

# Environmental Consistency Determines the Rate of Motor Adaptation

Luis Nicolas Gonzalez Castro,<sup>1,2,3,5</sup> Alkis M. Hadjiosif,<sup>1,5</sup> Matthew A. Hemphill,<sup>1</sup> and Maurice A. Smith<sup>1,4,\*</sup>

<sup>1</sup>Neuromotor Control Laboratory, Harvard School of Engineering and Applied Sciences, Cambridge, MA 02138, USA

<sup>2</sup>Harvard-MIT Division of Health Sciences and Technology, Cambridge, MA 02138, USA

<sup>3</sup>Harvard Medical School, Boston, MA 02115, USA

<sup>4</sup>Center for Brain Science, Harvard University, Cambridge, MA 02138, USA

## Summary

**Background:** The motor system has the remarkable ability not only to learn but also to learn how fast it should learn. However, the mechanisms behind this ability are not well understood. Previous studies have posited that the rate of adaptation in a given environment is determined by Bayesian sensorimotor integration based on the amount of variability in the state of the environment. However, experimental results have failed to support several predictions of this theory.

**Results:** We show that the rate at which the motor system adapts to changes in the environment is primarily determined not by the degree to which environmental change occurs but by the degree to which the changes that do occur persist from one movement to the next, i.e., the consistency of the environment. We demonstrate a striking double dissociation whereby feedback response strength is predicted by environmental variability rather than consistency, whereas adaptation rate is predicted by environmental consistency rather than variability. We proceed to elucidate the role of stimulus repetition in speeding up adaptation and find that repetition can greatly potentiate the effect of consistency, although unlike consistency, repetition alone does not increase adaptation rate. By leveraging this understanding, we demonstrate that the rate of motor adaptation can be modulated over a range that encompasses a 20-fold increase from lowest to highest.

**Conclusions:** Understanding the mechanisms that determine the rate of motor adaptation could lead to the principled design of improved procedures for motor training and rehabilitation. Regimens designed to control environmental consistency and repetition during training might yield faster, more robust motor learning.

## Introduction

The human motor system has the remarkable ability to not only adapt its output to reduce motor errors but also to adapt the rate at which this adaptation occurs [1–5]. However, the mechanisms by which adaptation rates change are still unclear.

Previous studies that have examined this phenomenon have posited that adaptation rates are determined by optimal estimation based on the sensory information available to guide

learning [1, 2, 5, 6]. The idea is that the adjustments that occur in motor output on the basis of noisy sensory information about the environment during motor adaptation represent an ongoing assessment of the motor system's belief about the state of the environment that was most recently experienced. Accordingly, Bayesian inference, in which the relative amplitudes of different types of noise determine how optimal estimates are made, has been suggested as a framework for understanding modulations in learning rate. However, several predictions of this theory have not been borne out experimentally [1, 2, 6]. Here we suggest that motor adaptation rates are primarily determined not by estimation of the state of the environment most recently experienced but by prediction, based on the tendency of changes in the environment to persist from one movement to the next, of the state most likely to be experienced during the next movement.

The contrast between prediction and estimation is made clear in the Kalman filter, a statistically optimal model widely used in the analysis of linear systems [7]. This model uses a two-step process for incorporating new information: estimation and prediction. First, in the *estimation step*, one computes a gain factor  $K$  (the Kalman gain) by using Bayesian inference to determine the statistically optimal weighting for updating the estimate of the current state on the basis of new sensory information including the relative levels of state noise  $\sigma_x$  and sensory input noise  $\sigma_u$ . In the equations below,  $x(n^-)$  and  $x(n^+)$  signify the state estimate before and after this update, respectively, and the error,  $e(n)$ , is the difference between new sensory information and the value predicted for it,  $x(n^-)$ . Second, in the *prediction step*, one multiplies a prediction factor  $A$  (the state transition gain) that models how the system evolves or decays from one state to the next by the result of the estimation step to make a prediction about the next state on the basis of the current estimate.

*Estimation:*  $x(n^+) = x(n^-) + Ke(n)$ , where  $K = \sigma_x^2 / [\sigma_u^2 + \sigma_x^2]$

*Prediction:*  $x([n+1]^-) = Ax(n^+)$

Putting these two steps together yields a form that closely resembles a learning rule commonly hypothesized to model motor adaptation [8–12]:

*Estimation and prediction step:*

$$x([n+1]^-) = Ax(n^-) + AKe(n)$$

*Trial-to-trial learning rule for motor adaptation:*

$$x(n+1) = Ax(n) + Be(n)$$

Here,  $A$  is a retention or prediction factor, and  $B$  is the apparent trial-to-trial learning rate. Thus, for a linear system that is capable of estimation and prediction and that is statistically optimal in the face of noise, the apparent, experimentally measurable learning rate ( $B$ ) should be determined by a combination of a next-trial prediction factor ( $A$ ) that characterizes environmental persistence and a same-trial Bayesian gain ( $K$ ) for estimating environmental change on the basis of new sensory information ( $B = AK$ ).

<sup>5</sup>Co-first authors

\*Correspondence: [mas@seas.harvard.edu](mailto:mas@seas.harvard.edu)

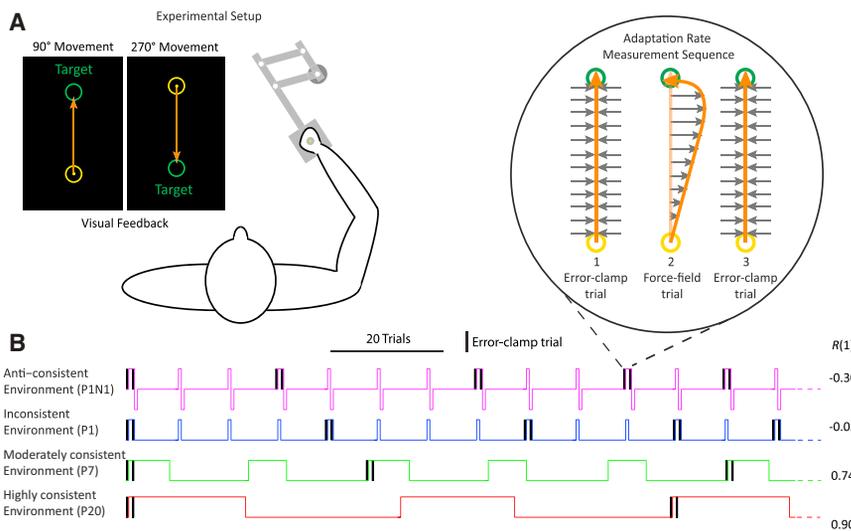


Figure 1. Experimental Paradigm

(A) Schematic of experimental setup. Subjects gripped the handle of a 2D robotic manipulandum and performed 10 cm reaching arm movements in two directions (90° and 270°) on the basis of targets presented on a vertical screen in front of them.

(B) Illustration of force field (FF) environments and adaptation-rate measurement sequence. Example sections of FF activation patterns for the anti-consistent, inconsistent, moderately consistent, and highly consistent learning environments are shown. In a subset of the FF cycles of each environment, we introduced error-clamp (EC) trials before and after the first FF trial of the cycle (vertical black bars and blowout area). The adaptation rate was calculated as the difference in FF compensation between the post- and pre-FF EC trials (see Equation 1 in the Experimental Procedures). The statistical consistency,  $R(1)$ , of each environment is shown on the right.

Adaptation rates are also known to be increased when the exposure to a particular perturbation is repeated, a phenomenon termed savings, and recent work has suggested that action repetition is what leads to the improved learning when savings is observed [13]. Thus, here we examine how motor adaptation rates are modulated by Bayesian estimation gain ( $K$ ), next-trial prediction ( $A$ ), and repetition. Although it is unclear whether a statistically optimal linear model such as the Kalman filter can provide insight into the mechanisms that determine how quickly the human motor system learns, because this learning is likely to be neither strictly linear nor statistically optimal, here we suggest that it can. In particular, we hypothesize that the motor system makes internal environmental-persistence estimates that are highly adaptable and readily modulate the rate at which motor learning proceeds. Accordingly, we demonstrate that experimental manipulation of the extent to which environmental change persists from one movement to the next can upregulate the rate of motor adaptation by 3-fold and downregulate it by 5-fold. In experiments where the persistence and the variability of environmental perturbations were systematically manipulated, we demonstrate a striking double dissociation supporting the hypothesis that prediction is fundamentally different from estimation and is modulated to a greater degree during motor learning. We find that the next-trial motor adaptation rate is predicted by environmental persistence but not variability, whereas the gain in strength of the same-trial feedback response is predicted by environmental variability but not persistence. Moreover, in experiments where persistence and repetition were systematically manipulated, we find that persistence, measured by environmental consistency defined as the lag-1 autocorrelation of the environment, alone can modulate adaptation rates, whereas repetition alone cannot. However, the combination of persistence and repetition greatly potentiates the effect of persistence alone, leading to even greater adaptation.

## Results

We began by systematically manipulating the persistence of the physical environment for movement by cycling different patterns of velocity-dependent force-field (FF) perturbations as illustrated in Figure 1B. After performing 200 baseline

reaching movements without perturbation (null trials), subjects were exposed to one of four different environments (anti-consistent, inconsistent, moderately consistent, and highly consistent environments) in which FF perturbation patterns with varying persistence were experienced. We operationally defined consistency as a statistical persistence measure based on the correlation coefficient between the FF strength in the current trial and that in the previous trial [the lag-1 autocorrelation,  $R(1)$ ]. In all experiments, subjects performed active point-to-point reaching movements in the horizontal plane in two alternating movement directions in the midline, toward and away from the chest, while holding the handle of a robotic manipulandum. Individuals experienced the environments in both movement directions, and, adaptation rates were measured for both; the results were then combined. However, the trial numbers we refer to are within each direction.

