps the most striking feature of the data was the similarity between the two traces. We found that, in general, the learning response appeared to be a shifted and scaled

version of the error-feedback response. This was best demonstrated by the EMG in the pectoralis, posterior deltoid, and triceps in their respective antagonist fields, where learning and error-feedback traces exhibited initial inhibition followed by excitation later in the movement (pectoralis for a CCW perturbation; posterior deltoid and triceps for a CW perturbation). Another clear example was presented by the bimodal excitatory pattern in the learning and feedback traces in the pectoralis for CW perturbations, where the second learning-excitation peak occurred just before movement termination.

For example, in the pectoralis muscle, the CW error-feedback response was an excitation that peaked at ~ 200 ms with respect

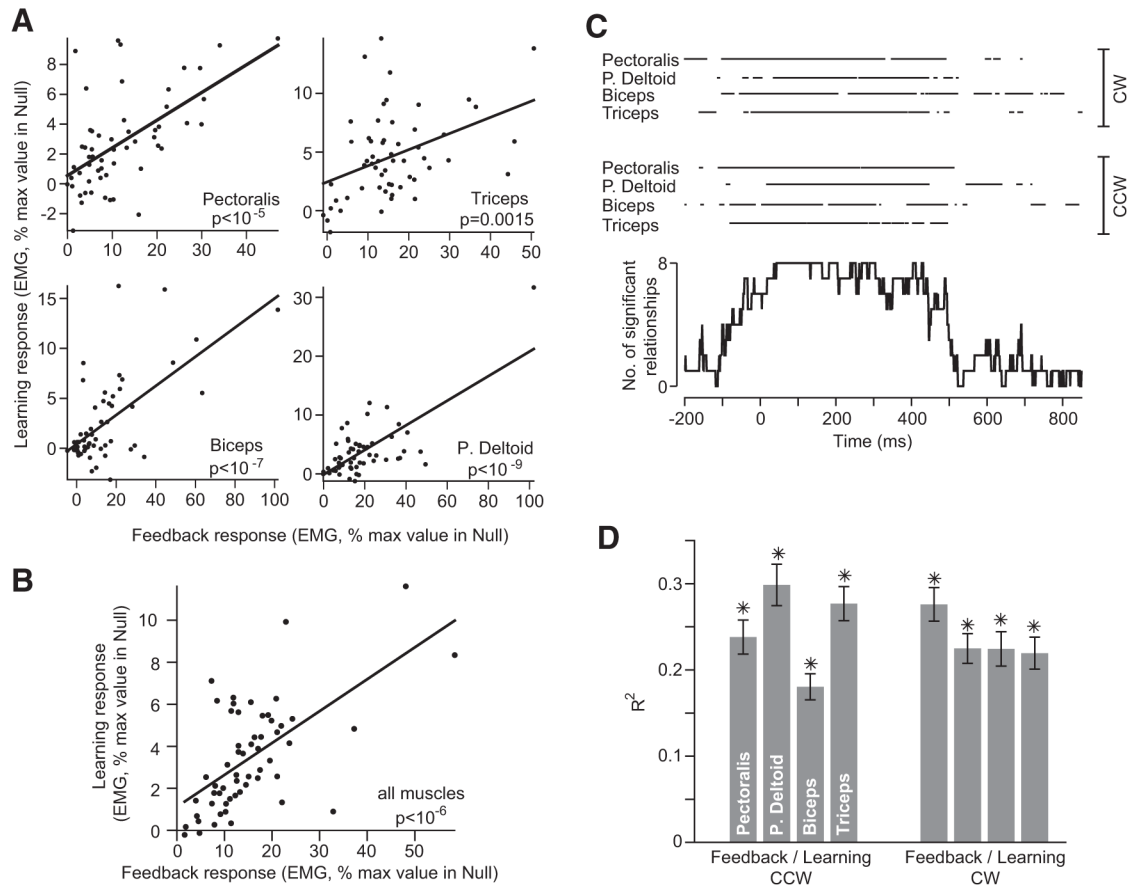


Figure 5. Learning and error-feedback responses across individuals. **A**, Feedback and learning responses in muscles that were agonists for a perturbation. The mean EMG learning response (from 2 perturbations) for each subject was computed over the interval between -200 and 500 ms and regressed onto the mean feedback response between 0 and 500 ms, for the first perturbation. **B**, The learning and feedback responses in each muscle for each subject were averaged, producing a single measure of feedback and learning in each subject. **C**, Learning is correlated to shifted feedback responses during movement on a fine temporal scale. Feedback responses (first perturbation) were shifted for each muscle–field orientation–subject trio independently to achieve temporal alignment with corresponding learning traces (cumulative learning from 2 perturbations), according to cross-correlation analysis. After alignment, learning at each time point was regressed onto the feedback response, across subjects, for each individual time point, independently for each muscle and field orientation (total of 8 combinations). The raster plots at top mark time points for which this regression was statistically significant ($p < 0.05$) and possessed positive slope (indicating a positive correlation). Each line oriented left to right shows a particular muscle in one of the field orientations. The continuous-time figure at bottom was constructed from these eight regression raster lines. At each time point, the number of significantly positive correlations (maximum of 8) was counted. Time at 0 ms refers to movement onset. The learning–feedback correlation appears to begin 100 ms before movement onset, and saturates for the entirety of the movement, which ends on average at 500 ms. **D**, We linearly regressed the learning response onto the feedback response over the time period from -100 to 500 ms. This regression was performed for each muscle and field orientation separately. The R^2 value of the regression is provided. The groups at left and right correspond to CCW and CW fields, respectively. Each bar represents a different muscle; from left to right, pectoralis, posterior deltoid, biceps, and triceps.

to reach initiation, followed by a second, smaller peak at ~ 600 ms. Learning also possessed two peaks of excitation, peaking at ~ 50 ms and again at ~ 400 ms. In the CCW perturbation, the error-feedback response was an inhibition that peaked at ~ 300 ms and an excitation that peaked at ~ 450 ms. Learning was also an inhibition followed by an excitation, but its timing had peaks at ~ 0 and 250 ms.

We had naively expected that only the early portion of the error-feedback response might resemble the corresponding learning response. However, we found that the learning and error-feedback responses appeared similar until near the cessation of the error-clamped reach (~ 500 ms), implying that both the short-latency and long-latency error-feedback responses were combined and shifted earlier in time to become the learning response.

