(Fig. 3). The resulting map is a function of the taxonomic cutoff used. At less than 99%, the model does not show any diversity or biogeographic patterns. This is expected because over 1400 years the genomes evolve by only 0.5%, corresponding to a 99% identity. However, at this time there are still provinces that have not yet experienced a takeover or coalescence event (Fig. 1A), and those are expected to continue to diverge beyond the 99% threshold for longer times. This distribution can change temporarily as a result of takeover events. For example, at times the Central Pacific and North Pacific provinces were distinguishable at 99.5% (0.5% difference, see Fig. 2A).

We conclude that neutral evolution (neutral mutations and genetic drift) coupled with dispersal limitation can produce substantial biogeographic patterns in the global surface ocean microbe population. Microbes evolve faster than the ocean circulation can disperse them, a feature that can also be seen in molecular observations (*10*). The patterns are dynamic. Provinces gradually emerge as subpopulations diverge by neutral evolution and periodically collapse due to coalescence. Neutral processes, along with environmental selection, must be considered in future research on microbial biogeography, and our results provide a quantitative benchmark for their potential role. Our results conflict with the notion that "everything is everywhere" (*11*, *32*) and may have important implications for how the oceans will respond to global change. Our model provides insights into the role of neutral evolution in shaping biogeographic patterns. The biology in the model is relatively simple, and future work may build on this by considering more spatial and temporal patterns (e.g., carrying capacity, as well as division and death rates based on ocean productivity) and more explicit representation of processes (e.g., recombination). Our modeling approach can theoretically be used to explore environmental selection as well. This will require relating genes to function, which is difficult but can be done for select genes or at the genome level (*23*, *33*).

REFERENCES AND NOTES

1. A. Ramette, J. M. Tiedje, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 2761–2766 (2007).
2. C. A. Hanson, J. A. Fuhrman, M. C. Horner-Devine, J. B. H. Martiny, *Nat. Rev. Microbiol.* **10**, 497–506 (2012).
3. R. T. Papke, D. M. Ward, *FEMS Microbiol. Ecol.* **48**, 293–303 (2004).
4. M. V. Brown, M. Ostrowski, J. J. Grzymski, F. M. Lauro, *Mar. Genomics* **15**, 17–28 (2014).
5. R. H. MacArthur, E. O. Wilson, *The Theory of Island Biogeography* (Princeton Univ. Press, Princeton, NJ, 1967).
6. S. P. Hubbell, *The Unified Neutral Theory of Biodiversity and Biogeography* (Princeton Univ. Press, Princeton, NJ, 2001).
7. D. B. Rusch *et al.*, *PLOS Biol.* **5**, e77 (2007).
8. B. K. Swan *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **110**, 11463–11468 (2013).
9. J.-F. Ghiglione *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 17633–17638 (2012).
10. A. C. Martiny, A. P. K. Tai, D. Veneziano, F. Primeau, S. W. Chisholm, *Environ. Microbiol.* **11**, 823–832 (2009).
11. S. M. Gibbons *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **110**, 4651–4655 (2013).
12. J. B. H. Martiny *et al.*, *Nat. Rev. Microbiol.* **4**, 102–112 (2006).
13. D. Wilkins, E. van Sebille, S. R. Rintoul, F. M. Lauro, R. Cavicchioli, *Nat. Commun.* **4**, 2457 (2013).
14. A. D. Barton, S. Dutkiewicz, G. Flierl, J. Bragg, M. J. Follows, *Science* **327**, 1509–1511 (2010).
15. R. Condit *et al.*, *Science* **295**, 666–669 (2002).
16. W. J. Sul, T. A. Oliver, H. W. Ducklow, L. A. Amaral-Zettler, M. L. Sogin, *Proc. Natl. Acad. Sci. U.S.A.* **110**, 2342–2347 (2013).
17. J. W. Drake, B. Charlesworth, D. Charlesworth, J. F. Crow, *Genetics* **148**, 1667–1686 (1998).
18. K. L. Vergin *et al.*, *Environ. Microbiol.* **9**, 2430–2440 (2007).
19. M. Vos, X. Didelot, *ISME J.* **3**, 199–208 (2009).
20. J.-U. Kreft *et al.*, *Proc. Natl. Acad. Sci. U.S.A.* **110**, 18027–18028 (2013).
21. F. L. Hellweger, V. Bucci, *Ecol. Modell.* **220**, 8–22 (2009).
22. P. G. Falkowski, C. D. Wirick, *Mar. Biol.* **65**, 69–75 (1981).
23. F. L. Hellweger, *Environ. Microbiol.* **11**, 1386–1394 (2009).
24. J. R. Clark, T. M. Lenton, H. T. Williams, S. J. Daines, *Limnol. Oceanogr.* **58**, 1008–1022 (2013).
25. W. B. Whitman, D. C. Coleman, W. J. Wiebe, *Proc. Natl. Acad. Sci. U.S.A.* **95**, 6578–6583 (1998).
26. M. Nei, W. H. Li, *Proc. Natl. Acad. Sci. U.S.A.* **76**, 5269–5273 (1979).
27. M. Slatkin, *Evolution* **47**, 264–279 (1993).
28. Materials and methods are available as supplementary materials on *Science* Online.
29. H. Ducklow, in *Microbial Ecology of the Oceans*, D. L. Kirchman, Ed. (Wiley-Liss, New York, 2000), chap. 4, pp. 85–120.
30. Y. Masumoto *et al.*, *J. Earth Simulator* **1**, 35–56 (2004).
31. J. M. Halley, Y. Iwasa, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 2316–2321 (2011).
32. P. Cermeño, P. G. Falkowski, *Science* **325**, 1539–1541 (2009).
33. A. M. Feist, M. J. Herrgård, I. Thiele, J. L. Reed, B. Ø. Palsson, *Nat. Rev. Microbiol.* **7**, 129–143 (2009).
34. S. J. Giovannoni *et al.*, *Science* **309**, 1242–1245 (2005).

ACKNOWLEDGMENTS

We thank R. Stepanauskas for interesting discussion that inspired this research. Three anonymous reviewers provided constructive criticism. F.L.H. and N.D.F. are supported by grants from the NSF and National Oceanic and Atmospheric Administration Massachusetts Institute of Technology SeaGrant. E.v.S. is supported by ARC via grant DE130101336. The source code is available at www.coe.neu.edu/~ferdi/Code/.

SUPPLEMENTARY MATERIALS

www.sciencemag.org/content/345/6202/1346/suppl/DC1
Materials and Methods
Supplementary Text
Figs. S1 to S9
Table S1
References (*35*–*43*)
Movie S1

4 April 2014; accepted 1 August 2014
10.1126/science.1254421

ERROR MEMORY

A memory of errors in sensorimotor learning

David J. Herzfeld,[1]* Pavan A. Vaswani,[2] Mollie K. Marko,[1] Reza Shadmehr[1]

The current view of motor learning suggests that when we revisit a task, the brain recalls the motor commands it previously learned. In this view, motor memory is a memory of motor commands, acquired through trial-and-error and reinforcement. Here we show that the brain controls how much it is willing to learn from the current error through a principled mechanism that depends on the history of past errors. This suggests that the brain stores a previously unknown form of memory, a memory of errors. A mathematical formulation of this idea provides insights into a host of puzzling experimental data, including savings and meta-learning, demonstrating that when we are better at a motor task, it is partly because the brain recognizes the errors it experienced before.

How does the brain alter behavior after experiencing an error? Classic theories assumed that the brain learns some fraction of the error regardless of its history or magnitude (*1*, *2*). However, recent experiments (*3*–*6*) demonstrate that the brain learns relatively more from small errors than large errors, and can modulate its error sensitivity (*7*–*9*).

Understanding error sensitivity is important because it may provide insight into the phenomena of "savings" and "meta-learning." Savings refers to the observation that when a person adapts to perturbation (A), and then the perturbation is removed (i.e., washout), they exhibit faster readaptation to (A) (*10*). Remarkably, savings of (A) is present even when washout is followed by adaptation to (-A), a perturbation in the opposite direction (*11*, *12*). Current error-dependent models of learning cannot account for these observations (*13*, *14*), nor explain meta-learning, where prior exposure to a random perturbation produces savings (*15*, *16*).

We begin with a standard model of motor learning (*17*–*20*) in which on trial n, a perturbation x is imposed on action u so that the sensory consequences are $y^{(n)} = u^{(n)} + x^{(n)}$. Based on their belief about the environment $\hat{x}^{(n)}$, the learner predicts the sensory consequences $\hat{y}^{(n)} = u^{(n)} + \hat{x}^{(n)}$, and then updates his or her belief from the prediction error $e^{(n)} = y^{(n)} - \hat{y}^{(n)}$. Such learning typically depends on a decay factor a, and error sensitivity η

$$\hat{x}^{(n+1)} = a\hat{x}^{(n)} + \eta^{(n)}e^{(n)} \quad (1)$$

Consider an environment in which the perturbations persist from trial to trial, and another

[1]Department of Biomedical Engineering, Laboratory for Computational Motor Control, Johns Hopkins University School of Medicine, Baltimore, MD 21205, USA. [2]Department of Neuroscience, Laboratory for Computational Motor Control, Johns Hopkins University School of Medicine, Baltimore, MD 21205, USA.
*Corresponding author: E-mail: dherzfe1@jhmi.edu

environment in which the perturbations switch (Fig. 1A). In a slowly switching environment, the brain should learn from error because the perturbations are likely to persist (learning from error in one trial will improve performance on the subsequent trial). However, in a rapidly switching environment, the brain should suppress learning from error because any learning will be detrimental to performance on subsequent trials.

Three groups of volunteers (n = 9 per group) made reaching movements while experiencing force perturbations from either a slow, medium, or rapidly switching environment (Fig. 1A). The mean of the perturbations was zero for all blocks (consisting of 30 trials). We measured error in a given trial and then computed the amount that was learned from that error (probe trials, purple bars, Fig. 1A). To quantify learning from error on trial n, we measured the change in force from the trial before to the trial after the perturbation, $f^{(n+1)}(t) - f^{(n-1)}(t)$ (Fig. 1C). In block 1, learning from error was similar in the three groups ($P > 0.99$), and in all probe trials the perturbation produced similar errors [Fig. 1B; repeated measure–analysis of variance (RM-ANOVA), effect of group $P > 0.8$, interaction, $P > 0.7$]. However, individuals who experienced the slowly switching environment increased their learning from error (Fig. 1C), whereas those who experienced the rapidly switching environment suppressed this learning.

We measured the force produced on a given trial and computed a coefficient representing percent ideal (Fig. 1D). RM-ANOVA indicated a significant block by group interaction ($P < 0.05$), suggesting that the history of perturbations altered the amount of learning from error. Posthoc tests showed that in the slowly switching environment, participants learned more from error than in the rapidly switching environment ($P < 0.03$). This change in error sensitivity developed gradually with training (Fig. 1D). The slowly switching environment induced an increase in error sensitivity (Fig. 1E; changes in sensitivity from the first half to second half of the experiment, ANOVA, $P < 0.05$).

Is control of error sensitivity local to the experienced errors? In experiment 2, participants performed rapid out-and-back movements for which no visual feedback was available during the outward part of the reach, with the aim of hitting a target at the turn-around point of their movement. An occasional perturbation altered the feedback regarding hand position at the turn-around point (Fig. 2A). We measured the relation between error $e^{(n)}$ and learning from error (change in reach extent).

Group 1 (n = 10) experienced a perturbation schedule that transitioned from slow, medium, to rapid switching (Fig. 2B), whereas group 2 (n = 10) experienced the reverse. In group 1, error sensitivity decreased, whereas in group 2, error sensitivity increased (Fig. 2C). We measured the mean error sensitivity in each environment, resulting in three measurements for each subject across the experiment. RM-ANOVA showed a significant main effect of group ($P < 0.005$) and block ($P < 0.001$) and group-by-block interaction

