

The role of error-based learning in movement and stillness

by
Scott T. Albert, B.S.

A dissertation submitted to Johns Hopkins University in conformity with
the requirements for the degree of Doctor of Philosophy

Baltimore, MD
March 2020

© 2020 Scott Albert
All rights reserved

Abstract

When people and other animals perform a movement that produces an unexpected outcome, they learn from the resulting error and retain a portion of this learning over time. Over many trials, individual events of learning accumulate into a motor memory. Understanding how these memories form and are retained over time, represents a fundamental question facing the neuroscience community. Curiously, for reaching movements of the arm, errors that occur solely during periods of movement cause changes to both the way we move and also the way we hold the arm still. Here, we explore the way the brain corrects for error after a single occurrence, how this response to error changes with experience, and finally, how these responses to error change the way we maintain stillness of the arm.

In Chapter 2, we consider the mechanisms that guide learning on the timescale of a single trial. How does the brain determine the patterns of muscle activity that will better compensate for future perturbations? Here we provide evidence that the brain uses its past corrections as a model for its future movement plans, and that this single trial learning may be dependent on the cerebellum.

The response to learning does not remain fixed over time. Rather, over the course of many movements, the brain modifies the way it learns. In cases where similar perturbations were experienced in the past, it accelerates its rate of learning (savings). In cases where dissimilar perturbations were experienced in the past, it slows its rate of learning (anterograde interference). And sometimes, even after memories appear forgotten, they reemerge in the absence of error (spontaneous recovery). In principle, these changes could occur because the brain alters its ability to retain memory, or the amount it is willing to change when it experiences an error. In Chapter 3, we describe a new algorithm that can extract these behavioral parameters more accurately than earlier techniques.

With this tool in hand, we first show that savings (faster rate of re-learning) is caused by an increase in the brain's sensitivity to error, specifically within fast motor learning processes. We also investigate if this may depend on a reward system inside the brain: the basal ganglia. Next, we show that reemergence of earlier memories (spontaneous recovery) is caused by the decay of fast learning processes. Finally, in Chapter 4, we show that anterograde interference (slowing of learning when a dissimilar perturbation is experienced) is caused by two different sources that operate on different timescales. The first is a lingering retention of the earlier, inappropriate memory. The second is a bonafide reduction in one's ability to learn from new errors.

Changes in the way we learn from error not only affect the rate at which we acquire new memories, but also the total amount our brain is capable of learning. In Chapter 5, we demonstrate that when errors are highly variable, learning is less complete. Remarkably, changes in the total amount of learning appear to be accomplished through changes in error sensitivity, but not our ability to retain memory that was already acquired. These changes in the total extent of learning appear to be primarily dependent on implicit adaptation, or learning that is acquired in the absence of any cognitive strategies.

Lastly, in Chapter 6, we revisit our initial observation, that adaptation changes not only the way we move, but also the way we hold still. In a series of experiments in human and non-human primates, we report a surprising relationship between movement and posture: on a within-trial basis, the commands that hold the arm and finger at a target location depend in part on the mathematical integration of the commands that moved the limb to that location.

Acknowledgements

This work has been made possible by grants from the National Institutes of Health (grant numbers NS078311, NS095706, F31NS095706), the National Science Foundation (grant number 1723967), and the Office of Naval Research (grant number N00014-15-1-2312).

Our studies could never have happened without the gracious time commitment of countless volunteers and participants. In particular, I am extremely grateful for the many Parkinson's Disease patients and stroke survivors who participated and contributed to this work. We hope that our careful analysis of your behavior will help us to understand your diseases, and one day, translate into meaningful treatments and therapies.

The data present in this work represents the collective output of many incredible investigators: Mark Churchland, Stuart Baker, Valeria Della-Maggiore, and John Krakauer. Each of these individuals has made an incredible mark on my development as a scientist. I hope that our partnerships will continue for the rest of my career and serve as an example of what is possible through scientific collaboration. In addition, this work was based on data collected by others and made publicly available to the scientific community and the students of the Computational Sensory-Motor Neuroscience School. In particular, I would like to thank Hugo Fernandes for graciously sharing his data, as well as Konrad Kording, Paul Schrater, and Gunnar Blohm for their efforts and leadership in organizing this course.

I consider myself so incredibly fortunate to have chosen Reza Shadmehr as my PhD advisor. As a convert from Chemical Engineering to Biomedical Engineering, I see it as fate that I landed in the laboratory of such a remarkably capable scientist, mentor, and friend. Under Reza's guidance I have learned the tools necessary to do rigorous scientific inquiry, think deeply about my observations, share my ideas with others in the community, appreciate the beauty of our nervous system, seek out the experiments that challenge our hypotheses the most, document findings with the highest level of clarity, and always seek to translate ideas into the universal language of mathematics. Any success that I find in the future will only have been possible because of our time together.

Equally important in my scientific development are my friends and colleagues in the Laboratory for Computational Motor Control. Throughout my years in lab I have greatly benefited from the truly collaborative environment present in the cubicles and around the lunch table of Traylor 416. So many of my ideas were molded and improved through our conversations. Every day, lab felt like home because of our friendship. I hope that I am always surrounded by such a gifted set of individuals in my place of work. Ehsan, talking with you about science and about life has become such an important part of my day-to-day in lab, and has made me a better scientist and a better person. Sarah, you were such an incredible mentor to me when I started in the lab, an amazing friend to me always, and a huge part of why I felt a sense of belonging at Hopkins. Simon, watching you as a scientist, a father, and a husband, has served as such an inspiration and role model for what I hope to become in the future. Kaveh, your unmatched patience with my nonstop questioning has helped me infinitely, and your unmovable sense of humor has made me smile every day in lab (I know the gates to the League of Legends will be open to you on the West Coast). Jay, your dependable friendship at lab and at Thread has made me feel so supported during these past few years in Baltimore. Paul, your vast technical knowledge and complete willingness to help others has made a huge impact on me and the lab. Tehrim, your inhuman mathematical brilliance has helped me think deeper about modeling and the meaning behind my data. Rob, your innate curiosity, deep intellect, and steadfast friendship has changed my time at Hopkins and in Baltimore. David your unmatched technical prowess and critical thinking has made a huge difference in the way I design my experiments and think about my data. Tom, your support during my tough times, your friendliness, and passion for science, has made a massive impact on my graduate career. Tara your

infectious laughter and supreme friendliness has really brightened Traylor 416 and made it such an incredible environment to do work. Jihoon, my thesis would be half the size without your tremendous work ethic and gifted thinking abilities. And even though I'm not entirely convinced that you play the bass, our partnership over the past few years has really kept the music playing.

John and David, your mentorship and guidance has challenged me to think more deeply about science and the truth behind the data. Our meetings have borne some of my most creative experiments and most exciting ideas. I could not ask for a thesis committee of higher quality.

I would also like to recognize the glue that holds the BME program together: Hong Lan. Hong is not only an incredible organizer and the hardest working employee at Johns Hopkins, but she is also one of the kindest and most caring people I have ever met. Her support both socially and scholastically has transformed my experience at Johns Hopkins.

I would be remiss if I did not thank all of the doctors and nurses at the Sidney Kimmel Cancer Center, in particular Nina Wagner-Johnston and Pierse Byrnes. Without them, I literally would not be here. And while I do not miss my time in the Weinberg building, it has taught me that no obstacle is insurmountable.

So much of my happiness was made possible by the friends and colleagues I made outside of the lab during my time in Baltimore. I am so thankful to have been a part of the Jazz Ensemble, the Concert Band, numerous soccer leagues at Hopkins, Patterson Park, and Fed Hill, and as a volunteer with Thread. In terms of this latter engagement, I would like to recognize Andrew Mallinoff, for his support, guidance, and friendship. Many of my fondest memories in Baltimore were made with the amazing people in each of these organizations. In addition, I would like to recognize David Saenz, for his mentorship, leadership, support, and friendship in both the best and toughest of times. I have learned so much about myself and the world around me, from the many friends I share with Mannat: Anushka, Baldeep, Erin and the Cooney's, Aliya, Mina, Akshay, Naira, and Gaurav.

And while I have definitely not been as great of a friend as they have been to me, I also would not be where I am without the support of my Lafayette crew: Chris, Leo, Nick, Rachel, Lauren, Rebecca, and Emily. I promise that one day I will become much better at returning your text messages. I am entirely overjoyed to see each of you thrive in the paths you have chosen. I would say that I wish I was back at Lafayette in the halls of Ramer, but I never want to take another exam in my life.

Baltimore would not be Baltimore if it were not for my two roommates: Nash and Bin. Although they graduated before me, and then left my here, I won't hold that against them. I do not know of any trio in Baltimore that have shaped each other's lives to the extent we did. I will forever and always miss Monday dinners at the Dizz, the music, the shows, dinner parties, debate, comradery, and knowing that I have two friends to come home to, even if one of them never opened his door. I cannot wait to see where your creativity brings you, and the continents and climes that will bring us all together again.

Mannat. Do you remember that my EM paper was accepted into the Journal of Neurophysiology when we were at the mall? When I told you and you looked back at me, your eyes and smile expressed a happiness, that if captured and placed in a bottle, could light up the entire sky for the rest of my life. Smiles like that make me feel like I can do anything. You are my complete and total partner, and you challenge me to be the best that I can be. I cannot imagine being here without you by my side, and without the love and support of the Malik family. You are and will forever be my Star Baker.

Mom, Dad, and Mel. The music at my Bar Mitzvah said it best – you are the wind beneath my wings. You have been here with me throughout it all. Whether it be on the phone, crammed into the McElderry House, sleeping on the couch downstairs or a cot in my room, all you can eat sushi, or holding onto a robot handle, you always lift me up during the highest of times and when I need you most. When I ask for your help you never respond with “why, what, how”, but always “when”. Every success I find, and every part of me, is because of you. I am sending all of my love to you, Oreo, Shadow, and Atticus.

Table of contents

Abstract.....	2
Acknowledgements.....	3
Table of contents	5
List of figures.....	10
Chapter 1. Introduction	13
1.1 Learning from a single error.....	14
1.2 The evolution of the state-space model of learning.....	15
1.3 Fitting models to behavioral data	18
1.4 The neural integrator	20
1.5 Self-published content and references.....	21
1.6 Specific aims.....	22
Chapter 2. The neural feedback response to error as a teaching signal for the motor learning system ...	24
2.1 Introduction	24
2.2 Materials and methods.....	25
2.2.1 Experiment.....	25
2.2.2 Data Recording and Analysis.....	27
2.3 Results.....	30
2.3.1 Measuring trial-to-trial feedback responses to error and learning.....	30
2.3.2 The learning response is a time-shifted copy of the feedback response to error.....	32
2.3.3 Across subject variability in feedback response	37
2.3.4 Within-subject variability in feedback response.....	39
2.3.5 Control studies	44
2.3.6 The effect of cerebellar stimulation on learning and feedback responses	49
2.4 Discussion.....	50
2.4.1 The feedback response to error is incorporated into the motor plan for the next movement	51
2.4.2 On the causality between feedback responses and subsequent learning	53
Chapter 3. Estimating properties of the fast and slow adaptive processes during sensorimotor adaptation.....	55
3.1 Introduction	55
3.2 Materials and methods.....	58

3.2.1 State-space model of learning	59
3.2.2 Experimental procedure	63
3.2.3 Epoch vs. trial-by-trial analysis of behavior	64
3.2.4 Simulating realistic behavior	65
3.2.5 Fitting EM and LMST to data	67
3.2.6 Measuring performance of the algorithms.....	68
3.2.7 Sources of noise	68
3.2.8 Power analysis.....	69
3.2.9 Control studies	71
3.3 Results.....	72
3.3.1 Fitting the two-state model to measured behavior.....	72
3.3.2 Fitting the two-state model to simulated data.....	75
3.3.3 A better tool for hypothesis testing.....	82
3.3.4 Modeling higher levels of noise	84
3.3.5 Different noises in the fast and slow adaptive processes.....	85
3.3.6 Other sources of noise	87
3.3.7 Changing the dynamics of the fast and slow adaptive processes.....	87
3.3.8 Changing the bounds on the parameter space.....	88
3.3.9 Trial-by-trial analysis of behavior.....	90
3.3.10 Savings is caused by an increase in error sensitivity of fast learning processes.....	93
3.3.11 Spontaneous recovery after the passage of time is caused by fast state decay	95
3.3.12 The expression of savings after impairment of the basal ganglia	96
3.4 Discussion.....	99
3.4.1 Design of a new algorithm	99
3.4.2 Evaluating the new algorithm	100
3.4.3 Limitations of the algorithm	101
3.4.4 Relationship between mathematical hidden states and neural substrates of learning	103
3.5 Appendix	104
3.5.1 Mathematical description of the generalized EM algorithm.....	104
3.5.2 Overview of the LMSE algorithm	117
3.5.3 Multiple target state-space model of learning	119
Chapter 4. The origins of anterograde interference in visuomotor adaptation.....	122
4.1 Introduction	122

4.2 Materials and methods.....	123
4.2.1 Experimental Paradigm.....	123
4.2.2 Experimental Procedure	124
4.2.3. Data post-processing	125
4.2.4 Model-free data analysis.....	125
4.2.5 State-Space Model.....	126
4.2.6 Validation of the single state-space model.....	127
4.2.7 Statistical assessment	129
4.3 Results.....	129
4.3.1 Memory of <i>A</i> decays in time, but persists even after 24 hours.....	130
4.3.2 Anterograde interference dissipates with increasing time separating <i>A</i> and <i>B</i>	132
4.3.3 Anterograde interference is caused by a decrease in sensitivity to error that recovers with time	133
4.3.4 Error sensitivity is independent of initial error size	137
4.3.5 Alternate hypotheses: a two-state model account of anterograde interference	137
4.3.6. Differences in the initial bias of two parallel adaptive states cannot explain interference ...	138
4.3.7 Error sensitivity of the slow state is impaired in a two-state model of anterograde interference	140
4.3.8 Relative implicit/explicit contributions are unlikely to explain differences in error sensitivity	141
4.3.9 The implicit adaptive system is strongly limited by anterograde interference	142
4.3.10 Reduced error sensitivity is likely the cause of anterograde interference in force field adaptation.....	145
4.4 Discussion.....	148
4.4.1 Anterograde interference differs from a lingering memory of a prior	148
4.4.2 Anterograde interference and memory stabilization	149
4.4.3 Anterograde interference results from an impairment in error sensitivity that recovers with time	150
Chapter 5. Asymptotic limits of sensorimotor adaptation	153
5.1 Introduction	153
5.2 Materials and methods.....	154
5.2.1 Visuomotor rotation	155
5.2.2 Force field adaptation	156
5.2.3 Statistics	158

5.2.4 Experiment 1	158
5.2.5 Experiment 2	158
5.2.6 Experiment 3	159
5.2.6 Experiment 4	159
5.2.7 Experiment 5	159
5.2.8 Re-analysis of prior work	160
5.2.9 State-space model of learning	160
5.2.10 Asymptotic properties of learning	161
5.2.11 Calculate of the retention factor	162
5.2.12 Calculation of error sensitivity	162
5.2.13 Memory of errors model.....	164
5.2.14 Simulation of the memory of errors model	164
5.3 Results.....	166
5.3.1 Perturbation variance limits sensorimotor adaptation	167
5.3.2 Perturbation variance limits the total extent of adaptation.....	170
5.3.3 Perturbation variance impairs the implicit component of learning	170
5.3.4 Perturbation variance reduces error sensitivity, but not forgetting rates	173
5.3.5 Perturbation variance reduced the ability to learn from small errors, not large errors	175
5.3.6 The consistency of past errors modulated error sensitivity	175
5.3.7 Error sensitivity changes throughout training according to the consistency of error	177
5.3.8 Asymptotic differences are caused by changes in implicit, not explicit adaptation.....	181
5.4 Discussion.....	185
5.4.1 Controlling the total extent of adaptation through modulation of error sensitivity.....	185
5.4.2 Plasticity in implicit learning systems	187
5.4.3. Alternate models.....	187
5.4.4 Neural basis of error sensitivity	188
Chapter 6. Postural control of arm and fingers through integration of movement commands	190
6.1 Introduction	190
6.2 Methods.....	192
6.2.A The reach integrator hypothesis.....	192
6.2.B Monkey experiments	193
6.2.C Human experiments	198
6.3 Results.....	210

6.3.1 Muscle activity during the hold period following reaching	211
6.3.2 Hold period activity for finger movements.....	214
6.3.3 Change in reach period commands alters hold period commands	217
6.3.4 Move and hold correlations are only weakly influenced by initial hold activity	222
6.3.5 The postural field during the hold period	223
6.3.6 Adaptation of the integration gain	224
6.3.7 Differential contributions of the corticospinal tract to reaching and holding.....	227
6.3.8 Holding forces are not specific to a particular holding location	231
6.3.9 The change in holding force, not final holding force, is predicted by the integral of reach forces.....	231
6.3.10 Reach forces predict hold forces independent of spatial orientation of movement	235
6.3.11 Cortically-driven feedback responses to error are also integrated	239
6.4 Discussion.....	240
6.4.1 A hypothetical architecture for control of arm posture	242
6.4.2 Limitations.....	243
Chapter 7. Conclusions and future directions	245
7.1 Mechanisms of error-based learning.....	245
Section 7.2 Models of error-based learning	248
Section 7.3 Cortical and subcortical control of holding still	250
Reference List.....	254
Curriculum Vitae: Scott T. Albert.....	270

List of figures

Figure 2.1. Measuring feedback and learning from EMG.	31
Figure 2.2. Learning and error feedback responses for each muscle.	34
Figure 2.3. Temporal shifts relating learning and feedback responses.	35
Figure 2.4. Learning resembles time-shifted copies of the error feedback response.	36
Figure 2.5. Learning and error feedback responses across individuals.	38
Figure 2.6. The error feedback response is predictive of the learning response within a subject.	39
Figure 2.7. Agonist learning magnitude correlates to the feedback response gain.	42
Figure 2.8. Antagonist learning magnitude correlates to the feedback response gain.	43
Figure 2.9. Error-clamp trials provide accurate approximations to unperturbed movements.	44
Figure 2.10. The difference between consecutive feedback responses is largely due to single trial learning.	45
Figure 2.11. Learning from one error is similar to learning from the next error.	46
Figure 2.12. The feedback response gain varied randomly with the progression of the experiment.	48
Figure 2.13. Cerebellar stimulation alters the amount of learning.	49
Table 3.1. Two-state model parameters.	66
Table 3.2. Upper and lower bounds for two-state model parameters.	67
Figure 3.1. EM and LMSE algorithms uncover different hidden processes.	73
Table 3.3. Parameter standard deviation for EM and LMSE.	74
Figure 3.2. Comparison of parameter values uncovered by EM and LMSE as fitted to experimental data.	77
Figure 3.3. Fitting EM and LMSE to simulated behavior.	80
Figure 3.4. Performance of EM and LMSE algorithms.	81
Figure 3.5. Parameter estimation errors for EM and LMSE.	82
Figure 3.6. Power analysis for EM and LMSE.	84
Figure 3.7. Sensitivity analysis for state and motor noise.	86
Figure 3.8. Sensitivity analysis for the dynamics of the fast and slow states.	89
Figure 3.9. Comparison of EM and LMSE in a restricted parameter space.	91
Figure 3.10. Comparison of EM and LMSE on a trial-by-trial analysis of the data.	92
Figure 3.11. Savings is caused by an increase in error sensitivity of the fast-adaptive process.	94
Figure 3.12. Spontaneous recovery after a break in time is caused by loss of the fast state.	96

Figure 3.13. Damage to the basal ganglia may prevent the expression of savings.	98
Table 3.4. Interpretation of model variables for common learning paradigms.	106
Figure 4.1. Paradigm to measure anterograde interference.	132
Figure 4.2. Effect of prior learning on the initial level of performance and the ability to learn.	133
Figure 4.3. Validation of the single state-space model for interference dataset.	135
Figure 4.4. State-space model fit to interference dataset.	136
Figure 4.5. Initial biases in a two-state learning system cannot account for anterograde interference.	140
Figure 4.6. Error sensitivity of the slow state of learning is impaired with anterograde interference.	142
Figure 4.7. Reaction times during adaptation to the countermanding perturbation.	143
Figure 4.8. Anterograde interference in implicit adaptation.	145
Figure 4.9. Anterograde interference in force field adaptation.	148
Figure 5.1. Perturbation variance impairs sensorimotor adaptation.	169
Figure 5.2. Perturbation variance decreases the total amount of implicit adaptation.	173
Figure 5.3. Perturbation variance decreases error sensitivity, not decay rates.	175
Figure 5.4. The consistency of trial-to-trial errors coincides with changes in error sensitivity.	177
Figure 5.5. The memory of errors model predicts that error sensitivity changes over the course of training.	179
Figure 5.6. Perturbation variability reduces the total amount of implicit, but not explicit adaptation.	183
Figure 5.7. Implicit and explicit adaptation compete for error.	184
Table 6.1. Measures of impairment in stroke patients.	199
Figure 6.1. Integral of muscle activity during the reach correlates with subsequent activity during the hold period.	213
Figure 6.2. The integral of muscle activity during finger flexion correlates with subsequent activity during the hold period.	216
Figure 6.3. The integral of reaching forces correlates with forces produced during the subsequent hold period.	221
Figure 6.4. Holding forces are an integration of moving forces, not a continuation.	223
Figure 6.5. The null point of the postural field is set by the integral of reaching forces.	225
Figure 6.6. Adaptation of the integration gain.	226
Figure 6.7. Cortical reaching commands are integrated in a subcortical area.	228
Figure 6.8. Muscles that do not integrate also do not contribute to posture.	230
Figure 6.9. Holding forces are sustained across long time intervals.	231
Figure 6.10. Moving and holding forces during exposure to dual force fields.	233

Figure 6.11. Holding forces at the same target position are specific to reaches initiated from different start positions. 234

Figure 6.12. The integral of reach force predicts the change in reach force, not the terminal reach force. 236

Figure 6.13. Moving and holding forces for reach movements of different angles. 237

Figure 6.14. The relationship between reaching and holding forces is consistent across different arm postures. 239

Figure 6.15. The visual feedback response to error is integrated. 240

Chapter 1. Introduction

Near the base of the skull sits a large structure known as the cerebellum. This relatively small neural volume houses nearly 50% of all of the neurons in the brain. While all of the roles of the cerebellum remain to be elucidated, it is clear that its crystalline circuitry is critical for the process of motor control and motor adaptation¹⁻⁵. Motor control refers to the measurements and theories that describe the way the nervous system generates movements of the body. Motor adaptation refers to the process by which people and other animals adjust their motor system due to changes in their environment, or perturbations of their movements⁶. While these two disciplines are fatefully intertwined, they are often studied in separation. But as patients who suffer from disorders of the cerebellum know all too well, there can be no accurate movements without learning, and no learning without movement.

In Chapter 2, we begin with the smallest unit of motor adaptation: a single reaching movement. We explore mechanisms by which the brain uses its muscles to respond and then predict the occurrence of a perturbation. Next, we extrapolate outwards from the response to a single error, to the response that emerges after many repeated errors. We consider various hallmarks of learning that appear not only in the motor domain, but also in our cognition: savings⁷, anterograde interference⁸, spontaneous recovery^{8,9}, and asymptotic performance¹⁰. Each of these phenomena share a common element – they all depend on the way we form and retain new memories. In Chapter 3, we develop and describe a new statistical tool that can extract various properties of memory formation. We demonstrate how this tool can be used to understand how learning rates increase (savings), how memories can unexpectedly reemerge (spontaneous recovery), and how some of these processes might be damaged when the reward systems of our brain degrade. In Chapter 4, we use our statistical tool to understand why the rate of learning slows (anterograde interference) and how this slowing can be alleviated by the passage of time.

At the heart of all of these phenomena is the concept of error sensitivity: the amount the brain learns from the experience of an error. In Chapters 3 and 4, we find that the plastic nature by which our experiences shape the acquisition of new memories is uniformly guided by error sensitivity. In Chapter 5, we demonstrate that our sensitivity to error not only determines the rate at which we learn, but also the total extent to which our brain is capable of learning. Furthermore, we test how implicit and explicit components of memory^{11,12} contribute to our total capacity for adaptation.

Along the way, we will find that the errors that we experience during a movement alter not only the way that we move, but also the way that we hold still. In Chapter 6, at the conclusion of this work,

we will explore a potential mechanism that might explain this curious interdependence, based not on our knowledge of the arm, but of the eye^{1,13}.

1.1 Learning from a single error

Movements we make are often affected by unintended errors. The occurrence of a single movement error engages learning mechanisms in the brain that change our future behavior. While we do not know how different areas of the brain contribute to error-based adaptation, there is at least one structure that appears critical for this process: the cerebellum⁶.

The cerebellum is not only needed for our ability to learn¹⁴⁻¹⁶, but more generally, our ability to produce accurate movements. That is, while we can move without a cerebellum, these movements suffer from ataxias (discoordination) that result in action tremor and dysmetria¹⁷. Why does disruption of the cerebellum cause these deficits?

The short answer to this question is that we do not yet know. But there are many conceptual frameworks that describe what the cerebellum might contribute to motor control and learning. One of the primary theories regards the cerebellum as a forward model, or state predictor⁶. In this view the cerebellum receives an efference copy of ongoing motor commands and uses an internal model to predict the future sensory outcome¹⁸. Another idea is that the cerebellum may do the reverse – convert an intended sensory state into the commands necessary to move there. Or, perhaps, the cerebellum houses both these forward and inverse models^{19,20}.

In any case, mismatch between the predicted sensory state of the body (potentially computed by the cerebellum) and the actual sensory state (measured in the periphery) yields a sensory prediction error⁶. The presence of these errors may teach the brain that its model of the plant and environment are incorrect and need to be adjusted. These adjustments to the cerebellum's internal models might then result in motor adaptation¹⁴.

How might this occur at the neurophysiological level? The climbing fibers from the inferior olive which produce complex spikes and thus LTD at the Purkinje cell-parallel fiber synapse (Marr-Albus-Ito hypothesis) appear to respond to different types of prediction errors, that is any unexpected change in the body's sensory state, even in the absence of movement²¹. In the case of associative learning, this LTD might modulate the strength or presence of reflexes that can be invoked in response to many types of sensory modalities²¹. In the case of voluntary movement, the consistent pairing of these unexpected sensory states (i.e., errors in the movement trajectory) with goal-directed movements, e.g., a saccadic eye movement²², would result in task-specific behavioral modifications that eliminate error.

What form might this adaptation take? A longstanding theory posits that the motor corrections we produce online to correct for an error may serve as a teaching signal for the motor learning system²³. In Chapter 2, we examine this possibility by recording the activity of several muscles during reaching movements of the arm. While we find ample evidence that learning responses resemble time-shifted copies of feedback responses to error²⁴, we also should note that it is unlikely to be the case that the motor correction alone is responsible for learning. That is, movements like saccades or brisk reaches of the arm¹⁴ occur too quickly for the presence of a feedback response, yet there is still learning. Therefore, a more holistic view of the cerebellar learning system might be that the presence of sensory prediction errors is necessary for adaptation, but the presence of a motor correction could modify the learned response to these errors when it is available to the cerebellar circuit. A more abstract argument might be that even when movements are too brisk for online correction, the subsequent movement (e.g., corrective saccade) could serve as an appropriate corrective response, or that corrections are computed implicitly by the nervous system and used for learning even when movements are too brief to allow for their expression.

To evaluate each of these theories it will be extremely critical to record from and make sense of the output neurons of the cerebellum: the deep cerebellar nuclei. It may be the case, that the output of the cerebellum is not a prediction at all, but rather a motor correction. For example, activation of the interposed nucleus prematurely reduces the velocity of movements executed away from the body, but increases the velocity of movements executed towards the body. In other words, stimulation of the interposed nucleus produces a correction of velocity in extrinsic space. This result is consistent with a different interpretation of cerebellar architecture, one in which the output of the cerebellum is a corrective motor signal. In the saccadic example^{22,25}, the cerebellar cortex may output its prediction for the motor commands it believes are necessary for a given sensory context (input from the granule cells) which is then compared with an efference copy of the ongoing motor response (encoded by mossy fiber input) at the level of the nucleus neurons, who then by way of a difference operator yield an estimate for the necessary motor correction prior to the arrival of sensory feedback.

1.2 The evolution of the state-space model of learning

In the field of motor control, mathematical models of learning have revolutionized our approach to understanding the dynamics of memory. Of these, the state-space model has emerged as a powerful and simple tool with which to interpret behavior and predict human performance. This model began

with the consideration of two fundamental components of adaptation: learning from error, and memory decay²⁶. That is, when an individual is confronted with a perturbation, errors in their movement are transduced by the nervous system into updated plans for movement. At the heart of this process of error-based learning is the concept of error-sensitivity: the amount that the brain learns from the experience of an error. However, sensorimotor adaptation, like our declarative memories, is also subject to forgetting: decay that accumulates over time²⁷. It is through the balancing of one's error sensitivity and rate of forgetting that the nervous system regulates the rate at which it acquires new memories and also the total extent to which it is capable of adapting^{10,28}.

As it is often said, "All models are wrong, but some are useful" – George Box. True to this idea, the state-space model of error-based learning has undergone several rounds of revision over the past 20 years as it has been confronted with ever-expanding bodies of data and behavioral phenomena. These constant revisions represent the true power of this simplistic approach: it generates easily testable predictions about behavior that guide new experiments. Here we will detail various modifications to this model upon which our current work is built.

The constant updating of the state-space model has been driven by three ubiquitous memory-related phenomena: savings, interference, and spontaneous recovery. Savings refers to an increase in the rate of learning that occurs when similar perturbations have been experienced in the past^{7,8,29–37}. Interference comes in two forms, retrograde and anterograde. In retrograde interference, adaptation to perturbation A, a countermanding perturbation B, and then perturbation A shows an impairment in the ability to recall the original A memory due to exposure to B^{38–41}. In anterograde interference, exposure to perturbation A directly impairs learning in B^{8,30,39,42}. And finally, spontaneous recovery refers to a reemergence of prior memory in the absence of any associated cues^{8,9,33,43–46}.

Whereas the original state-space model described adaptation using a single timescale of learning, in 2006 a new model was proposed that offered to explain each of these three hallmarks of adaptation: the two-state model of learning⁸. The two-state model posits that adaptation is composed of at least two different adaptive-processes that operate on different timescales. The first is a slow process of learning that exhibits a low sensitivity to error, but a robust ability to retain memory over time. This state is contrasted by the fast process of learning, which learns quickly from error but poorly retains its state over time. This model of adaptation was the first to provide a simple account of spontaneous recovery; here spontaneous recovery occurs when the fast and slow state counterbalance one another, and then the fast state is lost due to the passage of time or absence of error. In addition to this recovery, the two-state model also provided an account of savings and anterograde interference.

Since then, context-dependent models of adaptation have been proposed with additional states of learning, to account for the brain's ability to simultaneously adapt to conflicting perturbations^{47,48}. Furthermore, recent work has focused not only on the learning dynamics of these multiple adaptive states, but also their abilities to retain memories, focusing on differences between trial and time-based decay²⁹.

While multi-rate models remain the best descriptors of spontaneous recovery, the idea that the two-state model alone describes savings and anterograde interference has come under considerable scrutiny. In terms of savings, it does not appear that multidimensional time-invariant processes alone can lead to the established patterns of learning upon re-exposure to a perturbation^{7,35,36,49}. Instead, recent studies have suggested that savings is caused by a physical change to the error-based learning system: an increase in error sensitivity. Interestingly, changes in error sensitivity have recently been implicated in anterograde interference⁵⁰ and other phenomena such as meta-learning⁷.

Therefore, it appears that both the original state-space model as well as the two-state model were incorrect in their assumption that one's sensitivity to error is a fixed quantity: fixed across error size and time. We now know that each of these assumptions is untrue. First, error sensitivity is not constant across all error sizes, but declines as errors grow larger^{51,52}. Second, even for a given error size, one's sensitivity to error is not fixed over time but changes with experience. Many studies have elucidated various mechanisms and theories for how error sensitivity is modulated over time. Some have implicated the dopaminergic system in the alteration of learning rates through model-free or reinforcement-based learning mechanisms^{53,54}. One particularly influential class of models relies on a Bayesian framework^{33,55-60}; in the presence of motor noise that disrupts our movement plans, planning noise that disrupts the way we learn and select motor commands, and process noise that causes trial-by-trial variations in external disturbances, the Kalman filter provides a normative framework for how one should modify their rate of learning over time.

One debate that is central to our understanding of savings is whether faster rates of relearning are caused by a recall of previous actions⁶¹, or errors. In support of the latter possibility, Herzfeld and colleagues⁷ provided compelling evidence that sensitivity to error depends on the consistency of past errors. Errors that are experienced consistently and unlikely to change sign are followed by an increase in sensitivity. Errors that are volatile and likely to switch are followed by a decrease in sensitivity. More impressively, the nervous system can simultaneously increase sensitivity to certain consistent errors, while decreasing sensitivity to other inconsistent errors.

This memory of errors model forms the basis of our work in Chapter 5, where we consider how the variance of a perturbation alters not only the rate of learning, but also the total extent of learning. In keeping with the general theme of constant improvement, we make a critical adjustment to the original memory of errors model: we add decay. That is, just as memories are formed and stabilized through a combination of learning and decay, so too might the maintenance of our sensitivity to error. The state-space model of adaptation was designed to capture changes in behavior that occur due to implicit error-based learning mechanisms. However, more recently, many studies have pointed out that motor memories are derived from both implicit and explicit (strategy-based) learning mechanisms^{12,62–66}. In Chapter 5, we dissect these two processes to determine which contributes to improvements in asymptotic performance. We find evidence for the first time that implicit learning processes modify their response to error, and are not inflexible as other authors have suggested^{31,32,52,64,66–69}.