To quantify the temporal shifts that related the EMG measures of learning and error feedback, we computed their cross-correlation and found that the two traces were maximally correlated when the feedback response was shifted earlier in time by

123 ± 61 ms (mean \pm SD across all muscles and conditions; Fig. 3). To combine the data across various muscles, we labeled each muscle as agonist or antagonist for each perturbation. For example, a CW perturbation produced an initial excitation in the pectoralis and biceps, but inhibition in the posterior deltoid and triceps. Therefore, the pectoralis and biceps were agonists in responding to a CW perturbation. We found that the temporal shift from the feedback response to the learning response across muscles was larger when the muscle acted as an antagonist (137 ± 80 ms), responding in the direction of the perturbation, compared with when the muscle acted as an agonist (109 ± 78 ms, $p = 0.042$), responding in the direction opposite the perturbation. Similarly, the optimal shift was larger for Learning 2 (145 ± 75 ms) than Learning 1 (101 ± 80 ms, $p < 10^{-3}$), indicating that additional perturbations not only induced additional learning, but also caused this learning to be expressed earlier in time.

To better visualize the temporal relationship between learning and error-feedback responses, we plotted the time-shifted error-feedback response together with the learning response for the

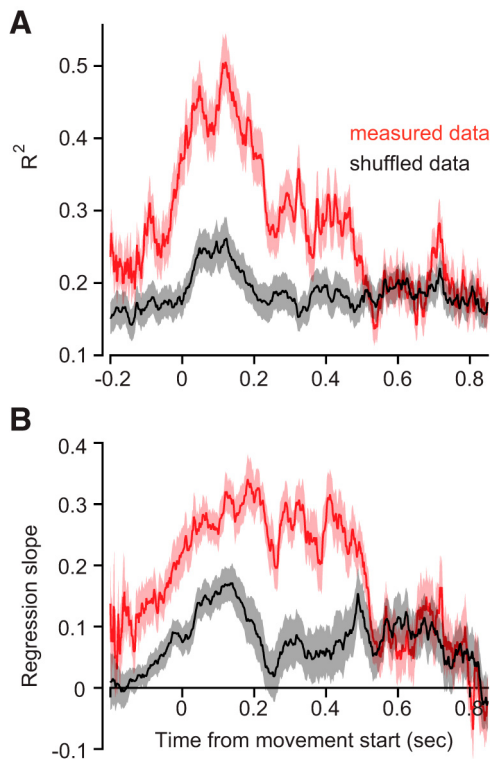


Figure 6. The error-feedback response is predictive of the learning response within a subject. For **A** and **B**, 0 s refers to movement onset. **A**, The feedback response accounts for within-subject variation in the learning response during movement. Regressions were performed within subject across the eight muscle–field orientation combinations (i.e., 8 points in each regression). Error feedback responses (first perturbation) were shifted to achieve temporal alignment with learning responses (cumulative learning from 2 consecutive perturbations) according to their cross-correlation. At each time point, learning across the muscle–field orientation pairs was regressed onto the corresponding error-feedback response. The R^2 value for the linear regression is provided in the figure (red). The red shaded error bars indicate ± 1 SEM. To quantify the baseline random correlation inherent in these signals (black), the regressions were repeated, this time randomly shifting the feedback response. The random shifts were sampled from a uniform distribution (0–400 ms). The black shaded error bars indicate ± 1 SD across 200 repetitions of the analysis. Comparison of the red and black traces indicates that the feedback response encodes variability in the learning response until movement termination at ~ 500 ms. **B**, The slopes of the regressions described in **A** are provided. This slope represents the scaling factor relating learning and feedback. It appears that $\sim 25\%$ of the feedback response magnitude was incorporated into the learning response at each point during the movement.

larger-amplitude Learning 2 traces (Fig. 4). The peaks and troughs in the error-feedback response appeared to be consistent with the features of the learning response throughout the duration of the movement. In addition, across the various muscles and perturbation orientations, the scaling factor relating the magnitudes of the learning and error-feedback response (reflected in the scaling factor relating the left and right y -axes of Fig. 4) was consistent at $\sim 25\%$, suggesting that approximately a quarter of the feedback response in all muscles became the learned response.

Across-subject variability in feedback response

These results indicated that the neural feedback responses in each muscle may be a strong predictor of the learning response in that muscle. To better quantify this relationship, we focused on agonist muscles and asked whether those subjects that demonstrated more agonist muscle activity during their feedback response also expressed more learning. For each subject, we computed the mean activity in the learning and feedback responses of agonist

muscles over the periods from -200 to 500 ms and from 0 to 500 ms, respectively. We then asked whether subjects that had shown a greater learning response also produced a greater feedback response. Figure 5A plots the magnitudes of the feedback response and learning response for each subject in each muscle. We found a statistically significant, positive correlation between the sizes of the two responses across all muscles. When we averaged the learning and feedback responses across all muscles in each subject, we found a very strong relationship (Fig. 5B). This indicated that subjects that had shown a larger feedback response were very likely to also show a greater learning response.

We performed a similar across-subject analysis but at a much finer temporal resolution to determine the length of the time over which learning and feedback were positively correlated (i.e., how much of the learning response could be explained by feedback as a function of time in the movement). To optimally align feedback responses with the learning time courses, we first shifted each feedback response by the shift determined via cross-correlation. Next, learning for a given muscle and field orientation was regressed onto the corresponding feedback response for that muscle and field orientation, independently at every time point (Fig. 5C). We found that both the agonist and antagonist muscles possessed significantly positive correlations between learning and feedback for the entire movement period. Specifically, this correlation began ~ 100 ms before the error-clamped movement onset and saturated for the entire movement period (500 ms on average), falling to baseline levels after the reach terminated. To determine the level of correlation between these signals, we linearly regressed the learning time course onto the feedback response (Fig. 4, aligned traces) over the time interval between -100 and 500 ms (Fig. 5D). We found that the R^2 values for these regressions were similar across muscles and field orientations, varying within the range of 0.18 ± 0.02 to 0.31 ± 0.02 , with a mean value of ~ 0.25 .

Within-subject variability in feedback response

Our across-subject analyses indicated that learning and feedback were correlated during preparation and execution of the error-clamped reach. We next asked whether, within subject, the strength of the correlation between learning and feedback responses varied during movement duration. In other words, we quantified the extent to which the variability in the feedback response accounted for the variability in the learning response as the movement progressed. To quantify this relationship, at each time point we used within-subject linear regression to compare the feedback responses across muscles and perturbation orientations with the corresponding learning responses. To generate a statistical comparison, we generated a null hypothesis by computing the correlation between the two signals when the time shift was randomly sampled from a uniform distribution between 0 and 400 ms. We found that the variability in the feedback response, within a subject, accounted for a maximum of $\sim 50\%$ of the variability in the learning response, peaking slightly after movement onset (Fig. 6A). However, the correlation between learning and feedback remained significantly above control levels up until movement termination (~ 0.5 s). Interestingly, the scaling factor describing the learning–feedback regression remained relatively stable within the range of 20 – 30% during the reaching movement, indicating that $\sim 25\%$ of the error-feedback response became the learning response generated on the following trial (Fig. 6B).