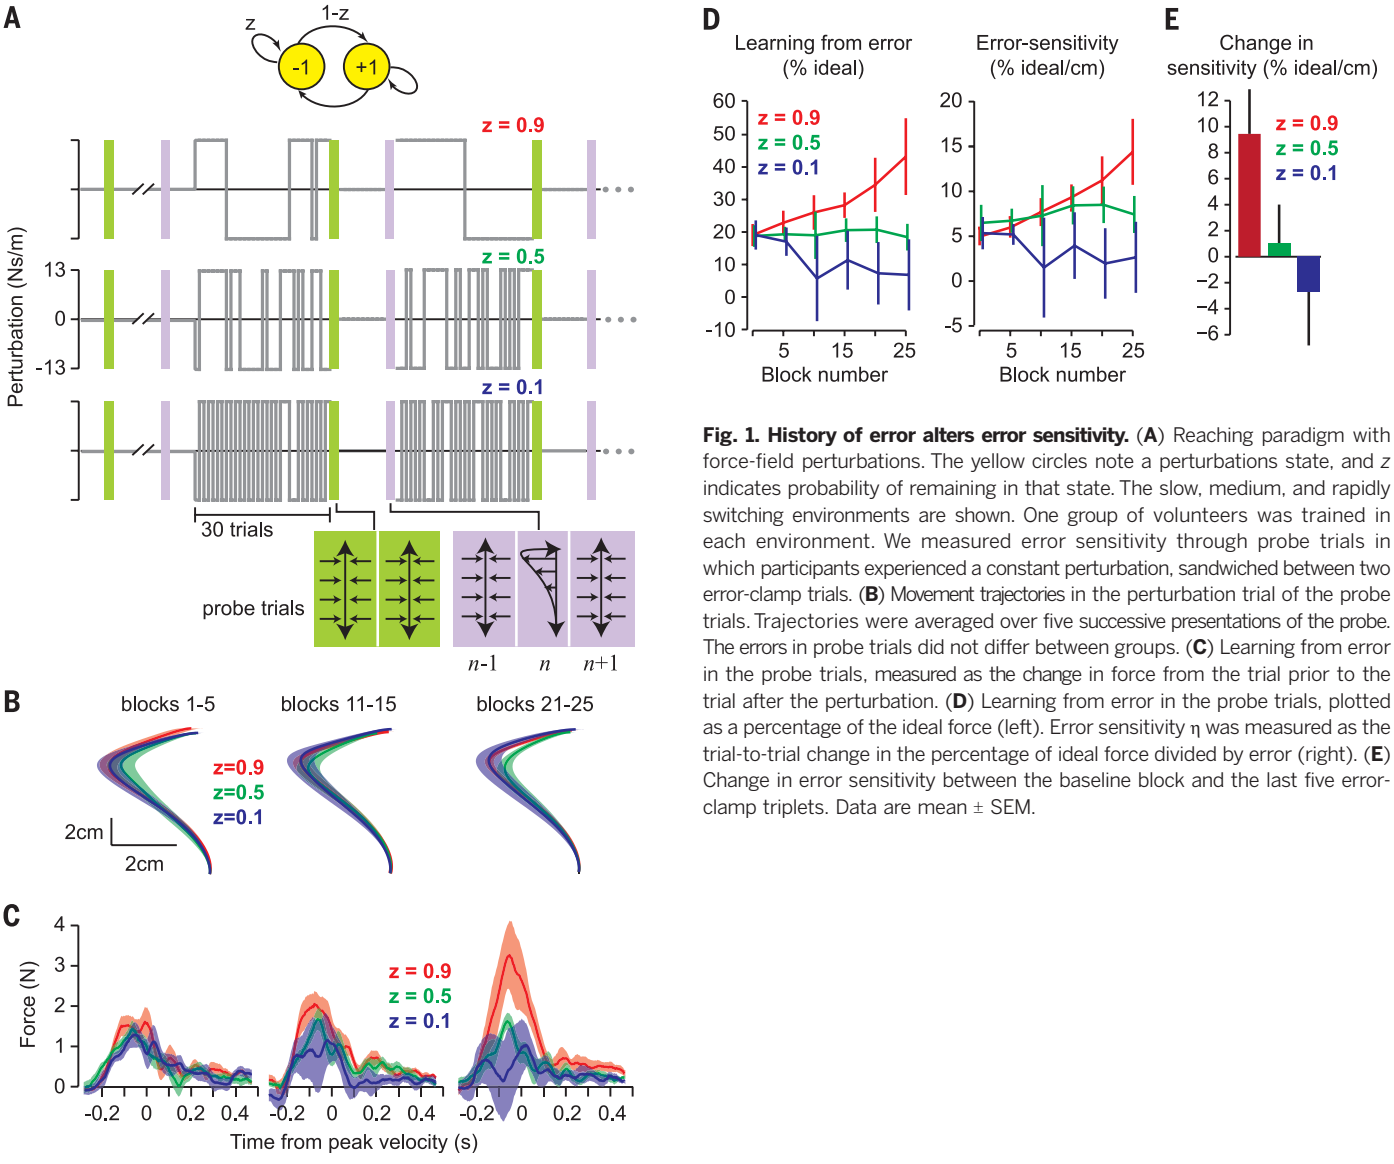

Fig. 1. History of error alters error sensitivity. (**A**) Reaching paradigm with force-field perturbations. The yellow circles note a perturbations state, and z indicates probability of remaining in that state. The slow, medium, and rapidly switching environments are shown. One group of volunteers was trained in each environment. We measured error sensitivity through probe trials in which participants experienced a constant perturbation, sandwiched between two error-clamp trials. (**B**) Movement trajectories in the perturbation trial of the probe trials. Trajectories were averaged over five successive presentations of the probe. The errors in probe trials did not differ between groups. (**C**) Learning from error in the probe trials, measured as the change in force from the trial prior to the trial after the perturbation. (**D**) Learning from error in the probe trials, plotted as a percentage of the ideal force (left). Error sensitivity η was measured as the trial-to-trial change in the percentage of ideal force divided by error (right). (**E**) Change in error sensitivity between the baseline block and the last five error-clamp triplets. Data are mean ± SEM.

($P < 0.001$). As the statistics of the perturbation changed, so did the error sensitivity.

We measured learning from error as a function of error in each environment (Fig. 2D). A given error produced greater learning when that error was experienced in a slowly switching environment (Fig. 2D, red line) (RM-ANOVA main effects of error size $P < 10^{-4}$, and environment $P < 0.001$, posthoc between slow versus medium or fast, $P < 0.001$). We quantified error sensitivity at each error size (Fig. 2E) and found that error sensitivity had not changed globally, but had changed predominantly for smaller error sizes. RM-ANOVA of the absolute sensitivities between 0.25 and 2 cm showed a significant main effect of environment ($P < 10^{-4}$), as well as a significant environment by error size interaction ($P < 0.05$). We found a significant difference in error sensitivity across environments for an error size of 0.25 cm ($P < 0.05$), but no significant difference for an error size of 2 cm ($P > 0.1$). Interestingly, the small error sizes for which the participants had shown the largest change in error sensitivity were also the most frequent errors (Fig. 2F). This hinted that control of error sensitivity was a function of error.

Current models of sensorimotor learning assume that error sensitivity $\eta^{(n)}$ is independent of error $e^{(n)}$. This is true for state-space models of learning (*18*, *21*–*24*), as well as Kalman filter models of learning (*5*, *25*–*27*). However, suppose that sensory prediction errors are encoded in the nervous system with a set of basis elements, where each basis element g_i has a preferred error $\breve{e}_i$. Further, suppose that error sensitivity is determined by a population coding

$$\eta(e^{(n)}) = \sum_i w_i g_i(e^{(n)})$$
$$g_i(e^{(n)}) = \exp\frac{-(e^{(n)} - \breve{e}_i)^2}{2\sigma^2} \quad (2)$$

On trial $n - 1$, the motor command $u^{(n-1)}$ produces an error $e^{(n-1)}$, as illustrated in the top part of Fig. 3A. The nervous system learns from this error and produces motor command $u^{(n)}$ on the subsequent trial, resulting in $e^{(n)}$. In a slowly switching environment (top part of Fig. 3A), $e^{(n)}$ has the same sign as $e^{(n-1)}$. In this case, error sensitivity should increase around $e^{(n-1)}$ (Fig. 3B, red line). By contrast, in a rapidly switching environment (Fig. 3A, bottom), $e^{(n)}$ has a different sign than $e^{(n-1)}$. In this case, error sensitivity should decrease

$$\mathbf{w}^{(n+1)} = \mathbf{w}^{(n)} + \beta \text{sign}(e^{(n-1)}e^{(n)}) \frac{\mathbf{g}(e^{(n-1)})}{\mathbf{g}^T(e^{(n-1)})\mathbf{g}(e^{(n-1)})} \quad (3)$$

In Eq. (3), $\mathbf{w} = [w_1 \quad w_2 \quad \cdots \quad w_N]^T$, $\mathbf{g} = [g_1 \quad g_2 \quad \cdots \quad g_N]^T$, and superscript T is the transpose operator. This rule is similar to the RPROP algorithm, a heuristic for adjusting the learning rate of machines (*28*), but has the unique feature of assuming that error sensitivity is via population coding of the error space.

Equations 1 to 3 represent a learner that stores two kinds of memory: a memory of the state of environment ($\hat{x}$, Eq. 1) and a memory of errors ($\mathbf{w}$, Eq. 3). We simulated the model (Fig. 3C, gray line) and found that in the slowly switching environment, error sensitivity increased in the neighborhood of the experienced errors, whereas in the rapidly changing environment, error sensitivity decreased (Fig. 3D).

Our model made a critical prediction: If the brain controlled error sensitivity via memory of errors, then it should be possible to simultaneously increase sensitivity for one error, while decreasing it for another. In experiment 3, we considered an isometric task in which participants ($n = 16$) produced a force to match a target (16 N) in the face of a perturbation. The perturbations were designed so that, according to our model, individuals would increase their sensitivity to –4 N errors, while simultaneously decreasing their sensitivity to +8 N errors.

In the baseline block, we probed sensitivity to +8 N and –4 N perturbations (probe 1, Fig. 3E). The resulting learning from error is plotted in Fig. 3F (probe 1). At baseline, participants responded to the +8 N and –4 N perturbations by learning a fraction of each error (Fig. 3F). We next produced 20 repetitions of a rapidly switching environment in which the perturbations were ±8 N (Fig. 3E, inset). After a period of washout, we then produced 15 repetitions of a slowly switching environment in which the perturbations were 0 N or +4 N. The critical aspect of our design was that the participants were never exposed to a –4 N perturbation. They nevertheless experienced –4 N errors (because removal of a learned +4 N perturbation results in a –4 N error).

The 8 N environment induced a decrease in sensitivity to a +8 N error, and subsequent exposure to the +4 N environment resulted in an increase in sensitivity to a –4 N error [Fig. 3G; RM-ANOVA showed a significant main effect of perturbation ($P < 0.03$) as well as a perturbation by block

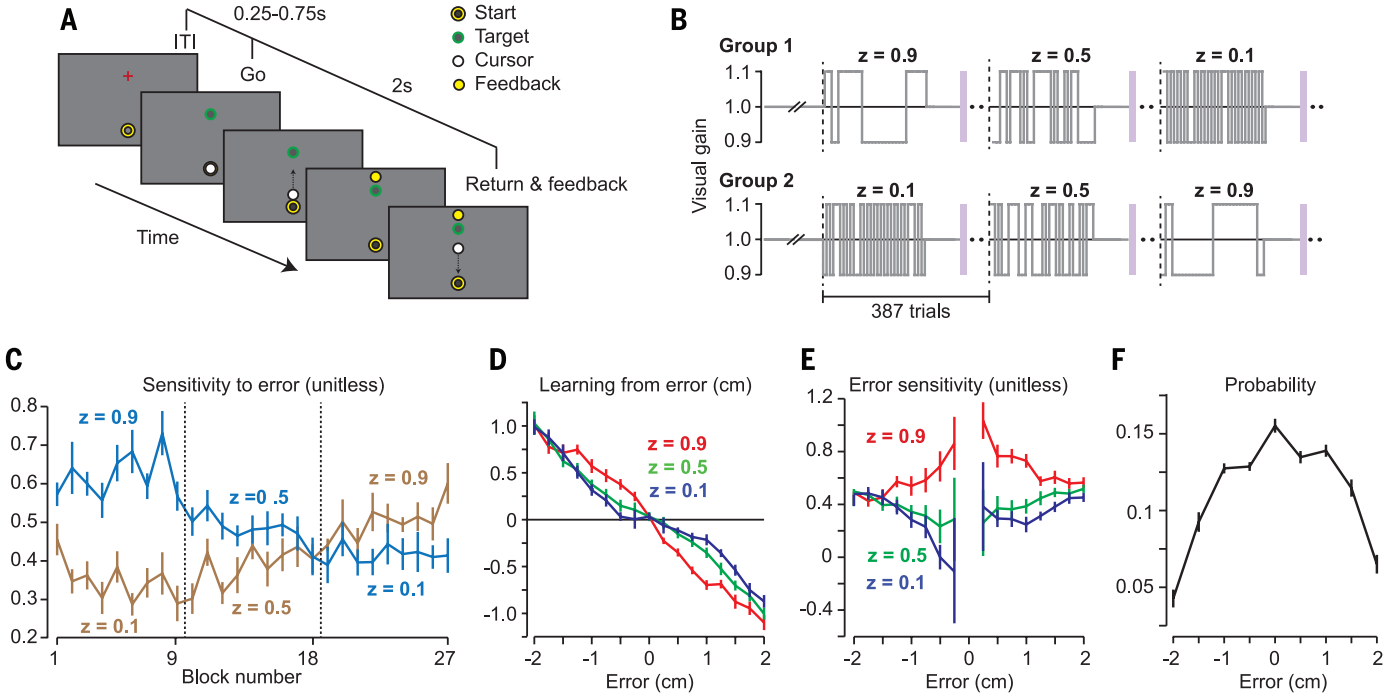

Fig. 2. Error sensitivity is a local function of experienced errors. (**A**) Paradigm with visuomotor gain perturbations. (**B**) Perturbation schedule. Dashed lines indicate changes in the statistics of the environment. (**C**) Error sensitivity averaged over all error sizes measured over each environment block. (**D**) Learning from error measured at various error sizes. (**E**) Error sensitivity as a function of error magnitude. (**F**) Probability of error.

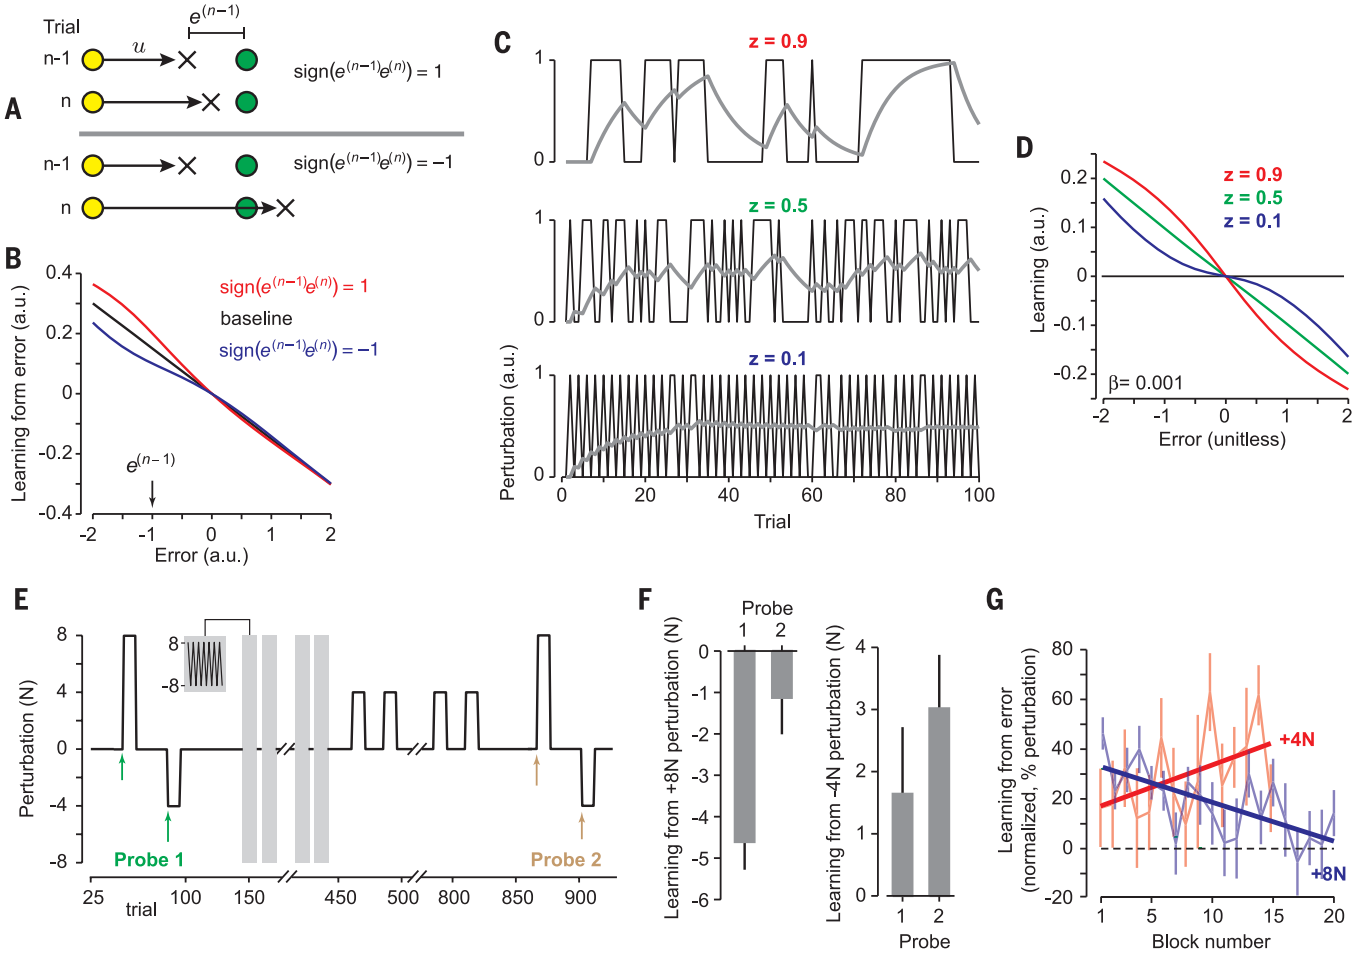

Fig. 3. Theoretical model and experiment 3. (**A**) On trial $n-1$, the motor command $u^{(n-1)}$ is generated, resulting in error $e^{(n-1)} = -1$. If the error in trial n is of the same sign as $e^{(n-1)}$, then error sensitivity should increase (top). However, if the error experienced in trial n has a different sign than $e^{(n-1)}$, then error sensitivity should decrease (bottom). (**B**) Learning from error after experience of two consecutive errors from (A). Error sensitivity around $e^{(n-1)}$ increases if $\text{sign}(e^{(n-1)}e^{(n)}) = 1$ and decreases otherwise. (**C**) Model performance for slow, medium, and rapidly switching environments (gray line represents $\hat{x}^{(n)}$). Learning from error (**D**) is increased in the slowly switching environment and decreased in the rapidly switching environment. (**E**) Experiment 3 perturbation protocol. (**F**) Single-trial learning from a +8 N perturbation and a −4 N perturbation in probes 1 and 2. Learning is increased for the −4 N perturbation, while simultaneously decreased for a +8 N perturbation. (**G**) Learning from error normalized by the perturbation magnitude (4 or 8 N) in the first trial of each repetition of the rapidly (blue) and slowly switching (red) environments. Learning increased in the slowly switching (4 N) environment but decreased when the perturbation was rapidly switching (8 N). Error bars are SEM.

interaction, $P < 0.01$]. The critical question, however, was whether both of these changes in sensitivity were simultaneously present. After the slowly switching block of perturbations, we again probed sensitivity to +8 N and −4 N errors (probe 2, Fig. 3E). Compared to the baseline block (probe 1), learning from a +8 N error had decreased ($P < 0.005$), while simultaneously, learning from a −4 N error had increased ($P < 0.05$) (Fig. 3F). When we ran our model on the same sequence of errors that participants had experienced, the change in error sensitivity predicted by the model was highly correlated with the change observed in our participants ($R^2 = 0.65$; $P < 10^{-8}$, fig. S1), suggesting that history of error induced changes in error sensitivity in the region of the experienced errors.