1.3 Fitting models to behavioral data

Spontaneous recovery of prior memory has been observed in several contexts: saccade paradigms^{9,61}, reach paradigms^{8,43,44}, vestibular paradigms⁴⁵, and classical conditioning paradigms⁴⁶. This recovery appears consistent with a mathematical model of learning where experience of error engages two (or more) independent learning processes: a fast process that learns strongly from error but forgets rapidly, and a slow process that learns weakly from error but exhibits robust retention^{8,55}. While the two-state model of adaptation provides an accurate account of spontaneous recovery, it fails to explain other equally ubiquitous phenomena, such as savings⁴⁹. Rather, savings appears to be caused by a change in the underlying process of adaptation: an increase in error sensitivity^{7,35–37}.

To understand how the process of adaptation changes over time, we often need to fit state-space models to behavior at different periods of time. However, when our models posit that behavior is composed of numerous hidden states (like the two-state model) it is not a simple task to extract these hidden states from measured behavior in a robust manner, especially at the level of individual subjects. The most common technique of fitting state-space models to measured behavior is a form of nonlinear optimization called least-mean-square-error estimation (LMSE). This algorithm begins with a state-space model of learning, and then searches the model's parameter space in order to minimize the sum of squared differences between the observed behavior and the model's predictions. LMSE has been widely applied to analyze trial-by-trial changes of behavior during motor learning^{45,48,70–73}.

In Chapter 3, we demonstrate that the LMSE algorithm has an inherent problem: it maximizes a likelihood function that assumes all measured variability comes from the experimenter's observation noise. In other words, every noisy data point we see is because of some electrical anomaly in our recording system that corrupts our measurement of participant behavior. Of course, this is simply not true. Motor behavior exhibits considerable variability both in the execution of movement, but also in the process of movement planning and error-based learning^{58,74-76}. More accurately understanding the hidden processes that support adaptation requires a model fitting technique that accounts for these stochastic realities of motor learning and motor control.

Enter Expectation Maximization (EM)^{77,78}: this statistical algorithm can be used to identify parameter sets that maximize the likelihood of observing participant behavior given a model of adaptation with any hypothesized noise structure. Yet despite the power of this technique, few have attempted to apply it to sensorimotor datasets^{74,79}. The reason for its limited use may be due to constraints with earlier formulations of the algorithm: (1) there is no easy way to constrain the parameter search making it difficult to isolate desired subsets of parameters that may correspond to a desired model of learning (e.g., the two-state model) and (2) there are common experiment probes such as set breaks that make analytical specification of the algorithm's update steps intractable.

In Chapter 3, we eliminate these limitations, by generalizing the M-step of the EM algorithm using a technique called generalized EM⁸⁰. The beauty of this approach is that it can work for any model complexity. However, this comes at the cost of computational workload. We will demonstrate that the EM algorithm better identifies the parameters of two-state models *in silico*, and also provides more robust parameter estimation for individual participant datasets.

With that said, we expect over time that even EM, the frequentist's approach to parameter estimation will yield to more robust Bayesian techniques⁵⁸. The key idea therefore is that model fitting is not a trivial process that should be thrown together after careful collection of data: models we choose make assumptions about the participant that will ultimately shape one's beliefs about their behavior.

With our GEM tool, we interrogate the motor learning systems during the expression of savings in Chapter 3. In both the saccadic system and the reach system, we find that similar mechanisms cause changes in adaptation rate: an adjustment to the error sensitivity of fast adapting processes. In addition, we point out scenarios where the two-state model can satisfactorily account for spontaneous recovery due to a decay in the fast state of learning, and other instances of recovery where a new framework may be required.

1.4 The neural integrator

The study of error-based learning in the human psychophysics literature has been heavily influenced by the force field paradigm⁸¹. Here participants reach to a target location while holding a robotic arm that can apply forces to the arm to change reaching dynamics. The most common set of forces is a velocity-dependent curl field, or one that applies forces perpendicular to the hand's trajectory and proportional to its speed. In Chapter 2, we study the consequence of this type of perturbation on the timescale of a single trial. However, a puzzle exists without the learned response to this perturbation; even though the perturbation only occurs during movement of the limb, on the next trial, participants change the way they move their arm, but also the way they hold still⁸². Specifically, participants exert forces against the handle during periods of holding still that are entirely unnecessary to keep the hand within the target.

In principle, there are at least two possible explanations for these postural changes. The first is that they are caused by adaptation⁸², or more specifically, that the errors we experience during a movement mistakenly generalize to the way our brain updates its plan for holding still. The second is that the way we hold is determined in part by the way we move. To understand this idea, consider the common expression, $d = r \times t$, or distance is equal to rate multiplied by time. If we know the velocity of an object and how long it was moving, we know where that object is in space without having to measure its position. Now, if we think about velocity like a reaching movement, and a position like a posture, it may be the case that we can calculate our postures from our movements.

Remarkably, this architecture is true of the oculomotor system. To move the eye, activity in the superior colliculus¹ causes neurons in the reticular formation to burst⁸³. These burst neurons drive a pulse of activity that moves the eye⁸⁴ to the intended target. But to maintain gaze on the target, motor neurons for the eye must maintain a "step" of activity that counteracts elastic forces of the eye tissue that pull the eye back towards the center of the globe⁸⁵. From where does this step of activity arise? Because the burst neurons of the reticular formation only fire during movement, there must be another set of premotor neurons that drive the motor neurons during gaze-holding. However, stimulation of the burst neurons thickens the plot. When burst neurons are stimulated transiently, the eye not only makes a saccade, but also has a sustained hold period long after the saccade has ended⁸⁶. Why does stimulation of neurons that only care about moving also produce a period of hold still.

The answer to this question lies in integration. If the burst neurons encode the velocity of the eye movement, then a set of target neurons could integrate the velocity over time, yielding the step that maintains the eye still at the desired target. Could this integration be a property of the motoneurons themselves? This does not seem to be the case, as stimulation of the abducens nucleus produces only

movement of the eye, but no holding still⁸⁷. The answer to this question was found elsewhere in the brainstem, in the prepositus nucleus^{13,88-90}. Damage to these neurons produces a specific deficit in the ability to hold the eyes in gaze, but no effect on movement. And equally important, the prepositus was found to contain neurons which linearly encode the position of the eye, and maintain their activity during gaze holding long after reticular activity has ceased⁹¹.

In summary, the eye possesses an amazing property: integration. Neurons in the reticular formation produce phasic activity that moves the eye, which is also integrated in a separate region of the brainstem, the oculomotor integrator. The output of the integrator is a sustained position signal that holds the eye still on target after movement. It is also known that the performance of the integrator is maintained by the cerebellar flocculus^{92,93}, as performance of the saccadic system is maintained by the cerebellar vermis²². Additionally, since these initial discoveries, a neural integrator has also been discovered for movements of the head⁹⁴⁻⁹⁶ within the interstitial nucleus of Cajal (INC).

In Chapter 6, we ask if the reach system might also possess a neural integrator, thus potentially offering an explanation for why adaptation to our movements also changes the way that we hold still. We not only find compelling evidence for this possibility, but also the idea that the putative reach integrator is located in a subcortical area that remains unharmed after cortical stroke. Combined with the observations that inhibition of the motor cortex spares the ability to hold still⁹⁷, transient stimulation of the brainstem in the decerebrate cat yields sustained changes in muscle tone⁹⁸, and stimulation of the spinal cord in the frog produces equilibrium positions of the leg⁹⁹, it seems possible that the putative reach integrator could be housed in the brainstem, spinal cord, or is distributed throughout both regions.

1.5 Self-published content and references

This thesis draws upon both published and unpublished work. Certain passages and figures of Chapter 2 are derived from Albert and Shadmehr (2016)¹⁰⁰. Certain passages and figures of Chapter 3 are derived from Albert and Shadmehr (2018)¹⁰¹. Certain passages and figures of Chapter 4 are derived from Lerner and Albert, et al. (2019)⁵⁰. Certain passages and figures of Chapter 5 are derived from Albert et al. (2019)¹⁰². And certain passages and figures of Chapter 6 are derived from Albert et al. (2020)¹⁰³.

1.6 Specific aims

When we experience an error in the way we move, we change our movement on the next attempt. Here we address three primary questions regarding the process of adaptation (1) how does the brain learn the correct motor response that should be used to prevent future sensory errors, (2) how does this process of learning change over time, and (3) how does this process of movement adaptation change the way that we hold our body still after movement.

In Chapter 2, we address this first aim by considering the feedback error learning hypothesis. By recording muscle activity during perturbed and error-free movements, we extract the commands the brain produces to correct for a movement online, and what it learns from the error on the next trial. We find that the learned response resembles a time-shifted and scaled copy of the feedback response to error, on a muscle-by-muscle basis. On trials where the feedback response to a perturbation is large for a given muscle, learned change in that muscle's activity is also large on the next trial. Finally, we use transcranial direct current stimulation to determine the role of the cerebellum in single trial learning.

In Chapters 3-5 we address our second aim, moving from our investigation of single trial learning, to the way the learning process changes as errors are experienced time and time again. Our approach is model-based; we fit models to behavior in different environments to understand how the dynamics of learning (i.e., the model parameters) change over time. In Chapter 3, we start with the description of a statistical algorithm that can better estimate the parameters of error-based learning models. With this tool in hand, we consider two hallmarks of adaptation, savings and spontaneous recovery, and interpret their origins from the standpoint of a two-state model of learning. For both movements of the eye and the arm, we find that error sensitivity modulation causes savings, without any changes to the process of forgetting. In addition, we describe preliminary evidence that savings might rely on the basal ganglia during adaptation to force fields. Finally, we also provide evidence that spontaneous recovery is caused by the loss of the fast adapting system.

In Chapter 4 we continue studying plasticity in the motor learning system, but this time in the context of anterograde interference. We expose participants to two opposing perturbations *A* and *B*, after some intervening break in time. We discover that there are two independent channels that impair performance during the second exposure: retention of the *A* memory, and a reduced sensitivity to error in *B*. While the former remains strong at 24 hours, the latter has almost entirely recovered. Therefore, we find that the brain changes its sensitivity to error to increase its learning rate (savings) and also decrease its learning rate (interference).

In Chapter 5, we continue to study error sensitivity modulation, not in the context of learning rate, but in the context of the total extent of adaptation. We revisit earlier ideas that the saturation point of motor adaptation is caused by the balancing of learning and forgetting. In an environment where perturbations are variable, we show that the nervous system greatly decreases its total extent of motor adaptation, by changing its error sensitivity, but not its rate of forgetting. We demonstrate that memory of errors models can account for these observations, with the addition of a forgetting term which saturates error sensitivity over time. Finally, we also use various probes to interrogate whether these changes in adaptation extent are mediated by implicit or explicit learning processes. We find strong evidence for the modulation of implicit adaptation, and offer a new framework to consider why changes in implicit learning might have been overlooked in earlier studies. In addition, we show that changes in implicit adaptation are not specific to the presentation of stochastic stimuli, but also occur in the other hallmarks of learning: savings and interference.

In Chapter 6 we move on to our third and final aim: understanding how the process of reach adaptation is related to our ability to maintain stillness of the arm. Drawing parallels to the oculomotor system we propose that part of the muscle commands used to maintain a posture are obtained through the integration of movement commands. To argue for the plausibility of this integration hypothesis we consider the muscle activity of different reaching movements of the macaque. We then test the more critical predictions of the integration model: changes to the way we move change the way we hold. We investigate this possibility first for simple finger movements of the macaque, and then for the control of reaching in the human. We find that the commands for movement appear are integrated into the commands for holding still, even to the extent that our movement commands can trick our holding system to produce a holding posture outside of desired visual targets. Finally, we analyze these same principles in stroke patients and healthy age-matched controls, finding that the cortical lesions that damage the movement controller appear to spare the putative reach integrator.

Chapter 2. The neural feedback response to error as a teaching signal for the motor learning system

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

2.1 Introduction

When we hold an object in our hand, the mass of the object alters the dynamics of our arm, changing the relationship between the motor commands sent from the brain to the muscles of the arm, and the resulting motion of the hand⁸¹. If the object is unfamiliar to us, our movement will exhibit errors, producing a sensation in our proprioceptive and visual organs. That is, the brain experiences errors in sensory coordinates. However, to improve performance, the brain must transform the sensory representation of error into better motor commands in muscle coordinates. How does the transformation from sensory coordinates of error to muscle coordinates of motor commands take place? That is, what signal serves as the teacher for the motor system?

Sensing error engages the proprioceptive and visual organs, but following a delay it also engages sensorimotor feedback pathways, producing reflexive and voluntary corrections that start as early as 50ms into the reach, continuing until the hand arrives at its goal. These corrections represent a sensorimotor transformation that takes error in sensory coordinates and produces a feedback response in muscle coordinates. The feedback response is a sequence of motor commands that can, in principle,

act as a template, providing the brain with an example of how to partially compensate for the unexpected dynamics^{23,24,104–106}. However, testing this hypothesis is difficult because on any given movement, the motor commands are a mixture of what the brain correctly predicted about the dynamics of the object, and what the feedback pathways added in response to the unexpected motion of the arm. To determine the relationship between error and the learning that resulted from error, one must dissociate the motor commands that reflect a process of prediction, from the motor commands that reflect a process of within-movement feedback correction.

Here, we approached this problem by using an important tool: error-clamp trials²⁷. An error-clamp trial makes it possible to reliably guide the movement precisely along a reproducible trajectory. To measure the feedback response to error, we measured the motor commands sent to various muscles of the arm in an error-clamp trial, and then re-measured the commands when novel dynamics (a force field) introduced errors in the reaching movement. By comparing the time-course of signals in the perturbation trial to the preceding error-clamp trial, we obtained a proxy for the neural feedback response to error. Following the perturbation trial, we again introduced an error-clamp trial. The change in the motor commands that occurred from the first error-clamp to the second error-clamp was a proxy for the learning that has occurred following the experience of error. We found that the learned motor commands were a scaled version of the feedback generated commands, but shifted earlier in time. This suggested that the sensorimotor transformation that was provided by the feedback system, from sensory coordinates of error to muscle coordinates of action, acted as a teacher for the motor system, instructing it on how to improve its commands on the next movement.

2.2 Materials and methods

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

2.2.1 Experiment

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

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

Our overall objective was to ask whether the feedback system that corrected for a perturbation during a movement produced a neural signal that became the teacher for the motor system, instructing it on how to predictively cancel the perturbation on the following trial. To test our hypothesis, we first measured the neuro-motor activity in a given muscle (EMG) during an unperturbed movement (termed error-clamp trial 1, EC1). On the next trial, we perturbed the reaching movement via a force field. The difference in EMG between the perturbed trial and the preceding error-clamp trial was our proxy for the feedback-generated response to the perturbation. On the next trial, with 50% probability the reach was in an error-clamp trial (EC2), or another perturbation trial. If a second perturbation trial occurred, the following trial was always an error-clamp trial. The difference in activity between EC2 and EC1 was our proxy for learning, indicating the change in neuro-motor activity due to experience of error in the preceding trial (or a pair of errors in the case of two consecutive perturbation trials).

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

The experiment (Fig. 2.1A) began with a block of 120 null field trials (not shown in Fig. 2.1A). This was followed by two consecutive blocks (labeled blocks A and B) of 263 trials each (one block is shown in Fig. 2.1A). Blocks A and B began with 23 null field trials. Following this, one or two perturbation trials were sandwiched between pairs of error-clamp trials. Each type of perturbation (CW, CCW) and number of consecutive perturbations (one or two) was assayed 10 times, for a total of 40 triplet/quartet perturbations per block. The orientation and number of consecutive perturbations were pseudo-randomly selected and counterbalanced so that subjects experienced an equal number of CW and CCW perturbations. In between each [error-clamp, perturbation, error-clamp] progression, either 2 or 3 null field trials were presented. The paradigm ensured that we could assess learning multiple times without accumulation of learning of either type of perturbation.

2.2.2 Data Recording and Analysis

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

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

The EMG signal was band-pass filtered (10 – 250 Hz) using a 4th order Butterworth filter and full-wave rectified. The filtered and rectified signal was smoothed by scaling the EMG amplitude at each time point by the root-mean-square of the signal in a 40 ms window centered at that time point. Following this pre-processing, we performed a within-subject, within-muscle normalization of each EMG trace by dividing the EMG amplitude at each time point by the average maximum EMG amplitude produced during the initial null trials of Blocks A and B (46 trials in total are included in this average). In

other words, following this normalization, the units of EMG activity for each muscle of a given subject were with respect to the average maximum value recorded in that muscle during an unperturbed reaching movement of the same subject.

To compute the neural correlates of learning from error, we compared the EMG activity recorded in the error-clamp trial following the perturbation (EC2) to the error-clamp trial preceding the perturbation (EC1), for each triplet (or quartet) progression. This difference (EC2 - EC1) represents the trial-to-trial change in the EMG following experience of an error. If the intervening trial was a single perturbation, we termed this change as learning1. If the intervening trials were two perturbations, we termed this change as learning2.

To compute the neural feedback response to error, we first focused on triplet progressions (a perturbation trial between two error-clamp trials) and compared the activity measured in the perturbation trial (P1) with the activity measured in the preceding error-clamp trial (P1 - EC1), and termed this difference feedback1 response. This difference represents how muscle activity was modified to counteract the perturbation during a perturbed movement, relative to an error-free reach. In quartet progressions (two perturbation trials between two error-clamp trials), we computed feedback responses in both the first and second perturbation trials. Importantly, for the feedback response to the second perturbation, we used EC2 from single trial perturbations to estimate the feedforward command produced by the brain after single trial learning, rather than EC1.

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

We asked at each movement of time into the reaching movement, how much the brain had learned from the feedback response. That is, we wished to answer whether there was greater learning from a specific part of the feedback signal (for example, its early part), or did the brain learn from the entire feedback signal. To answer this question, we first shifted each feedback response (independently

for each subject and muscle pair) by the optimal shift determined via cross-correlation. Next, we performed two separate analyses, one across-subject and one within-subject. In the former analysis, we looked at each muscle and field condition separately, and performed across-subject regressions of learning and feedback signals. Learning for a given muscle and field orientation was regressed onto the corresponding feedback response for that muscle and field orientation, independently at every time point. We identified time points for which these fits were statistically significant ($p < 0.05$) and possessed positive slope, signifying that learning and feedback were positively correlated at that point in time, across subjects. To determine the level of correlation between learning and feedback responses within each muscle, we linearly regressed the learning response onto the feedback response over the interval [-100, 500] ms and considered the R^2 value describing this regression, for CCW and CW fields separately.

We next performed with-in subject regressions of learning and feedback signals at each time point. For these regressions, we collapsed across muscles and field orientations. As we had recorded 4 muscles and 2 field orientations, each regression included a total of 8 feedback-learning data points. We considered both the within-subject R^2 metric for this regression (which represents how much of the variation of the learning response is explained by the feedback response for that time point) as well as the slope of the regression (which represents the scaling factor relating feedback and learning). As a control condition, we quantified the baseline correlation between learning and feedback for each subject and each muscle from a data set in which the feedback response was randomly shifted with respect to the learning response. We drew these random shifts from a uniform distribution [0 – 400] ms, and shifted each of the 8 feedback responses independently. The within-subject regression analysis described above was performed on the randomly shifted subject data set, and repeated 200 times, each time resampling shifts from the uniform distribution.

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

Finally, to determine how the relationship between feedback and learning might vary with temporal variation in the magnitude of the feedback response, we considered the fact that in some trials, a subject might produce a strong feedback response to the perturbation, whereas in other trials the same subject might produce a weak response. Did the variability in the feedback response

correspond to the variability in the learning response? To answer this question, for each subject, and each muscle, we separated the data into two classes that corresponded to high and low feedback gains. We will refer to these classes as the 'large' and 'small' feedback responses, respectively. To construct these two labels, we considered each muscle and each subject separately and focused on the feedback1 EMG traces. For agonist muscle feedback responses, we computed the mean feedback1 response over the interval [150, 450] ms relative to movement start. This interval was selected because it best captured differences in the gain of early agonist activity. For antagonist muscle feedback responses, we selected a longer averaging window defined by the range [0, 600] ms. This wider interval was selected so as to include both early inhibition in antagonist responses during the perturbation as well as excitation that occurred near movement termination. A perturbation trial was labeled as high feedback response if its feedback1 EMG trace exceeded the median feedback1 response observed in that muscle and that subject over the appropriate time interval (similarly for the low feedback response label). We computed the mean feedback responses for these two labeled data sets, and then for each labeled feedback trace we computed the learning trace that immediately followed. In addition, we considered kinematic correlates of these responses, corresponding to the maximum perpendicular displacement during the perturbation (for feedback) and maximum error-clamp force production (for learning). We used t-tests to determine if there existed a difference between these kinematic parameters, and expressed their difference as ratios (high feedback trials / low feedback trials).

2.3 Results

We asked whether the feedback response that corrected for a perturbation during a reach produced a signal that acted as a teacher for the motor system, instructing it on how to predictively cancel the perturbation on the following trial.

2.3.1 Measuring trial-to-trial feedback responses to error and learning

Our experiment employed triplet or quartet progressions of error-clamp, perturbation, error-clamp trials, as illustrated in Fig. 2.1A. The average hand paths for the error-clamp trials and perturbation trials are provided for a typical subject in Fig. 2.1B. The time-course of the perturbation-induced displacement perpendicular to the direction of motion is shown in Fig. 2.1C (at left). Following experience of this error, the nervous system altered the motor commands that it produced on the very next trial. To visualize this change, we compared the forces that were produced in the error-clamp trial preceding the

perturbation, to the forces that were produced in the error-clamp trial following the perturbation. The change in the motor commands produced a force pattern that was opposite in direction to that of the displacement (Fig. 2.1C, right).




Figure 2.1. Measuring feedback and learning from EMG. A. Subjects performed a center out reaching task to a single target. Each block began with null field trials. Next, random clockwise (CW) and counterclockwise (CCW) velocity dependent curl field perturbations were applied to the arm. Perturbations were applied either once or twice in a row. Each perturbation or pair of perturbations was sandwiched by error-clamp trials. B. Reach trajectories in the error-clamp trials before and after the perturbation (EC1 and EC2). Error bars are ± 1 SEM in 25 ms intervals. C. Kinematic correlates of error and learning from a single CW perturbation. Left: the time-course for the perpendicular hand displacement during CW perturbations. Right: the learned force production from the single error is shown (forces in EC2 - forces in EC1). Forces are normalized relative to the ideal force that would be produced for complete adaptation. D. Triceps EMG activity during error-clamp and perturbation trials. Left: triceps is active during the movement in EC1. As a result of the feedback response to the perturbation, triceps activity is suppressed early in the CW field, and enhanced near movement termination (P1, red). The brain changes the triceps activity for the subsequent error-clamp trial (EC2, blue). Right: the error feedback response in green (P1 - EC1). The learning response is the trial-to-trial

change in the motor command from EC1 to EC2 (purple, EC2 - EC1). Learning appears to be a time-shifted copy of the feedback response.

By analyzing the temporal patterns of muscle activity in the error-clamp and perturbation conditions, we obtained neural correlates of feedback response to error, as well as trial-to-trial learning. Example traces of EMG activity in the triceps for a typical participant are shown in the left column of Fig. 2.1D. In the error-clamp trial that preceded the perturbation (EC1), the triceps gradually increased its activity, peaking mid-movement at approximately 200ms. In the perturbation (red trace, left column of Fig. 2.1D, labeled P1), the triceps activity was inhibited relative to EC1 for the majority of the reach, but then demonstrated a sharp excitation as the movement was terminated.

To compute the feedback response to error (green, right, Fig. 2.1D), we subtracted the EMG time-course in the error-clamp trial (EC1) from the EMG time-course in the perturbation trial (P1 - EC1). In this participant, the CW displacement produced a feedback response that included an early inhibition of triceps (green curve, right column, Fig. 2.1D), followed by a late excitation of the same muscle. To compute the learning response, we compared the trial-to-trial change in the EMG signal in the two error-clamp trials that sandwiched the perturbation trial, EC1 and EC2 (blue trace, left column, Fig. 2.1D). This difference (EC2 - EC1) represents the trial-to-trial change in the motor command as a result of experiencing a single trial of error. We observed that the learning response (purple curve, right column, Fig. 2.1D) appeared to be a scaled version of the feedback response, but shifted earlier in time.

2.3.2 The learning response is a time-shifted copy of the feedback response to error

Group averaged kinematic and EMG traces for CW and CCW perturbations are shown in Fig. 2.2. The kinematic and force data are shown in Fig. 2.2A, where we have plotted the error induced by the first perturbation (error 1), and the resulting trial-to-trial change in force produced in the subsequent error-clamp trial (learning1). In trials in which a second perturbation followed the first, the errors were smaller (Fig. 2.2A, error2 vs. error1, peak displacement, $p < 10^{-23}$ for both fields). Similarly, in trials in which a second perturbation followed the first, learning following two perturbations was larger (Fig. 2.2A, learning2 vs. learning1, peak force, $p < 10^{-8}$ for CW field, $p < 10^{-4}$ for CCW field).

Fig. 2.2B illustrates the EMG measures of error feedback and learning responses. Perhaps the most striking feature of the data was the similarity between the two traces. We found that in general, the learning response appeared to be a shifted and scaled version of the error feedback response. This was best demonstrated by the EMG in the pectoralis, posterior deltoid, and triceps in their respective antagonist fields, where learning and error feedback traces exhibited initial inhibition followed by

excitation later in the movement (pectoralis for a CCW perturbation; posterior deltoid and triceps for a CW perturbation). Another clear example was presented by the bimodal excitatory pattern in the learning and feedback traces in the pectoralis for CW perturbations, where the second learning excitation peak occurred just prior to movement termination.

For example, in the pectoralis muscle, the CW error feedback response was an excitation that peaked at around 200ms with respect to reach initiation, followed by a second, smaller peak at around 600ms. Learning also possessed two peaks of excitation, peaking at around 50ms and again at around 400ms. In the CCW perturbation, the error feedback response was an inhibition that peaked at around 300ms and an excitation that peaked around 450ms. Learning was also an inhibition followed by an excitation, but its timing had peaks at around 0ms and 250ms. We had naively expected that only the early portion of the error feedback response might resemble the corresponding learning response. However, we found that the learning and error feedback responses appeared similar until the end of the reach (approximately 500ms), implying that both the short- and long-latency error feedback responses were combined and shifted earlier in time to become the learning response.

To quantify the temporal shifts that related learning and error feedback, we computed their cross-correlation and found that the two traces were maximally correlated when the feedback response was shifted earlier in time by 123 ± 61 ms (mean \pm SD across all muscles and conditions) (Fig. 2.3). To combine the data across various muscles, for each perturbation we labeled each muscle as agonist or antagonist. For example, a CW perturbation produced an initial excitation in pectoralis and biceps, but inhibition in posterior deltoid and triceps. Therefore, pectoralis and biceps were agonists in responding to a CW perturbation. We found that the temporal shift from the feedback response to the learning response across muscles was larger when the muscle acted as an antagonist (137 ± 80 ms), responding in the direction of the perturbation, as compared to when the muscle acted as an agonist (109 ± 78 ms, $p=0.042$), responding in the direction opposite the perturbation. Similarly, the optimal shift was larger for learning2 (145 ± 75 ms) than learning1 (101 ± 80 ms, $p < 10^{-3}$), indicating that additional perturbations not only induced additional learning, but also caused this learning to be expressed earlier in time.

To better visualize the temporal relationship between learning and error feedback responses, we plotted the time-shifted error feedback response together with the learning response, for the larger amplitude learning2 traces (Fig. 2.4). The peaks and troughs in the error feedback response appeared to be consistent with the features of the learning response throughout the duration of the movement. In addition, across the various muscles and perturbation orientations, the scaling factor relating the magnitudes of the learning and error feedback response (reflected in the scaling factor relating the left

and right y-axes of Fig. 2.4) was consistent at approximately 25%, suggesting that about a quarter of the feedback response in all muscles became the learned response.




Figure 2.2. Learning and error feedback responses for each muscle. The 0ms time point denotes movement start. Error-bars represent ± 1 SEM. A. Kinematic correlates of learning and error. During

perturbation trials, subjects (n=57) experienced large perpendicular displacements, specific to the orientation of the force field (right axis). Displacement during the 2nd perturbation (Error 2, purple) was smaller than that for the 1st perturbation (Error 1, green) due to partial learned compensation for the force field. As a result of the experience of error, subjects produced lateral forces against the error-clamp channel, in accordance with the field orientation (left axis). The net change in force production after the experience of 2 consecutive perturbations (Learning 2, blue) was larger in magnitude than single trial learning (Learning 1, red), due to the accumulation of 2 single trial learning events. B. EMG correlates of learning and error feedback responses. Learning and error feedback signals for 4 muscles of the upper arm and trunk are provided. Feedback responses (right axis) were in general much larger in magnitude than learning responses (left axis), as indicated by the 25% scaling factor relating the left and right axes for all muscles. The error feedback response for the 1st perturbation (Feedback 1, green) is nearly identical to the error feedback response for the 2nd perturbation (Feedback 2, purple). Note that the reference error-free reach for Feedback 1, is the EC1 trial before the perturbation. For Feedback 2, it is the EC2 trial after single perturbations, which accounts for single trial learning. The learning signals were computed as the change in muscle activity (EC2 - EC1) during the error-clamp reach before and after the perturbation(s). Learning from a single perturbation is shown in red and the accumulated learning from 2 consecutive perturbations is shown in blue. As would be expected, more is learned from 2 perturbations than 1 perturbation.




Figure 2.3. Temporal shifts relating learning and feedback responses. Within-subject cross-correlations were used to determine the time shifts for which the feedback responses and learning responses were maximally correlated. EMG time-courses of learning from a single perturbation and from two consecutive perturbations were cross-correlated with EMG time-courses of feedback responses to the 1st perturbation. Positive values indicate that feedback responses lagged the learning response. Shifts were averaged across subjects and error-bars represent ± 1 SEM. Each group of 4 bars represents a particular learning-feedback condition. Each bar in these groups represents a different muscle (left to right: biceps, triceps, posterior deltoid, and pectoralis). From left to right, the field orientations for each 4-bar group are as follows: CW, CW, CCW, and CCW. From left to right, the number of perturbations experienced for each 4-bar group were as follows: 1, 2, 1, and 2. The right-most bars represent data collapsed across all perturbations and separated for muscles that were agonist (responded early to restore the trajectory) or antagonist.




Figure 2.4. Learning resembles time-shifted copies of the error feedback response. The accumulated EMG learning responses for 2 consecutive errors (purple, left axis) are compared to shifted feedback responses to the 1st perturbation (blue, right axis). The feedback responses were shifted independently for each muscle-field orientation-subject trio, in order to maximally align them with their corresponding learning responses, according to cross-correlation analysis of the learning-feedback time-courses. Clear correspondence between the learning time-course and the feedback time-course is illuminated in this shifted feedback space. Here 0ms, represents movement onset for the non-shifted learning error-clamp trials. The behavior for CW perturbations and CCW perturbations is shown in the left and right columns, respectively. Each figure displays a different muscle. Shaded error-bars represent ± 1 SEM at each recorded time point.

2.3.3 Across subject variability in feedback response

These results indicated that the neural feedback responses in each muscle may be a strong predictor of the learning response in that muscle. To better quantify this relationship, we focused on agonist muscles and asked if those subjects that demonstrated more agonist muscle activity during their feedback response also expressed more learning. For each subject we computed the mean activity in the learning and feedback responses of agonist muscles over the periods [-200 – 500] ms and [0 – 500] ms, respectively. We then asked whether subjects that had shown a greater learning response also produced a greater feedback response. Fig. 2.5A plots the magnitudes of the feedback response and learning response for each subject in each muscle. Fig. 2.5B, shows the same relationship, for the average feedback response and learning response. We found a statistically-significant, positive correlation between the sizes of the two responses across all muscles, indicating that subjects that had shown a larger feedback response were likely to also show a greater learning response.

We performed a similar across-subject analysis but at a much finer temporal resolution to determine the length of the time period over which learning and feedback were positively correlated (i.e. how much of the learning response could be explained by feedback as a function of time in the movement). We first shifted each feedback response by the shift determined via cross-correlation, to optimally align them with the learning time-courses. Next, learning for a given muscle and field orientation was regressed onto the corresponding feedback response for that muscle and field orientation, independently at every time point (Fig. 2.5C). We found that both the agonist and antagonist muscles possessed significantly positive correlations between learning and feedback for the entirety of the movement period. Specifically, this correlation began approximately 100 ms prior to the error-clamped movement onset and saturated for the entirety of the movement period (500 ms on average), falling to baseline levels after the reach terminated. We linearly regressed the learning time-course onto the feedback response (the aligned traces in Fig. 2.4) over this time interval [-100,500] ms, to determine the level of correlation between these signals (Fig. 2.5D). We found that the R^2 values for these regressions were similar across muscles and field orientations, varying within the range $[0.18 \pm 0.02, 0.31 \pm 0.02]$, with a mean value of approximately 0.25.