The across-subject and within-subject results are thus far congruent with the hypothesis that error-feedback signals are instructors of learning. We observed that feedback and learning

signals possessed a “scaled-and-shifted” relationship; the feedback response appeared to be scaled down in magnitude, shifted earlier in time, and added to the feedforward motor plan to achieve the learning response. However, our experiment employed only a single perturbation magnitude, leaving one to question the generality of this proposed learning–feedback relationship. To address this question, we considered that on some trials, the subject would strongly resist the perturbation, whereas in other trials the same subject might only weakly resist the perturbation. For each subject and each muscle, we labeled the perturbation trials as large or small feedback response, based on the magnitude of the corresponding Feedback 1 response time course. For a given subject and given muscle, the large-feedback trials were constructed from all perturbation trials in which the Feedback 1 response exceeded the median feedback response (see Materials and Methods). The small-feedback trials were labeled similarly, but for responses that fell below the median Feedback 1 response. As is implied by this description, we divided the trials for each subject based on their feedback responses, not learning responses.

The two feedback responses are shown for agonist muscles in Figure 7 (left column). Labeling our data in this manner revealed that a perturbation produced feedback responses that were highly variable. The large feedback responses (red traces) possessed peak magnitudes in excess of twice the magnitude of the low feedback responses (black traces). For the trials used to label the two classes of triceps and posterior deltoid responses, there was no statistically significant difference in the maximum perpendicular displacement ($p = 0.703$ and $p = 0.279$, respectively). For the biceps and pectoralis, we found a significant difference between these maximal errors ($p < 10^{-5}$ and $p < 10^{-8}$, respectively) but their difference (7% for biceps, 8% for pectoralis) was too small to be adequately explained by the twofold difference in the feedback-response gains. To summarize, perturbation trials with nearly identical kinematics showed significant differences in the underlying patterns of feedback muscle activations.

Did these large differences in the feedback responses correspond to differences in the learning signals observed during the subsequent error-clamp trial? We computed the learning responses that were induced by the large-feedback and small-feedback trials (Fig. 7, right column). Remarkably, the differences in the feedback gains were mirrored in the magnitudes of the learning time courses. A perturbation that produced a large feedback response was followed by a large learning response, as shown by the red and black traces in Figure 7. This result suggested that the agonist learning responses were highly sensitive to

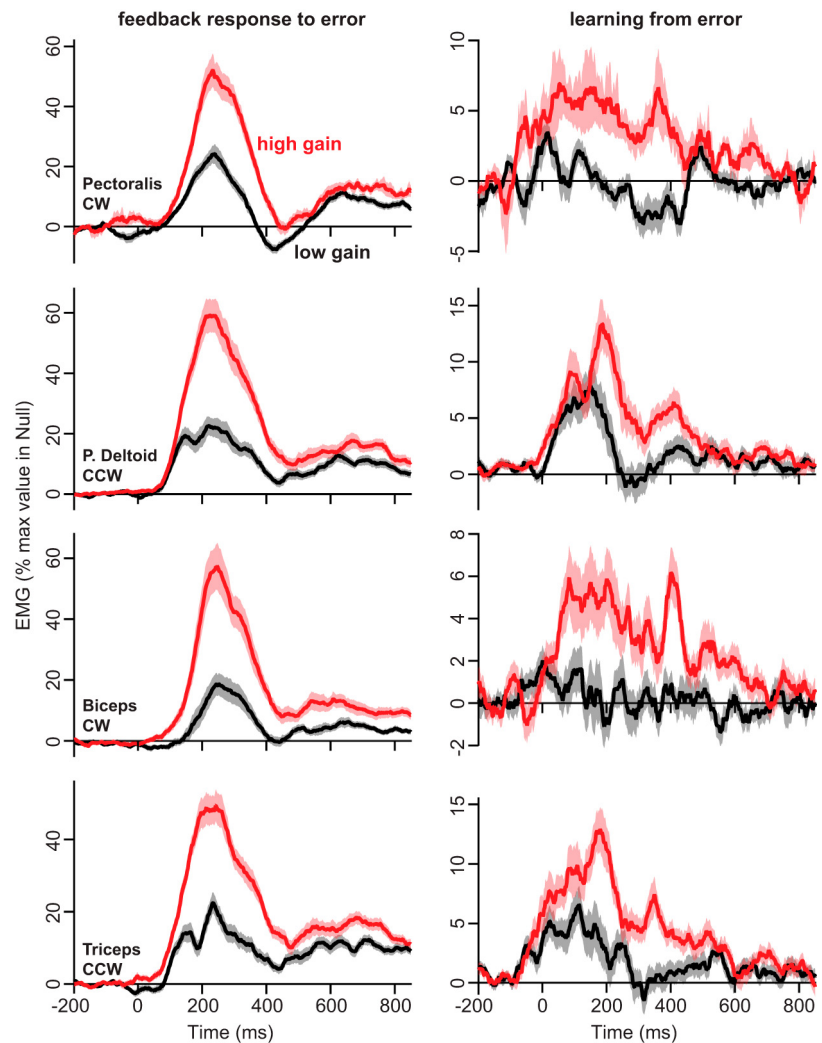


Figure 7. Agonist learning magnitude correlates to the feedback-response gain. Triplet trial progressions were classified into two groups of equal size based on the magnitude of the feedback response of agonist muscles. Large and small feedback groups were constructed from individual responses that fell above and below the median response over the period from 150 to 450 ms. In the left column, the mean time courses of the small feedback response group (black traces) are overlaid on the feedback responses for the large group (red traces). At right, the learning responses corresponding to the large (black traces) and small (red traces) feedback groups are shown. The difference in magnitude of the learning responses mirrored that of the corresponding feedback time courses. Importantly, these classifications were made solely based on the feedback responses, not the resultant learning. Each figure indicates a different muscle; from top to bottom, pectoralis, posterior deltoid, biceps, and triceps. Movement onset is indicated by the 0 ms time point. Errors bars indicate ± 1 SEM.

the gains of the feedback response. It did not appear that the size of the agonist learning responses was on average indicative of the force being produced during the channel trial (maximum peak force, $p > 0.05$ for all muscles).