This new model of learning provided insights on a wide range of puzzling experiments, including the phenomena of savings and meta-learning (fig. S2). It predicted that when one is better at a task than before, it is not because the brain recalled the motor commands, but because it recognized the errors—the errors for which error sensitivity had been altered. In addition, the model predicted that savings and meta-learning could be blocked by controlling the errors that are experienced during learning.

In experiment 4, volunteers participated in a visuomotor rotation experiment (Fig. 4A and fig. S4, $n = 10$ per group). The control group (ANA) experienced a +30° perturbation followed by extended washout and then relearning of +30°, a protocol that should produce savings (*13*). According to our model, savings occurs because during the initial exposure to (A), the stable sequence of perturbations increase error sensitivity, and these errors are revisited in the subsequent test of (A). If so, we should be able to block savings by presenting (A) gradually (GNA group, fig. S4), preventing prior exposure to errors that are visited at the onset of (A).

Furthermore, we should be able to produce savings in a very different way: Expose participants to perturbation (B) and then present sudden washout (Fig. 4A, BNA). During washout, they are exposed to a sequence of stable errors, which increase error sensitivity for those errors. Notably, the washout-induced aftereffects are errors that are also experienced during subsequent test of (A). If the meta-learning in BNA is due to errors that are experienced during washout of (B), we should be able to eliminate it by reducing the washout-induced errors. In $B_{wait}NA$, a wait period was inserted between −30° training and washout, reducing the size of aftereffects associated with the transition from −30° to washout (Fig. 4E). We also tested this idea in a different way: gradual washout (B) (group BGNA, fig. S4). In summary, the model predicted

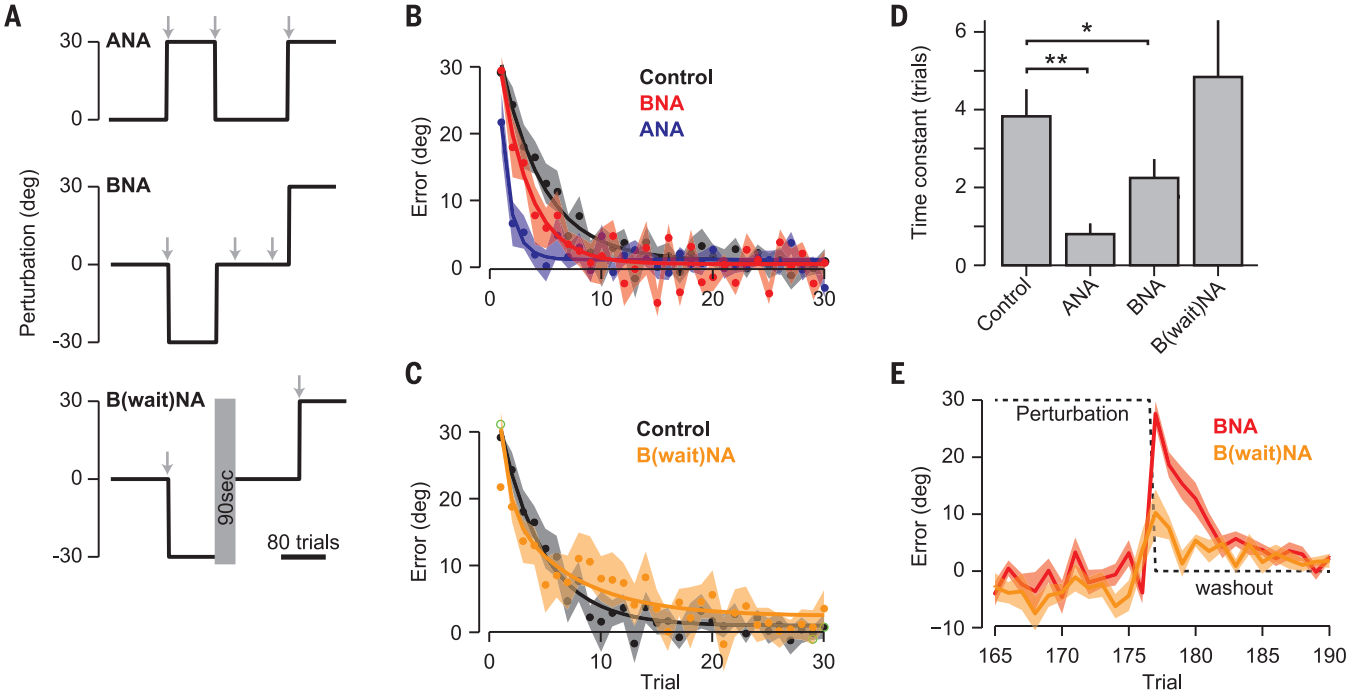

Fig. 4. Saving occurs only when previously experienced errors are revisited. (**A**) A visuomotor perturbation experiment. Gray arrows indicate 1- to 2-min set breaks. (**B**) Performance in the final +30° perturbation. ANA and BNA groups show savings, i.e., faster learning of the perturbation compared to control (naïve). Exponential fits are shown for the group data. (**C**) The $B_{wait}NA$ group does not exhibit savings. (**D**) Exponential time constants are compared to controls (*$P < 0.05$, **$P < 0.01$). A smaller time constant indicates faster learning. (**E**) Comparison of the errors (i.e., aftereffects) experienced by the BNA and $B_{wait}NA$ groups. The $B_{wait}NA$ group experienced smaller errors owing to the presence of the set break.

savings in ANA but not GNA, and meta-learning in BNA but not $B_{wait}NA$ and BGNA. Our experimental results confirmed these predictions (Fig. 4, C and D, and fig. S4).

We found that during learning, the brain controlled error sensitivity in a principled way: learning more from error when perturbations were likely to persist, and less when perturbations were likely to change. Error-sensitivity modulation was specific to the experienced errors, suggesting that training produced a memory of errors. This idea accounted for a host of puzzling observations, including saturation of error sensitivity (*5*, *6*, *29*), the phenomenon of meta-learning (*16*), examples of savings (*10*–*12*), and reinforced repetition (*15*).

The model predicted that meta-learning vanishes when a small delay or gradual washout alters the history of errors (Fig. 4A), demonstrating that savings depends crucially on the memory of errors that is accumulated during training. This memory of errors likely exists in parallel with the two traditional forms of motor memory, memory of perturbations (*23*) and memory of actions (*30*).

In our model, we chose to describe the learner as a process with a single time scale. However, data suggest that learning from error depends on a fast and a slow process with different error sensitivities (*23*, *25*, *31*). We speculate that the memory of errors exerts its influence through the error sensitivity of the fast process, and its manipulation through history of errors may be a useful strategy to speed recovery during rehabilitation (*32*).

REFERENCES AND NOTES

1. M. I. Jordan, D. E. Rumelhart, *Cogn. Sci.* **16**, 307–354 (1992).
2. M. Kawato, K. Furukawa, R. Suzuki, *Biol. Cybern.* **57**, 169–185 (1987).
3. F. R. Robinson, C. T. Noto, S. E. Bevans, *J. Neurophysiol.* **90**, 1235–1244 (2003).
4. R. Soetedjo, A. F. Fuchs, Y. Kojima, *J. Neurosci.* **29**, 15213–15222 (2009).
5. K. Wei, K. Körding, *J. Neurophysiol.* **101**, 655–664 (2009).
6. M. K. Marko, A. M. Haith, M. D. Harran, R. Shadmehr, *J. Neurophysiol.* **108**, 1752–1763 (2012).
7. M. A. Smith, R. Shadmehr, *Advances in Computational Motor Control* **3**, (2004).
8. L. N. Gonzalez Castro, A. M. Hadjiosif, M. A. Hemphill, M. A. Smith, *Curr. Biol.* **24**, 1050 (2014).
9. M. C. Trent, A. A. Ahmed, *Front. Comput. Neurosci.* **7**, 118 (2013).
10. Y. Kojima, Y. Iwamoto, K. Yoshida, *J. Neurosci.* **24**, 7531–7539 (2004).
11. L. A. Malone, E. V. Vasudevan, A. J. Bastian, *J. Neurosci.* **31**, 15136–15143 (2011).
12. A. M. Sarwary, L. P. Selen, W. P. Medendorp, *J. Neurophysiol.* **110**, 1269–1277 (2013).
13. E. Zarahn, G. D. Weston, J. Liang, P. Mazzoni, J. W. Krakauer, *J. Neurophysiol.* **100**, 2537–2548 (2008).
14. F. Mawase, L. Shmuelof, S. Bar-Haim, A. Karniel, *J. Neurophysiol.* **111**, 1444–1454 (2014).
15. D. A. Braun, A. Aertsen, D. M. Wolpert, C. Mehring, *Curr. Biol.* **19**, 352–357 (2009).
16. E. J. Turnham, D. A. Braun, D. M. Wolpert, *J. Neurophysiol.* **107**, 1111–1122 (2012).
17. N. J. Mackintosh, *Psychol. Rev.* **82**, 276–298 (1975).
18. K. A. Thoroughman, R. Shadmehr, *Nature* **407**, 742–747 (2000).
19. R. A. Rescorla, A. R. Wagner, in *Classical Conditioning*, A. H. Black, W. F. Prokasy, Eds. (Appleton-Century-Crofts, New York, 1972), vol. 2.
20. J. M. Pearce, G. Hall, *Psychol. Rev.* **87**, 532–552 (1980).
21. R. A. Scheidt, J. B. Dingwell, F. A. Mussa-Ivaldi, *J. Neurophysiol.* **86**, 971–985 (2001).
22. O. Donchin, J. T. Francis, R. Shadmehr, *J. Neurosci.* **23**, 9032–9045 (2003).
23. M. A. Smith, A. Ghazizadeh, R. Shadmehr, *PLOS Biol.* **4**, e179 (2006).
24. S. Cheng, P. N. Sabes, *Neural Comput.* **18**, 760–793 (2006).
25. K. P. Körding, J. B. Tenenbaum, R. Shadmehr, *Nat. Neurosci.* **10**, 779–786 (2007).
26. R. J. van Beers, *Neuron* **63**, 406–417 (2009).
27. R. J. van Beers, *PLOS ONE* **7**, e49373 (2012).
28. M. Riedmiller, H. Braun, in *IEEE International Conference on Neural Networks* (1993), vol. 1, pp. 586–591.
29. M. S. Fine, K. A. Thoroughman, *J. Neurophysiol.* **96**, 710–720 (2006).
30. V. S. Huang, A. Haith, P. Mazzoni, J. W. Krakauer, *Neuron* **70**, 787–801 (2011).
31. D. J. Herzfeld *et al.*, *Neuroimage* **98**, 147–158 (2014).
32. J. L. Patton, Y. J. Wei, P. Bajaj, R. A. Scheidt, *PLOS ONE* **8**, e46466 (2013).

ACKNOWLEDGMENTS

This work was supported by NIH grants T32EB003383, T32GM007057, R01NS078311, and 1F31NS079121. We thank J. Krakauer, J.-J. Orban de Xivry, and A. A. Ahmed for comments. D.J.H., P.A.V., M.K.M., and R.S. conceived and designed the experiments. D.J.H. performed the experiments, analyzed the data, and performed the simulations. R.S. and D.J.H. wrote the paper. The data reported in this paper and in the supplementary materials are available from the corresponding author.

SUPPLEMENTARY MATERIALS

www.sciencemag.org/content/345/6202/1349/suppl/DC1
Supplementary Text
Figs. S1 to S4
References (*33*–*40*)

10 March 2014; accepted 24 July 2014
Published online 14 August 2014;
10.1126/science.1253138

Corrected 9 September 2014; see below

www.sciencemag.org/cgi/content/full/science.1253138/DC1

Supplementary Materials for

A memory of errors in sensorimotor learning

David J. Herzfeld,* Pavan A. Vaswani, Mollie K. Marko, Reza Shadmehr

*Corresponding author: E-mail: dherzfe1@jhmi.edu

Published 14 August 2014 on *Science* Express
DOI: 10.1126/science.1253138

This PDF file includes:

Correction: A citation has been added to a recently published work.

In the main text, we formulated a new mathematical model of motor learning in which we hypothesized that the brain not only learned from errors that were experienced in each trial, but also accumulated them into a memory of errors. This memory of errors made it possible for the brain to recognize errors that it had experienced before, resulting in the ability to control error-sensitivity, i.e., control how much it was willing to learn from a given error. Here, we describe the steps that we took to test this mathematical model.

We performed four experiments to test the predictions of the model. In these experiments we considered various motor behaviors, including reaching and isometric paradigms, as well as various perturbations, including proprioceptive and visual modalities. A total of n = 113 right-handed volunteers participated in our experiments. The volunteers had no known neurological problems and were naïve to the purpose of the experiments. Protocols were approved by the Johns Hopkins School of Medicine Institutional Review Board. All subjects signed an informed consent form.

In addition to these experiments, we tested the model by considering a host of previously published results in various motor learning paradigms, including reaching, walking, and saccades. We chose these experiments because they represented puzzling data sets that were difficult to interpret within the current framework of motor learning. We used our model to simulate these experiments, and then compared the results with the experimental data. In sum, we tested whether a single idea, memory of errors, could provide a framework sufficient to unify the disparate data sets into one coherent body of results.

Experiment 1

In Exp. 1 we used a between-subject design to test the idea that the nervous system could modulate error-sensitivity. To do so, we used a constant perturbation in probe trials to produce an error in the movement and quantified how much the nervous system learned from this error.