Figure 2.5. Learning and error feedback responses across individuals. A. Feedback and learning responses in muscles that were agonists for a perturbation. The mean EMG learning response (from 2 perturbations) for each subject was computed over the interval [-200 to 500] ms and regressed onto the mean feedback response [0 to 500] ms, for the 1st perturbation. B. The learning and feedback responses were averaged across muscles, producing a single measure of feedback and learning in each subject. C. Feedback responses (1st perturbation) were shifted for each muscle-field orientation-subject trio independently to achieve temporal alignment with corresponding learning traces (cumulative learning from 2 perturbations), according to cross-correlation analysis. After alignment, learning at each time point was regressed onto the feedback response, across subjects, for each individual time point, independently for each muscle and field orientation (total of 8 combinations). The raster plots at top mark time points for which this regression was statistically significant ($p < 0.05$) and possessed positive slope (indicating a positive correlation). Each line that is oriented left to right shows a particular muscle in one of the field orientations. The continuous-time figure at bottom was constructed from these 8 regression raster lines. At each time point, the number of significantly positive correlations (maximum of 8) was counted. Time at 0ms refers to movement onset. The learning-feedback correlation appears to begin 100ms prior to movement onset, and saturates for the entirety of the movement, which ends on average at 500ms. D. We linearly regressed the learning response onto the feedback response over the time period [-100, 500] ms. This regression was performed for each muscle and field orientation separately. The R^2 value of the regression is provided. The groups at left and right correspond to CCW and CW fields, respectively. Each bar represents a different muscle: from left to right: pectoralis, posterior deltoid, biceps, and triceps.

2.3.4 Within-subject variability in feedback response

Our across-subject analyses indicated that learning and feedback were correlated during preparation and execution of the error-clamped reach. We next asked, if within-subject, the strength of the correlation between learning and feedback responses varied during movement duration. In other words, we quantified the extent to which the variability in the feedback response accounted for the variability in the learning response, as the movement progressed. To quantify this relationship, at each time point we used within-subject linear regression to compare the feedback responses across muscles and perturbation orientations with the corresponding learning responses. To generate a statistical comparison, we generated a null hypothesis by computing the correlation between the two signals when the time-shift was randomly sampled from a uniform distribution of [0-400] ms. We found that the variability in the feedback response, within a subject, accounted for a maximum of around 50% of the variability in the learning response, peaking slightly after movement onset (Fig. 2.6A). However, the correlation between learning and feedback remained significantly above control levels up until movement termination (around 0.5 seconds). Interestingly, the scaling factor describing the learning-feedback regression remained relatively stable within the range 20 – 30% during the reaching movement, indicating that approximately 25% of the error feedback response became the learning response generated on the following trial (Fig. 2.6B).




Figure 2.6. The error feedback response is predictive of the learning response within a subject. For A and B, 0sec refers to movement onset. A. The feedback response accounts for within-subject variation in the learning response during movement. Regressions were performed within-subject across the 8 muscle-field orientation combinations (i.e. 8 points in each regression). Error feedback responses (1st perturbation) were shifted to achieve temporal alignment with learning responses (cumulative learning from 2 consecutive perturbations) according to their cross-correlation. At each time point, learning across the muscle-field orientation pairs was regressed onto the corresponding error feedback response. The R^2 value for the linear regression is provided in the figure (red). The red shaded error bars

indicate ± 1 SEM. To quantify the baseline random correlation inherent in these signals (black), the regressions were repeated, this time randomly shifting the feedback response. The random shifts were sampled from a uniform distribution (0ms to 400ms). The black shaded error bars indicate ± 1 standard deviation across 200 repetitions of the analysis. Comparison of the red and black traces indicates that the feedback response encodes variability in the learning response up until movement termination at approximately 500ms. B. The slopes of the regressions described in Figure A are provided. This slope represents the scaling factor relating learning and feedback. It appears that approximately 25% of the feedback response magnitude was incorporated into the learning response at each point during the movement.

The across- and within-subject results are thus far congruent with the hypothesis that error feedback signals are instructors of learning. We observed that feedback and learning signals possessed a 'scaled-and-shifted' relationship; the feedback response appeared to be scaled down in magnitude, shifted earlier in time, and added to the feedforward motor plan to achieve the learning response. However, our experiment employed only a single perturbation magnitude, leaving one to question the generality of this proposed learning-feedback relationship. To address this question, we considered that on some trials, the subject would strongly resist the perturbation, whereas in other trials the same subject might only weakly resist the perturbation. For each subject and each muscle, we labeled the perturbation trials as large or small feedback response, based on the magnitude of the corresponding feedback1 response time-course. For a given subject and given muscle, the 'large feedback' trials were constructed from all perturbation trials in which the feedback1 response exceeded the median feedback response (see Methods). The 'small feedback' trials were labeled similarly, but for responses that fell below the median feedback1 response. As is implied by this description, we divided the trials for each subject based on their feedback responses, not learning responses.

The two feedback responses are shown for agonist muscles in Figure 2.7 (left column). Labeling our data in this manner revealed that a perturbation produced feedback responses that were highly variable. The large feedback responses (red traces) possessed peak magnitudes in excess of twice the magnitude of the low feedback responses (black traces). For the trials used to label the two classes of triceps and posterior deltoid responses, there was no statistically significant difference in the maximum perpendicular displacement ($p=0.703$ and $p=0.279$, respectively). For the biceps and pectoralis, we found a significant difference between these maximal errors ($p<10^{-5}$ and $p<10^{-8}$, respectively) but their difference (7% for biceps, 8% for pectoralis) was too small to be adequately explained by the two-fold difference in the feedback response gains. To summarize, perturbation trials with nearly identical kinematics showed significant differences in the underlying patterns of feedback muscle activations.

Did these large differences in the feedback responses correspond to differences in the learning signals observed during the subsequent error-clamp trial? We computed the learning responses that were induced by the large and small feedback trials (Fig. 2.7, right column). Remarkably, the differences in the feedback gains were mirrored in the magnitudes of the learning time-courses. A perturbation that produced a large feedback response was followed by a large learning response, as shown by the red and black traces in Fig. 2.7. This result suggested that the agonist learning responses were highly sensitive to the gains of the feedback response. It did not appear that the size of the agonist learning responses was on average indicative of the force being produced during the channel trial (maximum peak force, $p > 0.05$ for all muscles).

We next performed the same analysis for each muscle, in their respective antagonist field. The small and large antagonist feedback responses are shown in the left column of Figure 2.8. Our labeling method revealed two distinct differences between the high and low feedback responses. First, the early period of inhibition was attenuated in the large feedback responses (red traces) relative to the small feedback (black traces) responses. The second difference between the two feedback responses was characterized by enhanced late excitation of the large feedback response relative to the small feedback response. Again, as for the agonist responses, these differences were not reflective of some large difference in the underlying kinematics of the error. The only muscle for which we observed a statistically significant difference in the maximum perpendicular displacement was the triceps ($p = 0.019$), but this difference (only 3%) was too small to be explained by the differences in feedback response gains.¹⁰⁶

Once again, these differences in feedback responses were paralleled in learning (right column, Fig. 2.8). Learning traces for the small feedback response (black traces) diverged from those pertaining to the large feedback response (red traces) for the entirety of the reach. This divergence obeyed the differences we observed in the corresponding feedback responses. We found that the learning time-course resembled time-shifted and scaled replicas of the feedback response. One exception to this relationship was the activity of the triceps and posterior deltoid in the high feedback response, which lacked an early period of inhibition that would be expected from consideration of the corresponding feedback traces. We speculate that this inhibition was likely cancelled by co-contraction of agonist-antagonist pairs, which based on previous accounts¹⁰⁶ is often seen in the initial stages of force field learning. Indeed, in these two muscles, during the initial part of the reach, the excitation in the high feedback group's learning response overlapped with excitation in agonist muscles (*i.e.* we observed co-

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




Figure 2.7. Agonist learning magnitude correlates to the feedback response gain. Triplet trial progressions were classified into two equally-sized groups based on the magnitude of the feedback response of agonist muscles. Large and small feedback groups were constructed from individual responses that fell above and below the median response over the period [150, 450] ms. In the left column, the mean time-courses of the small feedback response group (black traces) are overlaid on the feedback responses for the large group (red traces). At right, the learning responses corresponding to the large (black traces) and small (red traces) feedback groups are shown. The difference in magnitude of the learning responses mirrored that of the corresponding feedback time-courses. Importantly, these classifications were made solely based on the feedback responses, not the resultant learning. Each figure indicates a different muscle: from top to bottom – pectoralis, posterior deltoid, biceps, and triceps. Movement onset is indicated by the 0ms time point. Errors bars indicate ± 1 SEM.




Figure 2.8. Antagonist learning magnitude correlates to the feedback response gain. Triplet trial progressions were classified into two equally-sized groups based on the magnitude of the feedback response of antagonist muscles. Large and small feedback groups were constructed from individual responses that fell above and below the median response over the period [0, 600] ms. In the left column, the mean time-courses of the small feedback response group (black traces) are overlaid on the feedback responses for the large group (red traces). At right, the learning responses corresponding to the large (black traces) and small (red traces) feedback groups are shown. The difference between the feedback responses appeared to correlate with differences in the sign of the learning responses. Importantly, these classifications were made solely based on the feedback responses, not the resultant learning. Each figure indicates a different muscle: from top to bottom – pectoralis, posterior deltoid, biceps, and triceps. Movement onset is indicated by the 0ms time point. Errors bars indicate ± 1 SEM.

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

2.3.5 Control studies

An assumption critical to our analysis was that the EMG patterns in error-clamp trials represent the motor output in an error-free movement. To test for this, we compared the EMG traces in EC1 with EMG traces recorded during baseline reaching conditions in the null field in each muscle (Fig. 2.9A). We computed the mean EC1 signal across blocks 1 and 2 of the experiment and compared it to the mean null field EMG signal of the 23 trial null periods that commenced each block. Indeed, the EC1 signal (red curves, Fig. 2.9A) appeared indistinguishable from that recorded in the null field periods (black curves, Fig. 2.9A). This analysis also suggested that on average, learning from the CW or CCW perturbations was washed out during the intervening null trials between consecutive perturbation periods, another assumption critical to our analyses.




Figure 2.9. Error-clamp trials provide accurate approximations to unperturbed movements. For A, B, and C, the 0sec or 0ms time points refer to movement onset. A. EMG during error-clamp trials before the perturbation is identical to null field EMG. The EMG activity of each muscle in the error-clamp trial preceding a movement (EC1, red) is contrasted with the EMG signal in the 46 null field trials that began blocks A and B (null, black, 23 trials at the start of each block). Each figure indicates a different muscle: top left – pectoralis, top right – triceps, bottom left – posterior deltoid, bottom right – biceps. B. Sufficient washout occurred between perturbations. We compared the EC1 activity after CW perturbations (red traces) to EC1 activity after CCW perturbations (black traces). We found that these activities were identical, indicating complete washout between consecutive perturbations. The layout of the figure is identical to A. C. The tangential kinematics of the error-clamp movements before and after the perturbation are identical. The tangential velocity in all of the error-clamp trials before the movement (EC1, black) was identical to that of the post-perturbation error-clamp (EC2, red). Errors bars indicate ± 1 SEM.




Figure 2.10. The difference between consecutive feedback responses is largely due to single trial learning. The difference in the EMG activity in the perturbation trials (purple, Pert. 2 - Pert. 1) is compared to single trial learning (red, Learning 1) as measured as the EMG change from the pre-perturbation error-clamp trial to the post-perturbation error-clamp trial. These signals should be similar, barring differences in the feedback responses for the 1st and 2nd perturbations. EMG activity in the CW and CCW fields is shown in the left and right columns, respectively. Each row of figures applies to a particular muscle. From top to bottom these muscles are as follows: pectoralis, posterior deltoid, biceps, and triceps. Here, 0ms refers to movement onset. Errors bars indicate ± 1 SEM.

To ensure that this apparent washout was not trivially caused by the cancellation of residual learning of the oppositely-oriented CW and CCW fields, we computed the mean EC1 signal

corresponding to trials that followed CCW perturbations and compared this to the mean EC1 signal of trials following CW perturbations (Fig. 2.9B). We found that these two groups of EC1 activities were identical, confirming that sufficient washout occurred between consecutive triplet/quartet progressions.




Figure 2.11. Learning from one error is similar to learning from the next error. The EMG learning signal induced by the 1st perturbation (red, 'learning from 1st perturbation') is compared to the learning induced by the 2nd perturbation (blue, 'learning from 2nd perturbation'). For the 1st perturbation, learning was calculated by subtracting the pre-perturbation error-clamp EMG signal from the post-perturbation error-clamp EMG signal. For the 2nd perturbation, learning was calculated by subtracting the post-perturbation error-clamp EMG signal after a single perturbation from the post-perturbation error-clamp EMG signal after two consecutive perturbations. EMG activities in the CW and CCW fields are shown in the left and right columns, respectively. Each row of figures applies to a particular muscle.

From top to bottom these muscles are as follows: pectoralis, posterior deltoid, biceps, and triceps. Here, 0ms refers to movement onset. Errors bars indicate ± 1 SEM.

To measure the learning response, we compared the EMG in EC2 with the EMG in EC1. This comparison requires that the kinematics of the two movements be identical. To check for this, we compared the tangential component of the reach in the two error-clamp trials (Fig. 2.9B), and found the two to be indistinguishable.

We wanted to determine if some simple linear transformation existed between the learning EMG signals and the learned force profiles. We found a weak, but significant ($p = 0.0186$) positive correlation between the magnitude of early agonist EMG activity (mean across the 4 muscles) over the period $[-100 - 200]$ ms and learned force production (Fig. 2.2A, left). This offered some evidence that larger EMG learning corresponded to larger learned forces. However, force production truly relates to changes in the net torque about a joint and therefore is determined by the balance between agonist and antagonist muscle activities; we found no significant ($p = 0.125$) correlation between antagonist muscle activity and force production, during this early period of the reach.

To determine the robustness of our estimate for the learning response, we considered an alternative approach by comparing the difference in the EMG recorded in the P2 and P1 trials. That is, $P2 - P1$ should resemble learning from a single error, provided that the feedback responses during the first and second perturbation trials are the same. Fortunately, the feedback1 and feedback2 responses of Fig. 2.2B are rather similar, though not identical. Therefore, we compared the $P2 - P1$ EMG signal with the learning1 signal, across subjects (Fig. 2.10). We confirmed that $P2 - P1$ resembled our estimate of the learning response (despite the fact that the measures relied on different comparisons). However, their correspondence was not exact, reflecting differences in feedback1 and feedback2. For example, the accentuated peak and trough in the CW biceps perturbation response difference (1st column, 3rd row of Fig. 2.10) corresponded precisely to a slight temporal shift relating the feedback1 and feedback2 responses of Fig. 2.2B (also 1st column, 3rd row).

Thus far we have analyzed learning from a single perturbation as well as cumulative learning from two perturbations. If the 'scale-and-shift' relationship between feedback and learning is a general learning rule, we should observe this phenomenon for the learning response from the 2nd perturbation experienced in quartet trial progressions. Given that the corresponding feedback responses were nearly identical during the 1st and 2nd perturbations (Fig. 2.2) we would predict that the single trial learning time-courses induced by these perturbations should be quite similar. To compute the learning that occurred solely due to experience of the 2nd perturbation we considered the difference EC2 (quartet) -

EC2 (triplet). This difference represents the learning that takes place in the feedforward command after the experience of P2. In Fig. 2.11 we have plotted our estimate of learning from the second perturbation in the quartets alongside our estimate of learning from the perturbation in the triplets. The two possess an extremely close resemblance. Thus, this analysis provided further evidence that supports our working hypothesis concerning the relationship between learning and feedback.

Finally, we wanted to ensure that the relationship between learning and feedback established by our analysis of high and low feedback trials (Figs. 2.7 and 2.8) was not trivially the result of some process that varied systematically during the progression of the experiment. For example, perhaps subjects became less sensitive and responsive to the triplet trial progressions due to the uncertain nature (i.e. frequency and orientation) of the force field perturbations. If such a phenomenon was responsible for the feedback response variability, we would expect that there would be some trend in which trials corresponded to the large and small feedback triplet trial progressions. However, we found no such trend in the trial orderings (Figure 2.12). The large and small feedback groups were constructed of trials that were sampled approximately uniformly across the experiment for both agonist and antagonist muscles. This finding suggests that the observed changes in feedback response gains was the result of random within-subject fluctuations in the gain of the neural feedback controller, rather than a systematic modulation due to passage of time.

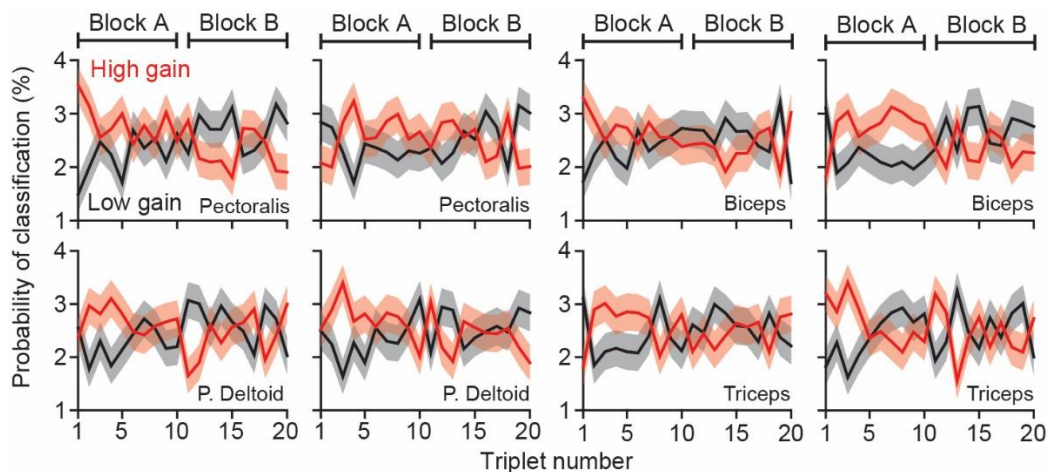


Figure 2.12. The feedback response gain varied randomly with the progression of the experiment. Triplet trial progressions were classified into two equally-sized groups based on the magnitude of the feedback response, for agonists and antagonists separately. Large (red) and small (black) feedback groups were constructed from individual responses that fell above and below the median response over a critical time window (agonists: 150 – 450ms; antagonists: 0 – 600ms). Here we show the probability of any particular triplet belonging to the large or small feedback group. The top and bottom rows show agonist and antagonist responses, respectively. Each figure indicates a different muscle: from left to

right – pectoralis, posterior deltoid, biceps, and triceps. No clear relationship between the triplet number and its assignment (large or small) existed for any of the muscle responses. Errors bars indicate ± 1 SEM.

2.3.6 The effect of cerebellar stimulation on learning and feedback responses

It is well known that the cerebellum is a critical site that supports motor adaptation. In the context of force field adaptation, stimulation of the cerebellar cortex appears to modify the amount of learning that takes place over a sequence of many trials⁴. Can we observe the cerebellar influence on adaptation after only a single trial? To answer this question, our experiment consisted of two phases. In the first phase, participants performed reaching movements in the absence of any stimulation. In the second phase of the experiment, we stimulated the cerebellum (2 mA, 25 min.), ipsilateral to the dominant arm (3 cm lateral to the inion) using two 5 x 5 cm saline-soaked sponges (Phoresor II device). Participants were divided into 3 groups: a sham stimulation group, a cathodal stimulation group, and an anodal stimulation group.

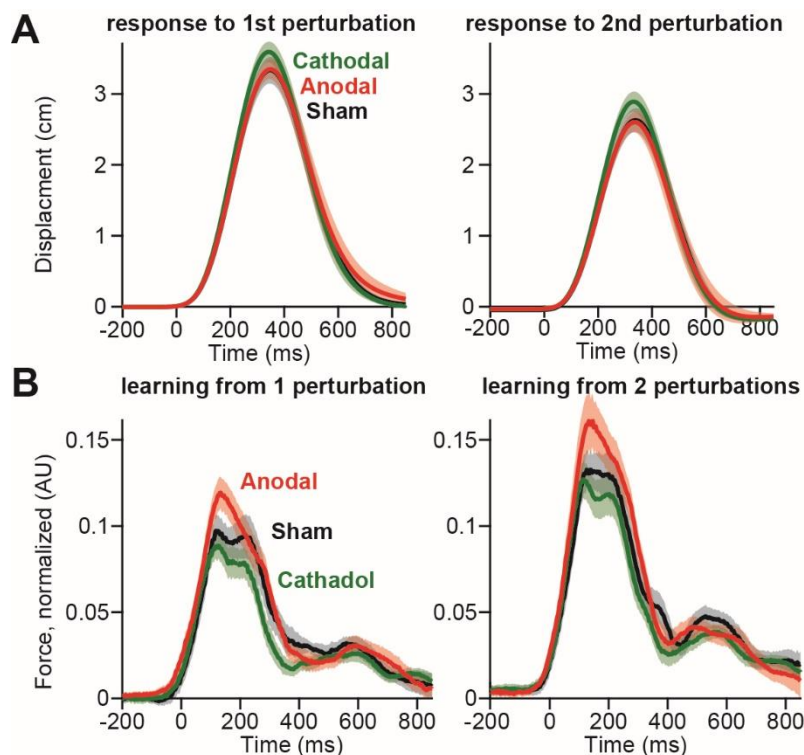


Figure 2.13. Cerebellar stimulation alters the amount of learning. A. The lateral displacement of the arm during the first perturbation (left) and second perturbation (right) is shown. Note the slight increase in error corresponding to cathodal stimulation. B. The amount of learning from one perturbation (left) or two consecutive perturbations (right) is shown. Note that anodal and cathodal stimulation appear to have a bidirectional effect on the magnitude of learning.

As before, perturbations were randomly encountered, either once, or twice in a row. Fig. 2.13 shows the feedback response to error (Fig. 2.13A) and the total amount of learning (Fig. 2.13B) that resulted from either one perturbation, or two perturbations. We do not provide the EMG analogues of these traces, as we did not have enough power to identify specific differences in muscle activity. Nevertheless, cathodal stimulation appeared to have a small effect on the feedback response to the error, slightly increasing the maximum error experienced on force field trials (Fig. 2.13A). Furthermore, a bidirectional effect of stimulation on the amount of learning was also observed. Cathodal stimulation appeared to decrease the amount of learning that resulted from either one or two perturbations, whereas anodal stimulation appeared to increase the amount of learning (Fig. 2.13B). These findings are consistent with the idea that modulation of the cerebellum alters the experience of learning from a single error. In terms of the multiple timescales of motor memory, this is most consistent with a role of the cerebellum in the fast timescale of adaptation, though it does not preclude the possibility of cerebellar involvement with slower timescales of memory⁸.

2.4 Discussion

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

During a reaching movement, the sensory encoding of error engages spinal and supra-spinal neural circuits that, following a delay, produce motor commands, partially correcting for the error. The motor response to the mismatched sensory feedback is termed an error feedback response²³. With practice, the gain of the sensory feedback response to error can be increased, resulting in more vigorous corrections to the repeatedly experienced errors¹⁰⁷. However, because of inherent delays in sensory feedback, the error feedback response alone cannot fully compensate for the errors. To solve this problem, theory has suggested that the feedback response may serve as a teaching signal for the brain, resulting in changes in the motor commands that are produced in a “feedforward” way²³.

2.4.1 The feedback response to error is incorporated into the motor plan for the next movement

In support of this hypothesis, several studies have demonstrated that on a trial-to-trial basis, aspects of the feedback response to error appear to be incorporated into the learned response^{24,104–106}. Specifically, these studies demonstrated that as the brain learns to compensate for an externally imposed perturbation, the early feedforward component of muscle activity during a reach grows to resemble the feedback-related muscle activity.

In these previous studies, the distinction between feedback and learning responses was made largely on the basis of timing of the response features, rather than separately isolating each component for the entirety of the movement trace. Here, we approached the problem by using error-clamp trials that sandwiched perturbation trials. This technique allowed us to more precisely examine the temporal relationship between learning and feedback, isolating the time-course of each response during the entire movement. We found that the error feedback response was a complex temporal pattern of activation/deactivation of each muscle, and included short- and long-latency feedback components that corrected for the perturbation and brought the hand to the target. Following experience of a single error, motor commands changed on the following trial. Similar to the feedback response, the learning response possessed complex temporal dynamics specific to each muscle, which persisted during the entirety of the movement. Remarkably, the time-course of the learning response appeared to be tightly correlated to the feedback response. In fact, the learning response included essentially all components of the feedback response, scaled by roughly 25% in magnitude (after 2 perturbations), and shifted approximately 125ms earlier in time. Considering that voluntary feedback corrections can be expected at around 150ms after the initiation of movement¹⁰⁸, or perhaps as early as 130ms in velocity-dependent curl fields²⁴, this shift appears to account for the delays in the salient voluntary correction component of the feedback response. It is not at this point clear if the magnitude of this shift is an invariant feature of the nervous system, or is drawn from the relative timing of the feedback controller's output and the onset of the movement.

One might have predicted that in the initial stages of training (*e.g.* after the experience of 1 or 2 errors), only earlier portions of the feedback response are learned. However, it appeared that the entirety of the feedback response affects the learning response (Figs. 2.5C and 2.6A); the fraction of the feedback response that transferred into the learned response was invariant throughout the trajectory of the movement (~25%, Fig. 2.6B). The consistency in the scaling factor, both across muscles (Fig. 2.4), and time (Fig. 2.6B), suggests that there exists some set point of error feedback sensitivity in muscle space.

We speculate that this set point may be intimately related to regulation of error sensitivity⁷ during motor learning. The amount learned from the experience of error may be in part modulated by the extent to which the feedback response is a reliable corrective signal, based on the task history and environment.

The motor commands produced by the sensory feedback system served as a template for the motor learning process, becoming the motor commands that were used to predictively compensate for the novel dynamics on the subsequent attempt. To determine the generality of this claim, we sorted our data within-subject based on the magnitude of the feedback responses. We found that for both agonists and antagonists, when the feedback response was large, so was the learning response (Figs. 2.7 and 8), suggesting that the trial-to-trial variability in the gain of the feedback system strongly affected the trial-to-trial variability in learning.

Our result calls for a re-evaluation of a previous claim that the nervous system produces nonspecific learning responses to single errors^{109,110}. These authors found that single trial perturbations with differing dynamics did not induce different kinematic correlates of learning. To reconcile our results with these prior findings, we speculate that the feedback system produced nonspecific, or saturated responses to the perturbations used in these studies, which would result in identical learning responses according to the 'scale-and-shift' hypothesis. We also note that kinematic similarity does not imply that the control signals in muscle space were identical; rather sophisticated differences might be present at the neuro-motor level, without corresponding distinct behavioral correlates, as is the case in Figures 2.7 and 2.8.

Previous investigators have observed co-contraction of agonist-antagonist pairs of muscles during force field adaptation, which increases the stiffness of the arm and likely stabilizes it against an unpredicted perturbation^{104-106,111}. We speculate that in our task, occasional co-contraction during the error-clamp trial following the perturbation may have resulted in the cancellation of learned antagonist inhibition (*e.g.* absence of initial inhibition in triceps and posterior deltoid in Fig. 2.8) thus partially masking the learning instructed by the error feedback signal. However, on the whole, given that the mean antagonist learning and feedback responses demonstrated clear decreases in muscle activity relative to an error-free reach (Figs. 2.2 and 2.4), rather than co-activation across agonist-antagonist pairs, we suspect that our task's perturbation infrequency might have partially disengaged this neural impedance controller. This relative absence of co-contraction may also explain a difference between our results and a general learning rule posed by Franklin et al. (2008)²⁴ where the authors found that antagonists increased their activity in response to error, in addition to their agonist counterparts. We

should note that apart from this difference, our proposed 'scale-and-shift' relationship between learning and feedback is quite similar to a computational architecture proposed by these authors.

2.4.2 On the causality between feedback responses and subsequent learning

Our results describe a correlation, and not a causal relationship. However, an earlier work provides some evidence for the idea that the feedback response is causally related to the learning response. Haith et al. (2011)¹¹² asked subjects to reach in a force field but not produce the voluntary corrective response associated with bringing the hand back to the target. They did this by having people reach to a line during perturbation trials (rather than to a point). They found that the learned response, measured via the forces that subjects produced in error-clamp point-to-point reaching trials, was significantly smaller as compared to when the perturbation trials were point-to-point. This result is consistent with our 'scale-and-shift' error feedback teaching hypothesis, as the absence of late voluntary corrections in the reach to a line movement would negatively influence the learning response late in the point-to-point movement.

We observed that people who produced a larger feedback response to error also learned more from the error that they had experienced at the level of muscle control signals (Figs. 2.5). This predicts that individuals who have reflexes that produce a stronger response to a given perturbation are likely to be able to adapt faster to that perturbation. Therefore, some of the between subject differences in rates of adaptation in force fields⁷⁵ may be due to between subject differences in their ability to correct for sensory prediction errors using reflexive and voluntary feedback pathways. However, because feedback response to error is in itself an adaptive process that benefits from experience^{105,113,114}, we do not know whether people who learn more do so because of inherently better feedback control, or because they are able to better tune their feedback control system to the range of perturbations.

Scaled and shifted feedback responses are not the sole progenitors of motor learning. More likely, feedback instruction is one of many mechanisms in a potpourri of neural motor learning/control strategies. For example, there exist motor learning paradigms where feedback corrections are not required; sensory prediction errors are sufficient to drive motor learning¹⁴. An example of this is saccade adaptation, where visual error detected at completion of the movement is sufficient for modulation of saccadic gain¹¹⁵. However, Wallman and Fuchs (1998)¹¹⁵ found that even in saccade adaptation, if subjects are allowed to correct their saccades with a second saccade that responds to the error, the rate of learning is faster than if this motor correction was not allowed. Therefore, the act of generating a corrective motor response appears to enhance the process of learning.

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

Chapter 3. Estimating properties of the fast and slow adaptive processes during sensorimotor adaptation

Experience of a prediction error recruits multiple motor learning processes: some that learn strongly from error but have weak retention, some that learn weakly from error but exhibit strong retention. These processes are not generally observable, but are inferred from their collective influence on behavior. Is there a robust way to uncover the hidden processes? A standard approach is to consider a state-space model where the hidden states change following experience of error, and then fit the model to the measured data by minimizing the squared error between measurement and model prediction. We found that this least-squares algorithm (LMSE) often yielded unrealistic predictions about the hidden states, possibly due to its neglect of the stochastic nature of error-based learning. We found that behavioral data during adaptation was better explained by a system in which both error-based learning and movement production were stochastic processes. To uncover the hidden states of learning, we developed a generalized Expectation Maximization (EM) algorithm. In simulation, we found that while LMSE tracked the measured data marginally better than EM, EM was far more accurate in unmasking the timecourses and properties of the hidden states of learning. In a power analysis designed to measure the effect of an intervention on sensorimotor learning, EM significantly reduced the number of subjects that were required for effective hypothesis testing. In summary, we developed a new approach for analysis of data in sensorimotor experiments. The new algorithm improved the ability to uncover the multiple processes that contribute to learning from error.

3.1 Introduction

When people and other animals perform a movement that produces an unexpected outcome, they learn from the resulting error and retain a portion of this learning over time. Analysis of behavior in numerous contexts, including saccade paradigms^{9,61}, reach paradigms^{8,43,44}, vestibular paradigms⁴⁵, and classical conditioning paradigms⁴⁶, has revealed an interesting behavioral property termed spontaneous recovery: following learning, washout, and then passage of time, behavior spontaneously reverts back to the previously learned state. That is, washout does not return memory to its baseline condition, but appears to engage a process that masks the previously acquired memory. With passage of time, this mask appears to lift, resulting in re-expression of the learned behavior.

Spontaneous recovery is consistent with a mathematical model of learning where experience of error engages two (or more) independent learning processes: a fast process that learns strongly from error but forgets rapidly, and a slow process that learns weakly from error but exhibits robust retention^{8,55}. It is possible that the putative learning processes represent interactions between distinct neural systems such as the cerebellum, the motor cortex, and the parietal cortex^{4,47,116}. The learning

processes may also be represented in behavior as explicit and implicit processes⁷², body versus world estimation³³, temporally-labile versus temporally-stabile processes²⁹, a memory of errors^{4,30}, and preparation-time dependent processes³². In all of these approaches, experience of error engages multiple hidden processes that act in parallel, each responding to error with their own characteristic learning and retention properties, and then combining their outputs to jointly influence behavior.

An example of a neural system that might implement such a learning model is the cerebellum and its principal cells, Purkinje cells (P-cells). Following experience of a visual error, some P-cells prefer that error^{25,117} and experience a strong modulation of their complex spikes. Experience of a complex spike in a P-cell produces plasticity among some of the synapses, resulting in a reduction in the simple spikes that the P-cell produces on the subsequent trial¹¹⁸. This resembles a learning process that adapts strongly from error. Interestingly, these neurons also exhibit rapid forgetting; displaying little or no retention (in terms of change in their simple spikes) after 10 seconds of time has passed since experience of error¹¹⁸. Other P-cells do not prefer that same error; for them, that error produces suppression of their complex spikes below baseline, resulting in a weak potentiation of their simple spiking rate. This resembles a learning process that adapts weakly from error. With repeated trials, the experience of these errors produces two timescales of change in the simple spikes: fast change in the P-cells that prefer the error, slow change in the P-cells that do not prefer the error¹¹⁸. Therefore, one potential neural mechanism for the multiple learning processes may be in the cerebellum, where the various P-cells learn differently from a given error based on their preference for the direction of that error.

Here, our aim was to build a mathematical tool that could, in principle, extract the hidden processes from observed behavior. A common tool currently used for analysis of behavioral data in motor learning is a form of nonlinear optimization called least-mean-square-error estimation (LMSE). This algorithm begins with a state-space model of learning, and then searches the model's parameter space in order to minimize the sum of squared differences between the observed behavior and the model's predictions. LMSE has been widely applied to analyze trial-by-trial changes of behavior during motor learning^{45,70-72}. However, we found that when we applied LMSE to behavioral data collected during a typical adaptation experiment, the algorithm fit the measured data well, but for many subjects, it yielded unrealistic predictions about the properties of the underlying hidden processes. We speculated that this problem was due to a fundamental limitation of LMSE: in the context of error-based learning where the errors we make are influenced by the movements we generate, the LMSE algorithm

is equivalent to a maximum likelihood estimator for a system that is ignorant of the stochastic nature of learning and moving.