During the anti-consistent environment (P1N1), a single positive FF trial (P1) was followed by a single negative FF trial (N1) and then by 11–13 washout (null) trials. Here the positive-negative subsequence led to a negative lag-1 autocorrelation [ $R(1) = -0.30$ ]. During the inconsistent environment (P1), subjects experienced a single positive FF trial followed by 10–12 washout trials, leading to a near-zero autocorrelation [ $R(1) = -0.05$ ]. Finally, during moderately and highly consistent environments [P7,  $R(1) = 0.74$ ; and P20,  $R(1) = 0.90$ ], subjects experienced blocks of 7 or 20 FF trials, followed by 15–18 or 28–32 washout trials, respectively. The lag-1 autocorrelation for these last two perturbation sequences is high because most perturbations predict subsequent ones, and it increases as the block length increases because the self-similarity of a sequence after a single-trial shift is higher for longer block lengths.

We designed the experiments so that we could directly compare the adaptation rates induced by each environment. To accomplish this, we estimated the single-trial adaptation rate induced by the first FF trial of each cycle by placing error-clamp (EC) measurement trials (black bars in Figure 1B) before and after the FF trial in order to determine the change in motor output associated with the FF exposure (blowout in Figure 1B) [14, 15]. We referred to this three-trial sequence as a measurement triplet and used the difference between the lateral force recorded on the two EC trials in each measurement triplet to measure the change in feedforward motor output induced by exposure to the intervening FF trial. This

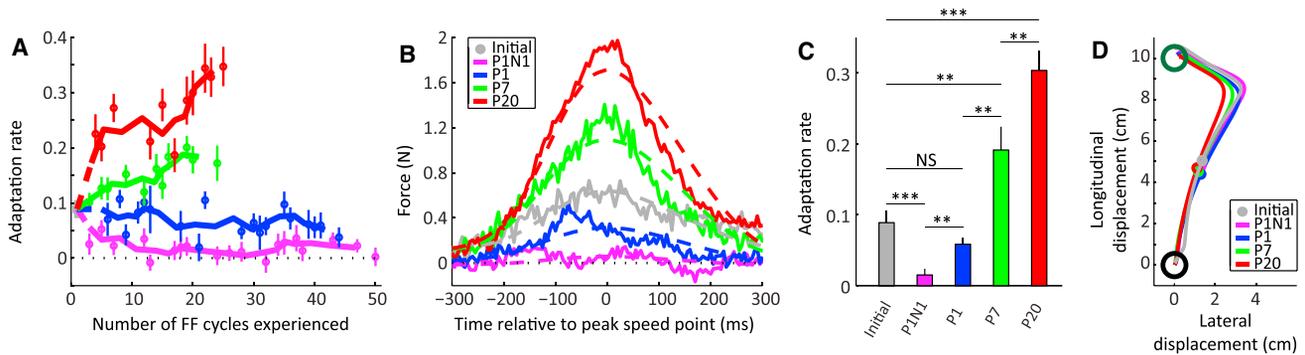


Figure 2. Upregulation and Downregulation of Motor Adaptation Rates by Environmental Consistency

(A) Average initial adaptation rate for the P1N1, P1, P7, and P20 environments and evolution of adaptation rate in these environments as a function of the number of FF cycles experienced. Circles indicate actual data from FF cycles where the adaptation rates were measured, and solid lines show a 3 point moving average. Error bars indicate SEM across all panels.

(B) Lateral-force profiles comprising the single-trial adaptive response for the different learning environments. The average learning-related change in lateral force across experiments is shown for the first cycle of the experiment (gray) and is compared to the average FF compensation observed in the last half of the measurement FF cycles in each different environment (colors).

(C) The average initial adaptation rate for all the different learning environments (gray) is compared to the last-half learning rates for each environment (colors). Significance levels are as follows: \* $p < 0.05$ , \*\* $p < 0.01$ , and \*\*\* $p < 0.001$ .

(D) Average perturbed hand paths seen during the first FF trial of the experiment (gray) and in the last half of the cycles in each environment (colors). The small circles near the midway point indicate the location of the peak speed point.

allowed us to estimate the single-trial adaptation rate identically for all environments by comparing the observed change in force to the change that would be necessary to fully compensate for the FF perturbation.

### Environmental Consistency Can Upregulate and Downregulate Motor Adaptation Rates

We found that all groups displayed similar adaptation rates on the initial FF trial (one-way ANOVA,  $F(3,66) = 0.87$ ,  $p = 0.46$ ;  $0.089 \pm 0.017$ , mean  $\pm$  SEM for all groups). However, subsequent exposure to different levels of environmental consistency resulted in single-trial adaptation rates that diverged over a range that encompasses a 20-fold increase from lowest to highest, as shown in Figure 2A. High environmental consistency resulted in increased adaptation rates, whereas low consistency resulted in decreases, as evidenced by the changes in lateral-force-output patterns associated with single-trial adaptation in the last half of the training period for each group (Figure 2B) and adaptation rates estimated from these changes in force (Figure 2C). Correspondingly, we found that the adaptation rates were significantly different between groups in the last half of the training period [one-way ANOVA,  $F(3,69) = 36.46$ ,  $p = 3.1 \times 10^{-14}$ ]. The highly consistent (P20) and moderately consistent (P7) groups both displayed significantly increased adaptation rates ( $0.303 \pm 0.028$ ,  $p = 1.5 \times 10^{-8}$  and  $0.191 \pm 0.033$ ,  $p = 0.0067$ , respectively) relative to the initial rate ( $0.089 \pm 0.017$ ), and the P20 group displayed significantly faster adaptation than the P7 group ( $p = 0.0075$ ). In contrast, the group exposed to the inconsistent environment (P1) displayed a small, nonsignificant decrease in adaptation rate ( $0.058 \pm 0.010$ ,  $p = 0.06$ ), and the anti-consistent group (P1N1) displayed a significantly decreased adaptation rate ( $0.015 \pm 0.009$ ) in comparison to both the initial rate ( $p = 0.0001$ ) and the rate displayed by the P1 group ( $p = 0.0014$ ). Because we randomized the duration of the washout period, the first perturbation trial in each cycle was somewhat unpredictable, and subjects showed no evidence of the ability to take advantage of the limited predictability that was present (see Figure S2 in the Supplemental Information available with

this article online). However, the second (negative) perturbation in P1N1 is perfectly predictable because it always follows the first, making the P1N1 environment inherently more predictable than P1, albeit in a negative manner, in line with its negative environmental consistency [ $R(1) = -0.30$ ]. Nevertheless, we observed a near-zero adaptation rate as opposed to a negative one in the P1N1 data, suggesting that the motor system is unable to make good use of this negative predictability, at least over the time course of our training paradigm. Overall, the data demonstrate that highly consistent environments upregulate adaptation rates by 3-fold, whereas anti-consistent environments downregulate adaptation rates by 5-fold. Together, these results indicate that adaptation rates can be modulated by more than an order of magnitude over the course of just a few minutes by changes in environmental consistency.

### Comparing Next-Trial Adaptive Responses with Same-Trial Feedback Responses

In order to determine whether the adaptation rate changes that we observe are due to changes in prediction versus estimation, we compared next-trial and same-trial perturbation responses. Our hypothesis was that a change in the prediction about the persistence of perturbations from one trial to the next would affect the next-trial adaptive response but not the same-trial feedback response. However, a change in the estimation of the size of perturbations would affect both the same-trial feedback response and the next-trial adaptive response.

An initial analysis based on data from the P1N1, P1, P7, and P20 environments yielded intriguing but inconclusive results, as detailed with statistics in the Supplemental Information. In short, we used the motor adaptation rate measured during the trial after each perturbation as a measure of next-trial response and the reduction in lateral error that was observed during the perturbation trial itself (with respect to the error observed on the very first perturbation trial) as a measure of the same-trial response—i.e., the feedback response strength. However, the results were unclear. We found that more highly

consistent environments (P7 and P20) that displayed increases in next-trial motor adaptation rates also showed increases in the strength of same-trial feedback responses. But the P7 and P20 environments also had the highest variability, obscuring whether consistency, variability, or both were responsible for the changes in same-trial and next-trial responses. Thus, variability and consistency were highly correlated across these four environments ( $r = 0.82$ ), preventing them from being dissociated. Interestingly, next-trial responses were somewhat more tightly coupled to changes in environmental consistency ( $R^2 = 0.94$ ) than to changes in variability ( $R^2 = 0.69$ ), and same-trial responses were more tightly coupled to changes in environmental variability ( $R^2 = 0.97$ ) than to changes in consistency ( $R^2 = 0.79$ ), hinting but only hinting at a double dissociation.

#### Next-Trial Adaptive Responses Are Driven by Environmental Consistency, Whereas Same-Trial Feedback Responses Are Driven by Environmental Variability

To clearly dissociate the effects of environmental consistency and variability on adaptation rate and feedback response strength, we designed two new environments to have high variability but low consistency. In the first, termed “random noise” (RN), the FF varied randomly from one trial to the next (see Figure 3A). The second environment, P1-long (P1L), was similar to P1 but had 2–4 rather than 10–12 null washout trials between the isolated FF trials in order to increase their frequency and thus increase environmental variability. RN’s variability ( $55.4 \text{ N}^2\text{s}^2/\text{m}^2$ ) was greater than that of all the other environments, whereas its consistency was essentially zero [ $R(1) = 0.02$ ] as a result of its random nature. Similarly, P1L’s variability was high ( $46.6 \text{ N}^2\text{s}^2/\text{m}^2$ ), whereas its consistency was low [ $R(1) = -0.45$ ]. Correspondingly, the inclusion of RN and P1L allows consistency and variability to be dissociated across the environments we studied by breaking the strong positive correlation between them; their correlation dropped from 0.82 to 0.38 ( $R^2$  of 0.67 versus 0.14).