We next performed the same analysis for each muscle in their respective antagonist field. The small and large antagonist feedback responses are shown in the left column of Figure 8. Our labeling method revealed two distinct differences between the high and low feedback responses. First, the early period of inhibition was attenuated in the large feedback responses (red traces) relative to the small feedback (black traces) responses. The second difference between the two feedback responses was characterized by enhanced late excitation of the large feedback response relative to the small feedback response. Again, as for the agonist responses, these differences were not reflective of some large difference in the underlying kinematics of the error. The only muscle for which we observed a statistically significant difference in

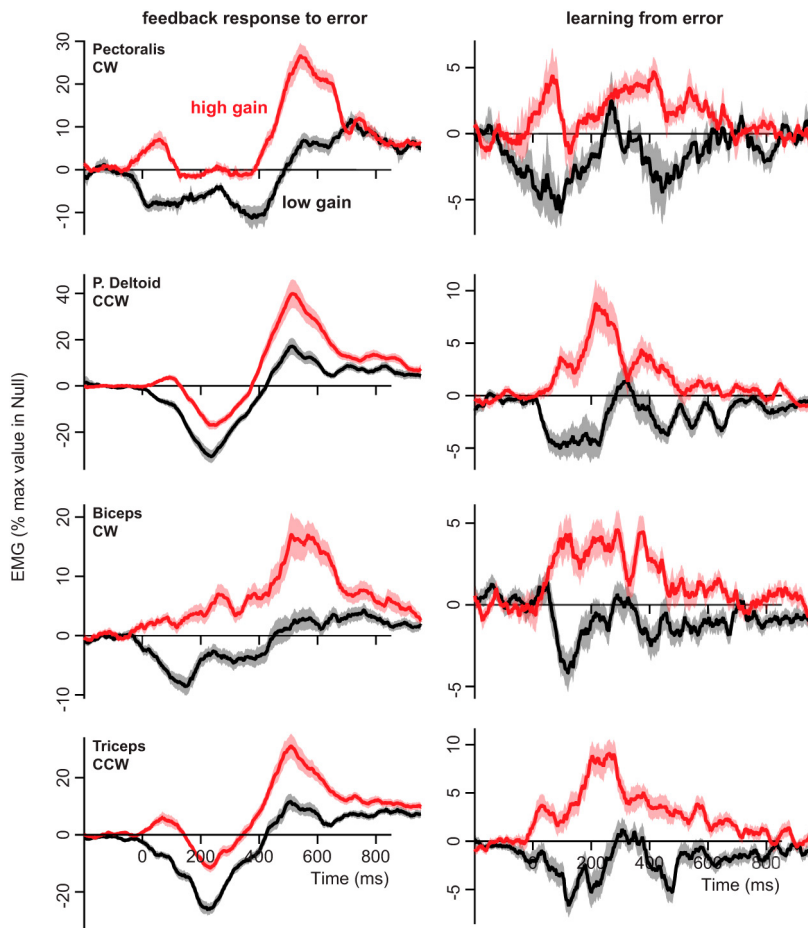


Figure 8. Antagonist learning magnitude correlates to the feedback-response gain. Triplet trial progressions were classified into two groups of equal size based on the magnitude of the feedback response of antagonist muscles. Large and small feedback groups were constructed from individual responses that fell above and below the median response over the period from 0 to 600 ms. In the left column, the mean time courses of the small feedback response group (black traces) are overlaid on the feedback responses for the large group (red traces). At right, the learning responses corresponding to the large (black traces) and small (red traces) feedback groups are shown. The difference between the feedback responses appeared to correlate with differences in the sign of the learning responses. Importantly, these classifications were made solely based on the feedback responses, not the resultant learning. Each figure indicates a different muscle; from top to bottom, pectoralis, posterior deltoid, biceps, and triceps. Movement onset is indicated by the 0 ms time point. Errors bars indicate ± 1 SEM.

the maximum perpendicular displacement was the triceps ($p = 0.019$), but this difference (only 3%) was too small to be explained by the differences in feedback-response gains.

Once again these differences in feedback responses were paralleled in learning (Fig. 8, right column). Learning traces for the small feedback response (black traces) diverged from those pertaining to the large feedback response (red traces) for the entirety of the reach. This divergence obeyed the differences we observed in the corresponding feedback responses. We found that the learning time course resembled time-shifted and scaled replicas of the feedback response. One exception to this relationship was the activity of the triceps and posterior deltoid in the high feedback response, which lacked an early period of inhibition that would be expected from consideration of the corresponding feedback traces. We speculate that this inhibition was likely cancelled by cocontraction of agonist–antagonist pairs, which, based on previous accounts (e.g., Milner and Franklin, 2005), is often seen in the initial stages of force-field learning. Indeed in these two muscles during the initial part of the reach, the excitation in the high feedback group's learning response overlapped with ex-

citation in agonist muscles (i.e., we observed cocontraction). Similar to the agonist learning responses, these differences in antagonist control signals did not correspond to differing levels of force production in the channel ($p > 0.05$ for all muscles).

In summary, for both agonists and antagonists muscles, despite constant perturbations, the occasion in which the feedback response was high often produced a learning response that was large, suggesting a strong coupling between the feedback response and the learned response.

Control studies

An assumption critical to our analysis was that the EMG patterns in error-clamp trials represent the motor output in an error-free movement. To test for this, we compared the EMG traces in EC1 with EMG traces recorded during baseline reaching in the null field in each muscle (Fig. 9A). We computed the mean EC1 signal across Blocks 1 and 2 of the experiment and compared it to the mean null field EMG signal of the 23 trial null periods that commenced each block. Indeed, the EC1 signal (Fig. 9A, red curves) appeared indistinguishable from that recorded in the null field periods (Fig. 9A, black curves). This analysis also suggested that on average, learning from the CW or CCW perturbations was washed out during the intervening null trials between consecutive perturbation periods, another assumption critical to our analyses. To ensure that this apparent washout was not trivially caused by the cancellation of residual learning of the

oppositely oriented CW and CCW fields, we computed the mean EC1 signal corresponding to trials that followed CCW perturbations and compared this to the mean EC1 signal of trials following CW perturbations (Fig. 9B). We found that these two groups of EC1 muscle activities were identical, further confirming that sufficient washout occurred between consecutive triplet/quartet progressions.

To measure the learning response, we compared the EMG in EC2 with the EMG in EC1. This comparison requires that the kinematics of the two movements be identical. To check for this, we compared the tangential component of the reach in the two error-clamp trials (Fig. 9B), and found the two to be indistinguishable.