We therefore wondered if an algorithm that considered both of these sources of stochasticity, noise in the system that learned from error, and noise in the system that produced the motor output, could improve our ability to estimate the hidden processes. We derived a canonical form of the two-state model that cast the learner in a framework where both learning from error, and the production of a movement, were stochastic processes^{74,76}. In this framework there was uncertainty in both the evolution of hidden states and the observation of movement⁵⁵.

To estimate parameters of this more general model of learning, we considered a maximum likelihood approach that was first applied to sensorimotor learning by Cheng and Sabes⁷⁴, called expectation maximization (EM). Unfortunately, it is difficult to constrain EM to enforce traditional two-state dynamics. In addition, previous descriptions of the algorithm assumed time-invariant state-space transitions⁷⁹. In contrast, a typical motor control experiment relies on behavioral probes such as error-clamp trials²⁷ and set breaks. The latter type of probe can make the state-space transitions time-dependent.

Here, we illustrate how a generalized EM algorithm can be used to estimate the hidden processes that may underlie a learning problem, even when constraints are applied to the model parameters, error-clamp trials are included in the experimental paradigm, and the generative model of learning varies in time due to the occurrence of set breaks. The result is a new mathematical toolbox.

We demonstrate that EM fits observed human reaching behavior similarly to LMSE, but uses an underlying parameter set whose likelihood is more likely to explain the observed behavior. To further evaluate EM, we consider several simulated sensorimotor learning paradigms. Unlike behavioral data measured in the laboratory, the fast and slow learning processes were explicitly known in the simulated data sets, allowing us to objectively quantify performance of EM.

Our work has two main results: 1) behavior during sensorimotor learning is better represented by a generative model in which both the generation of movement and learning from error are stochastic processes, and 2) in such a system, EM significantly improves the ability to uncover the hidden states of learning. The resulting algorithm has the practical implication of reducing the number of subjects that are needed for statistical testing of hypotheses.

3.2 Materials and methods

Our goal was to produce a mathematical toolbox that could robustly estimate the properties of a two-state learning process from data collected in a typical adaptation experiment. We employed a statistical algorithm known as EM. EM is an iterative parameter estimation technique that can be used for system identification in the presence of latent variables. As its name suggests, EM is composed of two separate steps. In the expectation step (E-step), a Kalman filter is used to provide the best estimate of the hidden states under the current estimate of the model parameters. In the maximization step (M-step), maximum-likelihood estimation is used to identify a set of model parameters that maximizes an objective function known as the expected complete log-likelihood function. The E- and M-steps together are guaranteed to identify model parameters that improve the likelihood of observing the measured data. The E- and the M-steps are iterated until the likelihood of observing the measured data converges.

Current application of EM in the sensorimotor literature is limited to linear time-invariant (LTI) systems⁷⁴ where the generative model assumes no constraints on the dynamics of the hidden learning processes. In this case, the E- and M-steps can be performed analytically via closed-form equations. Although closed-form equations simplify one's search for the optimal parameter set, performing this analytical formulation makes it difficult to enforce the conventional properties of a two-state model (e.g., the fast process forgets more rapidly than the slow process). In addition, closed-form expressions for the M-step of the algorithm cannot always be derived for complicated likelihood functions, as is the case when set breaks are included in the generative model. To rectify these issues, we considered a more general form of the EM algorithm, aptly named generalized EM⁷⁷. Our version of this algorithm is similar to previous descriptions of EM in the sensorimotor literature, differing only in the implementation of the M-step. In our algorithm, numerical techniques are used to search for the maximum value attained by the expected complete log-likelihood function within a constrained parameter space. Our implementation is described in Appendix 1.

To assess the performance of our approach, we tested EM against a different technique for fitting state-space models to behavioral data, LMSE. LMSE is a technique in which one identifies model parameters that minimize the mean-squared-error between the measured behavior and the model predicted behavior. In an error-based learning model, where movement result in errors that teach the motor learning system, LMSE is identical to a maximum likelihood estimator for a model that assumes that motor learning and movement production are deterministic processes, with uncertainty only arising in our measurement of the behavior of the subject, and is described in detail in Appendix 2.

We applied our EM and LMSE algorithms to data recorded from human subjects performing a visuomotor rotation task. Additionally, we simulated various paradigms and generated data sets in which the fast and slow states were known. Using these data, we asked how accurately EM and LMSE could uncover the true states. Finally, we used EM and LMSE to perform a power analysis, estimating how many subjects were needed in order to robustly test effectiveness of an intervention that modified error sensitivity and retention in a simulated population.

3.2.1 State-space model of learning

Here we derive a two-state model of learning. Our model possesses the same canonical form as previous models in the literature, but differs in a fundamental way: we mathematically formalize the inclusion of error-clamp trials and set breaks. This modification yields a more complicated time-varying form of the state-space equations, but also makes our approach compatible with general sensorimotor adaptation paradigms.

A learner is presented with a sequence of trials where she is instructed to make a movement towards a target. On trial n she is presented with the target $\mathbf{g}^{(n)}$. To achieve this target, she produces a movement $\mathbf{u}^{(n)}$, and observes the consequences of her action $\mathbf{h}^{(n)}$. The consequence of her action (Eq. 3.1) is determined by her movement as well as any external perturbation to her movement, denoted by $\mathbf{r}^{(n)}$.

$$\mathbf{h}^{(n)} = \mathbf{u}^{(n)} + \mathbf{r}^{(n)} \quad (3.1)$$

The learner adjusts her movement toward the target according to her estimate of the perturbation $\hat{\mathbf{r}}^{(n)}$. The movement she produces is altered by motor execution noise $\varepsilon_u^{(n)}$, which has a normal distribution with mean 0 and variance σ_u^2 . We have:

$$\mathbf{u}^{(n)} = \mathbf{g}^{(n)} - \hat{\mathbf{r}}^{(n)} + \varepsilon_u^{(n)} \quad (3.2)$$

In a two-state model, we assume that the learner estimates the perturbation via two independent states, referred to as the slow and fast states of learning. The values of the two states on trial n are represented by the vector $\mathbf{x}^{(n)} = \begin{bmatrix} x_s^{(n)} & x_f^{(n)} \end{bmatrix}^T$ where x_s and x_f are the scalar-valued slow and fast states. The learner's estimate of the perturbation is related to these states according to:

$$\hat{\mathbf{r}}^{(n)} = \mathbf{c}^T \mathbf{x}^{(n)} \quad (3.3)$$

Here $\mathbf{c} = [1 \ 1]^T$, meaning that the learner's estimate is equal to the sum of the fast and slow states.

Over time, the learner adjusts her estimate of the perturbation according to the errors she experiences. Error, denoted by $e^{(n)}$, is the difference between the observed outcome of the movement, and the target:

$$\begin{aligned} e^{(n)} &= h^{(n)} - g^{(n)} \\ &= r^{(n)} - \mathbf{c}^T \mathbf{x}^{(n)} + \varepsilon_u^{(n)} \end{aligned} \quad (3.4)$$

Note that this error could be further manipulated by the experimenter. In some cases, the experimenter can add additional noise to the observed movement to increase feedback uncertainty⁵⁶. One could also explicitly attempt to account for uncertainty in visual or proprioceptive transduction of error. We remark on these sources of noise further below.

A common experimental manipulation is the occurrence of an error-clamp trial²⁷. On these trials the learner is presented with an error that is independent of the movement she performed. Therefore, the error experienced by the learner can take different functional forms depending on the trial type according to:

$$e^{(n)} = \begin{cases} r^{(n)} - \mathbf{c}^T \mathbf{x}^{(n)} + \varepsilon_u^{(n)}, & \text{not an error-clamp trial} \\ e_c^{(n)}, & \text{error-clamp trial} \end{cases} \quad (3.5)$$

Here, the variable $e_c^{(n)}$ takes the value of the error imposed on trial n .

Two separate processes determine how the learner's estimate of the perturbation changes from one trial to the next: learning and forgetting. Together, learning and forgetting are captured by the state update equation.

$$\mathbf{x}^{(n+1)} = \mathbf{A}\mathbf{x}^{(n)} + \mathbf{b}e^{(n)} + \varepsilon_x^{(n)} \quad (3.6)$$

The forgetting process is controlled by the matrix \mathbf{A} which encodes the rate at which states decay in the absence of error due to the passage of time. If we assume that each state evolves independently, we

can represent \mathbf{A} as a diagonal matrix of the form $\mathbf{A} = \begin{bmatrix} a_s & 0 \\ 0 & a_f \end{bmatrix}$. Here a_s and a_f are retention factors

for the slow and fast states, respectively.

The learning process is controlled by the vector \mathbf{b} which encodes the learning rates of the fast and slow states. The parameter \mathbf{b} is a 2 x 1 vector of the form $\mathbf{b} = [b_s \ b_f]^T$, where b_s and b_f are

the error sensitivities of the slow and fast states, respectively. These error sensitivities determine the rate at which each state learns from error.

The entire process of updating the learner's estimate of the perturbation, like the process of generating a movement, is affected by noise, represented by ε_x . This noise source represents the combined effect of many sources of noise that accumulate in afferent pathways involved in learning from error. These include, but are not limited to, noise in the proprioceptive and visual transduction of error, noise in an error stimulus itself, noise in the synaptic mechanisms that contribute to learning from a given error, etc. We will refer to the collection of these processes as state update noise, and assume that it is distributed according to a multivariate normal distribution with mean $[0 \ 0]^T$ and variance-covariance matrix Q . In accordance with our assumption that the two states evolve independently, we require that their covariance be equal to zero, implying that Q is a 2 x 2 diagonal matrix. To simplify the model we assumed that the fast and slow state update variances were equal, yielding a variance-

covariance matrix of the form $Q = \begin{bmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_x^2 \end{bmatrix}$, where σ_x^2 represents the cumulative state update

variance described above. In a set of control analyses, we also considered a model of learning where the fast and slow states had different noise variances (see *Different noises in the fast and slow adaptive processes*).

Eq. (3.6) treats the errors experienced in error-clamp and non-error-clamp trials the same. That is, we assume that the learner does not differentiate from an error that was produced by her own behavior, and an error that was presented to her in an error-clamp trial. The validity of this assumption is currently under debate. Although blocks of error-clamp trials have been used extensively in the literature to assess decay properties of motor memory, two recent reports have found evidence that in some cases, error-clamp blocks appear to contain contextual cues that can affect the process of learning^{10,119}. In contrast, another report has found evidence that learning in error-clamp trials remains consistent with learning in non-error-clamp trials¹²⁰, or that differences in learning become evident only after long passages of time away from the task¹²¹. Given that this question remains unanswered, here we chose the simplest model wherein all errors were treated equally.

We do not directly observe the states of the fast and slow processes. Instead, we measure the movement of the subject on each trial. In many experiments, to normalize across different targets that

may be presented to the subject, it is useful to define the subject's movement relative to the target location, $y^{(n)}$.

$$\begin{aligned} y^{(n)} &= g^{(n)} - u^{(n)} \\ &= \mathbf{c}^T \mathbf{x}^{(n)} - \varepsilon_u^{(n)} \end{aligned} \quad (3.7)$$

Substitution of the motor action in Eq. (3.7) into our expression for error in Eq. (3.5) yields the following simplification:

$$e^{(n)} = \begin{cases} r^{(n)} - y^{(n)}, & \text{not an error-clamp trial} \\ e_c^{(n)}, & \text{error-clamp trial} \end{cases} \quad (3.8)$$

Finally, we account for set breaks in our generative model of learning, noting that set breaks result in significant forgetting of previously learned behavior⁹. Therefore, we imagined that set breaks could be modeled as additional decay that elapses after the conclusion of the trial preceding a set break:

$$\begin{aligned} \mathbf{x}^{(n+1)} &= D^{(n)} \left[\mathbf{A} \mathbf{x}^{(n)} + \mathbf{b} e^{(n)} + \varepsilon_x^{(n)} \right] \quad \varepsilon_x^{(n)} \sim N\left([0 \ 0]^T, Q\right) \\ \text{where } D^{(n)} &= \begin{cases} \mathbf{A}^d = \begin{bmatrix} a_s^d & 0 \\ 0 & a_f^d \end{bmatrix} & \text{trial } n \text{ is followed by a set break} \\ \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} & \text{trial } n \text{ is not followed by a set break} \end{cases} \end{aligned} \quad (3.9)$$

The parameter d in Eq. (3.9) is a decay factor that parametrizes elapsed time between trials in order to account for additional forgetting of the fast and slow states¹²² across a set break. A value of $d = 0$ means that a set break results in no further forgetting beyond that which accompanies an experimental inter-trial-interval (ITI). A positive value of d indicates that a set break results in more forgetting than an experimental ITI. Here we assumed that all set breaks are an equal length of time, though Eq. (3.9) could be easily modified to allow for set breaks of variable length by replacing d by the product of d with the ratio of the duration of a set break to the average ITI.

The form of Eq. (3.9) assumes that the rate of decay of the fast and slow states follows the retention properties that are observed trial-by-trial in the absence of set breaks. Our two-state model can now be represented as the following system of state-space equations that account for both error-clamp trials and set breaks:

$$\begin{aligned}
\mathbf{x}^{(n+1)} &= \mathbf{A}^{(n)} \mathbf{x}^{(n)} + \mathbf{b}^{(n)} e^{(n)} + \boldsymbol{\varepsilon}_x^{(n)} & \boldsymbol{\varepsilon}_x^{(n)} &\sim \mathcal{N}([0 \ 0]^T, \mathbf{Q}^{(n)}) \\
\mathbf{y}^{(n)} &= \mathbf{c}^T \mathbf{x}^{(n)} + \boldsymbol{\varepsilon}_u^{(n)} & \boldsymbol{\varepsilon}_u^{(n)} &\sim \mathcal{N}(0, \sigma_u^2) \\
\mathbf{A}^{(n)} &= \begin{cases} \mathbf{A} & \text{no set break} \\ \mathbf{A}^{d+1} & \text{set break} \end{cases} & \mathbf{Q}^{(n)} &= \begin{cases} \mathbf{Q} & \text{no set break} \\ \mathbf{A}^d \mathbf{Q} \mathbf{A}^{d^T} & \text{set break} \end{cases} \\
\mathbf{b}^{(n)} &= \begin{cases} \mathbf{b} & \text{no set break} \\ \mathbf{A}^d \mathbf{b} & \text{set break} \end{cases} & e^{(n)} &= \begin{cases} r^{(n)} - y^{(n)} & \text{not an error-clamp trial} \\ e_c^{(n)} & \text{error-clamp trial} \end{cases}
\end{aligned} \tag{3.10}$$

In Eq. (3.10), the ‘no set break’ condition indicates that trial n is not followed by a set break. The ‘set break’ condition indicates that trial n is followed by a set break. The ‘not an error-clamp trial’ condition indicates that trial n was not an error-clamp trial. The ‘error-clamp trial’ condition indicates that trial n was an error-clamp trial. Note that the sign of the motor noise in Eq. (3.10) was flipped; as this noise is Gaussian with mean zero, changing its sign describes an equivalent system.

3.2.2 Experimental procedure

We recruited $n=20$ healthy right-handed subjects (ages 17 – 59, 8 males) to perform a visuomotor adaptation study. All subjects signed a consent form approved by the Johns Hopkins University School of Medicine Institutional Review Board before participating in the experiment.

Subjects were seated in a chair and held the handle of a planar robotic manipulandum. The arm of the subject was obscured from view, and the position of the hand was represented by a white cursor projected onto the screen situated directly on top of the hand. The x and y positions of the manipulandum (i.e., the subject’s hand) were recorded at 200 Hz from optical encoders at a resolution of better than 0.1mm using custom C++ code. Subjects were instructed to move their hand from a starting circle through a target circle (radius = 1cm). The target circle was presented in one of 8 positions in the workspace, at a displacement of 10 cm. On some trials, the subject was provided no visual feedback of the cursor. We term these trials “no-feedback trials”. Apart from these no-feedback trials, subjects had continual visual feedback of their hand position during the outward reach. Our experiment employed a single perturbation condition. The perturbation was a 30° CCW rotation to the cursor position, about the starting position. The subject was awarded a point each time the cursor passed through the target within 190-290 ms following movement onset. Subjects were instructed to obtain as many points as possible during the experiment.

An epoch consisted of 8 trials, one to each of the 8 targets around the circle, chosen in a random sequence. The task began with a baseline period of 30 epochs where subjects reached without any perturbation (Fig. 3.1A). We interspersed 3 epochs of no-feedback trials within this baseline period, to familiarize subjects with this condition. The baseline period was followed by a block of 30 perturbation epochs. After the perturbation period, visual feedback was removed for 15 epochs. After this no-feedback period, feedback was reinstated and the perturbation was removed (washout trials) for 30 epochs.

Our analysis focused on the hand endpoint error, which was taken as the angular displacement of the hand from the target when the subject's hand displacement exceeded the 10 cm target displacement.

3.2.3 Epoch vs. trial-by-trial analysis of behavior

The two-state model described in Eq. (3.10) is readily applied to trial-by-trial data where only a single target is presented during the experiment. However, in our visuomotor rotation task, as in many experiments, the paradigm consisted of multiple targets. In paradigms with multiple targets, the learning process should be represented with different states for each target, that is, separate fast and slow states of learning for each target in the workspace. This expansion of the hidden state dimensionality is further complicated by generalization of learning across targets. That is, the error experienced when the subject moves towards one target generalizes to other targets in the workspace, resulting in differential amounts of learning across the workspace¹²³.

In our task, we did not directly probe this generalization function. Therefore, we attempted to minimize the effect of generalization on the learning process by averaging behavior across the 8 trials/targets visited in each epoch of the experiment. In this case, while the same target may not be visited from one trial to the next, the same set of targets is visited from one epoch to the next, reducing the effect of generalization on the recorded epoch-by-epoch behavior. Unless otherwise noted, all analysis of subject behavior in our visuomotor rotation task is based on application of the two-state model in Eq. (3.10) to epoch-by-epoch behavior.

However, we also considered two-state model fits to the raw trial-by-trial data. In order to apply our two-state model to the trial-by-trial data, we first had to decide how best to deal with generalization. We considered two extreme cases, (1) where learning generalized completely from one target to all other targets and (2) where there was no-generalization of learning to any other targets. We will refer to these trial-by-trial models as our full-generalization and no-generalization models,

respectively. In the full-generalization model, we assumed that the learner used a single fast and slow state to account for the perturbation for all targets as in Eq. (3.10). The fast and slow states were ignorant of the target location and fully generalized learning that occurred from one target to all other targets. Therefore, in the full-generalization model we applied Eq. (3.10) to the trial-by-trial data as if the same target had been visited on each trial.

In the no-generalization trial-by-trial model, learning from one target did not generalize to other targets. This trial-by-trial model had 16 states, one fast and one slow state for each of the 8 targets. All fast/slow states experienced forgetting on every trial according to a common fast/slow retention factor. However, if target k was visited on trial n , then only the fast and slow states associated with target k learned from the error experienced on that trial (all fast/slow states had a common fast/slow error sensitivity). Furthermore, all fast and slow states were subject to a non-zero state noise on each trial. In Appendix 4 we provide a complete description of this model and show how it can be extended to account for generalization that might extend beyond the width of a single target.

To summarize, the majority of our analysis of behavioral data averaged in epochs of 8 trials where all 8 targets are visited once in the epoch. In a set of control analyses (see *Trial-by-trial analysis of behavior*) we considered two trial-by-trial models that covered the extremes of generalization: one where learning generalized completely across targets and one where learning did not generalize at all across targets. For the simulations (see *Simulating realistic data*), we never performed averaging of the simulated behavior.

3.2.4 Simulating realistic behavior

To test our algorithm, we simulated realistic data using the two-state system described in Eq. (3.10). We simulated four paradigms (Fig. 3.3, top row) commonly encountered in the literature. These paradigms differed with respect to the inclusion or exclusion of error-clamp trials, set breaks, as well as the manner in which the perturbation was introduced (i.e., abruptly versus gradually).

All paradigms began with a baseline period in which the learner was simulated for many trials in the absence of any external perturbation to her movements. In Paradigm 1 (Fig. 3.3, Paradigm 1) this perturbation was followed by an abrupt introduction of a perturbation. We used a 30° perturbation to match our visuomotor rotation task. In Paradigm 2 (Fig. 3.3, Paradigm 2) we built upon Paradigm 1 by adding a prolonged error-clamp period that followed the perturbation, where error was completely eliminated. This type of intervention is a common way to isolate and measure retention of learned behavior. Paradigm 2 concluded with a washout period that allowed the simulated learner to return to

baseline behavior. Paradigm 2 closely resembled our visuomotor task, considering that the simulated error-clamps trials produced qualitatively similar behavior to that demonstrated by our subjects during no-feedback trials. In Paradigm 3, we simulated a learner in a gradual perturbation environment, followed again by an error-clamp period. In Paradigm 4 we built upon the other paradigms by adding set breaks. In this simulated experiment we followed a trial structure that is known to promote spontaneous recovery of behavior^{8,9,72}. After being exposed to a positive 30° perturbation, the sign of the perturbation is abruptly switched until the learner expresses approximately baseline behavior. Then a block of error-clamp trials is provided to test for spontaneous recovery (Fig. 3.3, Paradigm 4).

a_s	a_f	b_s	b_f	σ_x^2 (degrees ²)	σ_u^2 (degrees ²)	$x_s^{(1)}$ (degrees)	$x_f^{(1)}$ (degrees)	σ_1^2 (degrees ²)	d
0.985	0.556	0.097	0.213	1.694	1.037	0	0	0	8

Table 3.1. Two-state model parameters. This table provides the model parameters used for the simulated experiments in our primary analyses. We selected these parameters specifically to match the dynamics of learning observed in our experimental data (Fig. 3.1). To do this, we fit epoch-by-epoch single subject behavior with EM and LMSE. For the retention factors and error sensitivities, we computed the average parameter value across all 20 subjects, and both the EM and LMSE fits. We assumed both the initial slow and fast states were equal to 0, to represent a naïve learner. For the initial state variance, we used a value of 0, indicating that each simulated fast and slow state truly began at the value zero for all simulations. For the state and motor noise, we used the mean variances predicted by the EM algorithm, as LMSE does not separately measure these noise terms. Finally, note that the parameter d only applies for simulation of Paradigm 4 which included set breaks.

For each paradigm we simulated subject behavior 1000 times using fixed model parameter values. On each run we varied the seed for the random number generator, which resulted in different learning profiles due to motor and state noise. To simulate realistic data sets, we selected retention factors and error sensitivities that matched parameters estimated from our experimental data: we fit each subject’s reaching behavior using EM and LMSE (described below) and used the mean parameter values across subjects and algorithms for the simulated retention factors and error sensitivities. For the state and motor noise variances, we selected the mean values obtained using the EM algorithm, as LMSE does not provide an estimate of these two noise sources. The complete parameter set is reported in Table 3.1. Three additional parameters appear in Table 3.1 that were not discussed in our derivation of the two-state model. These parameters are related to the initial state of the learner (see Appendix 1). We modeled the initial fast and slow states of the learner as normally distributed random variables with

mean $x_s^{(1)}$, $x_f^{(1)}$, respectively, and common variance σ_1^2 . For our simulations, we considered a naïve learner who had an initial slow and fast state equal to zero.

3.2.5 Fitting EM and LMST to data

We fit our measured data collected in the visuomotor rotation experiment, and the simulated data, in an identical manner. Both the EM and LMSE algorithms were given the observed motor actions, y . Each algorithm then used its objective function to identify an estimate of the model parameters. This process is described in Appendix 1 for EM and Appendix 2 for LMSE.

To obtain the model parameters, both algorithms were treated identically: they were numerically constrained to search an identical parameter space using the function *fmincon* in MATLAB R2016a. Our constrained parameter space was defined by upper and lower bounds, as well as linear inequality constraints relating some of the parameters. The upper and lower bounds for each parameter are provided in Table 3.2. Linear inequality constraints were specified to enforce traditional two-state model dynamics according to:

$$\begin{aligned} a_s &\geq a_f + 0.001 \\ b_f &\geq b_s + 0.001 \end{aligned} \tag{3.11}$$

The first of these inequalities requires that the slow state be retained more strongly trial-by-trial than the fast state. The second of these inequalities requires that the fast state learn more rapidly from error than the slow state.

a_s	a_f	b_s	b_f	σ_x^2 (degrees ²)	σ_u^2 (degrees ²)	$x_s^{(1)}$ (degrees)	$x_f^{(1)}$ (degrees)	σ_1^2 (degrees ²)	d
[0,1.1]	[0,1.1]	[0,1]	[0,1]	[10 ⁻⁷ ,10]	[10 ⁻⁷ ,10]	[-30,30]	[-30,30]	[10 ⁻⁷ ,10]	[10 ⁻⁶ ,30]

Table 3.2. Upper and lower bounds for two-state model parameters. When fitting our behavioral data and simulated data, EM and LMSE searched the same bounded parameter space. Here we provide the upper and lower bounds used for each model fit.

EM is an iterative algorithm that attempts to increase the value of the likelihood function from one iteration to the next (see Appendix 1). For each EM fit, we performed 100 iterations of the algorithm. The EM algorithm is sensitive to the initial conditions used to initialize the first iteration (see

Appendix 1). For each simulation, we started the EM algorithm from 5 different initial guesses. For our experimental data, we used 10 different initial guesses for each subject. These initial guesses were randomly sampled from the constrained parameter space. We selected the parameter set with the greatest likelihood at the conclusion of the 100th iteration of the algorithm. Numerical implementation of the LMSE algorithm can also require different initial guesses for proper convergence of the *fmincon* algorithm. For LMSE, we seeded the *fmincon* search using 50 different starting parameter sets, to better ensure the identification of minimal squared error within the constrained parameter space. As for the EM algorithm, these initial starting parameter sets were also sampled randomly.

3.2.6 Measuring performance of the algorithms

We assessed how well EM and LMSE recovered the properties of the fast and slow states of learning in our simulated experiments, where the hidden states were explicitly known. After obtaining the parameter sets for EM and LMSE, we asked how well they predicted the time courses of the fast and slow states. To do this, we used each parameter set to simulate noise-free behavior. The noise-free version of our two-state model was obtained by removing the noise terms from Eq. (3.10):

$$\begin{aligned} \mathbf{x}^{(n+1)} &= \mathbf{A}^{(n)} \mathbf{x}^{(n)} + \mathbf{b}^{(n)} e^{(n)} \\ y^{(n)} &= \mathbf{c}^T \mathbf{x}^{(n)} \end{aligned} \quad (3.12)$$

This noise-free system is equivalent to the expected value of the hidden states and observed behavior at any point in time. We compared the noise-free time courses of the slow state, fast state, and overall behavior, predicted by EM and LMSE, to the actual time courses for each simulation. To determine how well the EM and LMSE time courses matched the actual time courses, we computed the root-mean-squared-error (RMSE) between the model fit and the actual data.

We also asked how well the EM and LMSE parameters sets matched the true parameter set. For this, we computed the absolute error between the fitted parameters and the underlying two-state model parameters used to simulate the data.

3.2.7 Sources of noise

In our model of learning, we considered two potential sources of noise, one in the generation of an action, and the other in learning from error. Our EM algorithm identifies a parameter set that maximizes a likelihood function that attributes the randomness in measured behavior to these two processes. In contrast, the likelihood function maximized by the LMSE parameter set attributes randomness in

measured behavior to the measurement of the behavior itself and assumes that the underlying learning system behaves deterministically (see Appendix 2). Therefore, the critical difference between EM and LMSE is the manner in which their likelihood functions account for variance in measured behavior.

To compare these likelihood models, we turned to our experimental data. We computed the corrected Akaike Information Criterion (AICc) for the likelihood models maximized by EM (Eq. A3.1.25) and LMSE (Eq. A3.2.4). AICc is a metric that can be used to compare the likelihoods associated with different models, discounted by the number of parameters contained by these models. The corrected AIC differs from the conventional AIC by further penalizing the number of parameters in the model. In this sense, it is a more conservative way to compare models that differ in parameter complexity. AICc is defined by the following equation:

$$AICc = 2k - 2\log_e \left(\{y\}_1^N | \theta \right) + \frac{2k(k+1)}{n-k-1}$$

Here k refers to the number of model parameters and n is the number of data points used to fit the model. The smaller the value of AICc, the greater evidence for the corresponding model. Our likelihood model for EM included a state noise and a motor noise, but omitted noise in the measurement of subject behavior. Our likelihood model for LMSE included noise in the measurement of subject behavior, but omitted state and motor noise in the underlying learning process. Therefore, because EM has two sources of noise and LMSE only one, the EM model had one greater parameter. We fit these models to the experimental data by searching for each model's maximum log-likelihood, and along with the number of model parameters, computed the AICc for each subject.

To compute the maximum log-likelihood of the LMSE model, we searched its incomplete (marginal) log-likelihood function directly (Eq. A3.2.4). We performed this search from 50 different initial points, using *fmincon* in MATLAB 2016a. We used the same search space and model constraints for EM and LMSE as described previously. Note that the likelihood model maximized by LMSE also neglects noise in the initial state of learning (see Appendix 2). We therefore excluded this initial state variance from both of our likelihood models.

3.2.8 Power analysis

One way to study the motor learning system is to compare how subjects perform in different experimental conditions, or at different points in time. For example, we may perform a savings experiment where we adapt subjects to a perturbation, wash out the adapted behavior, and then re-adapt subjects to the same perturbation³⁰. Typically, we find that subjects adapt to the perturbation

faster the second time. We can ask how this savings is expressed: did subjects experience an increase in retention, an increase in error sensitivity, or perhaps both? In other words, our analyses typically involve comparisons. In order to determine if an intervention resulted in a change in behavior, we can ask how the model parameters that describe learning changed as the result of the intervention. This statistical comparison is dependent on the variance in our estimates of two-state model parameters; as this variance increases, more subjects are required to obtain a statistically significant result.

Here we imagined that we performed an intervention that resulted in a change to the error sensitivity of a population, and a separate intervention that resulted in a change to the forgetting rate of the population. For the former, this is what is observed in savings paradigms^{7,30,124}. The latter has been observed in experiments that provide feedback in rewarding and punishing environments⁷¹. We performed power analyses to determine how well EM and LMSE could detect changes in these two-state model parameters at different effect sizes. Our power analysis considered two forms of experiments: within-subject experiments, and between-subject experiments. For our within-subject experiment, we imagined that a set of subjects performed Paradigm 2 at two different time points, A and B. For our between-subject experiment, we imagined that different sets of subjects performed Paradigm 2 in contrasting experimental conditions A and B, like a randomized control trial.

To generate data for the A condition, we simulated the behavior of 1000 subjects for Paradigm 2, where each subject's parameter vector was sampled from a multivariate normal distribution that we estimated from our experimental data (Table 3.1) by computing the mean and covariance matrix of the two-state model parameters estimated for our 20 individual subjects. To produce performance during Exposure B, we imagined that one of these parameters (i.e., a single entry in the parameter vector) had changed to a different value, and then re-simulated a new set of 1000 subjects. We fit each simulated behavior using EM and LMSE. From this subject pool we selected a certain number of subjects. For our within-subject experiment, we sampled the same subjects from Exposure A and Exposure B periods. For our between-subject experiment, we sampled subjects independently from the A and B distributions. We next identified the parameter values that EM and LMSE identified for the subjects in each exposure. For our within-subject experiment, we performed a paired t-test across the Exposure A and Exposure B parameter values to determine if either algorithm could be used to detect a statistically significant change in the parameter value, at a confidence level of 95%. For our between-subject experiment, we used a two-sample t-test. We repeated this analysis for a given number of subjects a total of 10,000 times, each time resampling subjects from our large 1000 subject pool. We tracked the percentage of times that EM and LMSE yielded a statistically significant difference in the parameter value across the

10,000 experiments. We used this percentage as a measure of how reliably each algorithm detected statistical differences for a given group size. We used these data to ask how many subjects would be required in order for EM and LMSE to detect a statistical difference for at least 85% of the simulated experiments.

The parameter values we selected for the Exposure B period were motivated by previous studies. For the slow state retention factor, we considered differences of -3% to +1.5% of the Exposure A retention factor, in agreement with the dynamic range seen across subjects adapting to visuomotor rotations with rewarding and punishing feedback⁷¹. We speculate that the difference in retention factor for the single state fits in this study is most reflective of the slow state of learning during the error-clamp period. For the fast state error sensitivity, we selected ± 15 , 25, and 50% to cover the large differences in error sensitivity observed for individuals adapting in punishing and rewarding environments⁷¹. This range in error sensitivity is also similar in magnitude to differences in error sensitivities of the slow and fast processes observed across the control and reporting groups in a study examining implicit and explicit components of learning in force field adaptation and visuomotor rotation learning⁷².