Volunteers (n = 27, 23.6 ± 4.3 years old, mean ± SD, 16 female) were asked to hold the handle of a robotic arm and make rapid out-and-back reaching movements to a target presented 10cm directly in front of them. In some trials, their reach was perturbed by a velocity dependent curl field, where force $\mathbf{f}$ was related to hand velocity as:

$$\mathbf{f} = \begin{pmatrix} 0 & b \\ -b & 0 \end{pmatrix} \begin{pmatrix} \dot{x} \\ \dot{y} \end{pmatrix} \tag{s1}$$

In Eq. (s1), $\dot{x}$ and $\dot{y}$ are the components of the subject's hand velocity. The force pushed the subject's hand perpendicular to the direction of movement on the outward reach, but was turned off on the

reach back. Perturbations were either clockwise (b =13 N.s/m) or counter-clockwise (b =-13 N.s/m). The subject's hand was occluded by an opaque horizontal screen located above the plane of the arm. An overhead projector displayed information about hand position and targets on this screen. Continuous feedback of the location of the hand was presented via a cursor (0.3cm in diameter). At the onset of each trial, subjects were presented with a start circle (1.0cm in diameter). Once the hand was placed inside this circle, following a randomly chosen inter-trial interval [0.25 - 0.75]sec a target circle (1.0cm in diameter) was displayed and an auditory tone was played. The display of the target and the sound of the tone served as the 'go' instruction. If the hand passed through the target circle in 300±30ms, the subject was rewarded with an animation of an explosion, an auditory tone, and a point added to their score. If the hand passed through the circle after 330ms, the target circle turned blue (indicating a movement that was too slow). Otherwise, if the subject's hand arrived at the target circle in less than 270ms, the circle turned red. The subjects were not required to have the turn-around point of their reach in the target circle; rather, they merely had to pass through the target to obtain reward. In some instances, the reach missed the target entirely. In this case, no reward or timing feedback was provided to the subject. Subjects were instructed to obtain as many points as possible.

Consider an environment in which the perturbations tend to switch slowly, i.e., persist from trial to trial, as compared to one in which the perturbations are rapidly switch. We can define this environment in terms of a Markov chain in which the perturbations can take on one of two states (Fig. 1A, top). For example, the perturbation can be + 1 or -1, where +1 refers to a force field that pushes the hand clockwise, and -1 refers to a field that pushes the hand counter-clockwise. The perturbation state can change from one trial to the next, and this change is governed by a transition probability z. In the slowly switching environment, the probability of staying in a given perturbation state is high ($z = 0.9$), whereas in the rapidly switching environment, this probability is low ($z = 0.1$). As a result, in the slowly switching environment the perturbations tend to repeat from one trial to the next, whereas in the rapidly switching environment the perturbations tend to change. We hypothesized that the brain would learn more from the error induced by the perturbation in the slowly switching environment because that perturbation was likely to persist (learning from error in one trial would improve performance on the subsequent trial). However, in the rapidly switching environment the brain would suppress learning from error because the perturbation that produced that error was likely to change (any learning would be detrimental to performance on the subsequent trial).

To produce such environments, we considered a perturbation schedule that was stochastic, as illustrated by the Markov chain in Fig. 1A. The perturbation state, labeled by variable b (indicating the

field produced by the robot) was a binomial, taking on one of two values [+13, -13]Ns./m. For example, suppose that on trial $n-1$ the perturbation state is $b^{(n-1)}$. Then the perturbation on trial n is determined by the following probabilities:

$$\begin{aligned} \Pr\left(b^{(n)} = b^{(n-1)}\right) &= z \\ \Pr\left(b^{(n)} = -b^{(n-1)}\right) &= 1-z \end{aligned} \tag{s2}$$

Eq. (s2) implies that if $z \approx 1$, then the perturbation is likely to repeat, i.e., the environment is slowly switching and persistent. However, if $z \approx 0$, then the perturbation is likely to change; the environment is rapidly switching.

We randomly divided our subjects into three groups and generated a single perturbation schedule for each group using Eq. (s2). One group (n = 9) experienced a slowly switching environment ($z = 0.9$), another group (n = 9) experienced a medium switching environment ($z = 0.5$), and a final group (n = 9) experienced a rapidly switching environment ($z = 0.1$). All groups began their training in a baseline block (156 trials). In the baseline block there were no perturbations, except for occasional probe trials in which we measured error-sensitivity, described below and illustrated in the inset of Fig. 1A. Following the baseline block, subjects experienced 5 blocks of perturbation trials (225 trials each). In each perturbation block there were 5 mini-blocks (learning block). Each mini-block included 30 perturbation trials, 10 washout trials, and 5 probe trials. In each mini-block the number of trials with clockwise or counter-clockwise perturbation was equal to 15. In this way, the mean of the perturbations in each block, as well as the mean of the perturbations in each mini-block, was zero for all subjects regardless of the environment. Furthermore, variances of the perturbations were identical across groups. The critical difference was the order of the perturbations. The experiment lasted about an hour. Subjects were allowed a 1-3 minute break between each block of trials.

Our objective was to estimate error-sensitivity during each block and ask whether this quantity changed as the subjects experienced the various environments. We approached the problem by considering a standard model of learning(*18, 22, 24*) in which on trial n, a perturbation x is imposed on action u so that the sensory consequences observed by the learner are $y^{(n)} = u^{(n)} + x^{(n)}$. On trial n, the learner predicts the sensory consequences $\hat{y}^{(n)} = u^{(n)} + \hat{x}^{(n)}$, and updates its belief about the state of the perturbation from the prediction error $e^{(n)} = y^{(n)} - \hat{y}^{(n)}$. Such learning typically depends on a decay factor a, and error-sensitivity η:

$$\hat{x}^{(n+1)} = a\hat{x}^{(n)} + \eta^{(n)} e^{(n)} \tag{s3}$$

If we assume that u_0 is the motor command generated in the null environment in which there are no perturbations, then the motor commands on a given trial is a proxy for the learner's estimate of the state of the perturbation:

$$u^{(n)} = u_0 - \hat{x}^{(n)} \tag{s4}$$

To measure error-sensitivity, we used probes that consisted of pairs and triplets of error-clamp trials(*6*). An error-clamp(*21*) is a trial in which the robot produces a channel with stiff walls along a line connecting the start position to the target, thereby reducing deviations from a straight line, eliminating error from that trial while allowing one to measure the forces that the subject produces against the channel walls. The error-clamp had the following properties: spring coefficient = 6000 N/m, damping coefficient = 250 N.s/m. On error-clamp trial n our proxy for motor output was $u^{(n)}$. To find $u^{(n)}$, we first regressed the measured force $f(t)$ that the subject had produced against channel walls onto the ideal force $f^*(t) = b\dot{y}(t)$ and found the parameters k_0 and k_1 that minimized the quantity $\left(f(t) - k_1 f^*(t) - k_0\right)^2$, and then set $u^{(n)} = k_1$. To measure error-sensitivity $\eta^{(n)}$, we first used Eq. (s3) to estimate a for each subject from all pairs of error-clamp trials that did not have a perturbation (Fig. 1A, green probe trials). As the subject did not experience an error in the first error-clamp, the forgetting factor was found by dividing the motor commands in the two trials: $a = \frac{u^{(n+1)}}{u^{(n)}}$. Next, we used this estimate of a to estimate $\eta^{(n)}$ from each triplet of error-clamp trials (Fig. 1A, lavender probe trials) in which there was a perturbation in the middle trial. The perturbation in this probe was always a counter-clockwise field (b =-13 N.s/m). As a result, we have:

$$\eta^{(n)} = \frac{u^{(n+1)} - a^2 u^{(n-1)}}{e^{(n)}} \tag{s5}$$

In the above equation, $e^{(n)}$ is the error on trial n, which we estimated by measuring the displacement of the hand from a straight line to the target at maximum velocity (this took place at 147±6.2ms, mean±SEM, into the movement). We estimated learning from error ($u^{(n+1)} - a^2 u^{(n-1)}$) and error-sensitivity by binning the data for 5 probe trials in each environment block.

Experiment 2

In Exp. 2 we designed a within-subject protocol to test the idea that a change in the history of perturbations would result in a change in error-sensitivity. Furthermore, in this protocol we had the

capacity to measure error-sensitivity on each trial. This allowed us to test whether changes in error-sensitivity were global, affecting learning from all error sizes, or local, specific to a range of error sizes.

We enrolled a new group of right-handed volunteers (n = 20, 24.1 ± 4.5 years old, mean ± SD, including 10 females) who were naïve to the purpose of the experiment. As in Exp. 1, subjects were asked to make rapid out-and-back reaching movements to a target at 10cm. However, unlike Exp. 1, there were no forces to perturb the movement (all movements were in error-clamp). Instead, we perturbed the visual feedback associated with position of the hand. At 100ms after reach onset, we removed the visual feedback, and then re-displayed hand position at the turn-around point of the reach by placing a stationary yellow dot at that location (Fig. 2A). In some trials the location of this dot was perturbed by either a 1.1x (magnifying) or a 0.9x (minifying) gain. We restored visual feedback of the hand after this turn-around point, but manipulated the location of the cursor using a gain so that it appeared that the subject had reached to the location indicated by the yellow dot. A trial was considered successful if the yellow dot fell within the target circle. Subjects were rewarded by a visual animation of an explosion, and the addition of a point to their score. They were instructed to maximize the total number of rewarded trials.

The difference between the visual feedback and the target position was our proxy for error. The brain responded to this error by changing the motor commands on the next trial, increasing or decreasing the extent of the reach. Because visual feedback was not available during the outward portion of the reach, the design of the experiment allowed us to measure error-sensitivity at every trial (the change in the magnitude of the reach divided by the experienced error). Suppose that on trial n, a perturbation was imposed, resulting in an error $e^{(n)}$ (defined as the difference between target position and the cursor position displayed to the subject to indicate their hand's turn-around point). We estimated the forgetting factor a in pairs of consecutive error-clamp trials via a technique identical to Exp. 1 (Fig. 2A, lavender trials).

Subjects were divided into two groups (n = 10 in each group). Both groups experienced a baseline block (100 trials, no perturbations). Following the baseline block, Group 1 (Fig. 2B) experienced three perturbation blocks, each 387 trials, composed of a slowly switching environment (z = 0.9), a medium switching environment (z = 0.5), and a rapidly switching environment (z = 0.1). Group 2 experienced the reverse sequence of environments. All subjects in each group experienced the same perturbation schedule. Each block was composed of 9 mini-blocks (30 perturbation trials, 10 no perturbations, and a probe triplet of trials). The mean of the perturbations within each mini-block, as well as the mean of the perturbations within each block, was zero.

We found that a change in perturbation statistics resulted in a change in error-sensitivity (Fig. 2C), and that the largest changes occurred where subjects experienced the majority of the errors (Fig. 2E and Fig. 2F). This suggests that error-sensitivity was a function of the experienced error.

A model of error sensitivity

When participants experience a prediction error, they update their motor command on the next trial to compensate for a fraction of that error. This can be mathematically described by a state space model (Eqs. s3 and s4), where $e^{(n)}$ is the error and a is a retention factor. Eq. (s3) describes a model in which the learner uses prediction error to form an estimate of the state of the environment, resulting in a memory of that state. Because our results from Exp. 2 suggested that error-sensitivity was a function of error, we constructed a new set of equations to account for a memory of errors (Eq. 2 and 3). In this model, the learner has a set of basis elements with which it encodes the error experienced on a given trial. In our simulations, we assumed that N basis functions with centers located at $\breve{e}_i$ that were uniformly distributed throughout a symmetric error space $\breve{e}_i \epsilon [-P, P], i = 1 \ldots N$. In addition, we assumed that at the beginning of the simulations all weights were equal (i.e. there was a constant error-sensitivity), $\mathbf{w}^{(0)}$. Therefore, our model that learned to represent the state of the environment had one parameter, a. Our model that controlled error-sensitivity had two parameters: σ, and β. In total, our model had 3 parameters.

In Fig. 3B we implemented this model with 10 basis elements, $\breve{e}$ equally spaced between -5 and 5, $\sigma = 1$, $a = 1$, and $\beta = 0.05$. In the model, sensitivity is a function of error size, and so any change in sensitivity is local to the errors experienced in the recent trials. Suppose that on trial $n-1$, the motor command produces error $e^{(n-1)} = -1$. If on the next trial the error $e^{(n)}$ is of the same sign as $e^{(n-1)}$, then $\text{sign}\left(e^{(n-1)} e^{(n)}\right) = 1$ and error-sensitivity is increased around the neighborhood of $e^{(n-1)}$ (Fig. 3A, top). As a result, learning from error is increased about $e = -1$, as illustrated by the red line in Fig. 3B. This means that if this error is ever experienced again, the system will learn more from it than before. On the other hand, if $e^{(n-1)}$ is of the opposite sign as $e^{(n)}$, $\text{sign}\left(e^{(n-1)} e^{(n)}\right) = -1$ (Fig. 3A, bottom), and sensitivity is decreased about $e = -1$, resulting in reduced learning from error around this neighborhood (as illustrated by the blue line in Fig. 3B).

In Fig. 3C we simulated the model in the slow (z = 0.9), medium (z = 0.5), and rapidly switching (z = 0.1) environments (identical parameters as in Fig. 3B except $\beta = 0.001$). In all cases, the model learns

from error on each trial, as illustrated by the gray line in Fig. 3C. However, the errors in the slowly switching environment tend to repeat, that is $E\left[\text{sign}\left(e^{(n-1)}e^{(n)}\right)\right]>0$, where $E[\ ]$ is the expected value operator. As a result, in the slowly switching environment error-sensitivity increases, producing an increase in learning from error (Fig. 3D, red-line). The errors in the medium switching environment have the following structure: $E\left[\text{sign}\left(e^{(n-1)}e^{(n)}\right)\right]\approx 0$. This produces little or no change in sensitivity, resulting in little or no change in learning from error, as illustrated by the green line in Fig. 3D. Finally, errors in the rapidly switching environment have the following structure: $E\left[\text{sign}\left(e^{(n-1)}e^{(n)}\right)\right]<0$, that is, error in one trial is usually of the opposite sign of the error in the previous trial. As a result, error-sensitivity decreases (Fig. 3D, blue line). The learning from error curves are qualitatively similar to those measured experimentally in Exp. 2 (Fig. 2D).

Experiment 3

In Exp. 3 we set out to test a critical prediction of the model: that by manipulating the history of errors that were experienced by the subject, we could simultaneously increase error-sensitivity for one range of errors, while decreasing it for another range.

We recruited a new group of right-handed volunteers (n = 16, 25.8 ± 2.6 years old, mean ± SD, including 6 female) who were naïve to the purpose of the experiment. They held a handle attached to a stationary force transducer. The handle was located 20-30cm in front of the subject, such that they could push against it comfortably while seated. The subject's hand was hidden from view by an opaque horizontal screen. Feedback regarding force generation was provided by an image projected on the screen. The objective of this isometric task was to produce a goal force of 16N.

At the onset of a trial, a start circle and a goal circle (both 0.75N in diameter) appeared. The goal circle was located approximately 15cm from the start circle. The screen was scaled such that a 15cm cursor displacement corresponded to 16N force. A cursor (0.3N diameter) appeared at trial onset. The displacement of the cursor corresponded to the total amount of force that the subject produced, multiplied by a scaling factor: $s\sqrt{f_x^2+f_y^2}$, where s corresponds to a scaling factor that maps units of force into screen displacement (15cm/16N).