We wanted to determine whether some simple linear transformation existed between the learning EMG signals and the learned force profiles. We found a weak, but significant ($p = 0.0186$), positive correlation between the magnitude of early agonist EMG activity (mean across the four muscles) over the period from -100 to 200 ms and learned force production (Fig. 2A, left). This offered some evidence that larger EMG learning corresponded to larger learned forces. However,

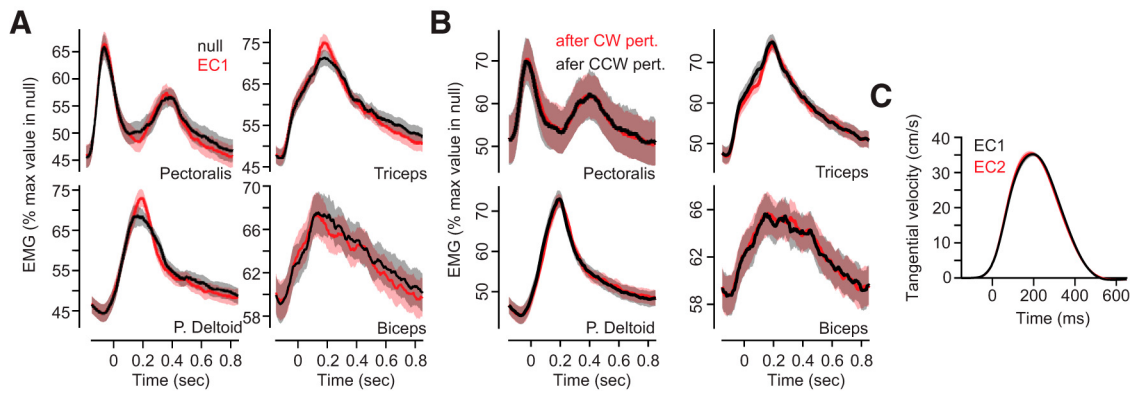


Figure 9. Error-clamp trials provide accurate approximations to unperturbed movements. For **A–C**, the 0 s or 0 ms time points refer to movement onset. **A**, EMG during error-clamp trials before the perturbation is identical to null field EMG. The EMG activity of each muscle in the error-clamp trial preceding a movement (EC1, red) is contrasted with the EMG signal in the 46 null field trials that began Blocks 1 and 2 (null, black, 23 trials at the start of each block). Each figure indicates a different muscle; top left, pectoralis; top right, triceps; bottom left, posterior deltoid; bottom right, biceps. **B**, Sufficient washout occurred between perturbations. We compared the EC1 activity after CW perturbations (red traces) to EC1 activity after CCW perturbations (black traces). We found that these activities were identical, indicating complete washout between consecutive perturbations. The layout of the figure is identical to **A**. **C**, The tangential kinematics of the error-clamp movements before and after the perturbation are identical. The tangential velocity in all of the error-clamp trials before the movement (EC1, black) was identical to that of the post-perturbation error clamp (EC2, red). Errors bars indicate ± 1 SEM.

force production truly relates to changes in the net torque about a joint and therefore is determined by the balance between agonist and antagonist muscle activities; we found no significant ($p = 0.125$) correlation between antagonist muscle activity and force production during this early period of the reach.

To determine the robustness of our estimate for the learning response, we considered an alternative approach by comparing the difference in the EMG recorded in the P2 and P1 trials. That is, P2 minus P1 should resemble learning from a single error, provided that the feedback responses during the first and second perturbation trials are the same. Fortunately the Feedback 1 and Feedback 2 responses of Figure 2*B* are rather similar, though not identical. Therefore, we compared the P2 minus P1 EMG signal with the Learning 1 signal, across subjects (Fig. 10). We confirmed that P2 minus P1 resembled our estimate of the learning response (despite the fact that the measures relied on different comparisons). However, their correspondence was not exact, reflecting differences in Feedback 1 and Feedback 2. For example, the accentuated peak and trough in the CW biceps perturbation response difference (Fig. 10, first column, third row) corresponded precisely to a slight temporal shift relating the Feedback 1 and Feedback 2 responses of Figure 2*B* (first column, third row).

Thus far we have analyzed learning from a single perturbation as well as cumulative learning from two perturbations. If the scale-and-shift relationship between feedback and learning is a general learning rule, we should observe this phenomenon for the learning response from the second perturbation experienced in quartet trial progressions. Given that the corresponding feedback responses were nearly identical during the first and second perturbations (Fig. 2), we would predict that the single-trial learning time courses induced by these perturbations should be quite similar. To compute the learning that occurred solely due to experience of the second perturbation, we considered the difference between the quartet EC2 and the triplet EC2. This difference represents the learning that takes place in the feedforward command after the experience of P2. In Figure 11, we have plotted our estimate of learning from the second perturbation in the quartets alongside our estimate of learning from the perturbation in the

triplets. The two possess an extremely close resemblance. Thus, this analysis provided further evidence that supports our working hypothesis concerning the relationship between learning and feedback.

Finally, we wanted to ensure that the relationship between learning and feedback established by our analysis of high and low feedback trials (Figs. 7, 8) was not trivially the result of some process that varied systematically during the progression of the experiment. For example, perhaps subjects became less sensitive and responsive to the triplet trial progressions due to the uncertain nature (i.e., frequency and orientation) of the force-field perturbations. If such a phenomenon was responsible for the feedback response variability, we would expect that there would be some trend in which trials corresponded to the large and small feedback triplet trial progressions. However, we found no such trend in the trial orderings (Fig. 12). The large and small feedback groups were constructed of trials that were sampled approximately uniformly across the experiment for both agonist and antagonist muscles. This finding suggests that the observed changes in feedback response gains was the result of random within-subject fluctuations in the gain of the neural feedback controller, rather than a systematic modulation due to passage of time.

Discussion

When we experience an error during a movement, the result is a sensory mismatch between the intended movement and the actual movement. This error is encoded in sensory coordinates. However, to improve our motor commands, the brain must transform the sensory representation of error to a motor representation of commands in muscle space. Our study sheds light on this process.

During a reaching movement, the sensory encoding of error engages spinal and supraspinal neural circuits that, following a delay, produce motor commands, partially correcting for the error. The motor response to the mismatched sensory feedback is termed an error-feedback response (Kawato et al., 1987). With practice, the gain of the sensory feedback response to error can be increased, resulting in more vigorous corrections to the repeatedly experienced errors (Ahmadi-Pajouh et al., 2012). However, because of inherent delays in