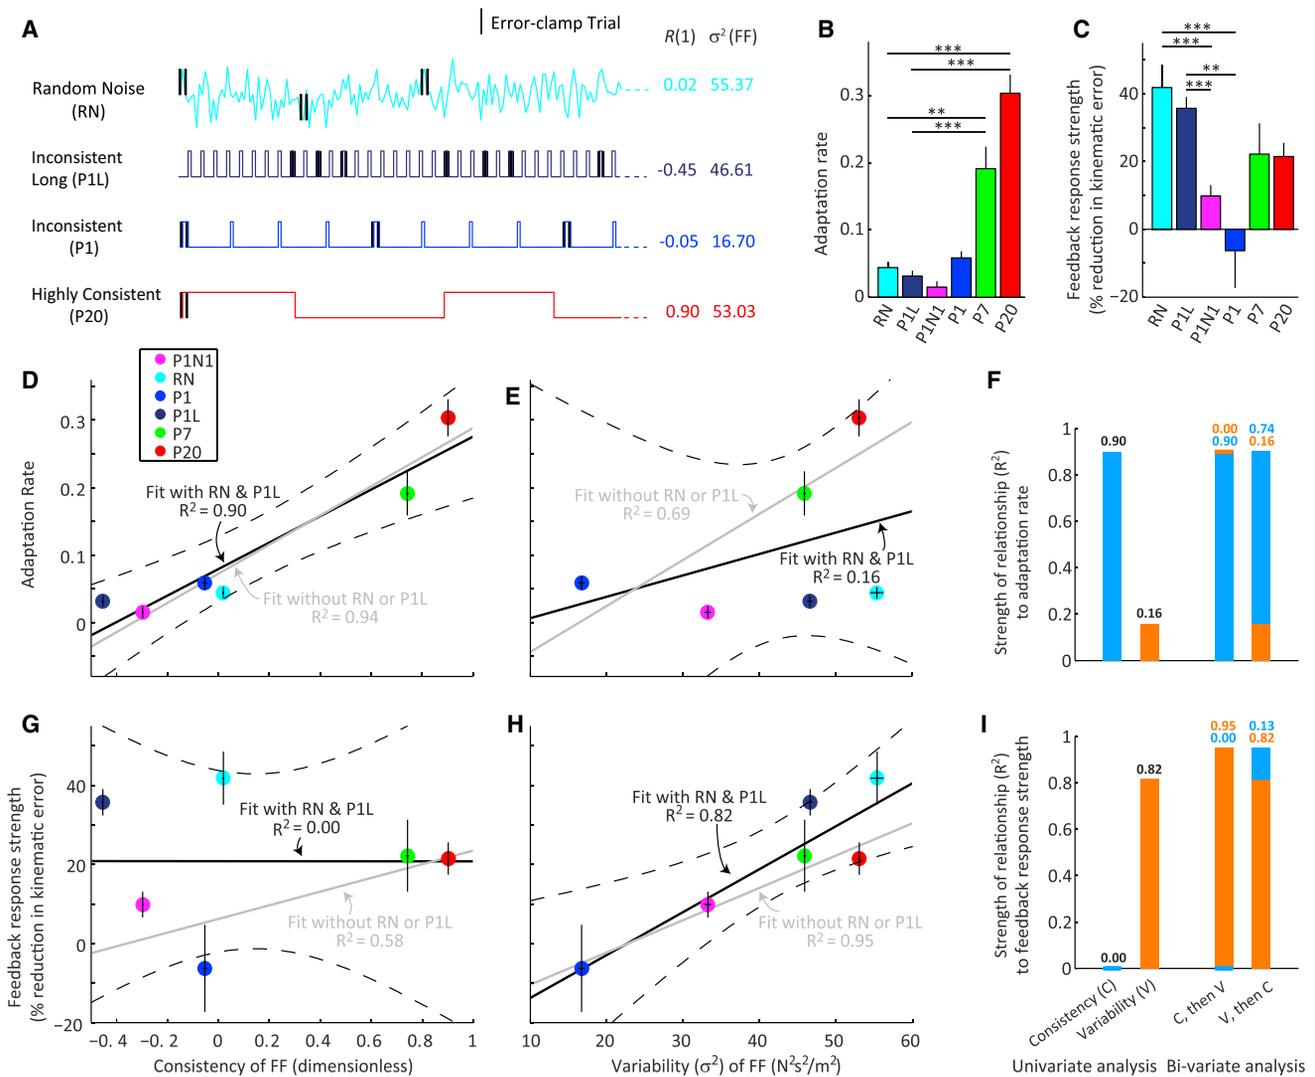
If adaptation rate depends on environmental variance, the RN environment should elicit adaptation-rate increases even greater than those seen in the P7 and P20 environments, and P1L should elicit adaptation-rate increases similar to those seen in P7. On the other hand, if adaptation rate depends on environmental consistency, neither RN nor P1L should elicit adaptation-rate increases. Correspondingly, the feedback response strength elicited by the P1L and RN environments should be high if it depends on environmental variability, but it should be low if it depends on consistency.

We found that despite substantially increased variability ( $55.4$  versus  $16.7 \text{ N}^2\text{s}^2/\text{m}^2$  for RN versus P1 and  $46.6$  versus  $33.2 \text{ N}^2\text{s}^2/\text{m}^2$  for P1L versus P1N1; see Figures 3B and 3D–3E), the RN and P1L environments, like P1 and P1N1, elicited low adaptation rates ( $0.044 \pm 0.008$  versus  $0.058 \pm 0.010$ ,  $p = 0.86$  for RN versus P1 and  $0.031 \pm 0.008$  versus  $0.015 \pm 0.009$ ,  $p = 0.09$  for P1L versus P1N1). However, this is in line with similar consistency ( $0.02$  versus  $-0.05$  for RN versus P1 and  $-0.45$  versus  $-0.30$  for P1L versus P1N1). In contrast, the RN and P1L environments displayed significantly lower adaptation rates ( $p < 0.001$  for all pairwise comparisons) than both the P7 and P20 environments ( $0.191 \pm 0.033$  and  $0.303 \pm 0.028$ ), in accordance with substantially lower consistency ( $0.02$  and  $-0.45$  for RN and P1L versus  $0.74$  and  $0.90$  for P7 and P20, despite similar variability ( $55.4$  and  $46.6$  versus  $45.9$  and  $53.0 \text{ N}^2\text{s}^2/\text{m}^2$ ). These results provide compelling

evidence for the idea that environmental consistency rather than variability determines adaptation rates. Correspondingly, the  $R^2$  value for the relationship between adaptation rate and environmental consistency remains essentially unchanged at 90% when the RN and P1L data are included [ $F(1,4) = 35.32$ ,  $p = 0.004$ ; see Figure 3D], and similar results are observed when the consistency of the kinematic error rather than the consistency of the FF environment is considered (see Figure S3). However, the  $R^2$  value between adaptation rate and variability drops precipitously from 69% to 16% [ $F(1,4) = 0.76$ ,  $p = 0.43$ ] when we include the RN and P1L experiment data as illustrated in Figure 3E. Moreover, a bivariate regression of consistency and variability onto adaptation rate reveals a significant effect of consistency but not variability [ $F(1,3) = 22.3$ ,  $p = 0.018$  for consistency and  $F(1,3) = 0.068$ ,  $p = 0.81$  for variability]. This bivariate regression accounts for 90% of the variance in the full data set, a 0% improvement over consistency alone but a 74% improvement over variability alone, as shown in Figure 3F. See Table S2 for the regression coefficients and additional detail.

When we instead examined changes in feedback response strength, we found that RN and P1L elicited levels ( $41.8\% \pm 6.7\%$  and  $35.7\% \pm 3.4\%$ ) that were far greater than those of P1 and P1N1 ( $-6.3\% \pm 11.0\%$  and  $9.8\% \pm 3.2\%$ ,  $p < 0.0015$  for all pairwise comparisons; see Figures 3C and 3G–3H) and more like those observed in the P7 and P20 environments ( $22.2\% \pm 9.1\%$  and  $21.5\% \pm 4.1\%$ , respectively). This is in line with levels of environmental variability in RN and P1L ( $55.4$  and  $46.6 \text{ N}^2\text{s}^2/\text{m}^2$ , respectively), which were substantially higher than environmental variability in P1 and P1N1 ( $16.7$  and  $33.3 \text{ N}^2\text{s}^2/\text{m}^2$ , respectively) but similar to that in P7 and P20 ( $45.9$  and  $53.0 \text{ N}^2\text{s}^2/\text{m}^2$ , respectively). However, this result is at odds with similar consistency for RN and P1 ( $0.02$  versus  $-0.05$ ) and for P1L and P1N1 ( $-0.45$  versus  $-0.30$ ) but reduced consistency for RN and P1L compared to P7 and P20 ( $0.02$  and  $-0.45$  versus  $+0.74$  and  $+0.90$ ), suggesting that, unlike next-trial adaptation, same-trial feedback response strength depends on environmental variability rather than consistency. Correspondingly, the  $R^2$  value for the relationship between feedback response strength and environmental variability remains high at 82% when the RN and P1L data are included [ $F(1,4) = 17.88$ ,  $p = 0.013$ , see Figure 3H]. However, the  $R^2$  value between feedback response strength and environmental consistency drops sharply from 58% to 0% [ $F(1,4) = 0.00$ ,  $p = 0.99$ ] when we include the RN and P1L data, as illustrated in Figure 3G. Moreover, a bivariate regression of consistency and variability onto feedback response strength reveals a significant effect of variability but not of consistency [ $F(1,3) = 59.03$ ,  $p = 0.0046$  for variability and  $F(1,3) = 8.34$ ,  $p = 0.063$  for consistency]. This bivariate regression accounts for 95% of the variance in the full data set, a 95% improvement over consistency alone but only a 13% improvement over variability alone, as shown in Figure 3I. See Table S3 for the regression coefficients and additional detail.