As the subject began pushing toward the target, we removed visual feedback when the cursor position reached 1/5 of the way to the target (>3.20N), and then placed a yellow dot on the screen (0.5N diameter) in the location corresponding to the maximum force that they produced. In some trials we

perturbed the location of this dot by adding an offset x. Visual feedback of the cursor was then restored as the force produced by the subject returned to zero. We scaled the position of the cursor continuously during the return so that it appeared that the subject had produced the force signified by the perturbed dot. Once the cursor had returned to the starting circle, the maximum force dot remained on the screen for 0.5s before the cursor, maximum force dot, and target circle disappeared. The subject then waited for an inter-trial-interval to elapse (randomly chosen between [0.25, 0.75]sec) before the next trial began.

A trial was successful if the yellow dot, corresponding to the subject's maximum force (plus the perturbation) landed inside the goal target. Feedback of a successful movement was indicated by an animation of an explosion and a point added to the score. If the subject failed to produce a force greater than 3.20N within 1.5 seconds of the go cue, the trial was aborted. Subjects were instructed to maximize the number of points.

The perturbation schedule is shown in Fig. 3E. The perturbations were designed so that, in theory, subjects would increase their sensitivity to +4N and -4N errors (despite the fact that they never experienced a -4N perturbation), while simultaneously decreasing their sensitivity to +8N and -8N errors. The experiment began with a baseline block (50 trials, no perturbations). Following the baseline block, we probed sensitivity to +8N and -4N perturbations (labeled as Probe 1 in Fig. 3E). We then exposed subjects to 20 repetitions of an alternating [+8N, -8N] environment, followed by 15 repetitions of a stable [0N, +4N] environment. Finally, we again probed sensitivity to +8N and -4N perturbations (labeled as Probe 2 in Fig. 3E). The experiment was divided into 6 blocks (105-120 trials each). The experiment took approximately 45 minutes to complete.

How well could the model explain behavior in this experiment? In Figure S1A we have plotted learning from error as a function of error size for all subjects across the entire data set. Remarkably, we found that the learning from error was not monotonic. Rather, subjects learned significantly more from a ±4N error than ±8N error (paired t-test, $t(15) = 7.76$, $p < 0.001$). Figure S1B plots the difference between learning from error and the regression line (our proxy for an unbiased learner). This reflects the change in learning that has been caused by the changes in error-sensitivity. We ran our model on the same sequence of errors that each subject experienced in Exp. 3 and have plotted the predicted change in Figure S1B (scaled by a multiplicative coefficient to convert to units of Newtons). The correlation between the predicted and observed values was $R^2 = 0.65$, $p < 10^{-8}$. Therefore, as the model

had predicted, we were able to use the history of errors to simultaneously increase learning at one error size, and decrease learning at another.

Modeling of previous experiments

Let us use this model to shed light on a set of puzzling observations in the field of motor learning, in paradigms such as reaching, walking, and saccades. For each simulation, we chose model parameters such that the errors experienced by the model were similar to the errors reported in the respective papers. However, similar qualitative results would be obtained for other parameter values.

1. Why does learning from error saturate with large perturbations?

Fine and Thoroughman (*29*) examined reaching movements that were perturbed by a force-pulse. The perturbations were drawn from a discrete uniform distribution: [±6, ±12, ±18]N and presented in 80% of the trials. They examined learning (change in motor commands from the trial prior to the trial after the perturbation) and noted that this measure did not grow linearly with perturbation size, but saturated (black bars, Fig. S2A). We ran our model on a perturbation schedule with the same distribution (50 iterations, 360 trials each, 80% perturbations, discrete perturbations drawn uniformly from [±6, ±12, ±18]N). In our model, error was encoded by 50 Gaussian bases distributed throughout an error space between ±30N, with a standard deviation of 7N. The initial weights of this network were set so that the error-sensitivity was roughly at 20%, a value typical for force field tasks(*22*). We allowed the weights to change with $\beta = 0.005$. After simulating the motor output for each trial, we estimated the learning following a particular perturbation by measuring the mean trial-to-trial change in motor output for each of the discrete perturbations. The results are shown with the red line in Figure S2A. The correspondence between model and experimental data is R^2 = 0.99±0.005 (mean±SD of 50 simulations).

Wei and Körding (*5*) measured learning from error in a visual displacement task (Fig. S2B, black bars). Subjects were asked to reach in the horizontal direction, and the location of the visual feedback was displaced in the vertical direction. Participants experienced 900 trials in which the perturbation on each trial was drawn from a discrete uniform distribution: [0, ±1, ±2, ±4, ±8]cm, and learning was measured by changes in motor commands in the vertical direction. We simulated the results for 50 randomly generated perturbation schedules with the same characteristics. We distributed 50 Gaussian bases throughout an error space between ±10cm, with a standard deviation of 1cm. The initial sensitivity was set at 5%. We allowed the weights to change with $\beta = 0.005$ (i.e., unchanged from the

above simulations). The model results are shown by the red line in Figure S2B. The correspondence between model and experimental data is $R^2 = 0.70 \pm 0.17$ (mean ± SD of 50 simulations).

Why did the learning saturate for large perturbations? The model explained that this was because the perturbations that were used in these experiments were drawn from a bounded uniform distribution. With such a distribution, error-sensitivity declines (and as a consequence, learning from error saturates) for the large errors produced by the perturbations near the bounds. This is because after experiencing an error from a perturbation near one of the bounds, it is much more likely that the next perturbation will produce a change in the sign of the error than not. To illustrate this, in the Figure S3A (top subplot) we have plotted the quantity $E\left[\text{sign}\left(e^{(n)}e^{(n+1)}\right)\right]$ as a function of the perturbation size on trial *n*. We find that the expected value is positive for small perturbations and negative for large perturbations. As a consequence of this uniform bounded distribution of perturbations, error-sensitivity decreases for large perturbations, as illustrated by the learning from error curve in Figure S3A (right subplot). It is also important to note that because the mean of the perturbation distribution is zero, the positive expected values around the mean perturbation produce little or no changes in learning from error function (Figure S3A, right subplot). That is, even though sensitivity increases for small errors ($e \approx 0$), learning from error is small for these errors. The dominant effect is the reduced error-sensitivity for errors other than those produced by the mean of the perturbation.

2. Why does error-sensitivity depend on the mean of the perturbation distribution?

Fine and Thoroughman (*40*) performed a force-field adaptation study in which the perturbations were similar to their previous study(*29*). However, subjects practiced in one of three environments: unbiased, in which the magnitude of the force-field was drawn from a zero mean discrete uniform distribution [+36, +24, +12, -12, -24, -36]Ns/m, weakly biased [+18, +9, -9, -18, -27, -36]Ns/m, or strongly biased [-6, -12, -18, -24, -36]Ns/m (Fig. S2C). If error-sensitivity is independent of error history, the three groups of data points should have the same slope. However, the authors found that the slope of the learning curve vs. perturbation magnitude was greater (more steep) for the strongly biased distribution, and smaller for the unbiased distribution. We ran our model on the same distributions (50 simulation runs per distribution) and have plotted the results in Fig. S2C. We used the same choice of model parameters for this simulation as we did for Fine and Thoroughman (*29*), except that the error space was enlarged to accommodate the larger perturbations ($N = 50$, $P = 50\,\text{N}$, $\sigma = 7\,\text{N}$, $\beta = 0.005$). The model results are shown in Fig. S2C. The correspondence between the model and experimental data is

R^2 = 0.97 ± 0.008 (mean ± SD). Note that the model SEM bars are very small in Fig. S2C and are obscured by the circle data markers.

In Figure S3B we have plotted the quantity $E\left[\text{sign}\left(e^{(n)}e^{(n+1)}\right)\right]$ as a function of perturbation size for this experiment. Each function has a peak near the mean of the corresponding distribution, but there is a clear asymmetry in the function associated with the strongly biased distribution, whereas the function is symmetric for the unbiased distribution. When the distribution of the perturbation is unbiased, the function $E\left[\text{sign}\left(e^{(n)}e^{(n+1)}\right)\right]$ is symmetric and generally negative for non-zero errors, suppressing learning from errors that arise from both positive and negative perturbations. In contrast, when the distribution is strongly biased, the expected value is asymmetric, only suppressing learning from errors that arise from large negative perturbations. In both cases, sensitivity is reduced to perturbation near the bounds of the uniform distribution, as in (*5, 29*). However, because of the distribution of errors in the strongly biased group, the expected change in sensitivity is asymmetric. As a consequence, learning from perturbation (Figure S2C) is shallow for the unbiased distribution since sensitivity is reduced for most perturbation magnitudes, but steep for the strongly biased distribution.

3. Why does error-sensitivity depend on the sequential order of the perturbation distributions?

In the above experiments, each group of subjects experienced one distribution of perturbations. Let us now consider what happens when different group of subjects experience a given set of distributions in distinct sequence.

Semrau, Daitch and Thoroughman (*33*) performed a visuomotor rotation experiment. Similar to Fine and Thoroughman (40), perturbations were presented in 80% of the trials. However, subjects were exposed to a sequence of three different perturbation distributions: unbiased [±30, ±20, ±10, 0]°, weakly biased [+30, +22.5, +15, +7.5, 0, -7.5 -15]°, or strongly biased [30, 25, 20, 15, 10, 5, 0]°. The order of the environments was counterbalanced across the two groups of participants. The authors found that there was an order effect: learning in a given environment depended on the specific order in which the environments were experienced. Subjects that experienced the strongly bias distribution first showed a steep learning function (Fig. S2D, top subplot, blue line), whereas participants that experienced the strongly biased perturbation last (Fig. S2E, top subplot, blue line) showed a shallow learning function.

We constructed a perturbation schedule which copied this design and simulated our model with the same order of perturbation distributions. As before, the error space was encoded by 50 Gaussian bases throughout a 40° error space with an initial sensitivity of 20% ($\sigma = 10\,°$, $\beta = 0.005$, $a = 0.8$). The

model's results are shown in Figures S2D and S2E. The correspondence between model and experimental data is $R^2 = 0.86$ (Figure S2D, $p < 10^{-8}$) and $R^2 = 0.93$ (Figure S2E, $p < 10^{-10}$).

The group that experienced the strongly biased distribution last had already experienced the weakly biased and unbiased distributions. Each of these prior experiences reduced error sensitivity for the perturbations near the bounds. Because all three distributions shared one of the bounds, the reduction in the error-sensitivity was particularly strong near this bound for the group that had experienced the unbiased and weakly biased distributions before the strongly biased distribution. As a consequence, the group with the prior experience showed a shallower learning curve than the group without the prior experience.

4. Model explains data attributed to structural learning

Let us now consider a remarkable observation termed structural learning. A key experiment is that of Turnham, Braun and Wolpert (*16*), in which subjects perform a visuomotor rotation task. The participants in the random condition trained in an environment in which the perturbation on the odd trials was drawn randomly and without replacement from a discrete sequence uniformly spanning the range [-60^o to $+60^o$]. Importantly, the even trials had a perturbation that was always the same as in the previous odd trial. That is, perturbations repeated twice in a row. [In the experiment, subjects were tested in 8 directions, termed a 'cycle'. Here, we considered a simpler version of this experiment in which there is only one direction and a trial represents a cycle.] An example of such a random distribution of perturbations is shown in Figure S2F (left sub-figure, black line). Following this training, the subjects were tested in a series of constant visuomotor perturbations of [+30, -30, +30]°. The authors found that as compared to a control group, subjects that had this prior training in the random perturbations showed faster learning in the subsequent constant perturbations (Fig. S2F, top-right). That is, the experience of the random perturbations appeared to facilitate learning of a constant perturbation, a phenomenon that the authors interpreted as structural learning. A second group completed a gradually imposed perturbation which spanned the range [-60° to +60°]. Importantly, subjects in this experiment did not fully compensate for the errors when the gradually imposed perturbations became large (c.f. Fig. 3 of (*16*)), resulting in persistent and repeating errors (see Figure S3C, right subplot).

We constructed similar perturbation schedules as the authors of this experiment. We used simulation parameters that were identical to those used in the visuomotor rotation experiment described above for Semrau and colleagues except that we increased the size of the error range to

accommodate the larger perturbations. As before, we used 50 Gaussian bases to span the error range [-150, 150]°, with an initial sensitivity of 20%. We assumed that each of these bases had a standard deviation of 10°. Finally, we set $\alpha = 0.8$ in Eq. (1) and $\beta = 0.005$. Our model reproduced the basic observation that following training in the random sequence, learning was faster than control in the series of constant perturbation of [+30, -30, +30]° (Fig. S2F, bottom-right). In addition, the model reproduced the result that learning was faster than control following the gradual perturbation. Finally, the model reproduced the observation that the gradual group performed slightly better than the random group in both the test of +30 and test of -30 conditions.

According to our model, the key fact in these experiments was the repetition of errors in the initial training, and then re-experiencing of these errors in the subsequent testing. In the random group, the repetition came about because every even trial had the same perturbation as the previous odd trial, resulting in sequence of errors that were likely to have the same sign, up-regulating error-sensitivity for the error experienced in the odd trial. In the gradual group, the repetition came about because the subjects could not adapt as fast as the gradually imposed perturbation, resulting in errors that accumulated and repeated near the end of the gradual perturbation (see Figure S3C, right subplot for data from the original experiment, and Figure S2F, bottom subplot for simulation data).

The model explained that in the random group, the faster than control learning that was observed was not due to a memory of perturbations (as the mean of the perturbation sequence was zero), but due to the accumulation of memory of errors. The +30° constant perturbation produced a +30° error, for which error-sensitivity had increased due to exposure to the previous 'random' perturbation. This is shown by the quantity $E\left[\text{sign}\left(e^{(n)}e^{(n+1)}\right)\right]$, plotted for various error sizes in Figure S3C (left subplot, labeled 'paired random'). Note that the $E\left[\text{sign}\left(e^{(n)}e^{(n+1)}\right)\right] > 0$ for all errors, resulting increased error-sensitivity for all errors. The model explained that this increase in error-sensitivity would not have occurred if the perturbations on the even trials were unrelated to the perturbations on the previous odd trial (Figure S3C, left subplot, labeled 'single random') where $E\left[\text{sign}\left(e^{(n)}e^{(n+1)}\right)\right] \approx 0$ for all errors. As a result, we found that error-sensitivity remained at near baseline, and the model did not show savings, as illustrated in Figure S3D.

The model explained that in the gradual group, the faster learning that was observed (with respect to control) for +30° perturbation was because at the end of gradual training subjects experienced repeated exposure to +30° errors (Figure S3C, right sub-plot). The repeated +30° errors resulted in

increased error-sensitivity for this error, which accounted for the fact that when they were tested with a +30° perturbation, the resulting +30° errors produced faster learning than in control.

In Figure S3F we have summarized the results of the various simulations. Learning from error increased in the paired-random perturbations, as well as in the gradual condition. As a result, the model suggested that data attributed to structural learning could be explained by memory of errors. The random and gradual conditions had resulted in a memory of errors for which error-sensitivity had increased, and during testing the subjects experienced errors similar to those they had experienced before. It was the repetition of errors, followed by subsequent revisiting of these errors that resulted in faster learning than control.