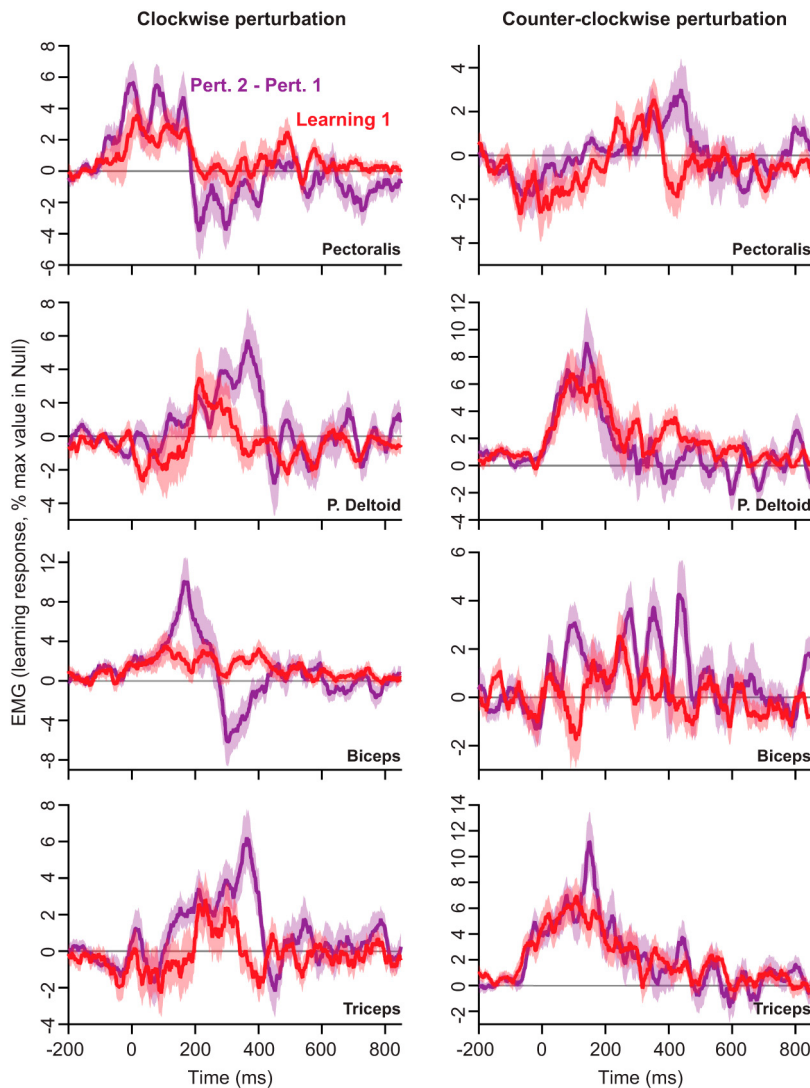


Figure 10. The change in EMG activity from perturbation 1 to perturbation 2 is largely due to single-trial learning. The difference in the EMG activity in the perturbation trials (purple, Pert. 2 – Pert. 1) is compared with single-trial learning (red, Learning 1) as measured as the EMG change from the preperturbation error-clamp trial to the postperturbation error-clamp trial. One would expect these signals to be similar, barring differences in the error-feedback responses for the first and second perturbations. EMG activity in the CW and CCW fields is shown in the left and right columns, respectively. Each row of figures applies to a particular muscle. From top to bottom, these muscles are as follows: pectoralis, posterior deltoid, biceps, and triceps. Here, 0 ms refers to movement onset. Error bars indicate ± 1 SEM.

sensory feedback, the error-feedback response alone cannot fully compensate for the errors.

To solve this problem, theory has suggested that the feedback response may serve as a teaching signal for the brain, resulting in changes in the motor commands that are produced in a “feed-forward” way (Kawato et al., 1987). In support of this hypothesis, several studies have demonstrated that on a trial-to-trial basis, aspects of the feedback response to error appear to be incorporated into the learned response (Thoroughman and Shadmehr, 1999; Franklin et al., 2003, 2008; Milner and Franklin, 2005). Specifically, these studies demonstrated that as the brain learns to compensate for an externally imposed perturbation, the early feedforward component of muscle activity during a reach grows to resemble the feedback-related muscle activity.

In these previous studies, the distinction between feedback and learning responses was made largely on the basis of timing of the response features, rather than separately isolating each com-

ponent for the entirety of the movement trace. Here, we approached the problem by using error-clamp trials before and after perturbation trials. This technique enabled us to more precisely examine the temporal relationship between learning and feedback, isolating the time course of each response during the entire movement. We found that the error-feedback response was a complex temporal pattern of activation/deactivation of each muscle, and included short-latency and long-latency feedback components that corrected for the perturbation and brought the hand to the target. Following experience of a single error, motor commands changed on the following trial. Similar to the feedback response, the learning response possessed complex temporal dynamics specific to each muscle, which persisted during the entire movement. Remarkably, the time course of the learning response appeared to be tightly correlated to the feedback response. In fact, the learning response included essentially all components of the feedback response, scaled by $\sim 25\%$ in magnitude (after two perturbations), and shifted ~ 125 ms earlier in time. Considering that voluntary feedback corrections can be expected at ~ 150 ms after the initiation of movement (Strick, 1978), or perhaps as early as 130 ms in velocity-dependent curl fields (Franklin et al., 2008), this shift appears to account for the delays in the salient voluntary correction component of the feedback response. It is not at this point clear whether the magnitude of this shift is an invariant feature of the nervous system, or is drawn from the relative timing of the feedback controller’s output and the onset of the movement.

One might have predicted that in the initial stages of training (e.g., after the experience of one or two errors), only earlier portions of the feedback response are learned. However, it appeared that the entire feedback response affects the learning response (Figs. 5C, 6A); the fraction of the feedback response that transferred into the learned response was invariant throughout the trajectory of the movement ($\sim 25\%$; Fig. 6B). The consistency in the scaling factor, both across muscles (Fig. 4) and time (Fig. 6B), suggests that there exists some set point of error-feedback sensitivity in muscle space. We speculate that this set point may be intimately related to regulation of error sensitivity (Herzfeld et al., 2014) during motor learning. The amount learned from the experience of error may be in part modulated by the extent to which the feedback response is a reliable corrective signal based on the task history and environment.

The motor commands produced by the sensory feedback system served as a template for the motor learning process, becoming the motor commands that were used to predictively compensate for the novel dynamics on the subsequent attempt. To determine the generality of this claim, we sorted our data within

subject based on the magnitude of the feedback responses. We found that for both agonists and antagonists, when the feedback response was large, so was the learning response (Figs. 7, 8), suggesting that the trial-to-trial variability in the gain of the feedback system strongly affected the trial-to-trial variability in learning.

Our result calls for a re-evaluation of a previous claim that the nervous system produces nonspecific learning responses to single errors (Fine and Thoroughman, 2006; Wei et al., 2010). These authors found that single-trial perturbations with differing dynamics did not induce different kinematic correlates of learning. To reconcile our results with these prior findings, we speculate that the feedback system produced nonspecific, or saturated, responses to the perturbations used in these studies, which would result in identical learning responses according to the scale-and-shift hypothesis. We also note that kinematic similarity does not imply that the control signals in muscle space were identical; rather, sophisticated differences might be present at the neuro-motor level, without corresponding distinct behavioral correlates, as is the case in Figures 7 and 8.