Taken together, these findings indicate a striking double dissociation, in which the motor adaptation rate is determined by environmental consistency rather than variability but the feedback response strength is determined by environmental variability rather than consistency (Figures 3F and 3I). This finding is in line with a recent study by Yousif and Diedrichsen [16], who showed that adaptive changes in feedback responses could be observed in both consistent and inconsistent environments, whereas changes in feedforward adaptation were only present in consistent environments. The



**Figure 3. Environmental Consistency Modulates Next-Trial Motor Adaptation Rates but not Same-Trial Feedback Control**

(A) Schematic depiction of learning environments, including the random noise (RN) environment. Note that the RN environment was designed to have high variability as in the highly consistent environment (P20) but at the same time a near zero consistency as in the inconsistent learning environment (P1). (B and C) Adaptation rate and feedback response strength (percent reduction in kinematic error) in the last half of the P1N1, P1, RN, P7, and P20 environments. Statistical comparisons between RN and the other environments are shown ( $*p < 0.05$ ,  $**p < 0.01$ , and  $***p < 0.001$ ). Error bars indicate SEM in all panels. (D and E) The relationships between adaptation rate and environmental consistency (D) or variability (E) across experiments. Both before (gray solid line) and after inclusion of the RN and P1L data (black solid line), environmental consistency explains a large fraction of the variance in motor adaptation rate:  $R^2 > 89\%$  in both cases. However, inclusion of the RN and P1L data reduces the ability of environmental variability to explain changes in adaptation rate; it reduces the  $R^2$  value from 69% to only 16%. The gray dotted horizontal line indicates the initial adaptation rate before exposure to the different environments. (G and H) The relationship between feedback response strength and environmental consistency (G) or variability (H) across experiments in a format similar to that in (D) and (E). Both before and after inclusion of the RN data, environmental variability explains a large fraction of the variance in feedback response strength:  $R^2 = 82\%$  when the RN and P1L data are included and  $R^2 = 95\%$  without these data. However, inclusion of the RN and P1L data reduces the ability of environmental consistency to explain changes in feedback response strength; it reduces the  $R^2$  value from 58% to 0%. (F and I) Summary of the strength of the relationships ( $R^2$ ) between adaptation rate (F) or feedback response strength (I) and consistency or variability. The first pair of bars summarizes the univariate regression analyses shown in (D), (E), (G), and (H). The third and fourth bars summarize the corresponding bivariate regressions. Here the third bar shows the improvement in  $R^2$  when a univariate analysis based on consistency (C, blue) is augmented by variability (V, orange). The fourth bar shows the improvement when a univariate analysis based on V is augmented by C. Full results of the bivariate analysis are shown in Tables S2 and S3.

positive relationship between environmental variability and feedback response strength is in line with several previous studies [17–21] and consistent with the idea that the variance of an environment informs the variance of the prior expectation for Bayesian integration [22–25]. Moreover, the extremely tight coupling we find between feedback response strength and

environmental variability indicates that the weak relationships we observe (1) between feedback response strength and environmental consistency and (2) between feedforward adaptation rate and environmental variability cannot simply result from the imprecise or noisy characterization of feedback response strength or environmental variability. The specific

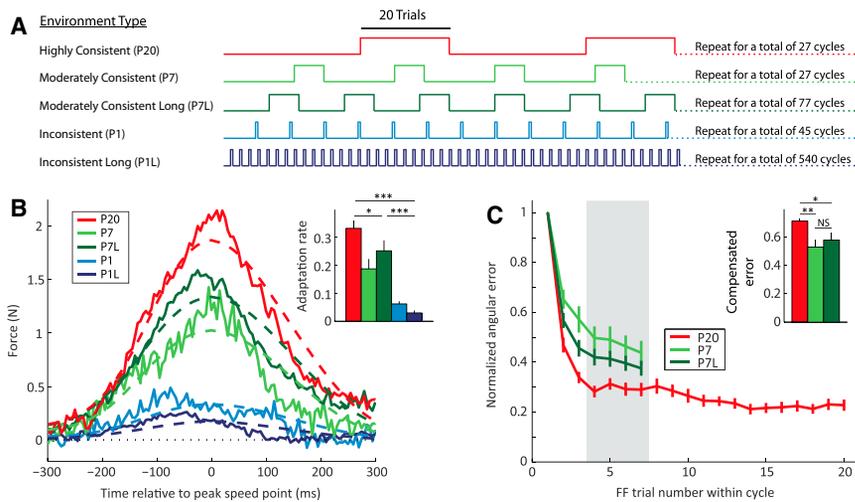


Figure 4. Increased Repetition Does Not Result in Increased Adaptation Rates in Blocked Environments

(A) Comparison of the P20 (red), P7 (light green), P7L (dark green), P1 (light blue), and P1L (dark blue) environments. Notice that although the number of FF cycles is the same (27) for both the P20 and P7 learning environments, the number of FF trials experienced is different, 540 versus 189. Correspondingly, the number of FF cycles for the P1 learning environment is 45 (i.e., only 45 FF trials). The P7L and P1L environments were designed to match the number of FF trials of the P20 environment (540) and thus assess the influence of the number of FF trials on the observed adaptation-rate increases.

(B) Lateral force profiles comprising the single-trial adaptive response in the P20, P7, P7L, P1, and P1L environments. Note that the subjects in the P7L environment, despite experiencing the same number of FF trials as those in the P20 environment, compensated less than those exposed

to the P20 environment. Subjects exposed to the P1L environment compensated even less. The inset shows the average adaptation rates in the P20, P7, P7L, P1, and P1L experiments in the last third of each environment. Error bars indicate SEM. Notice that subjects in the P20 environment exhibited adaptation rates that were significantly greater than those of subjects in all other environments.

(C) Mean angular error, a second measure of motor adaptation, in late FF cycles in the P20, P7, and P7L environments. The error in trials 4–7 (gray shaded region) of the FF cycles was significantly lower in the P20 environment than in the P7 or P7L environments. The inset quantifies these differences. Error bars indicate SEM. Significance rates are as follows: \* $p < 0.05$ , \*\* $p < 0.01$ , and \*\*\* $p < 0.001$ .

effect of environmental consistency on adaptation rate suggests that consistency drives changes in the motor system's predictions about the persistence of environmental perturbations from one trial to the next, which would affect next-trial adaptation, without driving changes in the motor system's estimation of perturbation size, which would affect same-trial feedback responses.

### Increases in Motor Adaptation Rate Are Not Explained by the Amount of Environmental Exposure

We next examined whether the improved adaptation levels we observed could be attributed to savings, which refers to the faster relearning of a previously learned task [10, 26–29]. Although our P7 and P20 experiments presented the same FF repeatedly, the durations of individual FF blocks were rather short (20 trials at most) in comparison to the paradigms in which savings has been demonstrated for motor adaptation [13, 26, 29], and the partial adaptation achieved in our experiments as a result of the relatively short FF exposures precludes the reinforcement of successful performance that has been suggested to be a key ingredient for the instantiation of savings [13]. In fact, subjects only learn 50%–70% of the force field by the last trial of the training blocks in the P7 and P20 environments where adaptation rate increases were observed (see Figure 4C). Correspondingly, we thought it unlikely that savings could be responsible for the faster adaptation observed in the P7 and P20 experiments. Nevertheless, we investigated two experimental manipulations to carefully examine the possibility.

Because the amount of savings should depend on the amount of training, we examined whether the total duration of the exposure to the FF environment could explain the observed increases in adaptation rates, as would be predicted by a savings mechanism. This possibility is compatible with the finding that the P20 environment elicited the largest adaptation rates because P20 exposed subjects to a greater number of FF trials than subjects in the other groups were exposed to. This was the case because the P1 and P7 experiments were designed to match or exceed the P20 experiment in the

number of cycles (45, 27, and 27, respectively) rather than FF trials (45, 189, and 540, respectively). We thus examined versions of the moderately consistent- and inconsistent environment experiments, P7-long (P7L) and P1-long (P1L), that were lengthened in order to match the number of FF trials experienced in the P20 environment (540). We increased the number of cycles from 27 to 77 for the P7L environment and from 45 to 540 for the P1L environment and reduced the number of washout trials somewhat to avoid experiment durations greater than 4 hr (see Figure 4A and Experimental Procedures).