3.2.9 Control studies

We performed a set of control studies where we tested EM and LMSE identification accuracy in situations of greater noise, different dynamics for the fast and slow processes, and different assumptions about the structure of our two-state model of learning. For each of these control studies, we followed the same general approach: we simulated our two-state model and fit simulated behavior with EM and LMSE to determine how well each algorithm identified the slow and fast states of learning. For certain control analyses we reanalyzed our subject data. Each of our control studies is discussed in a separate section of Results. In the section *Modeling higher levels of noise*, we tested EM and LMSE on simulated data sets with much greater levels of state and motor noise. In the section *Different noises in the fast and slow adaptive processes*, we discuss evidence for the existence of two different variances for the update of the fast and slow states and perform a sensitivity analysis where we test EM and LMSE on simulated data sets with different levels of fast and slow state update variances. In the section, *Other sources of noise*, we consider a way to model internal noise in the sensory observation of an error. In the section, *Changing the dynamics of the fast and slow adaptive processes*, we tested EM and LMSE on data sets where the dynamics of the fast and slow properties differed from the average processes measured in the subject behavior investigated in our primary analysis. In the section *Changing the bounds on the parameter space*, we tested the identification accuracy of EM and LMSE if we further restrict the

parameter space searched by both algorithms. And finally, in the section *Trial-by-trial analysis of behavior*, we show that fitting our state-space model to the trial-by-trial subject behavior as opposed to 8-trial epochs of subject behavior, does not have any effect on our primary conclusions. For brevity, we relegated further details of each of these control studies to the appropriate section in Results.

3.3 Results

Our aim was to design an algorithm that could uncover hidden processes that contribute to learning from error. We considered two algorithms: LMSE and EM. To perform our comparison, we asked volunteers ($n=20$) to participate in an adaptation experiment where they reached to 8 targets (Fig. 3.1A). The task consisted of a baseline period followed by a 30° visuomotor rotation to the cursor. After learning to compensate for the rotation, we removed visual feedback for an extended set of trials, and then reinstated feedback in the absence of any visual perturbation.

The data, represented in epochs of 8 trials, are shown in Fig. 3.1B for 4 subjects. At first, each subject reached to the target accurately with some noise (top row, gray). Upon introduction of the perturbation, subjects rapidly learned to counter the imposed rotation, learning about 80% of the total rotation within 30 epochs. Upon removing visual feedback, the adapted behavior decayed gradually towards baseline, and then rapidly washed out during the last 30 epochs of the experiment when feedback was reinstated.

3.3.1 Fitting the two-state model to measured behavior

We assumed that experience of error engaged two learning processes that differed in their sensitivity to error, as well as retention properties (Eq. 3.10). We fit this two-state model to the measured data from each subject to determine the properties of the hypothesized fast and slow states. To fit our model, we used two different algorithms, LMSE and EM. Each algorithm was provided with the noisy single subject epoch-by-epoch data (Fig. 3.1B, gray traces, top row). From this behavior, the algorithms estimated model parameters. We used these estimates to predict time courses of behavior (blue and red lines, Fig. 3.1B, top row), and the underlying slow and fast states (Fig. 3.1B, bottom row). The EM and LMSE time courses corresponded to the expected value of the behavior and hidden states, under each algorithm's estimates of the model parameters (Eq. 3.12).




Figure 3.1. EM and LMSE algorithms uncover different hidden processes. A. Subjects ($n=20$) participated in a reach adaptation task. There were 8 targets in total, each chosen pseudo randomly and presented once in epochs of 8 trials. Following a no-perturbation baseline period, a 30° counter-clockwise rotation was applied to the cursor representing the subject's hand position. After 30 epochs of this perturbation, visual feedback was removed for 15 epochs. Finally, visual feedback was reinstated during a washout block of 30 epochs. B. Single subject behavior. We fit the epoch-by-epoch data (reach direction) of each subject with EM (top row, blue) and LMSE (top row, red). Both provide good fits to the measured data. Each algorithm estimated the fast and slow processes that produced the measured behavior (bottom row). For subject S9, these time courses agreed across algorithms. For S16, EM and LMSE time courses exhibited reasonable two-state behavior, but had differing learning dynamics. For S2 and S1, the EM and LMSE predictions diverged completely.

We found that the two algorithms provided similar fits to the subject data (Fig. 3.1B, top row). However, differences between the algorithms emerged at the level of the fast and slow states of learning. For some subjects (Fig. 3.1B, bottom row, subject S9) EM and LMSE agreed quite well. However, for other subjects the predictions made by LMSE and EM differed considerably (Fig. 3.1B, bottom row, subjects S2 and S1). In these cases, the LMSE predictions often appeared to be unreasonable. For example, in subject S2, LMSE predicted large initial biases in the hidden states (Figure 3.1B, subject S2, bottom). These large initial states were accompanied by a near zero error sensitivity in the slow state of learning. This near zero error sensitivity allowed LMSE to use the slow state of learning as a bias, accounting for the dynamics of subject S2's behavior with the fast state alone. LMSE predicted that 5 of the 20 subjects possessed such an error-insensitive slow state. In contrast, the EM algorithm did not yield any such aberrant estimates. In another example, for subject S1, LMSE estimated a slow state retention factor that was greater than 1 (Fig. 3.1B, bottom row, subject S1). This led to a monotonically increasing slow state of learning (solid, red) that never decayed. LMSE estimated that 8 of

20 subjects possessed an unstable slow state retention factor. Again, the EM algorithm did not yield any such results.

Differences between EM and LMSE were further highlighted when the fits were averaged across subjects. That is, while the average fits to the behavior were nearly identical for EM and LMSE (Fig. 3.2A, top figure), the fast and slow state predictions differed considerably between the two algorithms. LMSE predicted larger contributions from the fast state than EM, and smaller contributions from the slow state. Relative to LMSE, EM predicted a smaller slow state retention factor (paired t-test, $t(19)=3.6$, $p<0.01$), a larger slow state error sensitivity (paired t-test, $t(19)=3.4$, $p<0.01$), a smaller fast state error sensitivity (paired t-test, $t(19)=2.6$, $p<0.05$), and a smaller fast state retention factor (paired t-test, $t(19)=2.7$, $p<0.05$). That is, the two algorithms made significantly different predictions regarding parameters of the system.

The different EM and LMSE model parameters led to contrasting levels of variance in the hidden state time courses; the LMSE time courses had considerably higher variability across subjects (compare sizes of error bars, Fig. 3.2A, bottom two figures). This elevated variability was largely driven by the outlying, likely errant, single subject fits we previously noted. The outlying fits also contributed to greater variance in the model parameters; the across subject standard deviation of the LMSE estimates exceeded that of EM for 5 of the 6 model parameters (all except fast state retention, Table 3.3).

These differences between the two algorithms produced a dilemma. Depending on the choice of algorithm, we obtained different descriptions of the fast and slow states of learning. Which estimate was closer to the truth?

To answer this question, we computed AICc to determine the evidence for the likelihood models maximized by EM and LMSE. The critical difference between EM and LMSE is that the likelihood function maximized by EM contains noise in both the state update process and generation of a movement, whereas the LMSE likelihood function considers noise only in the measurement of behavior, not in the underlying learning or moving processes (see Appendix 2).

<i>Paradigm and Algorithm</i>	<i>Two-state model parameters</i>						
	a_s	a_f	b_s	b_f	$x_s^{(1)}$	$x_f^{(1)}$	d
Paradigm 1 EM	0.0175	0.1878	0.0355	0.0669	1.5149	1.5345	-
Paradigm 1 LMSE	0.0330	0.2679	0.0472	0.0780	5.9684	6.5494	-

Paradigm 2 EM	0.0094	0.1667	0.0253	0.0549	1.5782	1.6078	-
Paradigm 2 LMSE	0.0225	0.2513	0.0409	0.0715	6.6765	7.4575	-
Paradigm 3 EM	0.0130	0.1966	0.0393	0.1042	1.5207	1.5892	-
Paradigm 3 LMSE	0.0368	0.2965	0.0601	0.2902	12.1512	12.6409	-
Paradigm 4 EM	0.0099	0.1057	0.0241	0.0300	1.3639	1.3469	6.2715
Paradigm 4 LMSE	0.0132	0.1521	0.0297	0.0354	3.8563	4.4528	10.0969
Human subjects	0.0127	0.3559	0.0616	0.1282	2.3391	2.7227	-
Human subjects	0.0339	0.3236	0.0942	0.2212	2.3771	4.5877	-

Table 3.3. Parameter standard deviation for EM and LMSE. Here we report the standard deviation of the two-state model parameter distributions shown in Fig. 3.5. For Paradigm 4, we also provide the standard deviation for the set break decay factor, though this is not shown in Fig. 3.5. Note that for all parameters and paradigms, the standard deviation of the LMSE distribution exceeded that of EM in simulation. On the bottom two rows we provide the standard deviation for each two-state model parameter across the 20 subjects that participated in our visuomotor rotation task.

We computed the differences in AICc across likelihood models for each subject, resulting in a within subject comparison of EM vs. LMSE (Fig. 3.2C). We found that the AICc for the model including state and motor noise (EM) was lower (better) than that of the model excluding state and motor noise (LMSE): paired t-test across subjects, $t(19)=3.4$, $p<0.01$. This suggested that despite having an additional parameter, the learning process is better described by the stochastic system considered by EM.

In summary, when we fit experimental data with each algorithm, we found that EM's estimates of the hidden states diverged from LMSE. Unlike LMSE, EM did not exhibit any aberrant fast and slow state predictions. EM also yielded states that were more consistent across subjects; they had smaller variances in the model parameters, leading to lower variance in the hidden state time courses across subjects. LMSE uses a likelihood function that is blind to state and motor noise. In contrast, EM's likelihood function takes these noise sources into account. The likelihood model maximized by EM possessed a lower AICc than LMSE. This implied that a stochastic model of learning and moving was a better descriptor of the experimental data.

3.3.2 Fitting the two-state model to simulated data

To better appreciate how well each algorithm could recover the hidden states where learning and moving are noisy processes, we performed simulations (Eq. 3.10) to produce realistic data sets using parameters obtained from our experiment. Unlike the behavior collected in the lab, the hidden states

were known in these simulated data sets, providing the opportunity to assess how well EM and LMSE could uncover the hidden states. Four paradigms were considered: abrupt perturbations that included or excluded error-clamp trials (Fig. 3.3A, Paradigms 1 and 2), a paradigm with a gradual perturbation (Fig. 3.3A, Paradigm 3), and a paradigm that demonstrated spontaneous recovery and decay of motor memory due to set breaks (Fig. 3.3A, Paradigm 4). These simulated paradigms included only a single target, and the resulting data were analyzed trial-by-trial.

Figure 3.3B shows the noise-free time courses of the simulated behavior (generated via Eq. 3.12), and the corresponding fast and slow states. Presence of error produced rapid changes in the fast state, but more gradual changes in the slow state when the perturbation was introduced abruptly (Paradigm 1). Complete loss of the fast state occurred after set breaks (Paradigm 4). Gradual introduction of the perturbation slowed the onset of the fast state of learning (Paradigm 3). Error-clamp trials produced rapid decay of the fast state, but produced only small changes in the slow state (Paradigm 2). Occasionally the two states had opposite signs (washout trials, Paradigm 2), which led to spontaneous recovery in error-clamp trials (Paradigm 4).




Figure 3.2. Comparison of parameter values uncovered by EM and LMSE as fitted to experimental data. A. Population behavior, represented by the average time course across all 20 subjects. At top, the average behavior (black) is shown overlaid with the average EM (blue) and average LMSE (red) fits. EM and LMSE had very similar fits to the behavior. However, the algorithms' predictions regarding the slow and fast states diverged. Error bars indicate ± 1 SEM. B. Model parameters. Bars indicate the mean value across the subjects. Error bars indicate ± 1 SEM. C. We compared the corrected AIC of two competing likelihood models: one with state and motor noise, and one without state and motor noise. AICc was lower for a model with state and motor noise for 14 of the 20 subjects (black lines) and larger for 6 (gray lines). A paired t-test across subjects indicated that a model with motor and state noise possessed a lower AICc (that is, a better fit) than a model without these noise sources.

Figures 3.3C-E show typical time courses of the simulated behavior and hidden states in the presence of noise (Eq. 3.10). Similar to our human data, we found that both EM and LMSE provided excellent fits to the simulated data (Fig. 3.3C). However, EM appeared more robust in uncovering the fast and slow states. In some conditions (Figs. 3.3D and 3.3E, Paradigms 1-3), LMSE produced estimates of the hidden states that diverged considerably from the truth. To compare each algorithm's accuracy in recovering the hidden states of learning, for each paradigm (1000 simulations for each of the four paradigms) we computed the root-mean-squared-error (RMSE) between three pairs of values: the true values of the simulated behavior, fast state, and slow state [y , x_f , x_s], and the predicted time courses obtained from EM and LMSE. We observed that across paradigms, LMSE fit the simulated behavior better than EM (Fig. 3.4, top row, parameter y ; paired t-test, $t(999) > 13.9$, $p < 10^{-5}$). This is expected, as the objective function of LMSE minimizes the RMSE between the observed behavior and model predicted behavior (see Appendix 2). With that said, LMSE offered only a modestly improved fit over EM (approximately 10% over EM, Fig. 3.4, bottom row, parameter y ; one-sample t-test, $t(999) > 31.5$, $p < 10^{-5}$).

While LMSE more closely tracked the observed data, EM produced far more accurate estimates of the fast and slow states (Fig. 3.4, x_f , x_s). In other words, LMSE was more prone to predicting errant fast and slow state time courses (Figs. 3.3D and 3.3E). LMSE performed particularly poorly in Paradigm 3, producing fast and slow state estimates that had errors exceeding those of EM by >275% (Fig. 3.4, Paradigm 3, bottom row; one-sample t-test, $t(999) > 25.0$, $p < 10^{-5}$). LMSE also performed poorly in Paradigm 1, where it produced state estimates that had errors exceeding those of EM by approximately 75-125% (Fig. 3.4, Paradigm 1, bottom row; one-sample t-test, $t(999) > 14.1$, $p < 10^{-5}$). Including error-clamp trials in Paradigm 2 improved LMSE performance marginally (Fig. 3.4, Paradigm 2). Finally, LMSE (and EM) performed best in Paradigm 4, the paradigm that included both error-clamp trials and set breaks (Fig. 3.4, Paradigm 4). In this case, EM was better than LMSE for the fast state (Fig. 3.4, Paradigm 4, bottom row; one-sample t-test, $t(999) = 11.8$, $p < 10^{-5}$), with no difference in estimation of the slow state (Fig. 3.4, Paradigm 4, bottom row; one-sample t-test, $t(999) = 0.9$, $p = 0.39$).

Overall, we found that LMSE performed substantially worse than EM in uncovering the hidden states. Why did LMSE fit the observed data well but was unable to robustly uncover the hidden states? We observed three modes of failure by LMSE.

In the first mode, LMSE identified retention factors for the slow state that were greater than 1 (Fig. 3.3D, Paradigm 1). Such retention factors resulted in unstable behavior of the slow state. Relative to EM, LMSE was significantly more prone to identifying unstable slow state retention factors. For

Paradigms 1 and 3, LMSE identified unstable slow state retention factors in 44.8% and 54% of the runs respectively, whereas for EM, this occurred in 15.2% and 6.4% of simulations. The inclusion of error-clamp trials in Paradigm 2 protected against this mode of failure, lowering the number of runs affected by unstable retention to 19.6% for LMSE and 1.3% for EM. Similarly, Paradigm 4 resulted in a very low frequency of error: 5.2% and 0.7% of simulations for LMSE and EM, respectively. This mode of failure was also demonstrated by LMSE fits to 8 of 20 subjects in our experiment (Fig. 3.1B, subject S1, bottom row).

In the second mode of failure, LMSE converged on a slow state of learning that was insensitive to error and possessed near complete retention, causing it to function as a behavioral bias (Fig. 3.3D, Paradigm 2). In this case, LMSE accounted for the dynamics of the observed data by relying solely on the fast state. These error-insensitive slow states of learning were also observed in LMSE fits to 5 of 20 subjects in our experiment (Fig. 3.1B, subject S2, bottom row).

In the final failure mode, LMSE predicted very large initial slow and fast states with similar magnitude and opposite sign (Figs. 3.3D and 3.3E, Paradigm 3). For Paradigms 1-4, LMSE identified hidden fast or slow states that differed from the true value (threshold of 10° or greater absolute error) in 8.1%, 12.2%, 25.8%, and 2.5% of simulations respectively. For EM, this never occurred. This mode of failure appeared more prevalent in simulation than in our behavioral data, with EM and LMSE both predicting only one subject with an initial state that exceeded 10° , or $1/3$ of the eventual perturbation.

Unsurprisingly, these three modes of failure were accompanied by larger error in the LMSE estimates of the model parameters. To quantify each algorithm's error in estimating model parameters we computed the absolute error for each parameter across all paradigms (Fig. 3.5). We found that for all parameters and in all paradigms, LMSE possessed a greater absolute error than EM (Fig. 3.5, bars labeled P1-4; paired t-test, $t(999) > 6.0$, $p < 10^{-5}$). EM also better estimated the set break decay parameter in Paradigm 4 than LMSE (errors of 4.99 ± 0.25 and 7.67 ± 0.43 for EM and LMSE respectively, paired t-test, $t(999) > 13.2$, $p < 10^{-5}$). The inclusion of error-clamp trials and set breaks improved the performance of both algorithms; with the exception of the initial fast and slow states, both algorithms had the lowest error in Paradigms 2 and 4.




Figure 3.3. Fitting EM and LMSE to simulated behavior. A. We simulated two-state models of learning in the context of four behavioral paradigms, with abrupt and gradual introduction of perturbations, and with error-clamp (EC) trials and set breaks. B. The expected value of the measured behavior (black), fast state (green), and slow state (blue) of learning. These time courses correspond to two-state model parameters extracted from our subject population (Table A3.1). C. For each of the four paradigms, behavior was simulated according to a two-state model of learning (Eq. 3.10). 1000 simulations were performed for each paradigm. The two-state model parameters were fixed for each simulation, solely the seed for the random number generator varied from simulation to simulation. Here we provide an example of a behavioral trajectory for each of the four paradigms. We fit each trajectory using EM and LMSE. D. The true slow state of learning along with EM and LMSE predictions. In the example simulations of Paradigms 1-3, LMSE failed to capture the slow state of learning. In Paradigm 4, both LMSE and EM closely tracked the true slow state. E. The true fast state of learning along with EM and

LMSE predictions. For Paradigms 1-3, LMSE predictions diverged from the true fast state trajectory. For Paradigm 4, both EM and LMSE tracked the true fast state time course.

Although error-clamp trials were also present in Paradigm 3, the gradual introduction of the perturbation hampered the response of the fast state of learning (Fig. 3.3, Paradigm 3) which likely impaired the ability of each algorithm (LMSE more than EM) to differentiate properties of the two learning processes. The opposite was true for Paradigm 4 where the two set breaks, opposite perturbations, and error-clamp trials resulted in several “excitations” of the fast state (Fig. 3.3B, Paradigm 4). These excitations significantly improved the ability to identify properties of the fast state, as evidenced by the markedly reduced error in fast state retention and error sensitivity (Fig. 3.5, Paradigm 4) for both algorithms. Error-clamp trials appeared to have a similarly dramatic effect on estimation of slow state retention and error sensitivity (Fig. 3.5, compare Paradigm 1 with Paradigm 2).




Figure 3.4. Performance of EM and LMSE algorithms. For each paradigm, 1000 simulations were performed, with fixed two-state model parameters, but a varying seed for the random number generator, altering noise. EM and LMSE were used to fit a two-state model to the simulated behavior. The EM and LMSE parameters were used to simulate noise-free time courses for behavior, slow state of learning, and fast state of learning. Next, we computed the RMSEs describing how well EM and LMSE recovered the hidden fast and slow states of learning, and the overall behavior. At top, the RMSE for the behavioral fit (y), slow state fit (x_s), and fast state fit (x_f) are shown. At bottom, a relative RMSE metric was computed to compare the RMSEs of EM and LMSE fits to the same simulated behavior; the RMSE for the LMSE algorithm was divided by that of EM, multiplied by 100, and then a factor of 100 was subtracted, to compute a percent increase of LMSE RMSE over that of EM. All bars in this figure represent the mean RMSE across 1000 simulations. Error bars represent 95% confidence intervals. LMSE improved upon the RMSE of the behavior fit by approximately 10% for all Paradigms (Rows 1-4, right). However, EM was superior in uncovering the slow and fast states. The largest difference was observed for Paradigm 3, followed by Paradigm 1, then Paradigm 2 and Paradigm 4.

In summary, in all paradigms EM was more accurate in uncovering the true parameters of the learning processes. Differences in hidden state recovery were driven by three failure modes in the LMSE algorithm, two of which were also prominent in LMSE fits to our human behavior data set. Inclusion of set breaks and error-clamp trials significantly improved the ability of both EM and LMSE to uncover the two-state model parameters.




Figure 3.5. Parameter estimation errors for EM and LMSE. For each simulation we computed the absolute value of the difference between each true parameter, and the parameter values predicted by EM and LMSE. All bars in this figure represent the mean absolute parameter error across all simulations. Error bars represent 95% confidence intervals. For all parameters and paradigms, EM had lower estimation error than LMSE.

3.3.3 A better tool for hypothesis testing

Sensorimotor tasks are occasionally designed to test the effectiveness of an intervention. Model fits provide a tool to ask whether the intervention significantly affected learning parameters such as error sensitivity or retention. Power analysis provides an estimate of how many subjects may be needed to detect a significant difference. The number of subjects needed to test a hypothesis depends on the noise properties of the data. In our experiment, EM parameter estimates had a lower variance than LMSE (evident by visual inspection of Fig. 3.2B, see Table 3.3 for numerical details). This lower variability in the parameter estimates has a practical implication: it should improve the ability to test hypotheses.

To explore this question, we considered a within-subject and a between-subject experimental design. A within-subject experiment is typical for the study of savings and anterograde interference, where the same subjects are exposed to similar or contrasting perturbations at two different time points. Between-subject designs are common when testing the effects of some intervention against a control group, where the subjects in both conditions differ. Here we imagined that these interventions may cause changes in learning rate, as is observed in savings paradigms⁷, and retention, as is observed in differing reward environments⁷¹.

We analyzed how many subjects would be required to achieve a particular level of confidence in the ability to detect differences in the learning parameter modified by the intervention. To do this, we created experiments by sampling subjects from a large pool. For our within-subject experiment, we sampled the same subjects in two different environments. For our between-subject experiment, we sampled subjects independently in the two different environments. We performed t-tests to see if there was a difference in the EM and LMSE learning parameters across the two simulated environments. We used a paired t-test to test for within-subject changes, and a two-sample t-test to test for between-subject changes. These simulated environments possessed the same trial structure (Paradigm 2) but differed in that subjects were simulated with a different error sensitivity or retention factor. We analyzed changes in each error sensitivity and retention factor separately, i.e., in each of our simulated experiments, only one parameter truly differed within the subject population across the two environments. We varied the number of subjects and determined the percentage of simulated experiments for which a statistically significant ($p < 0.05$) difference existed in the appropriate model parameter. Finally, we determined the number of subjects that would be required to detect a significant change in parameter value, for at least 85% repetitions of the experiment.

Figures 3.6A and 3.6B provide the results for our within-subject and between-subject analysis, respectively. Unsurprisingly, we found that a within-subject comparison required fewer subjects than a between-subject comparison, for both EM and LMSE. A within-subject test is more powerful, as it accounts for between-subject variability, increasing the power of the statistical comparison. For both the within-subject test and between-subject test we found that the number of subjects required to achieve an 85% detection rate generally decreased as the effect size (the magnitude of the parameter difference across the two environments) increased. This applied to both increases (solid lines) and decreases (dashed lines) in each parameter. For all effect sizes and parameters, EM required fewer subjects to reach the threshold detection level. Specifically, EM required approximately 20-95% and 30-95% fewer subjects than LMSE for the within-subject comparison and across-subject comparison.

In summary, we found that as compared to LMSE, EM reduced the number of subjects that would be needed to detect effects of an intervention that altered error sensitivity or retention.

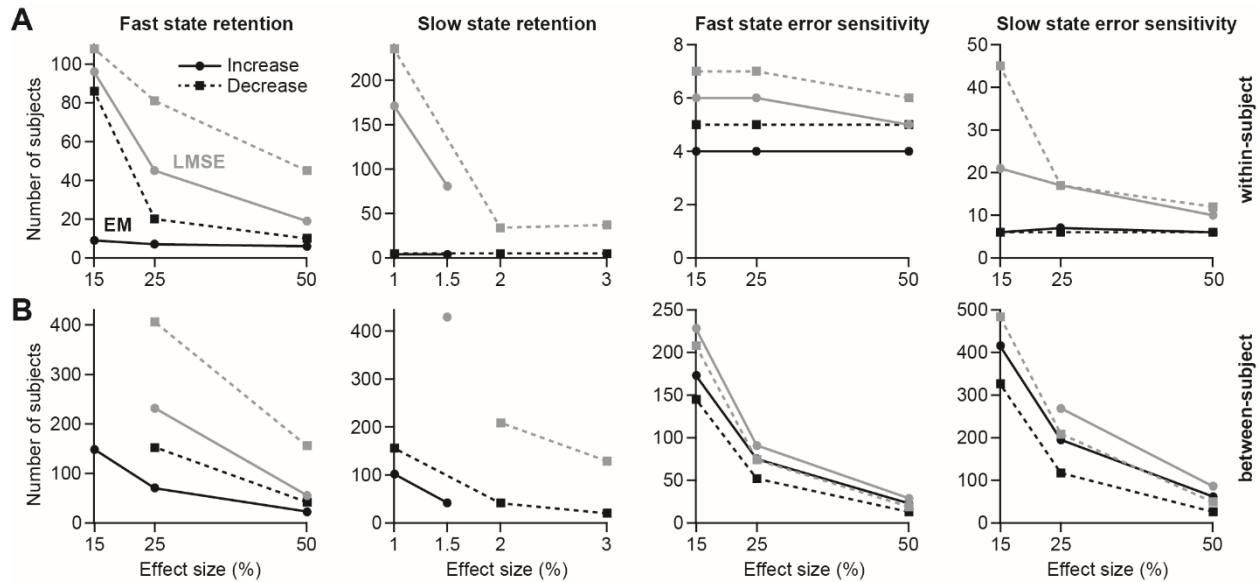


Figure 3.6. Power analysis for EM and LMSE. We simulated within-subject and between-subject experiments to determine the number of subjects that would be required to detect a change in learning parameters. We created a pool of 1000 simulated subjects by sampling two-state model parameters from a multivariate normal distribution. We created different distributions by scaling a single learning parameter for each of the 1000 subjects. We simulated behavior in Paradigm 2 and fit the data with EM and LMSE. We then sampled subjects to perform hypothesis testing. For within-subject tests, we sampled the same subjects from different parameter levels. For between-subject tests, we sampled subjects independently from different parameter levels. For each test, we performed a paired t-test (within-subject analysis) or two sample t-test (between-subject analysis) to determine if EM or LMSE detected a statistically significant difference in the learning parameter. We repeated this process for different random samples of our subject population (10,000 for each test). Finally, we determined the minimum number of subjects that would be required for each algorithm to detect a significant difference for 85% of our samples. Here we show the number of subjects required to reach an 85% detection rate for both EM (black) and LMSE (gray) as a function of the magnitude of the true parameter difference for each test (the effect size). We performed tests for both increases (solid lines with filled circles) and decreases (dashed lines with filled squares) in four two-state model parameters. The results for the within-subject analysis and between-subject analysis are shown in A and B, respectively. We only report results for which fewer than 500 subjects were required to reach the 85% detection rate. For LMSE, greater than 500 subjects were required for five different parameter-effect size pairs in the between-subject analysis (once for EM).

3.3.4 Modeling higher levels of noise

Subjects performing sensorimotor learning tasks exhibit different levels of noise in their motor behavior. Does the performance of either EM or LMSE worsen in the presence of higher levels of noise? To answer this question, we performed a sensitivity analysis by scaling the state and motor noise variances to different levels (0.5, 2, 4, 6, 8, and 10 times the values reported in Table 3.1). At each level, we

simulated the performance of 1000 subjects in Paradigm 2. We fit the simulated behavior with EM and LMSE, and computed the RMSEs for each algorithm's estimates of behavior and hidden states.

As expected, LMSE always provided a closer fit to the behavioral data (Fig. 3.7A). This result was expected because the objective function of LMSE minimizes the RMSE of the observed behavior. However, for all levels of noise, EM better isolated the hidden states of learning (Figs. 3.7B and 3.7C). The differences in the hidden state RMSE were approximately an order of magnitude larger than the difference in RMSE for the observed behavior. Therefore, irrespective of the noise level, it appeared that EM traded off small errors in fitting the observed behavior for larger improvements in uncovering the hidden fast and slow states. In summary, we expect that our conclusions about the relative performance of EM and LMSE would hold even at higher levels of state and motor noise.

3.3.5 Different noises in the fast and slow adaptive processes

We assumed that the fast and slow learning processes were affected by state noises with equal magnitude. However, faster processes that change more rapidly may also be accompanied by higher levels of noise. In fact, a relationship between learning rate and state variance would be expected from a Bayesian interpretation of learning⁵⁵. In light of these considerations, we reanalyzed our subject behavior with a two-state model of learning with two separate variances for fast and slow states of learning. In line with our intuition, we found that the variance of the fast state (mean \pm SEM, 2.74 ± 0.58 deg²) was greater (paired t-test, $t(19)=4.39$, $p<0.001$) than that of the slow state (mean \pm SEM, 0.63 ± 0.34 deg²).

We next asked if a model with separate variances for the slow and fast states was more likely to explain the data than a model with only one variance. To ask this question we computed the AICc for each subject for these two models. A within-subject paired t-test, across models, did not reveal a significant difference in the AICc for a model with a single state noise, or separate state noises for the slow and fast state (paired t-test, $t(19)=0.39$, $p=0.703$). This result indicates that in our experimental data, a model with two separate variances for the fast and slow states was not justified.

Despite this, we performed a sensitivity analysis to determine how well EM and LMSE would identify the fast and slow states of learning if each process was affected by different levels of state noise. In this analysis, we fixed the cumulative level of slow and fast state variance (i.e., we fixed the sum of the slow and fast state variances) and changed the fraction of the total state noise variance attributed to either the fast and slow state. We analyzed levels of 25%, 37.5%, 50% (i.e., equal contribution from the fast and slow state), 62.5%, and 75%. At each level, we simulated the behavior of

1000 subjects performing Paradigm 2. We fit the simulated behavior with EM and LMSE, and computed the RMSEs for each algorithm's estimates of behavior and hidden states of learning, as in Fig. 3.4. Regardless of noise levels tested, EM had greater error in the fitting of the measured behavior (Fig. 3.7D), but it remained superior to LMSE in the identification of the slow and fast states (Figs. 3.7E and 3.7F). Therefore, we expect the superior performance of EM would generalize to systems with different levels of variance in the fast and slow states.




Figure 3.7. Sensitivity analysis for state and motor noise. A, B, C. We scaled the state and motor noise variances by 0.5, 2, 4, 6, 8, and 10 times the values reported in Table A3.1 (measured from our subject population). At each noise level we performed 1000 simulations of Paradigm 2. We fit the simulated reaching behavior with EM and LMSE, generated EM and LMSE estimates of the behavior, fast, and slow states of learning, and finally computed the RMSE between the true time courses and model fits. The RMSEs for the behavior, slow, and fast state are shown in A, B, and C, respectively. Solid lines indicate the mean RMSE across all 1000 simulations at each noise level. The shaded error bars indicate 95% confidence intervals around the mean. D, E, F. We performed another analysis where we allowed the slow and fast processes to have different variances. We fixed the overall level of state noise (sum of the fast and slow state variances) and performed a sensitivity analysis where we assigned different fractions

of the overall state noise differentially to the slow and fast states. We tested levels where the slow (or fast) state had 25%, 37.5%, 50%, 62.5%, and 75% of the overall variance. For each level, we simulated 1000 simulations of Paradigm 2. We fit the simulated behavior using EM and LMSE and computed RMSEs for the behavior, slow process, and fast processes, as in A, B, and C. The RMSEs for the behavior, slow, and fast state are shown in D, E, and F, respectively. Solid lines indicate the mean RMSE across all 1000 simulations at each noise level. The shaded error bars indicate 95% confidence intervals around the mean.

3.3.6 Other sources of noise

We assumed that noise from all processes that contribute to learning from error could be combined into a cumulative state noise by adding the variances of each noise source together. However, there exist sources of noise that violate this assumption. Consider, for example, noise involved in the learner's observation of error. To correctly incorporate this source of noise, we could modify Eq. (3.6) by adding an observation noise that adds to the true error. In this model, the observation noise would be multiplied by error sensitivity, and therefore the state update noise in Eq. (3.6) would also depend on error sensitivity. We considered this more complete model, by simulating Eq. (3.6) with an added observation noise term. We performed 1000 simulations of this model for Paradigm 2 and attempted to recover the variances of the three different noise sources using EM. Unfortunately, we found that our algorithm was unable to estimate the variance of observation noise. For approximately 80% of simulations, EM's estimate of the observation noise variance converged to either an upper or lower bound of the parameter space, far from the true value we used in simulation. That is, our algorithm had no power to estimate this variance separately from the independent state and motor noises.

This limitation was caused by the multiplication of observation noise by error sensitivity. Because of this multiplication, the variance contributed by the observation noise is multiplied by the square of error sensitivity, and was therefore exceedingly small relative to the independent state noise variance. Therefore, while a more accurate model of learning might include this observation noise, we currently cannot estimate its magnitude using our algorithm.