5. Model explains savings following washout

A number of studies have considered the phenomenon of savings by conducting experiments in which subjects are trained with a constant perturbation, and then the perturbation is removed for a long duration, producing washout. Intriguingly, following this period of washout the re-learning of the perturbation is faster than control. Current state-space models in which the only memory is one of perturbations cannot account for such data(*13, 14*). Instead, some models have proposed that savings arises from a hypothetical ability of the brain to recognize a context and protect the memory from washout(*34*). In addition, other models have proposed that savings arises because during training certain motor commands are reinforced by repetition and reward(*30*). However, neither of these two hypotheses can explain the meta-learning results that we highlighted in Figure S2F. Let us show that the same memory of errors that accounted for meta-learning also readily accounts for these savings experiments.

Consider a typical scenario termed A, Null, A (ANA). In this simulation we exposed the model to 20 trials of a +1 (a.u.) perturbation, 20 trials of washout (0 perturbation) followed by 20 trials of relearning of the +1 perturbation (Figure S2G, top). We assumed an error region of ±3 consisting of 20 bases whose initial sensitivity was 20% ($\sigma = 0.25$, $\beta = 0.01$, $a = 1$). We found that despite washout, and the fact that memory of perturbation $\hat{x}$ had returned to zero, the model exhibited savings, i.e., faster rate of learning in the second exposure to the +1 perturbation. The reason for this savings, the model explained, was the fact that the previous errors had been experienced in a stable environment, enhancing error-sensitivity for the errors that were again experienced in the second exposure to +1.

In a second simulation, we exposed the model to A, Null, B, Null, A (ANBNA), as shown in Fig. S2G, lower subplot. The idea in this simulation was to illustrate that the washout of (A) produces after-

effects, which are errors that are subsequently revisited in (B). The model made a crucial prediction: that learning of (-1) would be faster than control, despite the fact that the model had never before been trained in a (-1) perturbation. This example of meta-learning is explained by our model via the fact that the after-effects following learning of (+1) produce enhanced error-sensitivity to the errors that are again experienced in the ensuing (-1) perturbation. We will present a test of this prediction in Experiment 4.

6. Model explains savings in a gait-adaptation experiment

Malone, Vasudevan and Bastian (*11*) considered an experiment with a sequence of perturbations similar to the one that we simulated in Figure S2G. In their split-belt gait adaptation task, subjects in the ANA group were asked to adapt to a perturbation in which the belt under the non-dominate leg was moving twice as fast as the dominant leg for 15 minutes. Subjects in this group then returned the next day and were exposed to a tied belt-condition (washout), followed by an additional re-learning period in the split belt condition. Subjects in the ANANA group were exposed to an additional cycle of 15 minutes of tied belt followed by 15 minutes of split-belt before leaving at the end of day 1. In the ANBNA group, subjects we exposed to 15 minutes of tied belt followed by 15 minutes of adaptation to a split-belt condition in which the non-dominant leg was moving half was fast as the dominate leg (opposite of the 'A' perturbation) before leaving after day 1. The results of this study are presented in Figure S2H (top). We simulated a similar perturbation schedule in which each of the 15 minute perturbation/null sessions was approximated by 100 trials of a perturbation with magnitude ±1 (adaptation) or 0 (washout). We distributed 25 bases throughout an error space of ±5. Initial sensitivity was set at 1% ($\sigma = 0.5$, $a = 0.9$, $\beta = 0.0005$) (Figure S2H, bottom). Consistent with the experimental data, our model exhibited the savings in ANA, and greater savings in ANANA.

According to the model, savings in ANA occurred because errors that were initially experienced in (A) were re-experienced in the second exposure to (A). The savings in ANANA occurred because the subject gained two prior exposures to (A) before the final test, resulting in greater increase in error-sensitivity than one prior exposure. Finally, the savings in ANBNA occurred because in addition to the errors in initial (A), subjects experienced similar errors upon the washout trials following (B). That is, in ANBNA subject also had two prior exposures to the errors of the (A) perturbation, despite the fact that they only experienced (A) once.

7. Model explains the limited range of savings

Kojima, Iwamoto and Yoshida (*10*) performed a saccade adaptation experiment in monkeys to quantify savings. In this experiment, a standard intra-saccade step paradigm(*35*) was used to produce errors that resulted in adaptation. After collecting 400-800 saccades during a 'gain-up' adaptation period (shown in Fig. S2I, top), the direction of the intra-saccade step was reversed until the animal was making saccades with an approximate gain of 1.0. The duration of this period of counter-adaptation was approximately the same as in the adaptation period. Finally, the monkey was exposed to a period of re-adaptation on the gain-up perturbation, but this block contained a larger number of trials than did the initial learning block (Figure S2I, top). The behavior showed clear evidence of savings, but the important observation was that the faster re-learning was present only in the first 100 or so trials, after which the learning curve returned to a rate similar to initial adaptation (red arrow, Figure S2I, top). Why was the faster learning present for only a limited number of trials?

To simulate this experiment, we constructed a perturbation schedule consisting of 300 trials of baseline movements, followed by adaptation to a 3.5° intra-saccade step over 750 trials. The learning of this perturbation was then washed out using the counter perturbation over 650 trials, before being re-exposed 1100 trials of the gain-up perturbation (Figure S2I, bottom). We distributed 50 bases in an error-region spanning ±6° with initial sensitivity of 0% ($\sigma = 0.25$ °, $\beta = 0.00001$, $a = 1.0$). Just as in the experimental data, our model also showed the fast initial re-learning, and then a return to slow learning following the first 100 or so trials.

The model explained this limited range of savings by noting that the inflection point occurred near the limit of the previously exposed errors. That is, saving was present only up until the errors that were previously experienced -- the errors for which error-sensitivity had been up-regulated.

8. Model explains savings that was attributed to reinforced repetition

A current hypothesis posits that in some conditions, savings may be the result of reinforcement of motor actions during the adaptation period. In their experiment, Huang, Haith, Mazzoni and Krakauer (*30*) constructed 4 different perturbation schedules for a visuomotor rotation task. In the first group, $Adp^{-}Rep^{-}$, subjects made movements with veridical feedback to targets drawn from a uniform distribution between 70° and 110°. This group then learned a constant 25° perturbation to a target located at 95°. In another group, $Adp^{+}Rep^{-}$, subjects moved to random targets between 70° and 110° (identical to the $Adp^{-}Rep^{-}$ group), except that the perturbations were randomly selected for each target from the uniform bounded distribution [0, 40]°. Therefore, this group adapted to the mean of the perturbation (20°), but did not repeat actions to a particular target. In the $Adp^{+}Rep^{+}$ group, subjects

were presented targets from a uniform distribution, but perturbations were chosen so that the correct solution to all rotations would be to move in the 70° direction. The final group, *Adp*$^{-}$*Rep*$^{+}$, made repeated movements to the 70° target in the absence of a perturbation. Each group experienced 80 trials of null movements, followed by 160 training trials, 80 trials of washout, and then a final test phase to the 25° perturbation for an additional 80 trials.

We simulated the same experiment for 50 random perturbation schedules for all groups except *Adp*$^{-}$*Rep*$^{-}$, which the authors found was not significantly different than the repetition control group, *Adp*$^{-}$*Rep*$^{+}$. We incorporated a model of movement generalization into our model which simulated the effects of generalization of motor commands to nearby targets (identical to (*36*)). We distributed 50 basis elements throughout an error-region of ±50° with an initial sensitivity of 10% ($\sigma = 10°$, $\beta = 0.05$, $\alpha = 1.0$). The authors found that only the *Adp*$^{+}$*Rep*$^{+}$ group showed savings after a washout block. Our model reproduced this result (Figure S2J), and explained that the reason was that the *Adp*$^{+}$*Rep*$^{+}$ group experienced errors that up-regulated error-sensitivity which were then re-visited during the test of savings.

We made two assumptions in this simulation. First, that motor commands generalize to nearby targets according to the model described by Tanaka, Sejnowski and Krakauer (*36*). Second, that the weights of the error-sensitivity bases are not target specific. That is, change in error-sensitivity is a function of error, and therefore generalizes fully to other targets. While we have not explicitly tested either of these predictions, we note that repetition of a rewarded action alone cannot account for the meta-learning results we have noted above.

Experiment 4

According to our model, savings and meta-learning are largely due to a memory of errors. If so, specific manipulations of the history of errors should affect the presence or absence of savings and meta-learning. In Exp. 4, we tested some of these critical predictions.

This experiment included n=50 subjects, 10 in each of the five groups (24.3 ± 5.4 years old, mean ± SD, including 20 females). The subjects were naïve to the purpose of the experiment. They held the handle of the robotic manipulandum. The subject's hand was hidden from view by an opaque horizontal screen. They were presented with a red circle (1cm diameter), which served as the start and end point for the trial. Subjects were asked to make a rapid shooting movement from the starting circle to a green target circle (0.5cm diameter) located 6cm directly in front of them. They were required to pass through the target within 150±50ms. Movements that fell outside this range were signaled by the target circle

changing color and a low frequency auditory tone. If the subject passed through the target circle within the time range, they were rewarded with an animation of an explosion and a point added to their score.

As subjects began moving towards the target (when the total velocity exceed 0.02m/s), we removed visual feedback for the remainder of the outward motion. When the reach exceeded 6cm eccentricity, a yellow dot (0.5cm diameter) was placed on the screen at that location. Visual feedback was withheld for the duration of the trial. To aid the subject in returning to the start position, when the hand was within 3cm of the starting circle, the position of the cursor was provided by a white dot (1mm diameter) which blinked at 1Hz (20% duty cycle). In some trials, we manipulated the location of the yellow dot by rotating it relative to the target using the perturbation schedules shown in Figures 4A and S4A.

We began with a control experiment to establish the basic idea that savings is present despite washout. In the ANA group (Fig. 4A), perturbation (A) was imposed, and then following a period of washout (N), perturbation (A) was again presented. We expected to observe a faster rate of learning in the second exposure to (A). According to the model, this faster learning was present because during the initial exposure to (A) the perturbation represented a stable environment, increasing sensitivity to the errors experienced in (A). Savings occurred because these errors were re-visited in the subsequent test of (A). If this is indeed the case, we should be able to prevent savings by presenting (A) so that the errors that are experienced during initial exposure are unlike the large errors that are experienced during re-exposure to (A). We tested this prediction with the GNA group (Figure S4A). In the GNA group, the perturbation was incremented gradually to (A), as opposed to a sudden presentation. As a result, subjects learned to make movements typical of learning to (A), but did not experience the same errors.

In the BNA group we attempted to produce meta-learning by exposing subjects to perturbation (B), and then null, followed by perturbation (A) (Figure 4A). In such a case, the model predicts that the savings occurs because the learner is exposed to errors during sudden washout of (B), i.e., the after-effects. The interesting idea is that the after-effects themselves present a sequence of stable errors, which increase error-sensitivity. Because these after-effect induced errors are re-experienced during subsequent learning of (A), subjects should show meta-learning, despite the fact that they have not previously experienced (A) before.

If the meta-learning in BNA is due to the errors that are experienced during washout of (B), then we should be able to block it by eliminating the errors in the washout. To do so, we considered two groups. In the B_{wait}NA group we installed a brief delay between (B) and the subsequent washout period

(Figure 4A). This wait period should reduce the after-effects in the subsequent washout trials(*23, 37*). Therefore, installation of a brief wait period would remove exposure to errors, the same errors that are the part of learning of (A). In this case, the model predicted that we should see no savings.

We followed this idea with a second group in which we eliminated the after-effects by introducing errors gradually in BGNA group (Figure S4). In BGNA, after exposure to (B) we gradually removed the perturbation so that there would be little or no errors that are similar to those that the subjects would experience during exposure to (A). In summary, the model predicted savings in ANA but not GNA, meta-learning in BNA but not B_{wait}NA and BGNA.

All groups experienced a baseline block of 90 null (no perturbation) trials. Each group experienced a training condition (90 or 120 trials) followed by 120 trials of washout (null, N), and a final phase in which we tested adaptation to an abruptly imposed 30° counterclockwise (CCW) perturbation over 90 trials (A). The ANA group was trained on a 30° CCW rotation over the course of 90 trials (A) before washout and testing. The BNA group experienced a 30° clockwise rotation (B), followed by washout testing. The B_{wait}NA group experienced the same perturbation schedule as BNA. However, subjects were asked to wait 1-2 minutes after training of (B). This delay was expected to reduce the adapted motor output(*38, 39*), resulting in reduced after-effects, i.e. smaller errors, when subjects experienced the washout condition. The GNA group experienced a 30° CCW rotation that was gradually imposed over 120 trials. The BGNA group experienced a 30° abruptly imposed CW rotation that was then gradually removed.

To test for savings, we fit an exponential to the performance in the test condition for each subject

$$k\exp(-t/\tau)+c \qquad \text{(s6)}$$

In Eq. (s6) the exponential time constant τ has units of trials. Therefore, savings compared to the initial learning of A by the ANA group is represented as a decrease in the value of τ.

Indeed we found that the subjects in the ANA and BNA groups learned significantly faster than control in the test of perturbation (A), whereas the GNA, BwaitNA, and BGNA groups learned at a rate that was no different than control.

Why does gradual perturbations sometimes produce savings, and sometimes not?

It is puzzling that in certain examples of a gradual perturbation there can be evidence of savings (Figure S2F), whereas in other examples of a gradual perturbation savings is precluded (Figure S4A, GNA group). Why?

The model explains that the critical factor is the history of errors during learning. In the experiment shown in Figure S2F in which savings is present, the gradual perturbation produced large errors, as shown in Figure S3C (right subplot). These residual errors at the end of the gradual perturbation were the same errors that were experienced by the subjects when they were tested in a +30 degree perturbation. In contrast, in the GNA group (Figure S4A), the errors at the end of the gradual perturbation were much smaller than the errors that are experienced at the onset of the +30 degree perturbation. As a consequence of these differing history of errors, one form of gradual training results in savings (Figure S2F), while another does not (Figure S4C).

Data collection and statistical analysis

Movement kinematics (position, velocity) and force information were recorded at 200 Hz. We were able to measure hand position at a resolution of better than 0.1mm, and force at a resolution of 1/80N. Statistical analyses were performed in SPSS 21 (IBM, NY). We used one-way ANOVAs (when there were more than 2 groups) or independent two-sided t-tests (when there were 2 groups) to compare the between-group differences at a single point in the experiment. We used paired t-tests to compare the results of two consecutive probes.

The standard statistical test used in adaptation studies is repeated-measure ANOVA (RM-ANOVA). In RM-ANOVA, the assumption is that the between subject variance of the measured variable is constant across measurements. However, in motor adaptation studies, the across-subject variance of the measured variable often changes as the experiment progresses, violating a primary assumption of a repeated measures ANOVA. To address this problem, we used the general linear model feature of SPSS (GLM-ANOVA) to test for main effects of trial, group, and group by trial interaction. Our analysis assumed a heterogeneous autoregressive correlation structure of the variance matrix, allowing for between-subject variance to change across repeated measurements.

In cases where there was a significant main effect of group or a group by trial interaction, we performed a series of post-hoc t-tests to determine which groups were significantly different. Post-hoc tests were corrected for multiple comparisons using the Dunn-Sidak approach. All figures and statistics are reported as ±SEM, unless otherwise noted.