The P1L and P7L environments both displayed significantly lower adaptation rates than P20 ( $0.030 \pm 0.009$  and  $0.252 \pm 0.038$  versus  $0.333 \pm 0.027$ ; mean  $\pm$  SEM for the last third of trials for P1L, P7L, and P20, respectively;  $p < 10^{-8}$  and  $p = 0.046$  for comparisons against P20). P1L also showed a lower adaptation rate than P1 ( $0.030 \pm 0.009$  versus  $0.062 \pm 0.010$ ,  $p = 0.011$ ), in line with even lower environmental consistency [R(1) of  $-0.45$  for P1L versus R(1) of  $-0.05$  for P1] but in contrast to increased training (540 cycles for P1L versus 45 cycles for P1). The kinematic-error learning curves displayed in Figure 4C provide additional evidence for an increased adaptation rate in the P20 group compared to P7 and P7L. Comparison of trials 4 through 7 in these normalized learning curves (the shaded region) indicates a significantly greater reduction in kinematic error in the P20 group than in P7L ( $p = 0.011$ ). These findings indicate that the P20 environment elicits greater adaptation rates than the inconsistent and moderately consistent environments when experiments were equated on the basis of either the number of cycles (P20 versus P1 or P7) or the number of trials (P20 versus P1L or P7L). Thus, the amount of exposure cannot account for the adaptation-rate differences observed between the inconsistent, moderately consistent, and highly consistent environments, as would be predicted by a savings mechanism.

### Specificity of Adaptation-Rate Increases for the Experienced Environment Dynamics

Although we found that differences in environmental consistency rather than differences in the amount of repetition

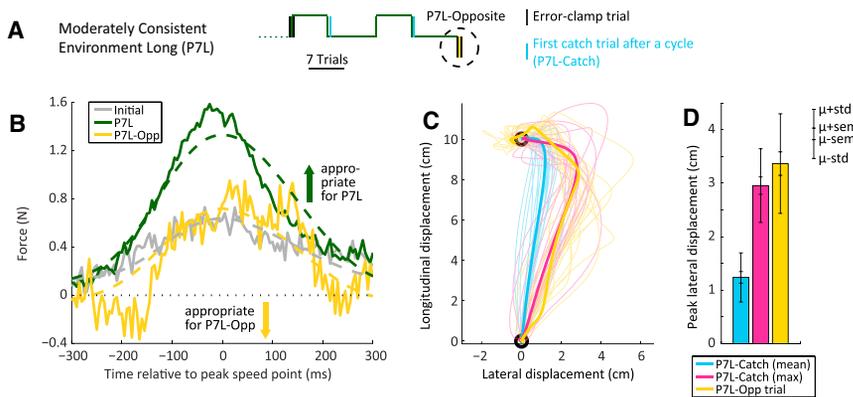


Figure 5. Specificity of Adaptation-Rate Increases for the Experienced Dynamics

(A) Illustration of the P7L-opposite (P7L-Opp, yellow) experiment, in which a single oppositely directed FF perturbation was presented at the end of the P7L environment. The adaptation rate for this trial was compared to the adaptation rate displayed in the last third of trials in P7L (green).

(B) Comparison of the average initial single-trial adaptive response (gray) with that observed for the last third of the P7L trials (green) and the P7L-opposite trial (yellow). Note that in the P7L-opposite trial, subjects produced a force compensation that was largely inappropriate for the FF experienced.

(C) Hand trajectories during the null trials immediately following FF blocks (P7L-catch; blue, aver-

age trajectories for each subject; pink, trajectories of trials where each subject experienced their maximum deviation) are compared to hand trajectories during the single P7L-opposite trial (yellow). Thin lines indicate data from individual subjects; thick lines indicate data averaged across subjects. Note that the P7L-catch trajectories that deviated the most are similar to the P7L-Opp ones, indicating that, although the FF experienced during the P7L-opposite trial was completely novel, the errors it elicited were not.

(D) Peak lateral errors experienced during the P7L-catch and P7L-opposite trials. Although on average each subject experienced weaker errors during P7L-catch trials (blue), the largest P7L-catch errors for each subject (pink) were similar to the errors experienced during P7L-opposite trials. Inner error bars indicate SEM, whereas outer error bars indicate the standard deviation across subjects.

explain the pattern of learning-rate changes observed across different environments (see Figure 4 above), we wondered whether repeated exposure to the same dynamics might interact with environmental consistency. In particular, does the relationship between adaptation rate and environmental consistency that we observe across environments require some degree of repetition for the trained dynamics? To investigate this possibility, we began by measuring the adaptation resulting from a single exposure to a completely novel FF—a FF opposite to the one that subjects were repeatedly trained on—after exposure to an environment that would increase adaptation rate. Specifically, we inserted an opposite FF trial at the end of the P7L experiment (P7L-opposite), as illustrated in Figure 5A. The idea was that only adaptation-rate upregulation that was entirely independent of repetition could result in full transfer of the adaptation-rate increases observed in P7L to the P7L-opposite trial. However, we found that the adaptive response to the P7L-opposite trial was in the same direction as the response for the previous FF exposures despite the fact that the P7L-opposite FF was opposite in direction to these exposures (Figure 5B). Thus, the P7L-opposite response was inappropriately directed, so that if an adaptation rate was calculated from it, that rate would be significantly negative in value ( $p < 0.001$ ); however, one might question whether an inappropriately directed response should be characterized by an adaptation rate. Instead, the adaptive responses elicited by P7L-opposite trials are more in line with *repetition of the adaptive responses* elicited during the P7L training than *transfer of the response gains* associated with this training, which were oppositely directed. However, it should be noted that the amplitude of the P7L-opposite responses falls short of amplitudes observed during P7L training (Figures 5B, 6C, and 6D), suggesting that the P7L-opposite responses might not be explained by repetition alone.

We next examined whether the inappropriately directed P7L-opposite responses were associated with the novelty of the experienced dynamics or of the associated motor errors. By design, the experienced dynamics were completely novel, opposite of the only FF previously experienced. In contrast, motor errors with the same direction as the ones experienced during the P7L-opposite trial would be expected at the

beginning of the washout period following each FF block. However, because subjects do not fully adapt by the end of each block (Figure 4C), the amplitude of the error on the first trial of washout (W1) might be substantially smaller than that experienced during the P7L-opposite trials. Indeed, we found that the mean W1 trajectory for each participant displayed motor errors that were aligned with, but 63% smaller than, the errors experienced during the P7L-opposite trials (blue versus yellow in Figures 5C and 5D, with no cases in which mean W1 errors were within 40% of the P7L-opposite errors). However, W1 trials displayed variable trajectories from one block to the next even within the same individual, and it turns out that the upper limit of this variability (measured as the maximally displaced W1 trial for each individual during the last third of P7L exposure) is largely indistinguishable (8.9% smaller, on average) from the P7L-opposite motor errors (pink versus yellow in Figure 5C and 5D; in 83% of cases, maximum W1 errors were within 25% of the P7L-opposite errors). Thus, whereas the experienced dynamics during P7L-opposite trials were completely novel, the motor errors that participants experienced were similar to errors sometimes experienced at the beginning of each washout block. However, the adaptive response following the maximally displaced W1 trials reduced, rather than increased, the subsequent displacement for all individuals in the P7L experiment, reflecting positive rather than negative adaptation rates for the W1 trials. Thus, the inappropriately directed P7L-opposite response cannot be attributed to the motor errors elicited on these trials and must instead be related to the novel dynamics experienced or the combination of the novel dynamics and the motor errors.

#### Repetition Contributes to, but Is Not Required for, Consistency-Driven Adaptation-Rate Upregulation

The adaptation-rate reductions observed in the P1L environment compared to P1 or to baseline indicate that repetition without consistency does not increase learning rates. However, the finding that P7L-opposite responses reflect repetition after exposure to an environment in which repetition and consistency are combined raises the possibility that the combination of consistency and repetition might be required for the learning-rate increases observed in high-consistency