3.3.7 Changing the dynamics of the fast and slow adaptive processes

Motor learning in different effectors (e.g. eye, arm, etc.) and paradigms (e.g. force field adaptation, visuomotor adaptation, etc.) occurs at different rates. Would the accuracy of EM or LMSE estimates of the hidden processes vary with the underlying learning and forgetting rates of the subject population? To investigate this question, we performed sensitivity analyses for four parameters: slow retention, fast retention, slow error sensitivity, and fast error sensitivity. We varied each of these parameters in turn

across the ranges used for our power analysis. For each parameter level, we performed 1000 simulations of our two-state model (Eq. 3.10) for Paradigm 2. We fit the simulated behavior with EM and LMSE, and computed the RMSEs for each algorithm's estimates of behavior and hidden states of learning, as in Fig. 3.4. As expected, LMSE continued to better fit the observed data in all simulations (Fig. 3.8, row 1). However, in every case EM was more accurate in uncovering the fast and slow states (Fig. 3.8, rows 2 and 3). This difference in hidden state RMSE was an order of magnitude larger than that of the observed behavior RMSE. Therefore, we expect that our conclusions about EM and LMSE would generalize to environments where learning exhibited different dynamics than those explored within our primary analysis.

3.3.8 Changing the bounds on the parameter space

Both algorithms were constrained to search an identical parameter space (Table 3.2). Despite this, we found that LMSE frequently exhibited modes of failure in the identification of two-state model parameters. One of these modes of failure involved the identification of slow state retention factors whose magnitude exceeded 1. Retention factors in this range can lead to unstable behavior of the slow state of learning. We doubt that a biological system would exhibit this unstable behavior. Could the LMSE algorithm be rescued by modifying the parameter space to prevent the identification of these unstable retention factors? To answer this question, we reanalyzed our primary simulations for Paradigms 1-4 (Figs. 3.3, 3.4, and 3.5) by refitting the EM and LMSE algorithm in a parameter space whose upper bound for the slow and fast state retention factor was equal to 1. We fit each simulated behavior with EM and LMSE, and computed the RMSEs for each algorithm's estimates of behavior and hidden states of learning (Fig. 3.9).

We found that restricting the parameter space had very little effect on the performance of EM (compare RMSE for the EM algorithm in Figs. 3.4 and 3.9). This is to be expected, considering EM did not identify many solutions with unstable retention factors for any of the 4 paradigms. In contrast, the performance of LMSE improved in terms of identifying the fast and slow states, specifically for Paradigms 1, 2, and 3 (compare RMSE for the LMSE algorithm in Figs. 3.4 and 3.9). This was also expected, considering LMSE predicted many behaviors with a slow state retention factor that exceeded 1. However, despite the improvement in LMSE identification of the hidden processes in Paradigms 1-3, LMSE error exceeded that of the EM algorithm (Fig. 3.9; Paradigms 1, 2, and 3; x_s , x_f). Therefore, restrictions on the parameter space that eliminated unstable behavior of the slow state of learning improved but did not rescue the performance of LMSE relative to the EM algorithm.




Figure 3.8. Sensitivity analysis for the dynamics of the fast and slow states. We performed sensitivity analyses to determine how well EM and LMSE could isolate that fast and slow states of learning, for two-state model parameters that differed from those observed for our visuomotor rotation subject population (Table A3.1). We analyzed one parameter at a time, fixing the remaining two-state model parameters to the values reported in Table A3.1. For each analysis, we scaled the two-state model parameter to several different values, corresponding to the effect sizes used in our power analysis in Fig. 3.6. At each parameter level we performed 1000 simulations of Paradigm 2. We fit the simulated reaching behavior with EM and LMSE, generated EM and LMSE estimates of the behavior, fast, and slow states of learning, and finally computed the RMSE between the true time courses and model fits. The top, middle, and bottom rows, show the RMSE for the behavior, slow state, and fast state fits, in that order. The shaded error bars indicate 95% confidence intervals. The parameters investigated are as follows: the fast state retention factor (shown in A), the slow state retention factor (shown in B), the fast state error sensitivity (shown in C), and the slow state error sensitivity (shown in D). For each analysis, EM identified slow and fast states of learning with lower RMSE than LMSE. These results indicate that the relative difference between EM and LMSE performance would generalize to other dynamics of learning.

3.3.9 Trial-by-trial analysis of behavior

Prior to fitting our state-space model to our behavioral data, we averaged single subject behavior across 8-trial epochs, where each of the 8 targets in our rotation task was visited once. We chose an epoch-based timescale for the analysis of subject behavior in order to minimize the effects of generalization on trial-based learning (see *Epoch and trial-based analysis of subject behavior*). However, an epoch-based timescale of motor behavior exhibits different dynamics than a trial-biased timescale; more is learned and forgotten in an epoch of 8 trials, and the variance in trial-by-trial behavior differs from that of epoch-by-epoch behavior. Therefore, the retention factors, error sensitivities, and noise variances that describe single subject data for our epoch-based analysis will differ from those of a trial-based analysis.

We performed a set of control analyses to confirm that these differences in the trial-by-trial data would not lead to changes in the identification accuracy of EM or LMSE. We reanalyzed our experimental data on a trial-by-trial basis with two different models of generalization. Because we did not measure the generalization function of each subject, we considered two extreme cases, (1) where subjects fully generalized learning from one target to all other targets (full-generalization model) and (2) where subjects had no-generalization of learning from one target to other targets (no-generalization model). For the full-generalization model, we applied Eq. (3.10) to our trial-by-trial subject behavior as if the same target was visited on each trial. For the no-generalization model, we extended the dimensionality of our state to include a fast and slow state for each target (a total of 16 states, 2 hidden states for each of the 8 targets). For a given trial, only the fast and slow states for that target experienced error-based learning consistent with the case of no-generalization, but all states experienced trial-by-trial forgetting. The full details of this model are described in Appendix 4.

The trial-by-trial analysis yielded strikingly similar results to our epoch-by-epoch analysis. As before (Fig. 3.2A), EM and LMSE provided similar fits of the measured data for both the full and no-generalization models (Fig. 3.10A, top and bottom rows, respectively). As in the epoch-by-epoch predictions in Fig. 3.2A, for both trial-by-trial models, EM estimated larger contributions from the slow state and LMSE estimated larger contributions from the fast state (Fig. 3.10B). The differences in these trajectories were driven by different estimates of the two-state model parameters (not shown in figure). As with our epoch-by-epoch model parameters (Fig. 3.2B), in our full-generalization model, EM estimated a smaller slow state retention factor than LMSE (paired t-test, $t(19)=3.20$, $p=0.0047$) and a larger slow state error sensitivity (paired t-test, $t(19)=3.65$, $p=0.0017$). For the no-generalization model these trends were not statistically significant. As before, the variance of the LMSE parameter estimates exceeded those of EM for all retention factors and error sensitivities of the no-generalization model, and

all parameters except the fast state retention factor of the full-generalization model. These increased variances were caused by errant LMSE slow and fast state predictions that resembled the failure modes noted in Fig. 3.1B for the epoch-by-epoch analysis. For 8 and 7 of the 20 subjects (no-generalization and full-generalization, respectively), LMSE estimated unstable retention factors that exceeded 1. For 3 and 6 of the 20 subjects (no-generalization and full-generalization, respectively), LMSE estimates of slow state error sensitivity tended towards zero yielding slow states of learning that were largely insensitive to error.




Figure 3.9. Comparison of EM and LMSE in a restricted parameter space. In our primary analysis, we found a preference for LMSE to assign slow retention factors that exceed 1, which led to unstable behavior of the predicted slow process. We asked whether LMSE could be rescued by modifying the parameter search space to prevent the identification of these unstable retention factors. To answer this question, we reanalyzed our simulations for Paradigms 1-4 (Figs. 3.3, 3.4, and 3.5) by refitting the EM and LMSE algorithm in a parameter space whose upper bounds for the slow and fast state retention factors were equal to 1. We used the EM and LMSE parameters to simulate noise-free time courses for behavior, slow state of learning, and fast state of learning. Next, we computed the RMSEs describing how well EM and LMSE recovered the hidden fast and slow states of learning, and the overall behavior for the same simulations depicted in Fig. 3.4. At top, the RMSE for the behavioral fit (y), slow state fit (x_s), and fast state fit (x_f) are shown. At bottom, we computed the percentage difference between the RMSEs for EM and LMSE. Positive values indicate a larger RMSE for the LMSE algorithm. Error bars represent 95% confidence intervals. We found that restricting the upper bound on the slow and fast state retention factors improved the RMSE of the LMSE fits to the hidden states (compare Fig. 3.4 with Fig. 3.9), but did not completely rescue LMSE predictions.

We repeated our AICc analysis on the likelihood models fit by EM and LMSE. Recall that the LMSE model excludes state and motor noise, and therefore possesses fewer parameters than EM. We found the same result as before (Fig. 3.2C): the likelihood used by EM possessed a lower AICc than the likelihood model used by LMSE (Fig. 3.10C; paired t-test for full-generalization, $t(19)=2.83$, $p=0.011$;

paired t-test for no-generalization, $t(19)=3.81$, $p=0.001$). Therefore, trial-by-trial analyses of our data assuming either complete generalization or no-generalization of learning both suggested that a model that includes state and motor noise in the process of error-based learning was more likely to explain measured behavior.

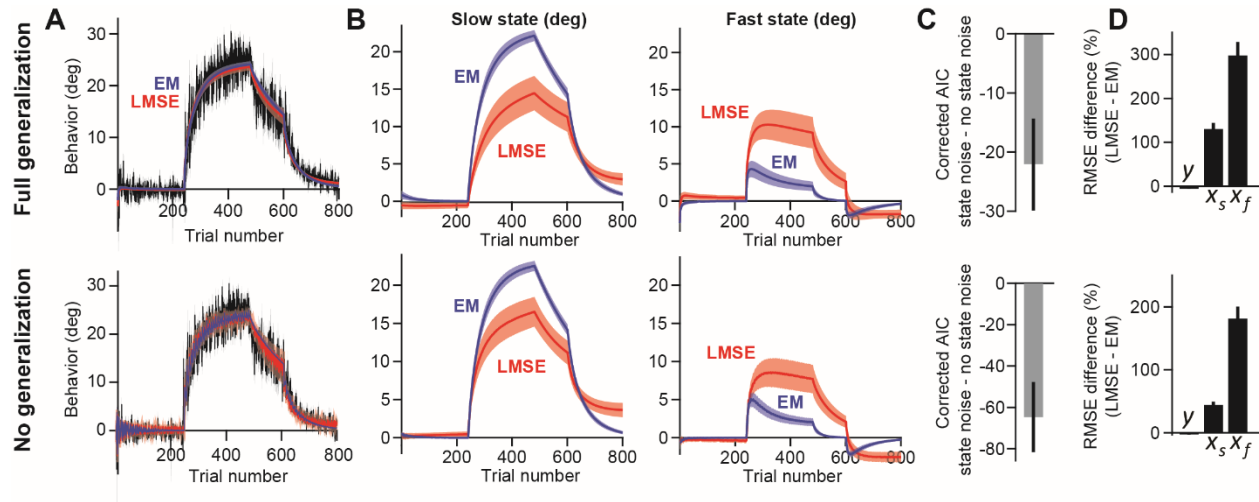


Figure 3.10. Comparison of EM and LMSE on a trial-by-trial analysis of the data. We collected the behavior of $n=20$ subjects in a visuomotor rotation task. We fit the two-state model to the trial-by-trial data recorded for individual subjects using EM and LMSE. We considered two trial-by-trial models that differed in terms of generalization. The full-generalization model (top row) consisted of a single fast and slow state whose learning completely generalizes across targets. The no-generalization model (bottom row) consisted of separate fast and slow states for each of the 8 targets, whose learning did not generalize across targets. A. Population behavior. We computed the average trial-to-trial behavior of the subject population. The average behavior (black) is shown overlaid with the average EM (blue) and average LMSE (red) fits. EM and LMSE had very similar fits to the behavior. B. Predicted fast and slow states. For both the full and no-generalization models, EM estimated larger contributions from the slow state of learning and smaller contributions from the fast state of learning. Error bars indicate ± 1 SEM. Here the average time courses across the 8 fast states and 8 slow states are shown for the no generalization model. C. We compared the corrected AIC of two competing likelihood models: one with state and motor noise and one without these noise sources. AICc was lower (better) for a model with state and motor noise. Here we provide the mean difference in AICc for both models (state and motor noise likelihood – no state and motor noise likelihood). D. We used the trial-by-trial parameters to perform a set of control simulations. We simulated Paradigm 2 a total of 1000 times, and fit each simulated data set with EM and LMSE. While LMSE fit the observed reaching behavior more closely (y), EM vastly outperformed LMSE in the identification of the hidden slow and fast states (x_s and x_f).

Finally, although in our primary analysis we did not perform any averaging of our simulated data sets, these simulations were indirectly affected by our epoch-based analysis of subject behavior because we used these parameter sets as a basis for the simulated data. For this reason, we used the trial-by-trial parameters (Fig. 3.10B) to perform a set of control simulations. We replicated our epoch-by-epoch

analysis by simulating Paradigm 2 (a total of 1000 times with our no-generalization and full-generalization state-space models). We found that while LMSE fit the observed reaching behavior more closely (Fig. 3.10D, y) improving on EM by approximately 5%, EM vastly outperformed LMSE in the identification of the hidden slow and fast states by approximately 50-300% (Fig. 3.10D, x_s and x_f).

In summary, in both the epoch-by-epoch and trial-by-trial data, EM and LMSE identified different slow and fast state trajectories. The likelihood model maximized by EM was more likely to explain both the epoch-by-epoch and trial-by-trial behavior. That is, in all cases the evidence pointed to a learning model that was stochastic in its adaptation to error. In simulation, EM was superior to LMSE in the identification of the slow and fast states when the parameter set was taken from trial-by-trial subject behavior, as well as from epoch-by-epoch subject behavior.

3.3.10 Savings is caused by an increase in error sensitivity of fast learning processes

One of the hallmarks of adaptation is savings: the observation that learning occurs more rapidly when similar perturbations were experienced in the past. It is well established that savings is caused by an increase in sensitivity to error^{7,37,49,125}, specifically of the fast adaptive process³⁷. Here, we aimed to determine how savings is expressed in two different modalities of learning: eye movements and arm movements. Unlike earlier studies, here we used GEM to obtain the maximum likelihood estimate for the learning parameters of a two-state model. In Experiment 1, participants ($n=20$) held the handle of a robotic arm and made reaching movements to 8 different targets. Participants were exposed to a visuomotor rotation that altered the position of the on-screen cursor by 30°. After adapting to the perturbation (Exposure 1), a no feedback block ensued, followed by a washout period. After a short set break, this perturbation scheduled was revisited once again (Exposure 2). We mirrored this paradigm for a separate set of participants who participated in a saccade adaptation paradigm ($n=16$). Both types of adaptation showed evidence of savings: a faster rate of learning in Exposure 2, relative to Exposure 1 (Fig. 11A, reaching at left, saccades at right; compare red with black).

To determine which of the adaptive processes contributed to this faster rate of re-learning, we fit our two-state model to each participant's behavior using our GEM algorithm. The algorithm revealed that memory retention of the slow and fast adaptive processes was no different during Exposures 1 and 2, for both reaching (Fig. 3.11B, reach) and saccades (Fig. 3.11B, eye). On the contrary, the cause of savings was a change in the fast state error sensitivity alone (Fig. 3.11B; paired t-test; reach, $p<0.05$; eye, $p<0.01$). Therefore, savings for both modalities of movement could be ascribed to sole changes in fast state error sensitivity.




Figure 3.11. Savings is caused by an increase in error sensitivity of the fast-adaptive process. A. Participants were tested in an arm movement ($n=20$, left) experiment and an eye movement ($n=16$, right) experiment. For reaching, participants reached to 8 different targets. Two exposures to the same perturbation (epochs 6-35) were followed by a no feedback period (epochs 36-50) and then a washout period (epochs 51-75). Learning appeared faster during Exposure 2 (red) than during Exposure 1 (black). A similar paradigm was used for saccade adaptation (right). Participants adapted to a 5-degree gain down perturbation in a memory-guided saccade adaptation task. Two exposures to the same perturbation (epochs 11-50) were followed by a no feedback period (epochs 51-70) and a washout period (epochs 71-100). As in reaching, learning during the second exposure (red) was faster than during the first exposure (black). B. We quantified these changes in learning using a two-state model of adaptation. We fit each the two-state model to each participant's behavior. At left, we show the retention factors for the slow and fast states for both reaching and saccades. We observed no differences in retention between Exposures 1 and 2. At right, we show the error sensitivity for the slow and fast states for both reaching and saccades. In both cases, fast state error sensitivity was increased during Exposure 2, but slow state error sensitivity remained the same. Statistics: no indicates means $p>0.05$, * means $p<0.05$, and ** means $p<0.01$.

Interestingly, for reach adaptation, the no feedback period (epoch 35-50) appeared to show a faster rate of decay during Exposure 2. If memory retention did not differ in either exposure, what was the cause of this change in forgetting? Interestingly, the model accurately predicted this trend, not because of any difference in decay rates, but because at the start of the no feedback period, the fast state of adaptation made up a larger fraction of the adapted behavior in Exposure 2 than in Exposure 1

due to its increase in error sensitivity. Because the fast state also decays more rapidly, increases in the relative contributions of the fast state will also increase the apparent rate of decay. Therefore, it is important to note that faster rates of forgetting do not necessarily mean that the underlying states of learning have changed their retention capabilities over time.

3.3.11 Spontaneous recovery after the passage of time is caused by fast state decay

A primary hallmark of behavior that is most consistent with two-state models of adaptation is spontaneous recovery⁸. Spontaneous recovery refers to the reemergence of an earlier behavior despite the absence of any cue. It is typically demonstrated in paradigms where participants are adapted to perturbation *A*, then exposed briefly to an opposite perturbation *B*, and then put in an error-free environment^{8,9,36}. Though our experiment was not designed to promote spontaneous recovery, we nevertheless observed of recovery of behavior after the set break separating the two exposures (Fig. 3.12A, compare behavior before and after the gray set break). To understand the cause of this recovery, we used GEM to estimate the underlying fast and slow states of adaptation throughout each period of the experiment (Fig. 3.12B, fast shown in green, slow shown in red). This decomposition indicated that spontaneous recovery of the memory of Exposure 1 occurred due to decay of the fast state. During the washout period, the slow state of adaptation remained positively biased, and the fast state became negatively biased. The set break caused an abrupt loss of the fast state, but spared the slow state (Fig. 3.12C, slow state comparison at left, $p > 0.05$; fast state comparison after set break at right ($p < 0.05$). This sudden loss of the fast state caused a small apparent recovery of the previous behavior. Thus, spontaneous recovery that emerges after a break in time, versus a sequence of error-free movements appears to have the same root cause: decay of the fast-adaptive process.

It should be noted that we also observed a spontaneous recovery of behavior in our saccade adaptation paradigm, but at a very peculiar point. During both the first and second exposure to the perturbation, the set break during the error-free period on epoch 60 (Fig. 3.11A, right), resulted in a different flavor of recovery. First, behavior appeared to decay as expected after the set break, but then slowly grew during the error-free period. Such recovery is not predicted by the two-state model, because participants were never exposed to an opposing perturbation that could counterbalance the sign of the fast and slow processes. We offer two potential explanations: (1) if saccades made during the set break are interpreted by the brain as zero perturbation washout trials, then such a spontaneous recovery could be predicted by the two state model or (2) perhaps certain instances of memory decay are not truly decay, but a period during which the brain switches between two contexts, in this case,

saccades made using the Eyelink apparatus, and those that are not. In terms of this latter possibility, the reemergence may represent the gradual recall of a prior context in which the system was adapted.




Figure 3.12. Spontaneous recovery after a break in time is caused by loss of the fast state. A. Behavior during the reaching paradigm described in Fig. 3.11 is shown side-by-side for the first (left, black) and second (right, red) exposures to the perturbation. A single set break separated each period. Behavior was observed to spontaneously increase after the set break (see black arrow on epoch 76). B. The GEM predictions for the fast (green) and slow (red) states are shown. Clearly, the model predicts that spontaneous recovery occurs due to the abrupt loss of the negatively-biased fast state after the passage of time. C. We quantified this loss by calculating the predicted slow (left) and fast (right) states before and after the set break. Statistics: no indication means $p > 0.05$; * means $p < 0.05$.

3.3.12 The expression of savings after impairment of the basal ganglia

Others have posited that reward systems in the brain contribute to savings through a reinforcement-based learning mechanism⁵⁴. Evidence for this is perhaps strongest for the saccadic system, where projections from the basal ganglia to the superior colliculus and then to the inferior olive are known to cause complex spikes in the cerebellar cortex¹²⁶ and the presence of reward appears to alter the rate of motor learning¹²⁷, and in some cases retention of memory⁷¹, though the latter results are mixed within the human psychophysics literature.

If reward circuitry modifies the rate of adaptation, it may be the case that individuals with damage to dopaminergic structures, like the basal ganglia, may exhibit deficits in the ability to save. Some studies have addressed this possibility, specifically in the context of visuomotor rotation^{128–131}. Uniformly, these authors have suggested that Parkinson’s Disease (PD) patients have an impairment in their expression of savings. With that said, a revisiting of some of this literature, does suggest some hints of savings, for example from the first exposure to the fourth exposure of the perturbation¹²⁸. Regardless, if taken at face value, the impaired savings in PD in this context could be attributed to deficits in either implicit adaptation, or more cognitive explicit strategies^{62,63,65,132}.

Here we aimed to determine how PD patients expressed savings in a different form of adaptation that is thought to primarily invoke implicit adaptation mechanisms: force field adaptation. Patients (n=8) and healthy controls (n=10) reached to a single target in a two-dimensional workspace. After a baseline period, an abrupt perturbation was applied to the arm, on four separate exposures. The maximum deviation of the hand for the controls and patients is shown in Fig. 3.13A and 3.13B. To quantify the presence of savings, we fit a decaying exponential to these timecourses. Healthy controls showed clear evidence of savings (Fig. 3.13C, left) with an increase in the rate of learning after the first exposure to the perturbation. In addition, healthy controls exhibited an increase in the total amount of adaptation with each subsequent exposure (Fig. 3.13D, left). The presence of savings was less clear in the PD group. There was no clear change in the rate of adaptation with subsequent exposure to the perturbation (Fig. 3.13C, right). In addition, there was no change in the total amount of adaptation (Fig. 3.13D, right). Therefore, these preliminary data suggested that savings may be impaired in PD even in force field adaptation.

There are considerable number of caveats that cloud this conclusion. First, the initial error experienced at the start of the perturbation (Fig. 3.13E) dropped significantly for the PD group after the first perturbation, but not for the control group. Considering the fast process is most likely to contribute to savings and is most active at the start of the perturbation, these differences in initial performance make it difficult to compare the PD and control groups. Likely, these changes were caused by an elevation in co-contraction in the limb. If true, the lack of savings may be caused by corruption of the motor plan by non-specific learning strategies (e.g., co-contraction), rather than a bonafide deficit in the patient’s response to error. Finally, we only collected performance during force field trials, but not on channel trials in an effort to reduce the total number of trials and prevent patient fatigue. In retrospect, this was likely a suboptimal choice. Without channel trials, we cannot disentangle non-specific learning

mechanisms like co-contraction, from perturbation specific error-based learning strategies. Therefore, more data is certainly needed to confirm any impairment in savings in PD.




Figure 3.13. Damage to the basal ganglia may prevent the expression of savings. A. Healthy age-matched controls (n=10) performed a standard force field adaptation task. They adapted to an abrupt force field on four separate exposures. Here we show the maximum displacement of the arm on each trial during each exposure. B. The same as in A but for a group of Parkinson’s Disease patients (n=8). We fit a decaying exponential model to behavior during each exposure to quantify the rate of learning (C), the total extent of adaptation (D), and the initial error at the start of each perturbation (E).

3.4 Discussion

State-space models were first applied to data in adaptation experiments following the observation that experience of a single error produced robust trial-by-trial changes in behavior^{26,133}. This provided the possibility to assay learning not only in typical scenarios where perturbations were sustained, but also in scenarios where the perturbations were random^{134,135}. The initial models assumed a single state; however, the observation of spontaneous recovery during saccade adaptation in monkeys⁶¹ suggested that experience of error engaged multiple learning processes. Smith et al.⁸ modified the state-space equations by proposing that the putative learning processes included a fast process that learned strongly from error but forgot rapidly, and a slow process that learned weakly from error but exhibited robust retention. Unfortunately, the task of identifying these processes was difficult because there was typically no direct way to measure them. Rather, their state had to be inferred from their collective influence on behavior. Here, we approached this estimation problem in the context of data measured in typical motor learning studies and designed an algorithm to uncover the hidden processes.

3.4.1 Design of a new algorithm

Previous attempts to estimate fast and slow processes of learning had predominantly relied on least-squares techniques^{45,70-72}. However, to our knowledge, robustness of the least-squares technique was not tested and compared with alternative algorithms. Here we used LMSE to fit a two-state model to behavioral data collected in a reach-adaptation task. We found that while LMSE fit the observed behavior well, for a subset of subjects it appeared to misidentify properties of the hidden processes; for these subjects LMSE estimated parameters that produced physically unrealistic trajectories of the fast and slow states. We hypothesized that LMSE occasionally yielded aberrant results because it incorrectly attributed any noise in the measured behavior to the measurement itself rather than the underlying learning process. In other words, LMSE was ignorant of any randomness in the state update process, which is affected by both state noise and motor noise (due to the process of error-based learning). To rectify this problem with LMSE, we developed a new algorithm based on EM.

Unlike LMSE, the EM algorithm is compatible with systems where both learning from error and generating a movement have independent noise sources. We developed a two-state model that represented the random processes involved in learning from error and production of movements. For such a system, our EM algorithm used a different, and in some sense, more complete, likelihood model than LMSE.

EM has received limited application in the sensorimotor literature⁷⁴. We speculate that this is because of two reasons. (1) Previous applications of the algorithm were restricted to experiments that could be described by time-invariant state-space models. That is, in previous descriptions of EM, one could not use the algorithm with modern sensorimotor experiments that include behavioral probes such as error-clamp trials²⁷ and set breaks, the second of which causes time-dependent changes to the state-space equations. Rather than using a closed form solution⁷⁴, here we used a numerical approach, allowing us to fit models to data that included error-clamp trials and set breaks. (2) Previous implementations could not specify bounds on the model parameters, and constrain relationships between model parameters. Here, we solved this problem by implementing a generalized EM algorithm that maximized the expected complete log-likelihood function numerically within a constrained parameter space.

3.4.2 Evaluating the new algorithm

We performed a visuomotor rotation experiment and fit the measured data with EM and found that parameters estimated by the algorithm differed from those of LMSE. EM appeared to eliminate the aberrant single subject fits observed for the LMSE algorithm. Additionally, EM parameters had lower variability across subjects, leading to reduced variance in the corresponding fast and slow state time courses. We computed the AICc for models that included (EM) or ignored (LMSE) noise in the process of motor learning. We found that the model with motor and state noise was more likely to explain the measured data in our experiment, than one where noise was attributed externally to the measurement of subject behavior. Therefore, we were able to make two conclusions: 1) experimental data suggested that equations that include noise in the learning process and moving process are a better descriptor of behavioral data than those that omit these noise sources; 2) EM, but not LMSE, is the appropriate algorithm to uncover parameters of a stochastic learning system.

To determine how well EM would perform over LMSE on a data set with state and motor noise, we performed simulations where the true trajectories of the hidden states were known. We tested EM and LMSE on simulated data generated in four learning paradigms, across a variety of noise conditions and parameter values. While in all cases EM performed slightly worse than LMSE in fitting the observed data, it consistently outperformed LMSE in the identification of the hidden states. Specifically, EM predicted hidden state time courses that were more closely matched to the true states, and identified model parameters that were more tightly distributed about the true values. For these reasons, we

expect the EM toolbox to provide a more robust method of fitting state-space models to single subject behavioral data.

In comparison to LMSE, EM identified learning parameters that had significantly reduced variance. For this reason, in simulated power analyses of typical experiments, we found that use of EM significantly reduced the number of subjects needed to statistically make within-subject and between-subject comparisons. Therefore, it appears that the new algorithm might allow for a more robust method of hypothesis testing.

3.4.3 Limitations of the algorithm

Our model assumed that two sources of noise, a state noise and a motor noise, affected the processes of learning and movement production, respectively. Other authors have also considered state noise in models of learning^{74,76}. However, at present we do not have a complete understanding of the properties of such noise. With that said, here we found evidence that a model with state noise was more likely to explain human behavior during a visuomotor rotation task. To show this, we repeated our analysis in Fig. 3.2C, for the comparison of a model with both state and motor noise and a separate model with only motor noise. For our subject population, the corrected AIC for the model with both noise sources was lower (better) than that with only a motor noise source (paired t-test, $t(19)=2.549$, $p<0.05$). Furthermore, the differentiation between state and motor noise is more harmonious with a Bayesian interpretation of motor learning^{55,56}. Additionally, the existence of state noise is also consistent with autocorrelations that arise between successive movements that can accumulate due to variability in planning a movement that is independent of variability in performing the movement itself¹³⁶.

The manner in which we described state and motor noise was not entirely accurate. For example, we assumed that the variance of motor noise was signal independent. While this assumption seems reasonable for visuomotor rotation learning where each movement has the same amplitude, for motor effectors like the eye, a better model of learning might account for scaling of noise with the amplitude of movement. Accounting for this signal dependency would require fundamental modifications to the E- and M-steps of our EM toolbox.

A useful modification to our model would be the inclusion of time-varying error sensitivity⁷. Such a modification could be incorporated by adding an additional parameter that determines the rate at which error sensitivity changes over time. Along these lines, the processes of learning and retention may possess nonlinearities not accounted for by our model. These modifications to the model would require derivation of a different expected complete log-likelihood function and an extended Kalman

filter (for nonlinear systems) and therefore fundamental modifications to our EM toolbox would be required to fit such nonlinear behavior.

Another extension could be the inclusion of more than two-states of learning. Preliminary evidence suggests that slow learning states can be subdivided into two component processes (yielding a total of 3 hidden states) with differing levels of susceptibility to temporal decay^{73,137}. Such a model of learning would be directly compatible with our EM approach with the addition of a retention factor, error sensitivity, and initial state.

Different model fits could be obtained by changing the restrictions on the parameter space we searched for both algorithms during the fitting process. In some cases, the experimenter can fix the initial states to zero to improve fits. However, the goal of our work was to identify an algorithm that could perform robustly for any perturbation sequence, independent of the subject's initial states and independent of the modeler's knowledge of the subject's initial states. For this reason, we did not attempt to improve EM and LMSE performance by fixing the initial states, as this constraint cannot be applied in general circumstances.

The form of the model can lead to correlations within the estimated parameters. For example, consider washout of learning. The rate of washout is determined by both forgetting and learning from error. Fast washout of learning can be explained by high error sensitivity and low retention factor. For this reason and others, there is a tendency to predict that these parameters are correlated. We quantified correlations within parameters estimated by EM and LMSE. For EM, there were two pairs of model variables with appreciable correlation (absolute value of correlation coefficient > 0.4), and for LMSE there were three. The two pairs shared by both algorithms were slow state retention-slow state error sensitivity (EM, $R=-0.455$; LMSE, $R=-0.628$) and fast state retention-slow state error sensitivity (EM, $R=-0.412$; LMSE, $R=-0.682$). The third pair exclusive to LMSE was fast state retention-slow state retention (LMSE, $R=0.536$). Therefore, both algorithms were affected by correlated parameter estimates, the magnitude of these correlations appeared smaller for the EM algorithm.

There are alternative approaches to uncovering hidden behavioral states. EM can be thought of as a frequentist's approach to mathematical estimation. It identifies the parameter set with the most likely solution to the problem. However, EM assumes a flat prior distribution over the parameters; in other words, the algorithm currently does not allow the modeler to use prior information regarding the probability distributions of the parameters of the learning system. These considerations can be accounted for within the context of Bayesian approaches to estimation. To our knowledge, this Bayesian

framework has not been applied to two-state models of adaptation; Bayesian techniques represent an exciting avenue that may further improve upon the robustness of the EM approach we pursued here⁵⁸.

3.4.4 Relationship between mathematical hidden states and neural substrates of learning

Why might it be useful to mathematically identify the hidden processes that underlie learning? In terms of behavior, various studies have posited that the fast and slow timescales of learning map onto dissociable components of a movement: during saccades, the early component of the movement exhibits properties that resemble influence of the slow process, whereas the later component of the same movement exhibits influence of the fast process¹³⁸. During reaching, fast processes appear to mirror more temporally-labile components of memory and slower processes appear more temporally-stable^{8,9,29,55}. Fast processes may relate to explicit or cognitive types of motor learning, where slower processes are supported by implicit, unconscious motor learning mechanisms^{11,72}. Fast states of learning may require larger amounts of preparation-time to be expressed than slow states of learning, which are present in behavior executed at low reaction times³².

In terms of neural substrates of learning, some authors have found that the neural basis of the fast process may depend on the cerebellum, as evidenced by the observation that non-invasive cerebellar stimulation can modulate learning from error^{4,116}, and damage to the cerebellum can spare slow processes². For example, people with cerebellar damage maintain the ability to modulate error sensitivity of the slow process¹³⁹. Imaging results suggest that for arm movements, both fast and slow adaptive processes may depend on the cerebellum⁴⁷, as well as regions of the cerebral cortex.