Previous investigators have observed cocontraction of agonist–antagonist pairs of muscles during force-field adaptation, which increases the stiffness of the arm and likely stabilizes it against an unpredicted perturbation (Thoroughman and Shadmehr, 1999; Osu et al., 2002; Franklin et al., 2003; Milner and Franklin, 2005). We speculate that in our task, occasional cocontraction during the error-clamp trial following the perturbation may have resulted in the cancellation of learned antagonist inhibition (e.g., absence of initial inhibition in triceps and posterior deltoid in Fig. 8), thus partially masking the learning instructed by the error-feedback signal. However, on the whole, given that the mean antagonist learning and feedback responses demonstrated clear decreases in muscle activity relative to an error-free reach (Figs. 2, 4), rather than coactivation across agonist–antagonist pairs, we suspect that our task's perturbation infrequency might have partially disengaged this neural impedance controller. This relative absence of cocontraction may also explain a difference between our results and a general learning rule posed by Franklin et al. (2008), where the authors found that antagonists, like their agonist counterparts, increased their activity in response to error. We should note that apart from this difference, our proposed scale-and-shift relationship between learning and feedback is quite similar to a computational architecture proposed by these authors.

Our results describe a correlation and not a causal relationship. However, an earlier work provides some evidence for the idea that the feedback response is causally related to the learning response. Haith et al. (2011) asked subjects to reach in a force

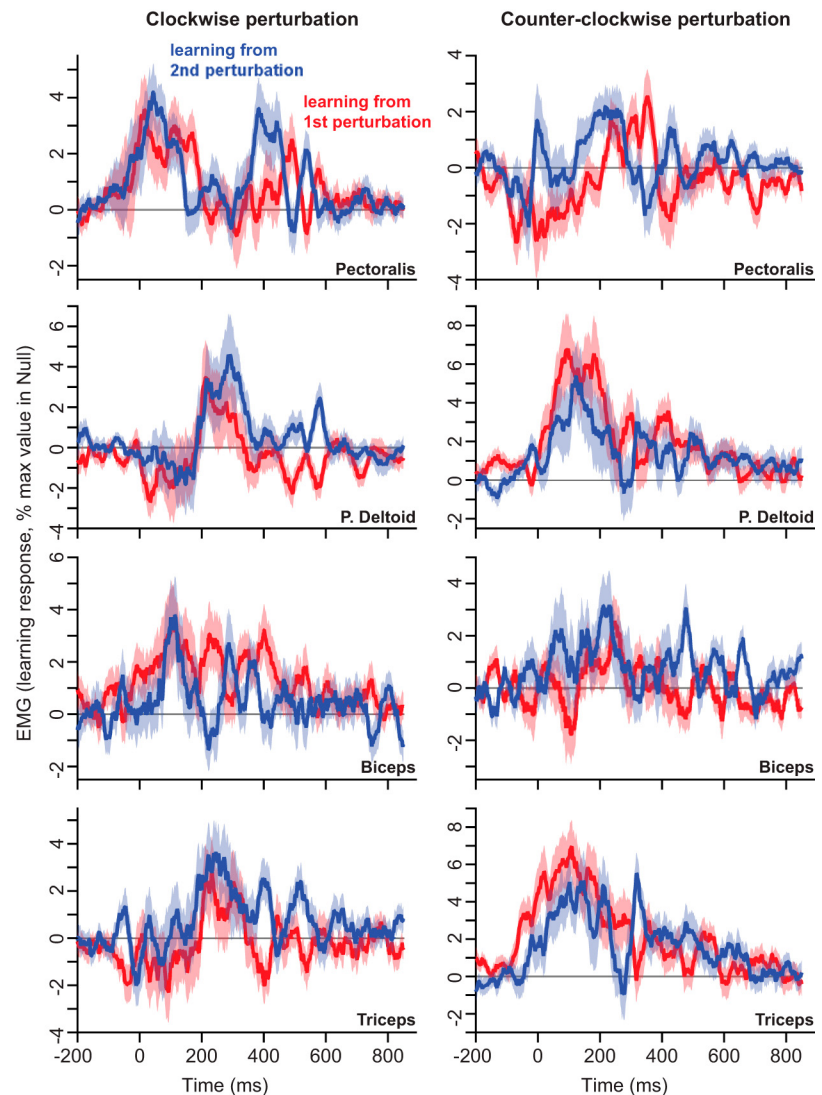


Figure 11. Single-trial learning is similar for first and second perturbations. The EMG learning signal induced by the first perturbation (red, learning from 1st perturbation) is compared with the learning induced by the second perturbation (blue, learning from 2nd perturbation). For the first perturbation, learning was calculated by subtracting the preperturbation error-clamp EMG signal from the postperturbation error-clamp EMG signal. For the second perturbation, learning was calculated by subtracting the postperturbation error-clamp EMG signal after a single perturbation from the postperturbation error-clamp EMG signal after two consecutive perturbations. EMG activities in the CW and CCW fields are shown in the left and right columns, respectively. Each row of figures applies to a particular muscle. From top to bottom these muscles are as follows: pectoralis, posterior deltoid, biceps, and triceps. Here, 0 ms refers to movement onset. Errors bars indicate ± 1 SEM.

field but not produce the voluntary corrective response associated with bringing the hand back to the target. They did this by having people reach to a line during perturbation trials (rather than to a point). They found that the learned response, measured via the forces that subjects produced in error-clamp point-to-point reaching trials, was significantly smaller compared with when the perturbation trials were point to point. This result is consistent with our scale-and-shift error-feedback teaching hypothesis, as the absence of late voluntary corrections in the reach-to-a-line movement would weaken the learning response late in the point-to-point movement.