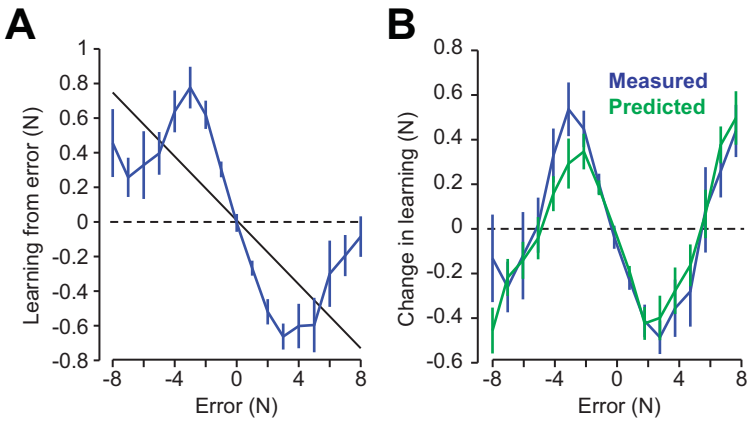

Figure S1. Comparison of model and experimental results for Exp. 3. **A**. Experimental results. Learning from error as a function of error size, measured across the experiment. Errors were binned across error sizes with a bin width of 1N. Black line represents the best-fit line, corresponding to a constant error-sensitivity across error magnitudes, our estimate of an unbiased learner. Error bars are mean±SEM. **B**. Difference between the measured learning from error and the unbiased learning curve. Predicted curves show the change in error sensitivity as predicted by the model, binned across error-magnitudes. Error bars represent mean±SEM across subjects.

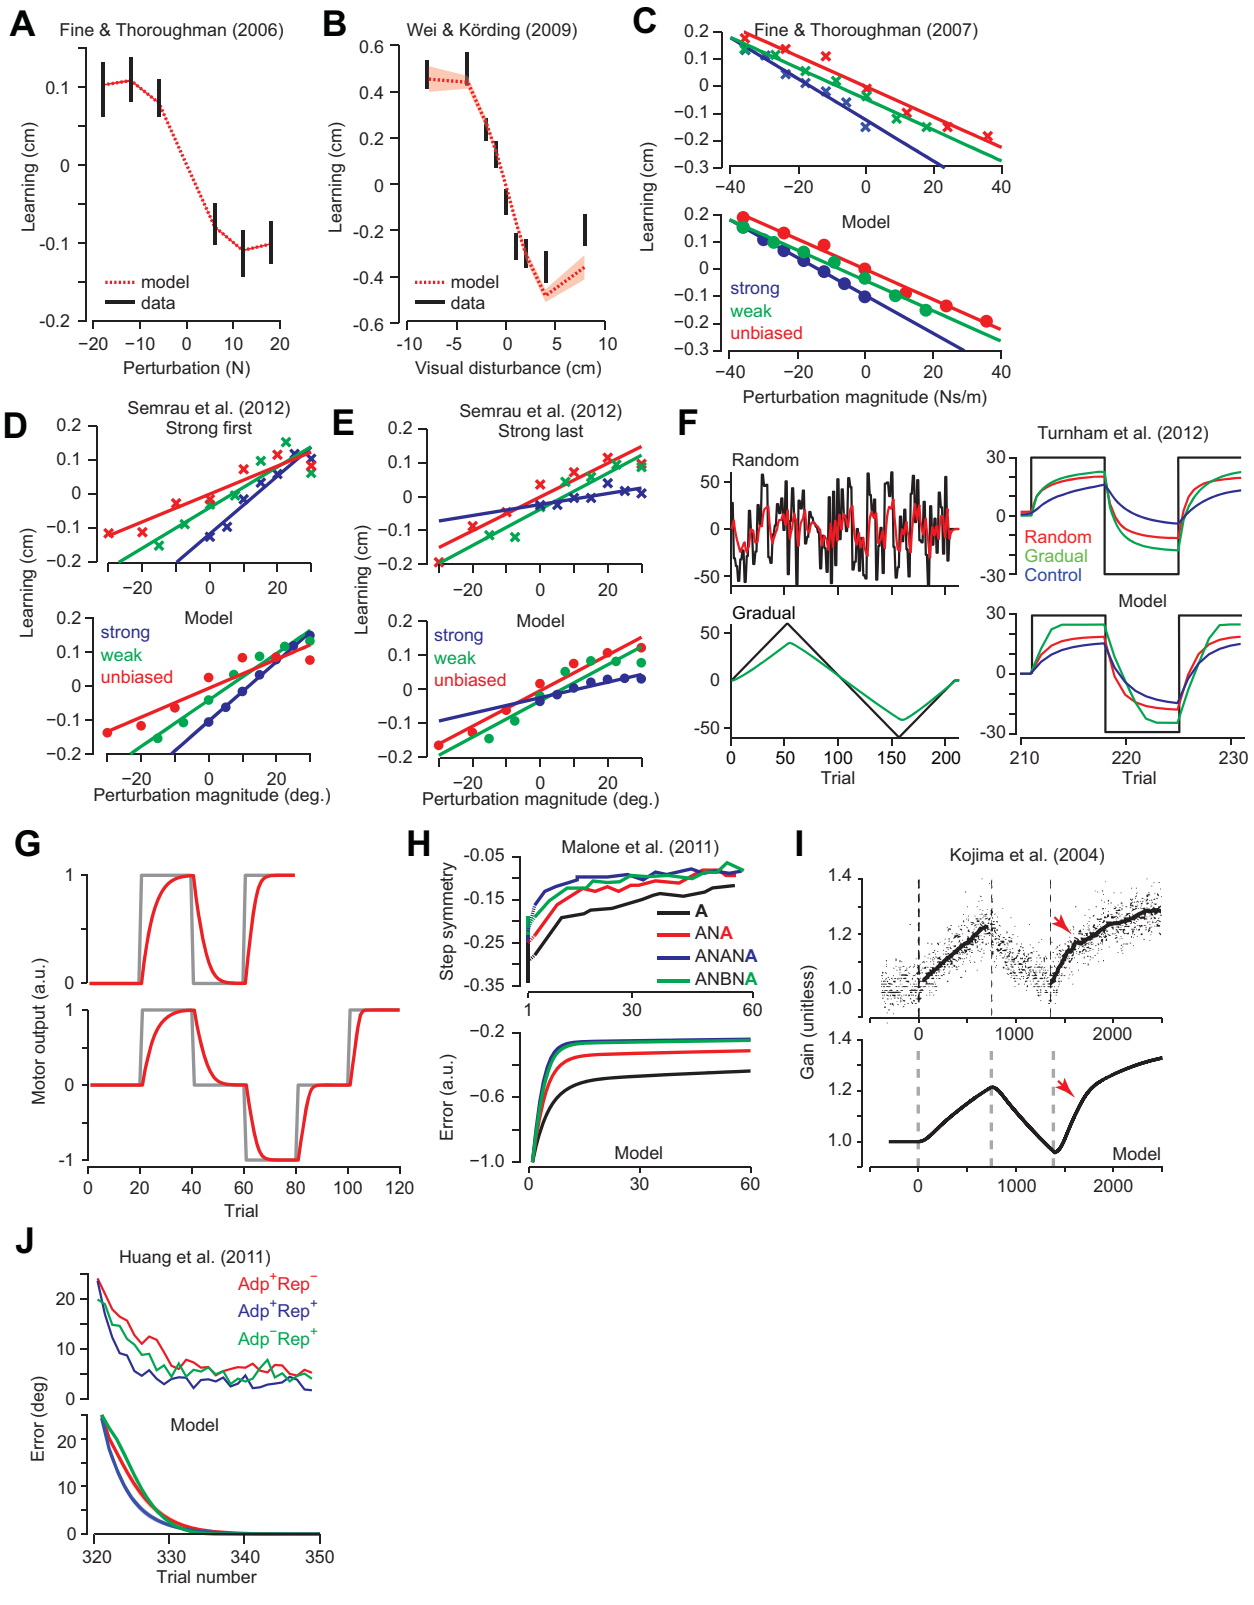
A
Fine & Thoroughman (2006)
Learning (cm)
model
data
Perturbation (N)
B
Wei & Körding (2009)
Learning (cm)
model
data
Visual disturbance (cm)
C
Fine & Thoroughman (2007)
Learning (cm)
Model
strong
weak
unbiased
Perturbation magnitude (Ns/m)
D
Semrau et al. (2012)
Strong first
Learning (cm)
Model
strong
weak
unbiased
Perturbation magnitude (deg.)
E
Semrau et al. (2012)
Strong last
Learning (cm)
Model
strong
weak
unbiased
Perturbation magnitude (deg.)
F
Random
Gradual
Trial
Turnham et al. (2012)
Random
Gradual
Control
Model
Trial
G
Motor output (a.u.)
Trial
H
Malone et al. (2011)
Step symmetry
A
ANA
ANANA
ANBNA
Error (a.u.)
Model
I
Kojima et al. (2004)
Gain (unitless)
Model
J
Huang et al. (2011)
Adp+Rep−
Adp+Rep+
Adp−Rep+
Error (deg)
Model
Trial number

Figure S2. Model accounts for a large body of experimental results. **A**. Mean adaptation to a series of discrete force-pulse perturbations. Data reported by Fine and Thoroughman (*29*) (black bars, mean±SEM) compared to model (red, mean±SEM). Model SE M is obscured by the line. **B**. Mean adaptation during a visual displacement task reported by Wei and Körding (*5*) (black bars) compared to model (red). **C**. Adaptation to discrete force-field perturbations during a reaching task. Volunteers experienced one of three environments in which perturbations were drawn from a distribution that was either unbiased, weakly biased (green x's) or strongly biased (blue x's) (top) (40). Model results shown in lower panel. **D-E**. Experimental results from a visuomotor rotation task (*33*). Participants experienced environments with statistics similar to sub-figure C over three separate days (top). Model results for the three environments in the order strong, weak, unbiased are shown in D (bottom); environments in the order unbiased, weak, strong are shown in E. **F**. Example of meta-learning (also termed structural learning). Results from a visuomotor rotation experiment (*16*). After exposure to a random or gradual environment (left), performance was assessed during learning of a [+30, -30, +30]° perturbation sequence (right). **G**. Examples of savings. Model results (red) in a perturbation schedule consisting of training (+1), washout (0), re-learning (+1) (grey, top). Model results for a schedule which includes learning of (-1) shows faster learning than control in (-1), as well as faster relearning in (+1) (bottom). **H**. Split-belt gait adaptation results(*11*) in which subjects were exposed to a similar sequence of perturbations as in G. Model results for similar perturbation schedule (bottom). **I.** Saccade adaptation experiment(*10*) (top) and model (bottom). Individual saccades are shown as black dots; black line shows 150 trial moving average. The red arrow denotes the end of facilitated learning. **J.** Example of data attributed to reinforced repetition. Model results for a paradigm similar to (*30*). Savings is present in the $Adp^{+}Rep^{+}$ group. In all plots, shaded error regions are model mean ± SEM, across randomly generated perturbation sequences.

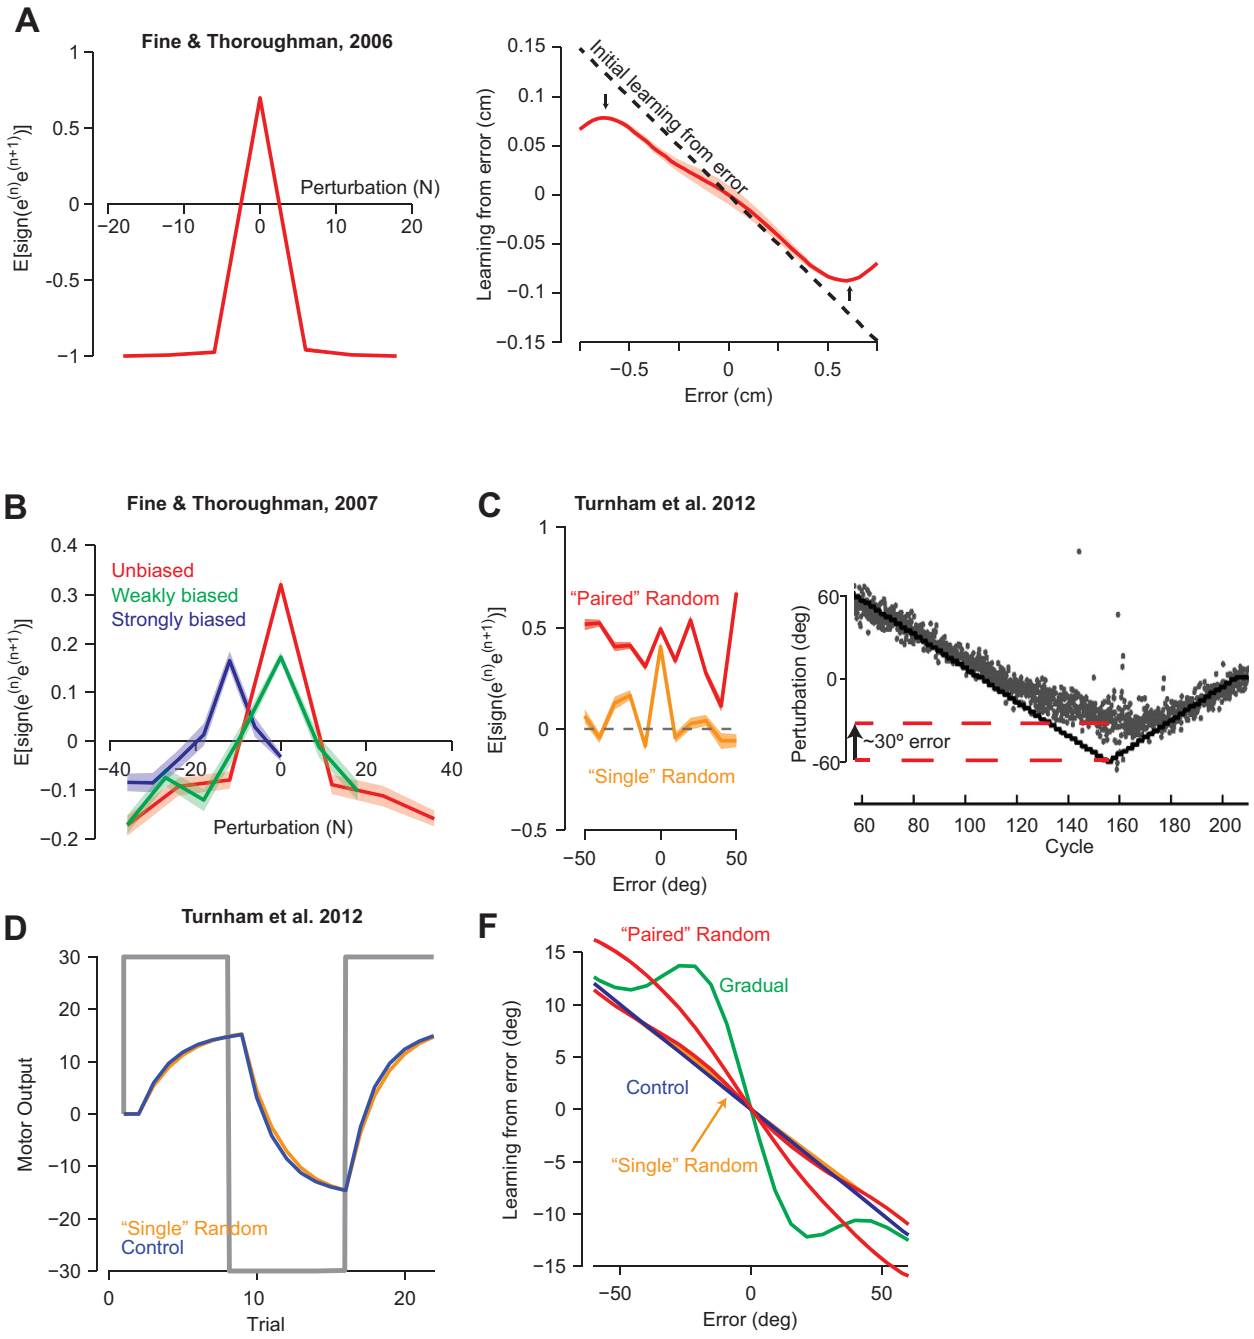
A
Fine & Thoroughman, 2006
E[sign(e(n)e(n+1))]
Perturbation (N)
Learning from error (cm)
Initial learning from error
Error (cm)
B
Fine & Thoroughman, 2007
Unbiased
Weakly biased
Strongly biased
Perturbation (N)
C
Turnham et al. 2012
"Paired" Random
"Single" Random
Error (deg)
Perturbation (deg)
~30° error
Cycle
D
Turnham et al. 2012
Motor Output
"Single" Random
Control
Trial
F
"Paired" Random
Gradual
Control
"Single" Random
Learning from error (deg)
Error (deg)