At the neuronal level, the existence of different timescales of memory may be present within the architecture of the cerebellum. A recent study found that the basic computational unit in the cerebellum may be micro-clusters of P-cells that share a common preference for error²⁵. Anatomical studies show that a given error is transmitted to the cerebellum via complex spikes that engage different micro-clusters of P-cells, placed in disparate regions of the cerebellum¹⁴⁰. This raises the possibility that a single error produces plasticity in multiple regions of the cerebellum, engaging distinct neural elements that can combine their outputs in service of adaptation¹. Indeed, P-cells in the flocculus exhibit a preference for error direction¹¹⁸. When a visual error is in the preferred direction of a P-cell, that cell produces complex spikes, which in turn results in depression of simple spikes on the subsequent trial. If the temporal distance between the two trials is large, these changes fade away, akin to a process of forgetting. In contrast, that same visual error is in the anti-preferred direction of another group of P-cells, resulting in reduction of complex spikes below baseline, which produces small

potentiation of the simple spikes on the subsequent trial. These two groups of P-cells exhibit different sensitivities to the same error, and exhibit forgetting with passage of time, two elements that appear quite similar to mathematical two-state models inferred from behavior. That is, the neural basis of multiple timescales of memory may be in part associated with the diversity of error-preferences in the P-cells of the cerebellum.

Comparison of mathematical estimation of these states of learning with these probes provides an opportunity to identify the neural substrates that mediate the multiple timescales of motor memory.

3.5 Appendix

This appendix contains four sections. The first section contains a mathematical description of our EM algorithm. Our form of the EM algorithm applies a different maximization step than previous descriptions in the sensorimotor literature⁷⁴. Our numerical implementation of the M-step would best classify our algorithm as a generalized EM algorithm. First, we provide a qualitative overview of EM and generalized EM. We then derive the equations required to apply EM to our two-state model. The toolbox that implements the generalized EM algorithm, along with supporting documentation is available at <http://shadmehrlab.org/Tools/tools.html>.

The second section provides a description of the least-squares (LMSE) algorithm. We discuss the general structure of LMSE and provide the equations we used for the algorithm. The third section provides tables and parameter values. The final section discusses an extended model of learning where multiple targets are presented to the subject.

3.5.1 Mathematical description of the generalized EM algorithm

Given a state-space model of behavior (Eq. 3.10), we ask how the parameters of the model can be estimated from a set of measured behavioral data. Suppose that our paradigm consists of N trials. The experimental design of our paradigm specifies the sequence of targets $\{g\}_1^N = g^{(1)}, g^{(2)}, \dots, g^{(N)}$, error-clamp trials, and perturbations $\{r\}_1^N = r^{(1)}, r^{(2)}, \dots, r^{(N)}$. During the experiment, we record the subject's motor outputs $\{u\}_1^N = u^{(1)}, u^{(2)}, \dots, u^{(N)}$ which, along with $\{g\}_1^N$, allows us to compute the subject's motor action on each trial relative to the target $\{y\}_1^N = y^{(1)}, y^{(2)}, \dots, y^{(N)}$. We provide a description of these variables in the context of three common sensorimotor learning paradigms (force field adaptation, visuomotor adaptation, and saccade adaptation) in the third section of this document within Table 3.4.

Our goal is to determine the parameter set $\theta = \{a_s, a_f, b_s, b_f, \sigma_x^2, \sigma_u^2, d\}$ that best explains the measured data. Note that the d parameter is only relevant for paradigms that include set breaks. A standard approach to parameter estimation is maximum likelihood estimation (MLE). MLE identifies the parameter set that maximizes the likelihood of observing the measured data, given a model parameter set. We will refer to this likelihood as the incomplete likelihood function. It is incomplete, as it does not include all the random variables of our system, i.e., it omits the hidden states of learning. Please note that other sources may refer to the incomplete likelihood function as the marginal likelihood function. Stated mathematically, MLE identifies the parameter set $\hat{\theta}$ according to $\hat{\theta} = \arg \max_{\theta} L(\{y\}_1^N | \theta)$. We provide a derivation of the incomplete likelihood for our two-state model in Section A1.4 of the Appendix. While many MLE problems can be solved by maximizing this function directly, for systems described by our two-state model (Eq. 3.10), this maximization has no closed-form solution, and can also be difficult to solve numerically.

Another approach to the MLE problem is an algorithm known as EM. Instead of finding the maximum likelihood estimator in one step by maximizing the incomplete likelihood function, EM iteratively increases the incomplete likelihood function by maximizing a different objective function known as the expected complete log-likelihood function. We will derive this function shortly. Central to EM is the complete likelihood function described by: $L_c = L(\{\mathbf{x}\}_1^N, \{y\}_1^N | \theta)$. We will later show that this complete likelihood is the product of several exponential terms. Therefore, it is simpler to work with the natural logarithm of the complete likelihood, l_c , where $l_c = \log_e [L(\{\mathbf{x}\}_1^N, \{y\}_1^N | \theta)]$.

As its name suggests, EM is composed of an expectation (E-step) and a maximization (M-step) step. The algorithm begins by guessing an initial parameter set θ_0 and then performs the E-step and M-step in order, in an iterative fashion. During the E-step, we consider the conditional expectation of the complete log-likelihood function, $E[l_c | \{y\}_1^N, \theta_t]$, where θ_t is the estimate of the parameter set obtained from the M-step of the previous EM iteration. We will refer to this expectation as the expected complete log-likelihood function. We will later show that a functional form of the expected complete log-likelihood function can be derived using the Kalman filter, concluding the E-step.

Model variable	Sensorimotor learning paradigm		
	Force field adaptation	Visuomotor adaptation	Saccade adaptation
$y^{(n)}$	The adaptation index describing the force profile	Heading angle of the reach relative to the target.	The saccadic endpoint relative to the cued target.
$r^{(n)}$	A value (typically 0 or 1) encoding the presence or absence of the force field. It can be standardized to some force field magnitude level and then take fractional values.	An external rotation to the cursor about the starting position of the reaching movement.	A displacement of the target position presented after the subject executes her primary saccade.
$h_c^{(n)}$	Error-clamp trials in force field tasks apply a stiff spring to eliminate error between subject forces and robot force. Should take the value 0.	The value of the clamped angular error between feedback of the subject's hand position and the target position. During no feedback trials, can take the value 0 (Kitago et al. 2013).	The value of the clamped angular error between the final target position and the endpoint of the subject's primary saccade.

Table 3.4. Interpretation of model variables for common learning paradigms. Our derivation of the two-state model used general language that could apply across different sensorimotor learning modalities. Here we provide a description of some key parameters in the context of force field adaptation, visuomotor rotation, and saccade adaptation, to assist the general reader.

In the M-step we compute an updated parameter set that maximizes the expected complete log-likelihood function according to $\theta_{t+1} = \arg \max_{\theta} E \left[l_c \left(\{y\}_1^N, \theta_t \right) \right]$. Critically⁷⁸, iteration of the E- and M-steps guarantees that the incomplete likelihood function increases with each update to the model parameters: $L \left(\{y\}_1^N | \theta_{t+1} \right) > L \left(\{y\}_1^N | \theta_t \right)$. The E-step and M-step are iterated until the incomplete likelihood function converges to a stationary point.

Previous descriptions of the EM algorithm^{74,79} have outlined the E-step and M-step for linear time-invariant dynamical systems similar, though not identical, to the form of Eq. (3.10). The previous implementations of EM assumed state equations that were time-invariant, representing experiments that had neither set-breaks nor error-clamp²⁷ trials. Under these assumptions, there existed closed-form solutions for the M-step. However, the introduction of set breaks in Eq. (3.10) introduces time-varying

nonlinearities to the two-state model that make the closed-form specification of the M-step difficult, if not impossible. Furthermore, our restrictions to the fast and slow state dynamics (i.e., constraints relating the fast and slow retention factors and error sensitivities) complicate our ability to identify a closed-form expression that globally solves the M-step. In such cases, a more general form of the EM algorithm known as generalized EM is useful⁷⁷. In generalized EM, rather than maximizing the expected complete log-likelihood function during the M-step, one selects θ_{t+1} such that:

$$E \left[\log_e \left[L \left(\{\mathbf{x}\}_1^N, \{\mathbf{y}\}_1^N \mid \theta_{t+1} \right) \right] \mid \{\mathbf{y}\}_1^N, \theta_t \right] > E \left[\log_e \left[L \left(\{\mathbf{x}\}_1^N, \{\mathbf{y}\}_1^N \mid \theta_t \right) \right] \mid \{\mathbf{y}\}_1^N, \theta_t \right]$$

In words, to increase the incomplete likelihood function using EM, it is sufficient to identify a parameter set that simply increases the expected complete log-likelihood over the value associated with the parameter set attained on the previous EM iteration.

In our generalized EM algorithm, we select an invariant parameter space. We numerically search this parameter space during each M-step of the algorithm to maximize the expected complete log-likelihood function. This maximization ensures that we satisfy the condition of the generalized EM algorithm above. That is, by identifying the maximal value of the expected complete log-likelihood function in a parameter space that does not change from one iteration to the next, we guarantee that the updated parameter set is better than (or at least as good as) the previous parameter set, which is also contained in the same parameter space. This is a generalized EM algorithm in that it does not globally maximize the expected complete log-likelihood function. Though more computationally intensive than the standard EM algorithm, our generalized EM algorithm has the benefits of allowing the modeler to specify hard parameter bounds as well as functional constraints on the relationship between model parameters. Therefore, when using our algorithm, the modeler can restrict the parameter space to obtain only physically relevant solutions, and appropriate two-state model dynamics. In the following section we describe the mathematics that define our generalized EM algorithm.

3.5.1.1 Expectation Step

The E-step requires derivation of the expected complete log-likelihood function. Note that the complete likelihood function can be factored given the Markov-form of Eq. (3.10) (Shadmehr and Mussa-ivaldi 2012). This factorization allows the complete likelihood function to be expressed as the following product.

$$L_c = \left[\prod_{n=1}^N L\left(y^{(n)} \mid \mathbf{x}^{(n)}, \theta\right) \right] \left[\prod_{n=1}^{N-1} L\left(\mathbf{x}^{(n+1)} \mid \mathbf{x}^{(n)}, y^{(n)}, \theta\right) \right] L\left(\mathbf{x}^{(1)}\right) \quad (\text{A3.1.1})$$

Eq. (A3.1.1) expresses the complete likelihood function in terms of three types of probability density functions. Our goal is to find general expressions for the likelihood functions on the right-hand-side of Eq. (A3.1.1). We can obtain the first likelihood $L\left(y^{(n)} \mid \mathbf{x}^{(n)}, \theta\right)$ directly from the observation equation of Eq. (3.10). This likelihood is the probability density function of a normal random variable, which is provided below:

$$L\left(y^{(n)} \mid \mathbf{c}^T \mathbf{x}^{(n)}, \theta\right) = N\left(\mathbf{c}^T \mathbf{x}^{(n)}, \sigma_u^2\right) \quad (\text{A3.1.2})$$

The second likelihood on the right-hand-side of Eq. (A3.1.1) can be obtained from the state update equation of Eq. (3.10), and is the probability density function for a multivariate normal random variable described by:

$$L\left(\mathbf{x}^{(n+1)} \mid \mathbf{x}^{(n)}, y^{(n)}, \theta\right) = N\left(\mathbf{A}^{(n)} \mathbf{x}^{(n)} + \mathbf{b}^{(n)} e^{(n)}, \mathbf{Q}^{(n)}\right) \quad (\text{A3.1.3})$$

To fully specify the complete likelihood of Eq. (A3.1.1), we must also obtain an expression for $L\left(\mathbf{x}^{(1)}\right)$, the probability density function for the initial state. We will assume that the initial state of the learner is itself a normal random variable:

$$L\left(\mathbf{x}^{(1)}\right) = N\left(\bar{\mathbf{x}}_1, \bar{\mathbf{V}}_1\right) \quad (\text{A3.1.4})$$

The mean of the normal random variable $\bar{\mathbf{x}}_1$ can be represented as $\bar{\mathbf{x}}_1 = \begin{bmatrix} x_s^{(1)} & x_f^{(1)} \end{bmatrix}^T$ which introduces two additional parameters to our state-space model, $x_s^{(1)}$ and $x_f^{(1)}$, the mean initial values of the slow and fast states, respectively. We will assume that we can represent the variance-covariance

matrix $\bar{\mathbf{V}}_1$ in the diagonal form $\bar{\mathbf{V}}_1 = \begin{bmatrix} \sigma_1^2 & 0 \\ 0 & \sigma_1^2 \end{bmatrix}$ which introduces the parameter σ_1^2 , the variance of the

initial states. Our full parameter set that we seek to identify now consists of ten variables, i.e.,

$\theta = \{a_s, a_f, b_s, b_f, \sigma_x^2, \sigma_u^2, d, x_s^{(1)}, x_f^{(1)}, \sigma_1^2\}$. Substitution of Eqs. (A3.1.2) to (A3.1.4) into Eq. (A3.1.1)

yields the following expression for the complete likelihood function.

$$\begin{aligned}
L_c &= \prod_{n=1}^N (2\pi\sigma_u^2)^{-1/2} \exp\left(-\frac{1}{2\sigma_u^2} \left(y^{(n)} - \mathbf{c}^T \mathbf{x}^{(n)}\right)^2\right) \\
&\times \prod_{n=1}^{N-1} (2\pi)^{-1} |Q^{(n)}|^{-1/2} \exp\left(-\frac{1}{2} \left(\mathbf{x}^{(n+1)} - A^{(n)} \mathbf{x}^{(n)} - e^{(n)} f^{(n)}\right)^T Q^{(n)-1} \left(\mathbf{x}^{(n+1)} - A^{(n)} \mathbf{x}^{(n)} - e^{(n)} f^{(n)}\right)\right) \\
&\times (2\pi)^{-1} |\bar{V}_1|^{-1/2} \exp\left(-\frac{1}{2} \left(\mathbf{x}^{(1)} - \bar{\mathbf{x}}_1\right)^T \bar{V}_1^{-1} \left(\mathbf{x}^{(1)} - \bar{\mathbf{x}}_1\right)\right)
\end{aligned} \tag{A3.1.5}$$

Due the various products of exponential functions in Eq. (A3.1.5), it is simpler to consider the natural logarithm of the likelihood function which we will refer to as l_c . Taking the natural logarithm of both sides of Eq. (A3.1.5), yields the following form for l_c .

$$\begin{aligned}
l_c &= -\frac{1}{2\sigma_u^2} \sum_{n=1}^N \left(y^{(n)} - \mathbf{c}^T \mathbf{x}^{(n)}\right)^2 \\
&\quad - \frac{1}{2} \sum_{n=1}^{N-1} \left(\mathbf{x}^{(n+1)} - A^{(n)} \mathbf{x}^{(n)} - \mathbf{b}^{(n)} e^{(n)}\right)^T Q^{(n)-1} \left(\mathbf{x}^{(n+1)} - A^{(n)} \mathbf{x}^{(n)} - \mathbf{b}^{(n)} e^{(n)}\right) \\
&\quad - \frac{1}{2} \left(\mathbf{x}^{(1)} - \bar{\mathbf{x}}_1\right)^T \bar{V}_1^{-1} \left(\mathbf{x}^{(1)} - \bar{\mathbf{x}}_1\right) \\
&\quad - \frac{1}{2} \log_e(|\bar{V}_1|) - \frac{N}{2} \log_e(\sigma_u^2) - \frac{3N}{2} \log_e(2\pi) - \frac{1}{2} \sum_{n=1}^{N-1} \log_e(|Q^{(n)}|)
\end{aligned} \tag{A3.1.6}$$

We proceed with the E-step by deriving an expression for $E\left[l_c \left\{ \mathbf{y}_1^N, \theta_t \right\}\right]$. This conditional expectation yields the expected complete log-likelihood function that we analyze in the M-step. To conserve space, we will represent the conditioned terms in the expectation using a \dots symbol. We can easily obtain $E\left[l_c \left\{ \dots \right\}\right]$ by expanding the quadratic terms in Eq. (A3.1.6) and then taking the conditional expectation. Doing so yields the following intermediate form of the expected complete log-likelihood function.

$$\begin{aligned}
E[l_c | \dots] = & -\frac{1}{2\sigma_u^2} \sum_{n=1}^N y^{(n)2} + \mathbf{c}^T E[\mathbf{x}^{(n)} \mathbf{x}^{(n)T} | \dots] \mathbf{c} - 2y^{(n)} \mathbf{c}^T E[\mathbf{x}^{(n)} | \dots] \\
& - \frac{1}{2} \sum_{n=1}^{N-1} \left[E[\mathbf{x}^{(n+1)T} \mathbf{Q}^{(n)-1} \mathbf{x}^{(n+1)} | \dots] - E[\mathbf{x}^{(n+1)T} \mathbf{Q}^{(n)-1} \mathbf{A}^{(n)} \mathbf{x}^{(n)} | \dots] - E[\mathbf{x}^{(n+1)T} | \dots] \mathbf{Q}^{(n)-1} \mathbf{b}^{(n)} e^{(n)} \right. \\
& \left. - E[\mathbf{x}^{(n)T} \mathbf{A}^{(n)T} \mathbf{Q}^{(n)-1} \mathbf{x}^{(n+1)} | \dots] + E[\mathbf{x}^{(n)T} \mathbf{A}^{(n)T} \mathbf{Q}^{(n)-1} \mathbf{A}^{(n)} \mathbf{x}^{(n)} | \dots] \right. \\
& \left. + E[\mathbf{x}^{(n)T} | \dots] \mathbf{A}^{(n)T} \mathbf{Q}^{(n)-1} \mathbf{b}^{(n)} e^{(n)} - e^{(n)} \mathbf{b}^{(n)T} \mathbf{Q}^{(n)-1} E[\mathbf{x}^{(n+1)} | \dots] \right. \\
& \left. + e^{(n)} \mathbf{b}^{(n)T} \mathbf{Q}^{(n)-1} \mathbf{A}^{(n)} E[\mathbf{x}^{(n)} | \dots] + e^{(n)} \mathbf{b}^{(n)T} \mathbf{Q}^{(n)-1} \mathbf{b}^{(n)} e^{(n)} \right] \\
& - \frac{1}{2} \left[E[\mathbf{x}^{(1)T} \bar{\mathbf{V}}_1^{-1} \mathbf{x}^{(1)} | \dots] - E[\mathbf{x}^{(1)T} | \dots] \bar{\mathbf{V}}_1^{-1} \bar{\mathbf{x}}_1 - \bar{\mathbf{x}}_1^T \bar{\mathbf{V}}_1^{-1} E[\mathbf{x}^{(1)} | \dots] + \bar{\mathbf{x}}_1^T \bar{\mathbf{V}}_1^{-1} \bar{\mathbf{x}}_1 \right] \\
& - \frac{1}{2} \log_e(\bar{\mathbf{V}}_1) - \frac{N}{2} \log_e(\sigma_u^2) - \frac{3N}{2} \log_e(2\pi) - \frac{1}{2} \sum_{n=1}^{N-1} \log_e(|\mathbf{Q}^{(n)}|)
\end{aligned}
\tag{A3.1.7}$$

As we can see from the above equation, the expected value operator only affects terms within Eq. (A3.1.7) that are functions of the hidden states. Our final step is to derive an alternative form for the expectation terms in Eq. (A3.1.7) that are quadratic functions of the unknown states. We note the following identity, which applies to any pair of multivariate random variables \mathbf{x} , \mathbf{y} and some matrix A of appropriate dimension.

$$E[\mathbf{x}^T \mathbf{A} \mathbf{y}] = E[\mathbf{x}]^T \mathbf{A} E[\mathbf{y}] + tr[\mathbf{A} \text{cov}(\mathbf{y}, \mathbf{x})]$$

Here $tr[\dots]$ is the trace operator. This identity allows us to express the quadratic terms of Eq. (A3.1.7) as a function of linear state expectations and covariances, which we can compute using the Kalman filter. Applying this identity to Eq. (A3.1.7) yields our final expression for the expected complete log-likelihood function.

$$\begin{aligned}
E[l_c] = & -\frac{1}{2\sigma_u^2} \sum_{n=1}^N y^{(n)2} + \mathbf{c}^T \left(\mathbf{V}^{n|N} + \hat{\mathbf{x}}^{n|N} \hat{\mathbf{x}}^{n|N^T} \right) \mathbf{c} - 2y^{(n)} \mathbf{c}^T \hat{\mathbf{x}}^{n|N} \\
& - \frac{1}{2} \sum_{n=1}^{N-1} \left[\begin{aligned}
& \hat{\mathbf{x}}^{n+1|N^T} \mathbf{Q}^{(n)-1} \hat{\mathbf{x}}^{n+1|N} + \text{tr} \left[\mathbf{Q}^{(n)-1} \mathbf{V}^{n+1|N} \right] - \hat{\mathbf{x}}^{n+1|N^T} \mathbf{Q}^{(n)-1} \mathbf{A}^{(n)} \hat{\mathbf{x}}^{n|N} - \text{tr} \left[\mathbf{Q}^{(n)-1} \mathbf{A}^{(n)} \mathbf{V}^{n+1, n|N^T} \right] \\
& - \hat{\mathbf{x}}^{n+1|N^T} \mathbf{Q}^{(n)-1} \mathbf{b}^{(n)} e^{(n)} - \hat{\mathbf{x}}^{n|N^T} \mathbf{A}^{(n)T} \mathbf{Q}^{(n)-1} \hat{\mathbf{x}}^{n+1|N} - \text{tr} \left[\mathbf{A}^{(n)T} \mathbf{Q}^{(n)-1} \mathbf{V}^{n+1, n|N} \right] \\
& + \hat{\mathbf{x}}^{n|N^T} \mathbf{A}^{(n)T} \mathbf{Q}^{(n)-1} \mathbf{A}^{(n)} \hat{\mathbf{x}}^{n|N} + \text{tr} \left[\mathbf{A}^{(n)T} \mathbf{Q}^{(n)-1} \mathbf{A}^{(n)} \mathbf{V}^{n|N} \right] + \hat{\mathbf{x}}^{n|N^T} \mathbf{A}^{(n)T} \mathbf{Q}^{(n)-1} \mathbf{b}^{(n)} e^{(n)} \\
& - e^{(n)} \mathbf{b}^{(n)T} \mathbf{Q}^{(n)-1} \hat{\mathbf{x}}^{n+1|N} + e^{(n)} \mathbf{b}^{(n)T} \mathbf{Q}^{(n)-1} \mathbf{A}^{(n)} \hat{\mathbf{x}}^{n|N} + e^{(n)} \mathbf{b}^{(n)T} \mathbf{Q}^{(n)-1} \mathbf{b}^{(n)} e^{(n)}
\end{aligned} \right] \\
& - \frac{1}{2} \left[\hat{\mathbf{x}}^{1|N^T} \bar{\mathbf{V}}_1^{-1} \hat{\mathbf{x}}^{1|N} + \text{tr} \left[\bar{\mathbf{V}}_1^{-1} \mathbf{V}^{1|N} \right] - \hat{\mathbf{x}}^{1|N^T} \bar{\mathbf{V}}_1^{-1} \bar{\mathbf{x}}_1 - \bar{\mathbf{x}}_1^T \bar{\mathbf{V}}_1^{-1} \hat{\mathbf{x}}^{1|N} + \bar{\mathbf{x}}_1^T \bar{\mathbf{V}}_1^{-1} \bar{\mathbf{x}}_1 \right] \\
& - \frac{1}{2} \log_e \left(|\bar{\mathbf{V}}_1| \right) - \frac{N}{2} \log_e \left(\sigma_u^2 \right) - \frac{3N}{2} \log_e \left(2\pi \right) - \frac{1}{2} \sum_{n=1}^{N-1} \log_e \left(|\mathbf{Q}^{(n)}| \right)
\end{aligned}
\tag{A3.1.8}$$

Eq. (A3.1.8) is the culminating result of the E-step. Note that the following shorthand notations have been applied in Eq. (A3.1.8).

$$\begin{aligned}
\hat{\mathbf{x}}^{n|N} &= E \left[\mathbf{x}^{(n)} \middle| \{ \mathbf{y} \}_1^N, \theta_t \right] \\
\mathbf{V}^{n|N} &= \text{var} \left(\mathbf{x}^{(n)} \middle| \{ \mathbf{y} \}_1^N, \theta_t \right) \\
\mathbf{V}^{n+1, n|N} &= \text{cov} \left(\mathbf{x}^{(n+1)}, \mathbf{x}^{(n)} \middle| \{ \mathbf{y} \}_1^N, \theta_t \right)
\end{aligned}
\tag{A3.1.9}$$

The shorthand quantities $\hat{\mathbf{x}}^{n|N}$, $\mathbf{V}^{n|N}$, and $\mathbf{V}^{n+1, n|N}$ can be computed using a smoothed Kalman filter⁷⁹. To summarize the Kalman smoother equations, we begin by evaluating the following posteriors which are computed using a forward pass of the standard Kalman filter.

$$\begin{aligned}
\hat{\mathbf{x}}^{n|n} &= E \left[\mathbf{x}^{(n)} \middle| \{ \mathbf{y} \}_1^n, \theta_t \right] \\
\mathbf{V}^{n|n} &= \text{var} \left(\mathbf{x}^{(n)} \middle| \{ \mathbf{y} \}_1^n, \theta_t \right)
\end{aligned}
\tag{A3.1.10}$$

The terms $\hat{\mathbf{x}}^{n|n}$ and $\mathbf{V}^{n|n}$ are our posterior state estimate and variance covariance matrix on the n -th trial, given our current parameter estimate and all our observations up to the n -th trial. Note that this expectation and variance are similar, but not equivalent to the desired $\hat{\mathbf{x}}^{n|N}$ and $\mathbf{V}^{n|N}$ which we will refer to as our smoothed Kalman estimates. These smoothed Kalman estimates are the expectation and

variance of the state on the n -th trial, given all the observations we have made. To first compute $\hat{\mathbf{x}}^{n|n}$ and $V^{n|n}$ we must calculate the Kalman gain $\mathbf{k}^{(n)}$ according to:

$$\mathbf{k}^{(n)} = \left(\mathbf{c}^T V^{n|n-1} \mathbf{c} + \sigma_u^2 \right)^{-1} V^{n|n-1} \mathbf{c} \quad (\text{A3.1.11})$$

To compute the Kalman gain we require the prior estimate for the variance-covariance matrix, denoted $V^{n|n-1}$. Using the Kalman gain we can compute $\hat{\mathbf{x}}^{n|n}$ and $V^{n|n}$ according to the following equations.

$$\hat{\mathbf{x}}^{n|n} = \hat{\mathbf{x}}^{n|n-1} + \mathbf{k}^{(n)} \left(y^{(n)} - \mathbf{c}^T \hat{\mathbf{x}}^{n|n-1} \right) \quad (\text{A3.1.12})$$

$$V^{n|n} = \left(\mathbf{I} - \mathbf{k}^{(n)} \mathbf{c}^T \right) V^{n|n-1} \quad (\text{A3.1.13})$$

Next we forward project these posteriors to obtain the prior estimates for the next trial:

$$\hat{\mathbf{x}}^{n+1|n} = \mathbf{A}^{(n)} \hat{\mathbf{x}}^{n|n} + \mathbf{b}^{(n)} e^{(n)} \quad (\text{A3.1.14})$$

$$V^{n+1|n} = \mathbf{A}^{(n)} V^{n|n} \mathbf{A}^{(n)T} + \mathbf{Q}^{(n)} \quad (\text{A3.1.15})$$

The forward Kalman filter proceeds by recursively iterating Eqs. (A3.1.11) to (A3.1.15) a total of N times to compute $\hat{\mathbf{x}}^{1|1}, \hat{\mathbf{x}}^{2|2}, \dots, \hat{\mathbf{x}}^{N|N}$ as well as $V^{1|1}, V^{2|2}, \dots, V^{N|N}$. The prior states and covariances will also be required for smoothing. We initialize this recursion with the priors $\hat{\mathbf{x}}^{1|0}$ and $V^{1|0}$. Here these priors are taken as the parameter estimates for $\bar{\mathbf{x}}_1$ and \bar{V}_1 obtained on the previous iteration of the algorithm.

That is, the prior expectation and variance are computed from the values of $x_s^{(1)}, x_f^{(1)}$, and σ_1^2 that were obtained from the M-step of the previous EM iteration.

To obtain the expectations and covariances required for EM (Eq. A3.1.9), we will now perform Kalman smoothing. The Kalman smoother uses backwards recursions to compute the means and variances of the probability distributions described in Eq. (A3.1.9). In other words, after obtaining the posterior state and variance-covariance matrix for all N time steps, we can recursively smooth our previous estimates. Our current implementation has been described previously^{74,79}. We note that the smoothed Kalman estimates on the final time step, N , were already computed in the final step of the forward Kalman filter. Therefore, computation of each quantity in Eq. (A3.1.9), begins with $N-1$. First, we compute the helper variable $J^{(n)}$ which functions similarly to a Kalman gain.

$$J^{(n)} = V^{n|n} \mathbf{A}^{(n)T} \left(V^{n+1|n} \right)^{-1} \quad (\text{A3.1.16})$$

With the computation of $J^{(n)}$ we can now compute $V^{n|N}$, our smoothed variance-covariance matrix.

$$V^{n|N} = V^{n|n} + J^{(n)} \left(V^{n+1|N} - V^{n+1|n} \right) J^{(n)T} \quad (\text{A3.1.17})$$

We also need to compute our smoothed state estimates.

$$\hat{\mathbf{x}}^{n|N} = \hat{\mathbf{x}}^{n|n} + J^{(n)} \left(\hat{\mathbf{x}}^{n+1|N} - \hat{\mathbf{x}}^{n+1|n} \right) \quad (\text{A3.1.18})$$

Recursion of Eqs. (A3.1.16) to (A3.1.18) computes $\hat{\mathbf{x}}^{1|N}, \hat{\mathbf{x}}^{2|N}, \dots, \hat{\mathbf{x}}^{N-1|N}$ and $V^{1|N}, V^{2|N}, \dots, V^{N-1|N}$. To complete the E-step, we also require a smoothed estimate for the covariance of consecutive states denoted by $V^{n+1,n|N}$. We can obtain this covariance using the following equation.

$$V^{n+1,n|N} = V^{n+1|N} J^{(n)} \quad (\text{A3.1.19})$$

Note that $V^{2,1|N}, V^{3,2|N}, \dots, V^{N,N-1|N}$ do not need to be computed in a recursive process, and can be calculated after recursion of Eqs. (A3.1.16) to (A3.1.18).

3.5.1.2 Maximization Step

In the M-step of the EM algorithm, the goal is to maximize the expected complete log-likelihood function that is derived in the E-step. For our two-state model, there exists no closed-form expression that globally maximizes the expected complete log-likelihood function. Therefore, we used a generalized M-step that numerically maximizes the expected complete log-likelihood function (Eq. A3.1.8) in a constrained parameter space. We maintained the same parameter space for each iteration of our generalized EM algorithm. As we described in the introduction to Appendix 1, maintaining this invariant parameter space is sufficient to guarantee convergence of the EM algorithm.

To perform our numerical maximization, we used *fmincon* in MATLAB R2016a. Because *fmincon* performs constrained minimization, we converted our maximization problem to a minimization problem by minimizing the negated expected complete log-likelihood in Eq. (A3.1.8) with respect to the two-state model parameter set $\theta = \{a_s, a_f, b_s, b_f, \sigma_x^2, \sigma_u^2, d, x_s^{(1)}, x_f^{(1)}, \sigma_1^2\}$. We constrained the parameter space for this numerical optimization in two ways. First, we specified lower and upper bounds for all the model parameters according to Eq. (A3.1.20).

$$\begin{aligned}
a_{s,\min} &\leq a_s \leq a_{s,\max} & a_{f,\min} &\leq a_f \leq a_{f,\max} \\
b_{s,\min} &\leq b_s \leq b_{s,\max} & b_{f,\min} &\leq b_f \leq b_{f,\max} \\
\sigma_{x,\min}^2 &\leq \sigma_x^2 \leq \sigma_{x,\max}^2 & \sigma_{u,\min}^2 &\leq \sigma_u^2 \leq \sigma_{u,\max}^2 \\
x_{s,\min}^{(1)} &\leq x_s^{(1)} \leq x_{s,\max}^{(1)} & x_{f,\min}^{(1)} &\leq x_f^{(1)} \leq x_{f,\max}^{(1)} \\
\sigma_{1,\min}^2 &\leq \sigma_1^2 \leq \sigma_{1,\max}^2 & d_{\min} &\leq d \leq d_{\max}
\end{aligned} \tag{A3.1.20}$$

The numerical values for the upper and lower bounds that specify Eq. A3.1.20 are provided in Table 3.2. We used identical bounds for our least-squares algorithm, as described in Appendix 2. The second way we constrained our parameter space is by enforcing conventional two-state model dynamics. Recall that the fast and slow states have the following properties; the fast state learns rapidly, but also forgets rapidly. The slow state learns slowly, but forgets slowly. To enforce these state dynamics, one can specify the following parameter constraints.

$$\begin{aligned}
a_s &\geq a_f + \Delta a \\
b_f &\geq b_s + \Delta b \\
\Delta a, \Delta b &> 0
\end{aligned} \tag{A3.1.21}$$

As defined in Eq. (A3.1.21), the slow state will have a greater retention factor than the fast state, and the fast state will have greater error sensitivity than the slow state. For Eq. A3.1.21, we used the value 0.001 for Δa and Δb . To summarize the generalized M-step, we numerically maximize Eq. (A3.1.8) using *fmincon* in a parameter space that is constrained by Eqs. (A3.1.20) and (A3.1.21).

Note that the selection of the upper and lower bounds in Table 3.2 will be specific to modeler preferences and the features of the behavioral data. In the current work, we selected upper bounds on the initial states that were equal in magnitude to the size of the perturbation. We felt this was a logical bound as it represents the maximum value that could be attained by the slow or fast state at any point during the adaptation timecourse, in the absence of noise. For our retention factors we specified an upper bound of 1.1. We selected this bound to be greater than 1 to demonstrate that the LMSE algorithm tended to identify unstable properties in the slow state of learning. In *Changing the bounds on the parameter space*, we detail a control analysis where we changed this upper bound to 1, to prevent the identification of unstable fast and slow retention factors. For our error sensitivities, we specified a lower and upper bound of 0 and 1, to prevent “negative” learning, or unstable learning, respectively. Finally, for all the variances of our noise terms (state, motor, and initial state) we specified an upper bound of 10 degrees². This parameter will be specific to the range and units of the behavioral data. Here

we selected this value to greatly exceed the variance of the residuals for the state-space model fit to any of our subject behaviors.