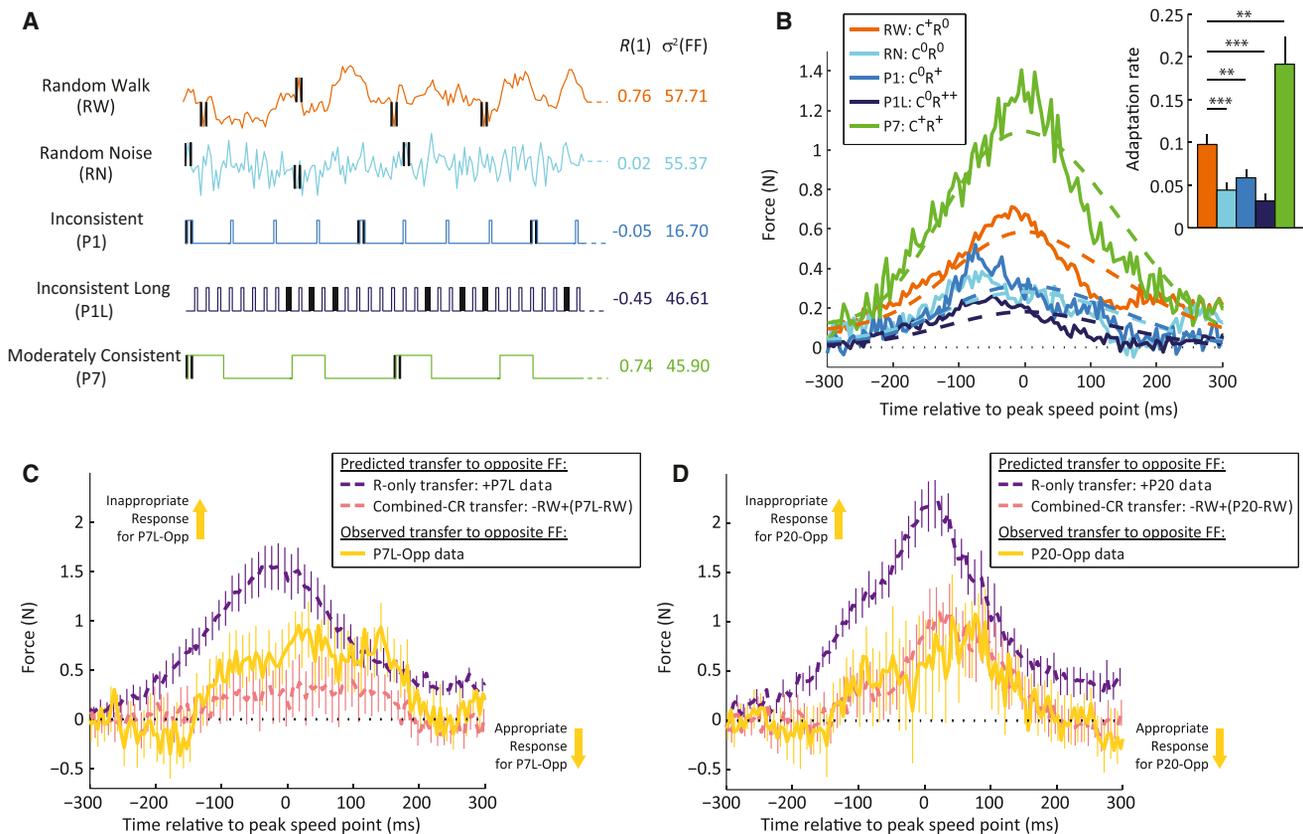


Figure 6. Synergistic Interaction between Repetition and Consistency for Learning-Rate Upregulation

(A) Illustration of the random walk (RW) environment (orange trace). Note that this environment was designed to have a variance that is similar to that of the random noise (RN) environment (57.71 versus 55.37) but at the same time to have a consistency similar to that of the moderately consistent (P7) environment (0.76 versus 0.74) without having its repetitive structure.

(B) Lateral-force profiles and average adaptation levels (inset) for the single-trial adaptive response in the second half of training for the RW, RN, P1, P1L, and P7 environments. Note that the force compensation in the RW environment (orange) is greater than that during low consistency (RN, P1, and P1L;  $p < 0.01$  in all cases), indicating that consistency, even without repetition, leads to higher adaptation rates. However, the learning in RW does not reach the level of that in P7 ( $p < 0.01$ ), indicating that repetition can enhance consistency-modulated learning-rate increases. Error bars indicate SEM in all panels. Significance levels are as follows: \* $p < 0.05$ , \*\* $p < 0.01$ , and \*\*\* $p < 0.001$ .

(C) Repetition and consistency-driven responses combine to produce the observed P7L-opposite response. A response predicted by the hypothesis that both consistency and repetition contribute distinct components to the adaptation (combined-CR transfer, pink dashed curve) matches the P7L-opposite response (yellow curve) much better than a response predicted by the hypothesis that consistency enables repetition-based learning (R-only transfer, purple dashed curve).

(D) Same as (C) but for the P20 experiment, illustrating how the P20-opposite response matches the response predicted by the combined-CR hypothesis rather than that predicted by the R-only hypothesis.

environments. Thus, we tested whether learning-rate upregulation can occur in a highly consistent environment without repetitious exposure to specific dynamics. In particular, we studied an environment based on a damped random walk (RW) where dynamics randomly varied without repetition from one trial to the next, but rather than varying around zero (as in RN), they varied around a fixed fraction of the previous trial's dynamics (see [Supplemental Experimental Procedures](#)). This resulted in a slowly drifting pattern (see [Figure 6A](#)), which we parameterized so that consistency [ $R(1) = 0.76$ ] was similar to P7 [ $R(1) = 0.74$ ] but much higher than RN [ $R(1) = 0.02$ ]; however, variability ( $57.7 \text{ N}^2\text{s}^2/\text{m}^2$ ) was approximately matched to both RN ( $55.4 \text{ N}^2\text{s}^2/\text{m}^2$ ) and P7 ( $45.9 \text{ N}^2\text{s}^2/\text{m}^2$ ).

We found that the RW environment elicited adaptation rates ( $0.097 \pm 0.012$ ) that were significantly higher than those for RN ( $p = 0.0004$ ) and the other environments involving low consistency, P1 ( $p = 0.0079$ ) and P1L ( $p = 0.00003$ ), as shown in [Figure 6B](#), in line with the higher consistency of the RW

environment. However, its adaptation rate was significantly lower than that observed in P7 ( $p = 0.0094$ ), where a similar consistency was combined with repetition. These findings indicate that environmental consistency alone can elicit learning-rate increases; however, when consistency is combined with repetition, even greater increases occur.

To examine the mechanism by which consistency and repetition combine in the P7 environment, we can re-examine the P7L-opposite data in light of the RW results. Here, we consider three possibilities: first, that repetition enables consistency-based learning (C-only hypothesis); second, that consistency enables repetition-based learning (R-only hypothesis); and third, that both consistency and repetition contribute distinct components to the adaptation (combined CR hypothesis). The P7L-opposite data provide a unique opportunity for examining these hypotheses because consistency-based gain modulations and the repetition of previously learned adaptive responses lead to opposite predictions, and these predictions

can be made without reference to the P7L-opposite data themselves—instead, the R-only prediction can be derived from P7L, and the C-only prediction from RW, and the combined CR prediction from a combination of the two. We can immediately reject the C-only hypothesis, that repetition further enables pure consistency-based learning, because consistency-based adaptation-rate transfer to the P7L-opposite trial would predict an appropriately directed response that is opposite to what we observe. The R-only hypothesis, that high consistency enables pure repetition-based learning, correctly predicts an inappropriately directed response for P7L-opposite; however, the amplitude of the predicted response (Figure 6C, purple) is considerably larger than that for P7L-opposite (yellow). In contrast, the combined CR hypothesis, where both consistency and repetition contribute distinct components to the adaptation, appears to explain both the amplitude and direction of the P7L-opposite response. Here, a flipped version of the RW data ( $-RW$ ) data provides an estimate of how the consistency-based component of learning should transfer to P7L-opposite, whereas the difference between the P7L and the RW data ( $P7L - RW$ ) provides an estimate of the repetition-based component of learning. Figure 6C shows that the prediction generated from combining these components [ $-RW + (P7L - RW)$ ; pink] explains the P7L-opposite data more closely than the R-only and C-only predictions would. However, it should be noted that, although certainly closer to the combined-CR prediction, the P7L-opposite data fall between the R-only and combined-CR predictions, suggesting that this analysis might not be definitive. Thus, to further examine the issue, we examined opposite-FF responses in a group of 16 subjects who experienced a single opposite-FF trial after exposure to the P20 environment. These P20-opposite data (Figure 6D) also match the prediction of the combined CR hypothesis considerably better than the R-only or C-only predictions. Taken together, these data suggest that environments that combine repetition and consistency elicit a combination of consistency-based and repetition-based adaptation. In the P7 and P20 conditions these contributions act in synergy; however, in the P7L-opposite and P20-opposite conditions, these contributions act in opposition.

## Discussion

We find that experimental manipulation of environmental consistency allows for adaptation-rate modulation that spans an order of magnitude ( $20\times$ ). In particular, we show that a highly consistent environment, in which external disturbances tend to persist from one trial to the next, elicits learning rates that are 3-fold greater than baseline, whereas an anti-consistent environment, in which these disturbances fluctuate in a negatively correlated manner, elicits learning rates that are 5-fold lower than baseline. Moreover, we find that whereas environmental variability determines feedback response strength but not motor adaptation rate, environmental consistency determines motor adaptation rate but not feedback response strength. This striking double dissociation gives insight into the specificity of mechanisms for same-trial control versus next-trial adaptation. However, consistency can be achieved through repetition, and the somewhat surprising results of the P7L-opposite experiment suggested a key role for repetition in the adaptive responses we observed. Thus, we closely examined the interactions between consistency and repetition. We find that even in the absence of repetition, increased

consistency elicits highly significant increases in learning (RW versus RN), but in the absence of consistency, increased repetition fails to elicit increased learning (P1L versus P1 versus RN), suggesting that consistency rather than repetition is primarily responsible for the learning-rate increases we observed. However, we find that consistency-modulated learning-rate increases are greatly magnified when repetition is combined with increased environmental consistency, indicating a strong synergistic effect between repetition and consistency.

The idea that the consistency of the environment modulates the rate of motor adaptation is in line with the idea that forming a *prediction* of the degree to which environmental change should persist from trial to trial is a critical step in determining the rate of adaptation. In an anticonsistent or inconsistent environment, the motor system would benefit from a low adaptation rate in order to avoid overlearning from current disturbances that do not positively predict future disturbances. However, in a highly consistent environment, the motor system would benefit from rapid adaptation because disturbances experienced during one trial would be highly predictive of future disturbances. Our results show that the motor system adjusts its rate of adaptation in a way that is consistent (pun intended) with these ideas.