We observed that people who produced a larger feedback response to error also learned more from the error that they had experienced at the level of muscle control signals (Fig. 5). This predicts that individuals who have reflexes that produce a stronger response to a given perturbation are likely to be able to adapt

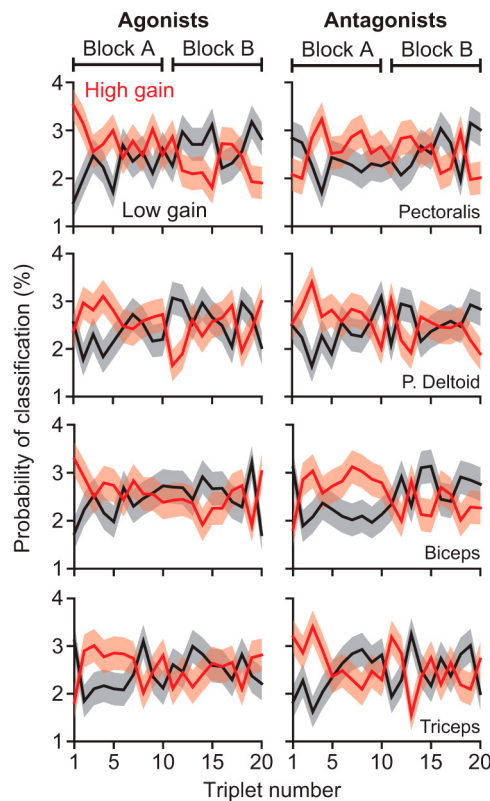


Figure 12. The feedback-response gain varied randomly with the progression of the experiment. Triplet trial progressions were classified into two groups of equal size based on the magnitude of the feedback response, for agonists and antagonists separately. Large (red) and small (black) feedback groups were constructed from individual responses that fell above and below the median response over a critical time window (agonists, 150–450 ms; antagonists, 0–600 ms). Here we show the probability of any particular triplet belonging to the large or small feedback group. The left and right columns show agonist and antagonist responses, respectively. Each figure indicates a different muscle; from top to bottom, pectoralis, posterior deltoid, biceps, and triceps. No clear relationship between the triplet number and its assignment (large or small) existed for any of the muscle responses. Error bars indicate ± 1 SEM.

faster to that perturbation. Therefore, some of the between-subject differences in rates of adaptation in force fields (Wu et al., 2014) may be due to between-subject differences in their ability to correct for sensory prediction errors using reflexive and voluntary feedback pathways. However, because feedback response to error is in itself an adaptive process that benefits from experience (Burdet et al., 2001; Franklin et al., 2003, 2007), we do not know whether people who learn more do so because of inherently better feedback control, or because they are able to better tune their feedback control system to the range of perturbations.

Scaled and shifted feedback responses are not the sole progenitors of motor learning. More likely, feedback instruction is one of many mechanisms in a potpourri of neural motor learning/control strategies. For example, there exist motor learning paradigms where feedback corrections are not required; sensory prediction errors are sufficient to drive motor learning (Tseng et al., 2007). An example of this is saccade adaptation, where visual error detected at completion of the movement is sufficient for modulation of saccadic gain (Wallman and Fuchs, 1998). However, Wallman and Fuchs (1998) found that even in saccade adaptation experiments, if subjects are allowed to correct their saccades with a second saccade that responds to the error, the rate of learning is faster than if this motor correction was not allowed.

Therefore, the act of generating a corrective motor response appears to enhance the process of learning.

In summary, our results demonstrate that the transformation from sensory representation of error to motor representation of commands produced by the feedback system serves as a teacher for the motor learning system. The patterns of muscle activity that comprise the feedback response to error are shifted earlier in time to become the learned response. Individuals who have well-tuned feedback systems that produce a larger feedback response to error, have access to a better teacher, resulting in more learning from a given error.

References

- Ahmadi-Pajouh MA, Towhidkhal F, Shadmehr R (2012) Preparing to reach: selecting an adaptive long-latency feedback controller. *J Neurosci* 32:9537–9545. [CrossRef Medline](#)
- Burdet E, Osu R, Franklin DW, Milner TE, Kawato M (2001) The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 414:446–449. [CrossRef Medline](#)
- Fine MS, Thoroughman KA (2006) Motor adaptation to single force pulses: sensitive to direction but insensitive to within-movement pulse placement and magnitude. *J Neurophysiol* 96:710–720. [CrossRef Medline](#)
- Franklin DW, Osu R, Burdet E, Kawato M, Milner TE (2003) Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. *J Neurophysiol* 90:3270–3282. [CrossRef Medline](#)
- Franklin DW, Liaw G, Milner TE, Osu R, Burdet E, Kawato M (2007) Endpoint stiffness of the arm is directionally tuned to instability in the environment. *J Neurosci* 27:7705–7716. [CrossRef Medline](#)
- Franklin DW, Burdet E, Tee KP, Osu R, Chew CM, Milner TE, Kawato M (2008) CNS learns stable, accurate, and efficient movements using a simple algorithm. *J Neurosci* 28:11165–11173. [CrossRef Medline](#)
- Haith A, Pekny S, Shadmehr R, Krakauer JW (2011) Evidence for model-free learning during force field adaptation. *Advances in Computational Motor Control (ACMC)* 2011.
- Herzfeld DJ, Vaswani PA, Marko MK, Shadmehr R (2014) A memory of errors in sensorimotor learning. *Science* 345:1349–1353. [CrossRef Medline](#)
- Huang VS, Haith A, Mazzoni P, Krakauer JW (2011) Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. *Neuron* 70:787–801. [CrossRef Medline](#)
- Kawato M, Furukawa K, Suzuki R (1987) A hierarchical neural-network model for control and learning of voluntary movement. *Biol Cybern* 57:169–185. [CrossRef Medline](#)
- Milner TE, Franklin DW (2005) Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. *J Physiol* 567:651–664. [CrossRef Medline](#)
- Osu R, Franklin DW, Kato H, Gomi H, Domen K, Yoshioka T, Kawato M (2002) Short- and long-term changes in joint co-contraction associated with motor learning as revealed from surface EMG. *J Neurophysiol* 88:991–1004. [Medline](#)
- Scheidt RA, Reinkensmeyer DJ, Conditt MA, Rymer WZ, Mussa-Ivaldi FA (2000) Persistence of motor adaptation during constrained, multi-joint, arm movements. *J Neurophysiol* 84:853–862. [Medline](#)
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208–3224. [Medline](#)
- Strick PL (1978) Cerebellar involvement in volitional muscle responses to load change. In: *Cerebral motor control in man: long loop mechanisms* (Desmedt JE, ed), pp 85–93. Basel: Karger.
- Thoroughman KA, Shadmehr R (1999) Electromyographic correlates of learning internal models of reaching movements. *J Neurosci* 19:8573–8588. [Medline](#)
- Tseng YW, Diedrichsen J, Krakauer JW, Shadmehr R, Bastian AJ (2007) Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J Neurophysiol* 98:54–62. [CrossRef Medline](#)
- Wallman J, Fuchs AF (1998) Saccadic gain modification: visual error drives motor adaptation. *J Neurophysiol* 80:2405–2416. [Medline](#)
- Wei K, Wert D, Körding K (2010) The nervous system uses nonspecific motor learning in response to random perturbations of varying nature. *J Neurophysiol* 104:3053–3063. [CrossRef Medline](#)
- Wu HG, Miyamoto YR, Gonzalez Castro LN, Ölveczky BP, Smith MA (2014) Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat Neurosci* 17:312–321. [CrossRef Medline](#)