Fig. S3. The model was simulated on each perturbation distribution in the various experiments and then the quantity $E\left[sign\left(e^{(n)} e^{(n+1)} \right) \right]$ was plotted as a function of perturbation magnitude in that distribution. **A**. Simulation results for the perturbation distribution in Fine and Thoroughman (*29*). In the right subplot the baseline learning curve is shown by the black line and he final learning curve is shown by the red line. **B**. Simulation results for the perturbation distributions in Fine and Thoroughman (40). **C**. Simulation results for the random perturbation sequence in Turnham, Braun and Wolpert (*16*) labeled as "paired" random. Additional results for a schedule in which perturbations are random on each trial rather than each pair is labeled as "single" random. Mean errors from the gradual training in (*16*), showing sustained errors of approximately 30°, similar to those tested at the end of the experiment. **D.** Simulation results for the "single" random perturbation. The "single" random schedule does not show savings whereas the "paired" random schedule from (*16*) shows savings (Fig. S2F). **F.** Learning from error as a function of error magnitudes for all groups from the simulation of (*16*). Learning from error increases for the gradual and paired-random perturbations, but not single-random perturbations. In all plots, shaded error regions are mean ± SEM, across randomly generated perturbation sequences.

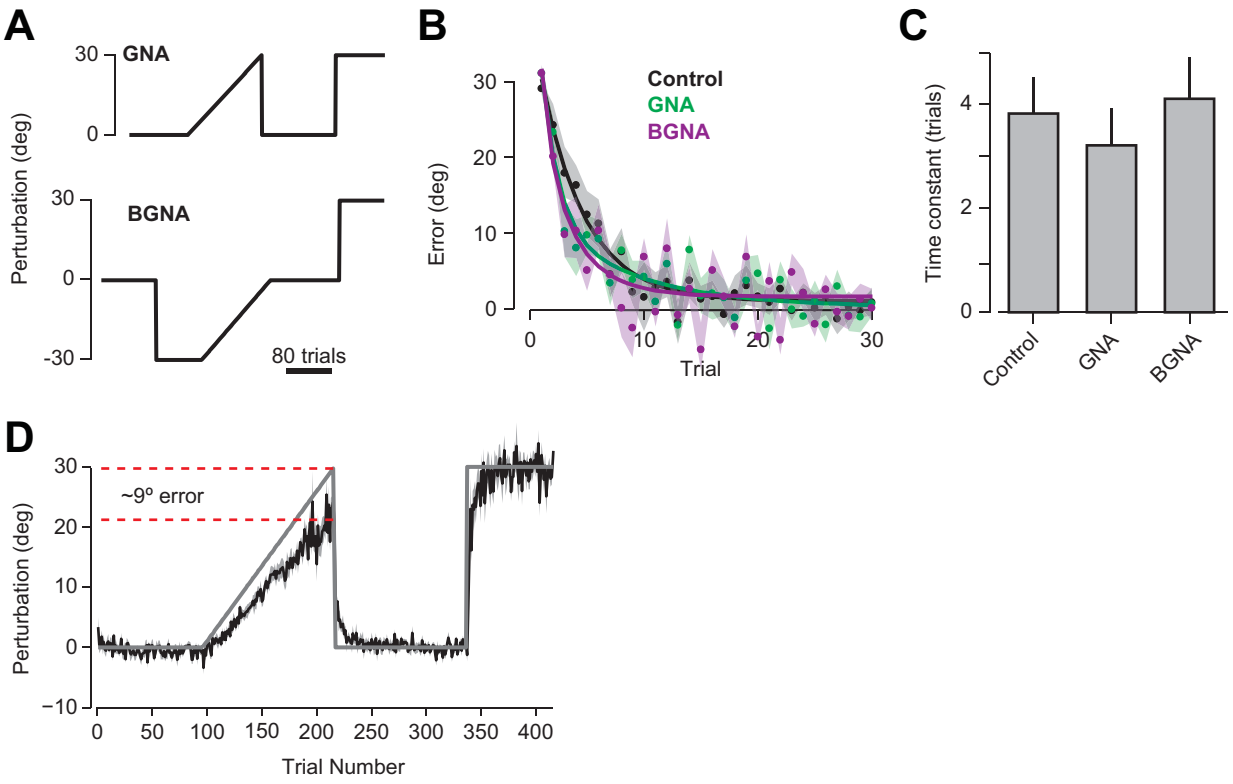

Figure S4. Additional groups in Experiment 4: saving and meta-learning occur only when previously experienced errors are re-visited. **A**. Perturbation protocols for groups GNA and BGNA. **B.** Performance during exposure to +30° perturbation. Exponential fits are shown for the group data. Performances of the GNA and BGNA are not different than control, demonstrating that these perturbation protocols blocked savings. **C**. Exponential time constants are not different than control. A lower time constant indicates faster learning. **D.** Time course of adaptation in the GNA group. Note that the magnitude of the error at the end of gradual learning is significantly smaller than 30°. Data are mean ± SEM across subjects.

References and Notes

1. M. I. Jordan, D. E. Rumelhart, Forward Models - Supervised Learning with a Distal Teacher. *Cogn. Sci.* **16**, 307–354 (1992). doi:10.1207/s15516709cog1603_1

2. M. Kawato, K. Furukawa, R. Suzuki, A hierarchical neural-network model for control and learning of voluntary movement. *Biol. Cybern.* **57**, 169–185 (1987). Medline doi:10.1007/BF00364149

3. F. R. Robinson, C. T. Noto, S. E. Bevans, Effect of visual error size on saccade adaptation in monkey. *J. Neurophysiol.* **90**, 1235–1244 (2003). Medline doi:10.1152/jn.00656.2002

4. R. Soetedjo, A. F. Fuchs, Y. Kojima, Subthreshold activation of the superior colliculus drives saccade motor learning. *J. Neurosci.* **29**, 15213–15222 (2009). Medline doi:10.1523/JNEUROSCI.4296-09.2009

5. K. Wei, K. Körding, Relevance of error: What drives motor adaptation? *J. Neurophysiol.* **101**, 655–664 (2009). Medline doi:10.1152/jn.90545.2008

6. M. K. Marko, A. M. Haith, M. D. Harran, R. Shadmehr, Sensitivity to prediction error in reach adaptation. *J. Neurophysiol.* **108**, 1752–1763 (2012). Medline doi:10.1152/jn.00177.2012

7. M. A. Smith, R. Shadmehr, Modulation of the rate of error-dependent learning by the statistical properties of the task. *Advances in Computational Motor Control* **3**, (2004).

8. L. N. Gonzalez Castro, A. M. Hadjiosif, M. A. Hemphill, M. A. Smith, Environmental consistency determines the rate of motor adaptation. *Curr. Biol.* **24**, 1050–1061 (2014). Medline doi:10.1016/j.cub.2014.03.049

9. M. C. Trent, A. A. Ahmed, Learning from the value of your mistakes: Evidence for a risk-sensitive process in movement adaptation. *Front. Comput. Neurosci.* **7**, 118 (2013). Medline doi:10.3389/fncom.2013.00118

10. Y. Kojima, Y. Iwamoto, K. Yoshida, Memory of learning facilitates saccadic adaptation in the monkey. *J. Neurosci.* **24**, 7531–7539 (2004). Medline doi:10.1523/JNEUROSCI.1741-04.2004

11. L. A. Malone, E. V. Vasudevan, A. J. Bastian, Motor adaptation training for faster relearning. *J. Neurosci.* **31**, 15136–15143 (2011). Medline doi:10.1523/JNEUROSCI.1367-11.2011

12. A. M. Sarwary, L. P. Selen, W. P. Medendorp, Vestibular benefits to task savings in motor adaptation. *J. Neurophysiol.* **110**, 1269–1277 (2013). Medline doi:10.1152/jn.00914.2012

13. E. Zarahn, G. D. Weston, J. Liang, P. Mazzoni, J. W. Krakauer, Explaining savings for visuomotor adaptation: Linear time-invariant state-space models are not sufficient. *J. Neurophysiol.* **100**, 2537–2548 (2008). Medline doi:10.1152/jn.90529.2008

14. F. Mawase, L. Shmuelof, S. Bar-Haim, A. Karniel, Savings in locomotor adaptation explained by changes in learning parameters following initial adaptation. *J. Neurophysiol.* **111**, 1444–1454 (2014). Medline doi:10.1152/jn.00734.2013

15. D. A. Braun, A. Aertsen, D. M. Wolpert, C. Mehring, Motor task variation induces structural learning. *Curr. Biol.* **19**, 352–357 (2009). Medline doi:10.1016/j.cub.2009.01.036

16. E. J. Turnham, D. A. Braun, D. M. Wolpert, Facilitation of learning induced by both random and gradual visuomotor task variation. *J. Neurophysiol.* **107**, 1111–1122 (2012). Medline doi:10.1152/jn.00635.2011

17. N. J. Mackintosh, Theory of attention - variations in associability of stimuli with reinforcement. *Psychol. Rev.* **82**, 276–298 (1975). doi:10.1037/h0076778

18. K. A. Thoroughman, R. Shadmehr, Learning of action through adaptive combination of motor primitives. *Nature* **407**, 742–747 (2000). Medline doi:10.1038/35037588

19. R. A. Rescorla, A. R. Wagner, *A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement*, in *Classical Conditioning*, A. H. Black, W. F. Prokasy, Eds. (Appleton-Century-Crofts, New York, 1972), vol. 2.

20. J. M. Pearce, G. Hall, A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol. Rev.* **87**, 532–552 (1980). Medline doi:10.1037/0033-295X.87.6.532

21. R. A. Scheidt, J. B. Dingwell, F. A. Mussa-Ivaldi, Learning to move amid uncertainty. *J. Neurophysiol.* **86**, 971–985 (2001). Medline

22. O. Donchin, J. T. Francis, R. Shadmehr, 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 (2003). Medline

23. M. A. Smith, A. Ghazizadeh, R. Shadmehr, Interacting adaptive processes with different timescales underlie short-term motor learning. *PLOS Biol.* **4**, e179 (2006). Medline doi:10.1371/journal.pbio.0040179

24. S. Cheng, P. N. Sabes, Modeling sensorimotor learning with linear dynamical systems. *Neural Comput.* **18**, 760–793 (2006). Medline doi:10.1162/neco.2006.18.4.760

25. K. P. Körding, J. B. Tenenbaum, R. Shadmehr, The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat. Neurosci.* **10**, 779–786 (2007). Medline doi:10.1038/nn1901

26. R. J. van Beers, Motor learning is optimally tuned to the properties of motor noise. *Neuron* **63**, 406–417 (2009). Medline doi:10.1016/j.neuron.2009.06.025

27. R. J. van Beers, How does our motor system determine its learning rate? *PLOS ONE* **7**, e49373 (2012). Medline doi:10.1371/journal.pone.0049373

28. M. Riedmiller, H. Braun, in *IEEE International Conference on Neural Networks* (1993), vol. 1, pp. 586–591.

29. M. S. Fine, K. A. Thoroughman, Motor adaptation to single force pulses: Sensitive to direction but insensitive to within-movement pulse placement and magnitude. *J. Neurophysiol.* **96**, 710–720 (2006). Medline doi:10.1152/jn.00215.2006

30. V. S. Huang, A. Haith, P. Mazzoni, J. W. Krakauer, Rethinking motor learning and savings in adaptation paradigms: Model-free memory for successful actions combines with internal models. *Neuron* **70**, 787–801 (2011). Medline doi:10.1016/j.neuron.2011.04.012

31. D. J. Herzfeld, D. Pastor, A. M. Haith, Y. Rossetti, R. Shadmehr, J. O'Shea, Contributions of the cerebellum and the motor cortex to acquisition and retention of motor memories. *Neuroimage* **98**, 147–158 (2014). Medline doi:10.1016/j.neuroimage.2014.04.076

32. J. L. Patton, Y. J. Wei, P. Bajaj, R. A. Scheidt, Visuomotor learning enhanced by augmenting instantaneous trajectory error feedback during reaching. *PLOS ONE* **8**, e46466 (2013). Medline doi:10.1371/journal.pone.0046466

33. J. A. Semrau, A. L. Daitch, K. A. Thoroughman, Environmental experience within and across testing days determines the strength of human visuomotor adaptation. *Exp. Brain Res.* **216**, 409–418 (2012). Medline doi:10.1007/s00221-011-2945-z

34. J. Y. Lee, N. Schweighofer, Dual adaptation supports a parallel architecture of motor memory. *J. Neurosci.* **29**, 10396–10404 (2009). Medline doi:10.1523/JNEUROSCI.1294-09.2009

35. S. C. McLaughlin, Parametric Adjustment in Saccadic Eye Movements. *Percept. Psychophys.* **2**, 359–362 (1967). doi:10.3758/BF03210071

36. H. Tanaka, T. J. Sejnowski, J. W. Krakauer, Adaptation to visuomotor rotation through interaction between posterior parietal and motor cortical areas. *J. Neurophysiol.* **102**, 2921–2932 (2009). Medline doi:10.1152/jn.90834.2008

37. S. E. Criscimagna-Hemminger, R. Shadmehr, Consolidation patterns of human motor memory. *J. Neurosci.* **28**, 9610–9618 (2008). Medline doi:10.1523/JNEUROSCI.3071-08.2008

38. A. M. Hadjiosif, M. A. Smith, Savings is restricted to the temporally labile component of motor adaptation. *Translational and Computational Motor Control*, (2013).

39. R. Shadmehr, T. Brashers-Krug, Functional stages in the formation of human long-term motor memory. *J. Neurosci.* **17**, 409–419 (1997). Medline

40. M. S. Fine, K. A. Thoroughman, Trial-by-trial transformation of error into sensorimotor adaptation changes with environmental dynamics. *J. Neurophysiol.* **98**, 1392–1404 (2007). Medline doi:10.1152/jn.00196.2007