## Interactions between Environmental Consistency and Repetition

Our results reveal a surprising synergy between consistency and repetition in the learning environment. The initial set of experiments, which employed a blocked design in which a single-amplitude FF alternated with a null environment (P1N1, P1, P7, and P20), demonstrated large increases in adaptation rate as consistency increased. However, the increased number of FF trials in each block resulted in increased repetition of the FF in the more highly consistent environments where the block length was increased (45 versus 189 versus 540 FF trial repetitions for the P1, P7, and P20 experiments, respectively). We thus performed the P1L and P7L experiments (540 FF trials each) to determine whether the additional repetition present in the P20 experiment was responsible for the increased adaptation rate it elicited. But in both cases the adaptation rates were smaller than those in the P20 environment ( $p < 0.0001$  for P1L and  $p = 0.046$  for P7L) and not larger than the corresponding environments with less repetition (P1L actually shows a *lower* learning rate than P1 [ $p = 0.011$ ]; P7L is not different than P7 [ $p = 0.22$ ]). This suggests that the extremely high learning rate observed in P20 was due to its increased environmental consistency, rather than to increased repetition relative to the original P1 and P7 environments, and that increased repetition alone does not improve learning rates, especially when consistency is low (P1 versus P1L). To further study the effect of repetition, we created two environments (RW and RN) in which the amplitude of the FF varied randomly over a continuous spectrum so that, strictly speaking, no repetition was present except in the probe trials themselves. This allowed us to examine the effect of environmental consistency in the absence of repetition. Here we found that the moderately consistent [ $R(1) = 0.76$ , similar to P7], zero-repetition RW environment elicited an adaptation rate that was significantly greater ( $p = 0.0004$ ) than that observed in zero-repetition RN environment in which there was low consistency [ $R(1) = 0.02$ ], as shown in Figure 6. However, the adaptation rate observed for RW was significantly lower ( $p = 0.009$ ) than for the P7 environment, to which

it was matched in consistency, suggesting that the combination of consistency and repetition leads to faster learning than consistency alone.

Together, our results show that increased environmental consistency improves adaptation rates, both when repetition is low (RW versus RN) and when repetition is high (P20 versus P7L; P7L versus P1L). In contrast, increased repetition fails to increase adaptation rates when consistency is low (P1 versus RN; P1L versus P1), but increased repetition does increase adaptation rates when consistency is moderate or high (either P7 or P7L versus RW). The substantial effect of repetition when consistency is high is in line with the inappropriately directed learning we observe during the P7L-opposite trial at the end of the P7L experiment, suggesting that repetition might lead to a stereotypic adaptive response when a large error is encountered, regardless of the direction of the error. This is in line with the idea that repetition-enabled recall of actions can lead to savings in motor adaptation [13]. In summary, we find that both consistency and repetition play a role in the modulation of adaptation rates. We show that repetition alone does not lead to an increase in adaptation rate but that consistency alone does lead to an increase in adaptation rate; however, together they lead to an increase in adaptation rate that is about double that elicited by consistency alone, indicating a synergistic effect between these two key features of the learning environment.

One interpretation of these results (the C-only hypothesis) is that the primary effect that consistency has on adaptation rate is potentiated by repetition. A very different interpretation (R-only) is that the primary effect that repetition has on adaptation rate is inhibited at low levels of environmental consistency. A third possibility (combined-CR) is that repetition and consistency contribute distinct components to the adaptive response but that the repetition-based component requires consistency. Interestingly, we show that these interpretations make distinct predictions for the P7L-opposite and P20-opposite data. Of particular interest is that the combined-CR hypothesis predicts responses that sum contributions from appropriately directed consistency-driven rate increases and inappropriately directed repetition-enabled recall. In contrast, the responses that the R-only hypothesis predicts for the P7L-opposite and P20-opposite data are identical to the responses in the P7L and P20 environments. We find that the predictions of the combined-CR model match both the P7L-opposite and P20-opposite data better than the predictions of the R-only model, suggesting that adaptive responses are characterized by an additive combination of consistency-driven learning-rate increases and repetition-enabled recall; however, additional work might be needed to definitively answer this important question.

#### Attempts to Manipulate Measurement Noise and State Noise to Alter Adaptation Rates

Previous studies that have investigated the modulation of motor adaptation rates have focused on the estimation of the identity of the experienced perturbations [1, 2, 5]. Theoretically, this estimation should be based on the appropriate weighting of the uncertainty about the state of the environment (*state noise*, i.e., variance on the *prior* expectation) relative to uncertainty about the fidelity of sensory information (*measurement noise*, i.e., variance on the *likelihood*).

According to optimal Bayesian estimation, the gain on the sensory information should decrease when the measurement noise rises, and it should increase when the state noise rises [1, 2, 5, 23]. Correspondingly, studies have attempted to

manipulate state noise and measurement noise in order to modulate the rate of adaptation [1, 2, 5]; however, these manipulations have not yielded consistent results. One set of approaches for manipulating state and measurement noise has involved perturbing the visual feedback about hand position by using a sequence that varied from one trial to the next either as white noise or a random walk [1, 2]. Here the idea was that measurement noise, like white noise, should be independent from one trial to the next, whereas state noise should accumulate like a random walk, and that the sensorimotor system would leverage this dichotomy by parsing the trial-to-trial variability in sensory information into independent and accumulating components to estimate the levels of measurement noise and state noise. However, there is no compelling reason to believe that this occurs. Instead, we suggest that the sensorimotor system's estimate of measurement noise should be affected by the quality of the sensory feedback rather than its pattern across trials. For example, when a child tracks the motion of a firefly at night, she can clearly identify the location of each flash, despite having little idea where each flash will appear. Thus, the random pattern of flashes does not dictate high measurement noise. Only a blurred view of a flash would introduce uncertainty about the firefly's current location. Correspondingly, experiments using randomized white-noise perturbations in an attempt to manipulate the sensorimotor system's estimate of measurement noise have consistently failed to show the predicted decreases in adaptation rate [1, 2]. In contrast, when studies have attempted to increase the motor system's estimate of measurement noise by degrading the intrinsic quality of visual feedback, decreases in adaptation rate have been consistently observed [2, 5].

Moreover, studies of sensorimotor integration for same-trial feedback responses have convincingly showed that a white-noise perturbation sequence leads to Bayesian integration in which state noise (the prior expectation) rather than measurement noise (the likelihood) is increased by the perturbation variance [22–25, 30]. In fact, state noise has been taken to fully reflect the variance of the perturbation sequence in these studies. This leaves us with the conclusion that the motor system does not parse trial-to-trial variability in sensory feedback into independent and accumulating components to differentially estimate the levels of measurement noise and state noise; instead, it takes the entire variability as an estimate of state noise.

But then, how can we explain that studies have found that accumulating random-walk perturbations result in faster adaptation rates than white noise perturbations [1, 2]? A key observation is the fact that random-walk perturbations display a lag-1 autocorrelation that approaches 1 for long exposures and are thus, in our parlance, highly consistent, whereas white-noise perturbations display a low-consistency, near-zero lag-1 autocorrelation. Therefore, changes in environmental consistency can provide a unifying explanation for the current set of results and for adaptation-rate increases previously observed with random-walk perturbations. This interpretation is also in line with the finding that in dynamic environments where the spatial complexity has a frequency that is likely to exceed that of the basis elements for adaptation, learning rates are consistently reduced [31] because high spatial complexity could make an environment appear inconsistent.

#### Implications for Retention of Motor Memories

The idea that forming a prediction of the degree to which environmental change should persist from trial to trial is a critical

step in determining the rate of adaptation also has implications for the retention of motor memories. The basic trial-to-trial learning rules presented in the Introduction (see the last two equations) suggest that when error feedback is withheld, retention of the current memory will, like the learning rate, depend on the prediction factor,  $A$ . Thus, the current results predict that consistency-modulated increases and decreases in learning will be accompanied, respectively, by increases and decreases in retention, i.e., decreases and increases in decay. Future work will investigate this idea.

### Implications for Rehabilitation

The findings described in the present study might enable the design of more effective rehabilitation protocols for stroke patients. The first four weeks after stroke is the time window during which rehabilitation has the greatest impact on the functional recovery of patients [32–34]. Therefore, an intervention that improves the rate of motor learning during this initial period might help patients achieve better functional recovery, possibly via the design of motor rehabilitation paradigms that upregulate motor learning through the use of highly consistent environments.

### Experimental Procedures

Subjects performed 10 cm point-to-point reaching movements in the 90° and 270° directions in the horizontal plane with their dominant hands while grasping the handle of a 2-link robotic manipulandum (Figure 1A). Subjects first performed 200 movements to obtain familiarity with the task. Afterward, in certain movements, the subjects' trajectories were disrupted by velocity-dependent curl FFs. We assessed the level of adaptation by using error-clamp (EC) trials (i.e., we measured the force pattern that subjects produced when their lateral errors were held to near-zero values in an error clamp) [10, 35–38]. We calculated adaptation rates by obtaining the difference between the adaptation coefficient ( $x$ ) [10, 35–38] for the EC trial following the first FF trial in a measurement cycle ( $EC_{Post}$ ) and the adaptation coefficient for the EC trial preceding this force-field trial ( $EC_{Pre}$ ):

$$\text{Adaptation rate} = x(EC_{Post}) - x(EC_{Pre}) \quad (1)$$

Learning environments with different levels of consistency (operationally defined as the lag-1 autocorrelation [ $R(1)$ ] of the environment) and variability allowed us to study the environmental modulation of adaptation rate.

$$\text{Consistency: } R(1) = E[(FF_n - \mu_{FF})(FF_{n+1} - \mu_{FF})] / \sigma_{FF}^2 \quad (2)$$

In the anti-consistent environment [P1N1;  $R(1) = -0.3$ ], 21 subjects experienced 50 cycles each with a single positive FF trial, followed by a single negative FF trial, followed by 11–13 washout (null) trials. In the inconsistent environment [P1;  $R(1) = -0.05$ ], 12 subjects experienced 45 cycles with a single positive FF trial, followed by ten to 12 washout trials. In the moderately consistent environment [P7;  $R(1) = 0.74$ ], 12 subjects experienced 27 cycles with seven positive FF trials, followed by 15–18 washout trials. In the highly consistent environment [P20;  $R(1) = 0.90$ ], 28 subjects experienced 27 cycles with 20 positive FF trials, followed by 28–32 washout trials. Two additional groups completed extended versions of the P1 and P7 experiments modified to include the same number of FF trials used in the P20 experiment [P1L;  $R(1) = -0.45$  and P7L;  $R(1) = 0.73$ ]. Twelve P1L subjects experienced 540 cycles with a single positive FF trial, followed by one to three washout trials, and 18 P7L subjects experienced 77 cycles with seven positive FF trials, followed by ten to 14 washout trials. In addition, we added a single negative FF trial, surrounded by EC trials, after the last FF cycle in P7L and in a subset of P20 (P7L- and P20-opposite) to assess the adaptation rate associated with this novel perturbation. To assess the adaptation rate during exposure to the learning environments, we randomly interspersed EC trials before and after the positive FF trial to form measurement triplets in a subset of the FF cycles in these experiments (40%, 44%, 44%, 44%, 22%, and 45% in P1N1, P1, P7, P20, P1L, and P7L, respectively).

Thirteen subjects experienced a random-noise environment [RN;  $R(1) = 0.02$ ] in which the FF trials varied randomly and magnitudes were determined according to a normal distribution with a standard deviation of

7.5 N/(m/s). Measurement triplets consisting of a single positive or negative FF trial, in which the FF had the same magnitude [15 N/(m/s)] as in the cycle experiments described above, were randomly interspersed after 3% of trials so that we could assess the adaptation rate. Finally, 23 subjects experienced a random-walk environment ([RW;  $R(1) = 0.76$ ]) where the FF trials followed a random walk with a carryover coefficient of 0.88 and a noise term with a standard deviation of 2.7 N/(m/s). As with RN, we used measurement triplets after 4% of trials to assess the adaptation rate. Please refer to the Supplemental Experimental Procedures for additional details.

### Supplemental Information

Supplemental Information includes four figures, three tables, and Supplemental Experimental Procedures and can be found with this article online at <http://dx.doi.org/10.1016/j.cub.2014.03.049>.

### Acknowledgments

We thank Jasmine Bailey for help with experiments. This work was supported in part by the McKnight Scholar Award, a Sloan Research Fellowship, and a grant from the National Institute on Aging (R01 AG041878) to M.A.S..

Received: December 1, 2012

Revised: February 27, 2014

Accepted: March 17, 2014

Published: May 1, 2014

### References

- Baddeley, R.J., Ingram, H.A., and Miall, R.C. (2003). System identification applied to a visuomotor task: near-optimal human performance in a noisy changing task. *J. Neurosci.* 23, 3066–3075.
- Burge, J., Ernst, M.O., and Banks, M.S. (2008). The statistical determinants of adaptation rate in human reaching. *J. Vis.* 8, 1–19.
- Körding, K.P., Tenenbaum, J.B., and Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat. Neurosci.* 10, 779–786.
- Malone, L.A., Vasudevan, E.V.L., and Bastian, A.J. (2011). Motor adaptation training for faster relearning. *J. Neurosci.* 31, 15136–15143.
- Wei, K., and Körding, K. (2010). Uncertainty of feedback and state estimation determines the speed of motor adaptation. *Front Comput Neurosci* 4, 11.
- van Beers, R.J. (2012). How does our motor system determine its learning rate? *PLoS ONE* 7, e49373.
- Kalman, R.E. (1960). A new approach to linear filtering and prediction problems. *J. Basic Eng.* 82, 35–45.
- Donchin, O., Francis, J.T., and Shadmehr, R. (2003). Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *J. Neurosci.* 23, 9032–9045.
- Sing, G.C., and Smith, M.A. (2010). Reduction in learning rates associated with anterograde interference results from interactions between different timescales in motor adaptation. *PLoS Comput. Biol.* 6, e1000893.
- Smith, M.A., Ghazizadeh, A., and Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol.* 4, e179.
- Smith, M.A., and Shadmehr, R. (2005). Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. *J. Neurophysiol.* 93, 2809–2821.
- Thoroughman, K.A., and Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature* 407, 742–747.
- Huang, V.S., Haith, A., Mazzoni, P., and Krakauer, J.W. (2011). Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. *Neuron* 70, 787–801.
- Sing, G.C., Orozco, S.P., and Smith, M.A. (2013). Limb motion dictates how motor learning arises from arbitrary environmental dynamics. *J. Neurophysiol.* 109, 2466–2482.
- Sing, G.C., Joiner, W.M., Nanayakkara, T., Brayanov, J.B., and Smith, M.A. (2009). Primitives for motor adaptation reflect correlated neural tuning to position and velocity. *Neuron* 64, 575–589.

16. Yousif, N., and Diedrichsen, J. (2012). Structural learning in feedforward and feedback control. *J. Neurophysiol.* *108*, 2373–2382.
17. Selen, L.P.J., Franklin, D.W., and Wolpert, D.M. (2009). Impedance control reduces instability that arises from motor noise. *J. Neurosci.* *29*, 12606–12616.
18. Mitrovic, D., Klanke, S., Osu, R., Kawato, M., and Vijayakumar, S. (2010). A computational model of limb impedance control based on principles of internal model uncertainty. *PLoS ONE* *5*, e13601.
19. Franklin, D.W., Burdet, E., Tee, K.P., Osu, R., Chew, C.-M., Milner, T.E., and Kawato, M. (2008). CNS learns stable, accurate, and efficient movements using a simple algorithm. *J. Neurosci.* *28*, 11165–11173.
20. Osu, R., Kamimura, N., Iwasaki, H., Nakano, E., Harris, C.M., Wada, Y., and Kawato, M. (2004). Optimal impedance control for task achievement in the presence of signal-dependent noise. *J. Neurophysiol.* *92*, 1199–1215.
21. Takahashi, C.D., Scheidt, R.A., and Reinkensmeyer, D.J. (2001). Impedance control and internal model formation when reaching in a randomly varying dynamical environment. *J. Neurophysiol.* *86*, 1047–1051.
22. Körding, K.P., Ku, S.P., and Wolpert, D.M. (2004). Bayesian integration in force estimation. *J. Neurophysiol.* *92*, 3161–3165.
23. Berniker, M., Voss, M., and Körding, K. (2010). Learning priors for Bayesian computations in the nervous system. *PLoS ONE* *5*, 5.
24. Chalk, M., Seitz, A.R., and Seriès, P. (2010). Rapidly learned stimulus expectations alter perception of motion. *J. Vis.* *10*, 2.
25. Seydell, A., Knill, D.C., and Trommershäuser, J. (2010). Adapting internal statistical models for interpreting visual cues to depth. *J. Vis.* *10*, 1–27.
26. Kojima, Y., Iwamoto, Y., and Yoshida, K. (2004). Memory of learning facilitates saccadic adaptation in the monkey. *J. Neurosci.* *24*, 7531–7539.
27. Krakauer, J.W., Pine, Z.M., Ghilardi, M.F., and Ghez, C. (2000). Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J. Neurosci.* *20*, 8916–8924.
28. Medina, J.F., Garcia, K.S., and Mauk, M.D. (2001). A mechanism for savings in the cerebellum. *J. Neurosci.* *21*, 4081–4089.
29. Zarah, E., Weston, G.D., Liang, J., Mazzoni, P., and Krakauer, J.W. (2008). Explaining savings for visuomotor adaptation: linear time-invariant state-space models are not sufficient. *J. Neurophysiol.* *100*, 2537–2548.
30. Körding, K.P., and Wolpert, D.M. (2004). Bayesian integration in sensorimotor learning. *Nature* *427*, 244–247.
31. Thoroughman, K.A., and Taylor, J.A. (2005). Rapid reshaping of human motor generalization. *J. Neurosci.* *25*, 8948–8953.
32. Krakauer, J.W. (2006). Motor learning: Its relevance to stroke recovery and neurorehabilitation. *Curr. Opin. Neurol.* *19*, 84–90.
33. Kwakkel, G., Kollen, B., and Lindeman, E. (2004). Understanding the pattern of functional recovery after stroke: facts and theories. *Restor. Neurol. Neurosci.* *22*, 281–299.
34. Kwakkel, G., Kollen, B.J., van der Grond, J., and Prevo, A.J.H. (2003). Probability of regaining dexterity in the flaccid upper limb: impact of severity of paresis and time since onset in acute stroke. *Stroke* *34*, 2181–2186.
35. Joiner, W.M., and Smith, M.A. (2008). Long-term retention explained by a model of short-term learning in the adaptive control of reaching. *J. Neurophysiol.* *100*, 2948–2955.
36. Scheidt, R.A., Reinkensmeyer, D.J., Conditt, M.A., Rymer, W.Z., and Mussa-Ivaldi, F.A. (2000). Persistence of motor adaptation during constrained, multi-joint, arm movements. *J. Neurophysiol.* *84*, 853–862.
37. Gonzalez Castro, L.N., Monsen, C.B., and Smith, M.A. (2011). The binding of learning to action in motor adaptation. *PLoS Comput. Biol.* *7*, e1002052.
38. Joiner, W.M., Brayanov, J.B., and Smith, M.A. (2013). The training schedule affects the stability, not the magnitude, of the interlimb transfer of learned dynamics. *J. Neurophysiol.* *110*, 984–998.