3.5.1.3 Algorithm summary

Here we offer a practical summary of the algorithm. The algorithm begins by specifying an initial guess for θ_0 , which are the initial values for the model parameters that will seed the algorithm.

1. Use the current parameter estimate θ_t to compute the posteriors $\hat{\mathbf{x}}^{1|1}, \hat{\mathbf{x}}^{2|2}, \dots, \hat{\mathbf{x}}^{N|N}$ and $V^{1|1}, V^{2|2}, \dots, V^{N|N}$ as well as the priors $\hat{\mathbf{x}}^{2|1}, \hat{\mathbf{x}}^{3|2}, \dots, \hat{\mathbf{x}}^{N|N-1}$ and $V^{2|1}, V^{3|2}, \dots, V^{N|N-1}$ by recursively applying Eqs. (A3.1.11) to (A3.1.15). These forward recursions are seeded using the current parameter estimates, $\hat{\mathbf{x}}^{1|0} = \bar{\mathbf{x}}_1$ and $V^{1|0} = \bar{V}_1$.
2. Use the posterior and prior estimates from Step 1 to compute the smoothed conditional expectations and variances $\hat{\mathbf{x}}^{1|N}, \hat{\mathbf{x}}^{2|N}, \dots, \hat{\mathbf{x}}^{N-1|N}$ and $V^{1|N}, V^{2|N}, \dots, V^{N-1|N}$ by recursively applying Eqs. (A3.1.16) to (A3.1.18) backwards in time. Compute the conditional covariances $V^{2,1|N}, V^{3,2|N}, \dots, V^{N,N-1|N}$ by applying Eq. (A3.1.19).
3. Numerically maximize (*e.g.*, *fmincon* in MATLAB) the expected complete log-likelihood function (Eq. A3.1.8) with respect to the model parameters subject to desired bounds (Eq. A3.1.20) and linear constraints (Eq. A3.1.21). The maximizing model parameters now become the parameter estimates for the current EM iteration.
4. Return to Step 1 and start the next EM iteration using the updated model parameters computed in Step 3 to perform the state estimation. Stop when the incomplete likelihood function has converged.

3.5.1.4 The incomplete likelihood function

The EM algorithm iteratively locates local extrema of the incomplete (marginal) likelihood function. This function can be evaluated at the conclusion of each iteration of the algorithm in order to track convergence. Here we provide a brief derivation of the form of the incomplete likelihood function $\mathcal{L}(\{\mathbf{y}\}_1^N | \theta)$. First, by successive application of the definition of conditional likelihood, we can factor the incomplete likelihood function as follows:

$$L(y^{(1)}, y^{(2)}, \dots, y^{(N)} | \theta) = L(y^{(1)} | \theta) \prod_{n=2}^N L(y^{(n)} | \theta, \{y\}_1^{n-1}) \quad (\text{A3.1.22})$$

This factored incomplete likelihood reveals a direct relationship between the complete likelihood and the Kalman filter; $L(y^{(n)} | \theta, \{y\}_1^{n-1})$ is a normal random variable with a mean and variance that can be computed from the priors obtained using the forward Kalman filter:

$$\begin{aligned} E[y^{(n)} | \theta, \{y\}_1^{n-1}] &= E[\mathbf{c}^T \mathbf{x}^{(n)} + \varepsilon_u^{(n)} | \theta, \{y\}_1^{n-1}] = \mathbf{c}^T \hat{\mathbf{x}}^{n|n-1} \\ \text{var}(y^{(n)} | \theta, \{y\}_1^{n-1}) &= \text{var}(\mathbf{c}^T \mathbf{x}^{(n)} + \varepsilon_u^{(n)} | \theta, \{y\}_1^{n-1}) = \mathbf{c}^T \mathbf{V}^{n|n-1} \mathbf{c} + \sigma_u^2 \end{aligned} \quad (\text{A3.1.23})$$

Therefore, the incomplete likelihood function can be expressed as follows:

$$L(y^{(1)}, y^{(2)}, \dots, y^{(N)} | \theta) = \prod_{n=1}^N N(\mathbf{c}^T \hat{\mathbf{x}}^{n|n-1}, \mathbf{c}^T \mathbf{V}^{n|n-1} \mathbf{c} + \sigma_u^2) \quad (\text{A3.1.24})$$

Given that Eq. (A3.1.24) is a product of exponentials, we consider the natural logarithm:

$$\begin{aligned} \log_e \left(L(\{y\}_1^N | \theta) \right) &= -\frac{1}{2} \sum_{n=1}^N \frac{1}{\Sigma^{(n)}} (y^{(n)} - \mu^{(n)})^2 - \frac{1}{2} \sum_{n=1}^N \log_e(\Sigma^{(n)}) - \frac{N}{2} \log_e(2\pi) \\ \text{where } \mu^{(n)} &= \mathbf{c}^T \hat{\mathbf{x}}^{n|n-1} \\ \Sigma^{(n)} &= \mathbf{c}^T \mathbf{V}^{n|n-1} \mathbf{c} + \sigma_u^2 \end{aligned} \quad (\text{A3.1.25})$$

In summary, to compute the incomplete log-likelihood associated with a given set of model parameters, we use Eqs. (A3.1.11-15) to compute the prior hidden state expectations and variances, and subsequently apply Eq. (A3.1.25).

3.5.1.5 Convergence

In this section we discuss issues related to the convergence of the generalized EM algorithm. In a constrained parameter space, we can expect our implementation of the generalized EM algorithm to converge to either a stationary point or a boundary of the constrained parameter space⁸⁰. As with any EM algorithm, we are not guaranteed that this stationary point is the desired global maximum of the incomplete log-likelihood function. The stationary point reached by an EM algorithm is determined by its initial conditions (i.e., the starting parameter guess). Therefore, it is imperative to perform the EM algorithm using different initial conditions. Here for each set of data, we used 5 or 10 initial conditions (for simulated and behavioral data, respectively) in an attempt to identify the parameter set that

resulted in the greatest incomplete log-likelihood. We found that using 50 different initial conditions for the algorithm did not meaningfully affect our results, and therefore chose a smaller number of initial conditions to make the computation time for our study more tractable.

For each initial condition, we performed a fixed number of EM iterations. We found that the number of iterations required to achieve convergence scaled with the size of the data set. In virtually all cases, we found that 100 iterations of the EM algorithm was more than sufficient to achieve convergence of the incomplete log-likelihood function. The only exception, was for our trial-by-trial control analysis, where we used 200 iterations, due the greater number of trials and slower convergence rate. Although we used a fixed number of iterations in this study, the modeler could set a convergence criterion that terminates the algorithm once the change in log-likelihood from one EM iteration to the next falls below some threshold.

3.5.2 Overview of the LMSE algorithm

The current standard technique used for fitting state-space models to motor learning data is one that selects the model parameters that minimize the squared-error between measured variables and model predictions. Here we offer a brief description of one form of this algorithm, the least-mean-squared-error (LMSE) technique, which generalizes to other least-squares techniques implemented within the literature. To use LMSE, we imagine a noise-free state-space analogue of Eq. (3.10). We previously described this system in Eq. (3.18) but reproduce it again below:

$$\begin{aligned} \mathbf{x}^{(n+1)} &= \mathbf{A}^{(n)} \mathbf{x}^{(n)} + \mathbf{b}^{(n)} e^{(n)} \\ \mathbf{y}^{(n)} &= \mathbf{c}^T \mathbf{x}^{(n)} \end{aligned} \tag{A3.2.1}$$

This noise free system is equivalent to the expected value of the states and observed behaviors predicted by our general two-state model of Eq. (3.10). To be clear, $e^{(n)}$ refers to the model prediction for the error in our noise-free system (i.e. the expected value of the error in Eq. (3.10) given our model parameters), not the errors actually measured during the experiment (with the exception of error-clamp trials where the expected value of the error is equal to the error imposed on that trial).

Note that Eq. (A3.2.1) describes a deterministic system; specification of the parameter set $\theta_{LMSE} = \{a_s, a_f, b_s, b_f, x_s^{(1)}, x_f^{(1)}, d\}$ entirely determines the progression of the slow and fast states for a given sequence of perturbations, error-clamp trials, and set breaks. The LMSE algorithm simply searches this 7-dimensional parameter space to identify the parameter set that satisfies the following optimization.

$$\theta_{LMSE} = \underset{\theta}{\operatorname{argmin}} \left\{ \frac{1}{N} \sum_{n=1}^N \left(y^{(n)} - \hat{y}^{(n)} \right)^2 \right\} \quad (\text{A3.2.2})$$

Here $y^{(n)}$ is the measurement on trial n and $\hat{y}^{(n)}$ is the model prediction, which is computed from the output of Eq. (A3.2.1). The argument of the *argmin* function of Equation (A3.2.2) is the mean-squared-error between the observation and model prediction for a given parameter set. For all LMSE fits in this study, we used the MATLAB function *fmincon* to identify the least-squares solutions. For each LMSE fit, the algorithm was seeded at 50 points scattered across the parameter space, to better ensure the identification of the minimum mean-squared-error within the search space. The parameter space used for the LMSE algorithm was identical to that of the EM algorithm and is reported in Table 3.2.

The least-squares algorithm can also be thought of as a maximum likelihood estimator, under certain conditions. For our two-state model, the least-squares algorithm maximizes the likelihood of a system where any randomness in the measured behavior is attributed to the measurement process itself (i.e. it is external to the underlying motor learning process). For this system, we assume that the actual motor action generated by the subject $y_{true}^{(n)}$ differs from the observed behavior $y^{(n)}$ due to some non-zero Gaussian measurement noise:

$$\begin{aligned} \mathbf{x}^{(n+1)} &= \mathbf{A}^{(n)} \mathbf{x}^{(n)} + \mathbf{b}^{(n)} e^{(n)} \\ y_{true}^{(n)} &= \mathbf{c}^T \mathbf{x}^{(n)} \\ y^{(n)} &= y_{true}^{(n)} + \boldsymbol{\varepsilon}_{measure}^{(n)} \quad \boldsymbol{\varepsilon}_{measure}^{(n)} \sim \mathcal{N}(\mathbf{0}, \sigma_m^2) \end{aligned} \quad (\text{A3.2.3})$$

where

$$e^{(n)} = \begin{cases} r^{(n)} - y_{true}^{(n)}, & \text{not an error-clamp trial} \\ e_c^{(n)}, & \text{error-clamp trial} \end{cases}$$

Here σ_m^2 represents the variance of our measurement. Note for the EM algorithm, we did not include this term and assumed that there was no measurement noise in our system (i.e. $y^{(n)} = y_{true}^{(n)}$). To reiterate, Eq. (A3.2.3) describes a system where the underlying learning process is deterministic. The only noise in our model is that in the experimental measurement of the behavior, which does not affect the true behavior of the learner.

To prove that the MLE for this system satisfies the least-squares optimization in Eq. (A3.2.2) we must compute its incomplete likelihood function. For this, we consider the factored form of the incomplete likelihood function of Eq. (A3.1.22). To specify this form, we must compute the conditional

means and variances described in Eq. (A3.1.23). Critically, for our system in Eq. (A3.2.3) the measured behavior of the subject is corrupted by a measurement noise that is independent on every trial.

Therefore, all of the measured behaviors are independently distributed and Eq. (A3.1.23) simplifies to

$$E\left[y^{(n)}\middle|\theta,\{y\}_1^{n-1}\right]=E\left[y^{(n)}\middle|\theta\right] \text{ and } \text{var}\left(y^{(n)}\middle|\theta,\{y\}_1^{n-1}\right)=\text{var}\left(y^{(n)}\middle|\theta\right)=\sigma_m^2. \text{ Applying these}$$

conditional means and variances to Eq. (A3.1.25) yields the following form of the incomplete log-likelihood:

$$\log_e\left(L\left(\{y\}_1^N\middle|\theta\right)\right)=-\frac{1}{2\sigma_m^2}\sum_{n=1}^N\left(y^{(n)}-E\left[y^{(n)}\middle|\theta\right]\right)^2-\frac{N}{2}\log_e\left(\sigma_m^2\right)-\frac{N}{2}\log_e(2\pi) \quad (\text{A3.2.4})$$

Note that $E\left[y^{(n)}\middle|\theta\right]$ in Eq. (A3.2.4) is also equal to $\hat{y}^{(n)}$ in the objective function for LMSE (Eq. A3.2.2).

Furthermore, note that maximization of Eq. (A3.2.4) implies minimization of the quantity

$$\sum_{n=1}^N\left(y^{(n)}-E\left[y^{(n)}\middle|\theta\right]\right)^2. \text{ This is the same quantity minimized by the LMSE algorithm in Eq. (A3.2.2).}$$

Therefore, the parameter set that solves our least-squares algorithm also maximizes the likelihood of a system without state and motor noise.

3.5.3 Multiple target state-space model of learning

In the state-space model outlined in our Methods, we assume that the learner possesses a single slow and fast state which are consistently engaged in the learning process on each trial. This model is most compatible with data sets where the same movement target is provided to the subject on each trial.

Here we consider a more complicated paradigm where the learner is presented with a sequence of trials where she is instructed to make a movement towards a target that can change from one trial to the next. We assume that the number of targets is finite, and equal to G . Note that for our visuomotor rotation paradigm, $G = 8$. We will now provide the necessary modifications to the equations described in our Methods for this multiple target model.

As for the single target case, the learner adjusts her movement towards each target according to her estimate of the perturbation. As before, this estimate depends on the state of a slow and fast adaptive process. The learner has a separate fast and slow state for each target, all of which are

included in the state vector $\mathbf{x}^{(n)} \in \mathbb{R}^{2G \times 1} = \left[s_1^{(n)} \quad \dots \quad s_G^{(n)} \quad f_1^{(n)} \quad \dots \quad f_G^{(n)} \right]^T$. Here s_k is the slow

state for target k and f_k is the fast state for target k . The learner's estimate of the perturbation depends on this state vector according to the following equation:

$$\hat{f}^{(n)} = \mathbf{c}^T \mathbf{x}^{(n)}$$

$$\text{where } \mathbf{c} \in \mathbb{R}^{2G \times 1} = \left[c_1^{(n)} \quad \dots \quad c_G^{(n)} \quad c_1^{(n)} \quad \dots \quad c_G^{(n)} \right]^T \quad (\text{A3.3.1})$$

$$\text{and } c_k^{(n)} = \begin{cases} 1, & \text{target } k \text{ is presented on trial } n \\ 0, & \text{otherwise} \end{cases}$$

Eq. (A3.3.1) formalizes the selection of the appropriate fast and slow state for the current target. On each trial it will contain two entries that are equal to one (all others are zero). In this way, the estimate of the perturbation is the sum of the fast and slow state that correspond to the target on trial n .

The evolution of the fast and slow states from one trial to the next depends both on forgetting and error-based learning according to Eq. (A3.3.2).

$$\mathbf{x}^{(n+1)} = \mathbf{A}\mathbf{x}^{(n)} + \mathbf{C}^{(n)}\mathbf{b}\mathbf{e}^{(n)} + \mathbf{C}^{(n)}\boldsymbol{\varepsilon}_{\text{learn}}^{(n)} + \boldsymbol{\varepsilon}_{\text{base}}^{(n)} \quad (\text{A3.3.2})$$

In the no-generalization model analysis described in our Results, we constrained the parameters of our modified state update equation, with the following four assumptions. Of course, any of these restrictions could be relaxed to allow for a more general model. First, the fast states of learning all exhibit the same forgetting properties, as do the slow states (i.e., they have the same slow and fast retention factors). Second, all fast states of learning all learn at the same rate, as do the slow states (i.e., they have the same slow and fast error sensitivities). Third, there is no-generalization of learning across targets. Therefore, the error experienced on trial n only engages the fast and slow processes that correspond to the target presented on trial n . And finally, all states experience a baseline level of state noise on each trial. The fast and slow states engaged in the learning process experience amplified noise on that trial.

To enforce all these rules, we made the following constraints on the parameters in Eq. (A3.3.2). To enforce that each state forgets from one trial to the next, with no difference in forgetting across targets, we represented \mathbf{A} as a $2G$ by $2G$ diagonal matrix of the form

$$\mathbf{A} \in \mathbb{R}^{2G \times 2G} = \text{diag}(a_s, \dots, a_s, a_f, \dots, a_f). \text{ To enforce that all slow states and all fast states learn at the}$$

same rate, we used a common error sensitivity vector $\mathbf{b} = \mathbb{R}^{2 \times 1} = \begin{bmatrix} b_s & b_f \end{bmatrix}^T$.

To account for generalization, we introduced the selector matrix $\mathbf{C}^{(n)}$. Suppose that target k is visited on trial n . For our no-generalization model, the selector matrix causes only the slow and fast

states corresponding to target k to learn on trial n . In this case, the selector matrix is a $2G \times 2$ matrix of

the form: $C^{(n)} \in \mathbb{R}^{2G \times 2} = \begin{bmatrix} c_1^{(n)} & \cdots & c_G^{(n)} & 0 & 0 & 0 \\ 0 & 0 & 0 & c_1^{(n)} & \cdots & c_G^{(n)} \end{bmatrix}^T$. It is critical to note that this selector

matrix could be used to encode generalization of learning across targets. In the case of generalization, the modeler could replace the zero elements in matrix C to cause a single error to differentially affect the update of each state.

And finally, to enforce our rule concerning the variance of the state update process, we separated state noise into two terms in Eq. (A3.3.2): a baseline state noise that affects the evolution of each state on every trial $\boldsymbol{\varepsilon}_{base}^{(n)}$ and an additional noise source that only affects the states that

experienced learning on a particular trial $\boldsymbol{\varepsilon}_{learn}^{(n)}$ (i.e., the states for the target presented on trial n). To enforce the latter property, we multiplied the learning noise by the selector matrix. Given the dimensions of our system, our baseline and learning noises had the following form: the baseline noise

$\boldsymbol{\varepsilon}_{base}^{(n)} \in \mathbb{R}^{2G \times 1}$ was unbiased with mean $[0 \ \cdots \ 0]^T$ and covariance matrix

$B \in \mathbb{R}^{2G \times 2G} = \sigma_{base}^2 I_{2G \times 2G}$ where σ_{base}^2 represents a common baseline variance for all states. The

learning noise $\boldsymbol{\varepsilon}_{learn}^{(n)}$ was unbiased with mean $[0 \ 0]^T$ and variance-covariance matrix $L = \sigma_{learn}^2 I_{2 \times 2}$,

where σ_{learn}^2 represents the state update variance associated with learning.

We can account for set breaks with the decay factor introduced in our model described in Eq. (3.10) of the manuscript. With the introduction of set breaks, our two-state model can now be represented as the following system of state-space equations that account for both error-clamp trials, set breaks, and multiple targets:

$$\begin{aligned}
 \mathbf{x}^{(n+1)} &= A^{(n)} \mathbf{x}^{(n)} + \mathbf{b}^{(n)} e^{(n)} + \boldsymbol{\varepsilon}_x^{(n)} & \boldsymbol{\varepsilon}_x^{(n)} &\sim N\left([0 \ \cdots \ 0]^T, Q^{(n)}\right) \\
 \mathbf{y}^{(n)} &= \mathbf{c}^T \mathbf{x}^{(n)} + \boldsymbol{\varepsilon}_u^{(n)} & \boldsymbol{\varepsilon}_u^{(n)} &\sim N\left(0, \sigma_u^2\right) \\
 A^{(n)} &= \begin{cases} A & \text{no set break} \\ A^{d+1} & \text{set break} \end{cases} & Q^{(n)} &= \begin{cases} C^{(n)} L C^{(n)T} + B & \text{no set break} \\ A^d \left(C^{(n)} L C^{(n)T} + B \right) A^{dT} & \text{set break} \end{cases} \quad (\text{A3.3.3}) \\
 \mathbf{b}^{(n)} &= \begin{cases} C^{(n)} \mathbf{b} & \text{no set break} \\ A^d C^{(n)} \mathbf{b} & \text{set break} \end{cases} & e^{(n)} &= \begin{cases} r^{(n)} - y^{(n)} & \text{not an error-clamp trial} \\ e_c^{(n)} & \text{error-clamp trial} \end{cases}
 \end{aligned}$$

Chapter 4. The origins of anterograde interference in visuomotor adaptation

Anterograde interference refers to the negative impact of prior learning on the propensity for future learning. There is currently no consensus on whether this phenomenon is transient or long-lasting, with studies pointing to an effect in the time scale of hours to days. These inconsistencies might be caused by the method employed to quantify performance, which often confounds changes in learning rate and retention. Here, we aimed to unveil the time course of anterograde interference by tracking its impact on visuomotor adaptation at different intervals throughout a 24h period. Our empirical and model-based approaches allowed us to measure the capacity for new learning separately from the influence of a previous memory. In agreement with previous reports, we found that prior learning persistently impaired the initial level of performance upon revisiting the task. However, despite this strong initial bias, learning capacity was impaired only when conflicting information was learned up to 1h apart, recovering thereafter with passage of time. These findings suggest that when adapting to conflicting perturbations, impairments in performance are driven by two distinct mechanisms: a long-lasting bias that acts as a prior and hinders initial performance, and a short-lasting anterograde interference that originates from a reduction in error-sensitivity.

4.1 Introduction

We gain robustness through adaptation: in the face of environmental and/or internal perturbations, adaptation allows us to maintain precise control of elementary movements like reaching and saccades. Like other types of learning, adaptation may lead to interference or facilitation depending on the level of congruency of sequentially learned materials. Facilitation of learning is commonly referred to as savings, a process by which subsequent exposure to the same perturbation results in faster learning^{44,141}. In contrast, successive adaptation to opposing perturbations, for example, rotation *A* followed by rotation *B* may lead to a deficit in the learning of *B*. This phenomenon, known as anterograde interference, has been reported in various adaptation paradigms^{38,142,143}. Yet, there is currently no consensus on whether anterograde interference is transient or long lasting. In fact, whereas some studies suggest that anterograde effects may last less than a few hours^{104,142}, others appear to point to a long-lasting impact in the time scale of days^{39,42}. It has even been suggested that anterograde interference may be stronger than retrograde interference^{39,42,143}, masking the effect of interest in retrograde protocols aimed at unveiling the time course of memory consolidation³⁹.

This lack of consensus may be partly due to the method employed for measuring interference⁸². Previous studies estimated the amount of interference of *A* on *B* predominantly based on the initial level of performance, computed by averaging through the first trials of the learning curve. This empirical

measure does not discriminate between changes in learning rate and retention. That is, initial performance in *B* is a mixture of how much the subject has retained what they learned in *A*, and how much they can learn from errors experienced in *B*. If anterograde interference arises from impairment in the ability to learn, one would expect that prior exposure to *A* would reduce the learning rate in *B*. Yet, with the exception of¹⁴³ no study that we are aware of has focused on the rate of learning as the fundamental measure of anterograde interference.

Here, we aimed to unveil the origins of anterograde interference by varying the time interval elapsed between adaptation to opposing rotations throughout a 24 h period. This approach allowed us to estimate how the passage of time affected two potential sources of performance impairment: (1) retention of an opposing prior memory versus (2) changes in the rate at which new information could be acquired. We recruited a large number of subjects ($n = 93$) in order to measure how adaptation changed from *A* to *B*, when the two events were separated by 5 min, 1 h, 6 h and 24 h. We used a trial-by-trial error-based model of learning^{8,9,74,101,135} to determine the impact of prior learning on three separate processes: (1) biases in performance due to the memory of *A*, (2) the rate of memory decay in *B*, and (3) the capacity of learning from error in *B*. These three processes are represented separately by three specific model parameters: (1) the initial state of the learner in *B*, (2) the retention factor and (3) the error sensitivity. In contrast with previous findings, our work shows that anterograde interference recovers gradually with passage of time. This recovery proceeds despite initial impairments in performance that originate from a lingering memory of *A* that persists over a much longer time scale.

4.2 Materials and methods

Ninety-three healthy volunteers (33 males; ages: mean \pm std. dev. 24 ± 4 years old) with no known history of neurological or psychiatric disorders were recruited from the School of Medicine of the University of Buenos Aires. All subjects were right-handed as assessed by the Edinburgh handedness inventory¹⁴⁴. The experimental procedure was approved by the local Ethics Committee and carried out according to the Declaration of Helsinki.

4.2.1 Experimental Paradigm

Subjects were seated in a comfortable chair and performed a center-out-back task using a joystick operated with the thumb and index fingers of their right hand. Visual information was presented on a computer screen. The right elbow laid comfortably on an armrest and the wrist laid on a structure that

fixed the joystick over a desktop. Subjects were told to maintain the same wrist posture across experimental sessions. Vision of the hand was occluded throughout the study.

At the beginning of each trial, we displayed one of 8 potential targets (0.4 cm diameter, placed 2 cm from the start point and concentrically located 45° from each other) on a computer screen. Joystick position was represented on the screen with a gray cursor of the same size as the target. The gain of the joystick was set to discourage subjects from correcting their movements online. Specifically, a displacement of 1.44 cm of the tip of the joystick moved the cursor on the screen by 2 cm. On average, movement time for correct trials was 125.5 ± 26.6 ms (mean \pm 1 std. dev.), providing little or no opportunity for within-movement corrections based on visual feedback. Participants were instructed to make a shooting movement through the target, as fast as possible, starting at target onset. There were 8 trials per cycle (one for each target) and 11 cycles per block. The order of target presentation was randomized within each cycle.

Two types of trials were presented throughout the experimental session (Fig. 4.1A). During *null* trials, participants performed shooting movements in the absence of a perturbation. During *perturbed* trials, a counterclockwise (CCW, labeled as perturbation A), or a clockwise (CW, labeled as perturbation B) visual rotation of 30° was applied to alter the trajectory of the cursor.

Feedback about the subject's movement was provided on each trial via the color of the cursor, which varied along a gradient between red (miss) and green (hit). Furthermore, subjects had a limited amount of time to complete the movement after the appearance of the target. If the elapsed time exceeded 900 ms, the trial was aborted and the cursor was turned red until the next trial. Target hits with error $< 2.5^\circ$ were rewarded with a simulated sound of an animated explosion. The total score (hit percentage) was displayed on the screen at the end of each block. Subjects were instructed to try to maximize this score throughout the experiment. The task was programmed using MATLAB's Psychophysics Toolbox, Version 3.

4.2.2 Experimental Procedure

Figure 4.1A illustrates the experimental design. Participants were randomly assigned to one of four experimental groups or a control group. The experimental groups (Fig. 4.1A) performed one block (11 cycles) of null trials followed by six blocks (66 cycles) of CCW perturbed trials (perturbation A). After a variable time interval, each group performed six blocks (66 cycles) of the CW perturbation (perturbation B). The four experimental groups were distinguished by the amount of time that separated the two rotations: 5 min ($n = 16$), 1 h ($n = 20$), 6 h ($n = 19$), and 24 h ($n = 18$). This variation in the period between

perturbations *A* and *B* allowed us to assess how the passage of time impacted the initial level of performance in *B* (first cycle), as well as on each subject's ability to adapt to *B*.

A group of subjects ($n = 20$) experienced only the *B* perturbation. This control group served two purposes. First, it was critical for our analysis of anterograde interference, serving as our benchmark for performance in *B* without any potential influence of learning in *A*. Second, given that subjects always learned *A* before *B*, this group was key in ruling out an order effect. Control subjects performed 1 block (11 cycles) of null trials followed by 6 blocks (66 cycles) of *B*.

4.2.3. Data post-processing

For each trial, the pointing angle was computed based on the angle of motion of the joystick relative to the line segment connecting the start and target positions. Trials in which pointing angles exceeded 120° or deviated by more than 45° from the median of the trials for each cycle were excluded from further analysis (1.6% of all trials). After this processing, the trial-by-trial data were converted to cycle-by-cycle time series by calculating the median pointing angle in each 8-trial cycle for each subject. Unless otherwise noted, the cycle-by-cycle data were used for each analysis reported in this work.

4.2.4 Model-free data analysis

We empirically quantified each subject's learning rate in *A* and *B* by fitting a single exponential function (Eq. 4.1) to the sequence of pointing angles measured in the *A* and *B* periods.

$$y(t) = \alpha \exp(-\beta t) + c \quad (\text{Eq. 4.1})$$

Here $y(t)$ represents the pointing angle on cycle t . The first cycle of the rotation was represented by $t = 0$. The exponential fits included three parameters. The parameters α and c determine the initial bias and the asymptote of the exponential, respectively. The parameter β represents the learning rate of the subject. We constrained the relationship between α and c to force the exponential fit to intersect subject behavior at time step $t = 0$. Therefore, the exponential function had only two free parameters; the third was fixed by the initial level of subject performance. We fit one exponential function to the 66 cycles of the *A* rotation and another one to the 66 cycles of the *B* rotation (Figure 4.1). Each period was fit using the *fmincon* function (MATLAB 2018a) to minimize the squared error between subject behavior and the exponential fit.

Although the exponential function closely approximates the decay of motor error during adaptation to a single perturbation, its learning rate parameter reflects a mixture of cycle-by-cycle forgetting and error-based learning. This potentially confounds our analysis of interference because during the *A* perturbation, the direction of forgetting (always towards baseline performance) opposes the direction of error-based learning. However, during the *B* perturbation, an initial bias in the performance of the experimental groups towards *A* causes forgetting and error-based learning to act in the same direction. This relationship switches once subjects pass the “zero point” of baseline performance: here retention and error-based learning again oppose one another. These considerations illustrate the difficulties inherent in using exponential fits to disambiguate the differential effects learning and forgetting may have on behavior.

4.2.5 State-Space Model

To better quantify subject performance in *A* and *B*, we used a state-space model that dissociates the effect of cycle-to-cycle learning from forgetting while appreciating initial biases in learning.

When people perform a movement that produces an unexpected result, they learn from their movement error and retain part of this learning over time. In other words, behavior during sensorimotor adaptation can be described as a process of error-based learning and trial-by-trial forgetting^{8,26,135}. State-space models of learning consider how the behavior of a learner changes due to trial-by-trial error-based learning, and decay of memory due to the passage of time (i.e., trials). To examine the anterograde interference of *A* on *B*, we fit a single module state-space model to the empirical data. This allowed us to ascribe any differences in performance during the *B* period to meaningful quantities: sensitivity to error, forgetting rate, and initial state.

We imagined that the state of the learner (the internal estimate of the visuomotor rotation) changed from one cycle to the next, due to error-based learning and partial forgetting, according to Eq. (4.2).

$$\mathbf{x}^{(t+1)} = \mathbf{a}\mathbf{x}^{(t)} + \mathbf{b}e^{(t)} + \varepsilon_x^{(t)} \quad (\text{Eq. 4.2})$$

Here $\mathbf{x}^{(t)}$ represents the state of the learner on cycle t . The parameter \mathbf{a} is a retention factor that encapsulates how well the subject retains a memory of the perturbation from one cycle to the next. The parameter \mathbf{b} represents sensitivity to error and determines the rate at which each subject learns from error. The error sensitivity is multiplied by the visual error $e^{(t)}$ between the pointing angle and target.

The change in state from one cycle to the next is corrupted by state noise $\varepsilon_x^{(t)}$ which we assumed to be Gaussian with mean zero and variance equal to σ_x^2 .

The internal state of the subject is not a measurable quantity. Rather, on each cycle, the motor output of the subject is measured. We imagine that the motor output directly reflects the internal state but is corrupted by motor execution noise according to Eq. (4.3).

$$y^{(t)} = x^{(t)} + \varepsilon_y^{(t)} \quad (\text{Eq. 4.3})$$

As with our exponential fit of Eq. (4.1), here $y^{(t)}$ represents the subject's pointing angle on cycle t . We assumed that the motor execution noise $\varepsilon_y^{(t)}$ corrupting the reaching movement was Gaussian with mean zero and variance equal to σ_y^2 .

We fit the state-space model to cycle-by-cycle single subject behavior using the Expectation Maximization (EM) algorithm¹⁰¹. The algorithm identified the parameter set that maximized the likelihood of observing each sequence of subject pointing angles given the parameters and structure of our state-space model. This parameter set contained 6 parameters: the retention factor a , error sensitivity b , state noise variance σ_x^2 , motor noise variance σ_y^2 , and two parameters describing the initial state of the learner. We modeled the initial state of the learner as a normally distributed random variable with mean \mathbf{X}_1 and variance σ_1^2 . The parameter \mathbf{X}_1 served as our estimate of the initial bias of the learner.

To fit the model, we started the EM algorithm from 5 different initial parameter sets, performed 100 iterations of the algorithm¹⁰¹, and selected the parameter set with the greatest likelihood. We fit our state-space model to single subject behavior separately for the A and B periods. For the A period, we fit the 77 cycles encompassing the first 11 null cycles and the following 66 CCW rotation cycles (Fig. 4.1). We fit the initial null trials along with the perturbation trials to increase confidence in the model parameters. For the B period, we fit the 66 cycles encompassing the CW rotation (Fig. 4.1).

4.2.6 Validation of the single state-space model

Our primary analysis assumed that learning could be represented using a single adaptive state. For a single state system, impairment in the learning rate in B requires that the learning system (*i.e.*, the

model parameters) has changed from the *A* to the *B* period. In contrast, two-state models of learning posit that adaptation is supported by two parallel learning processes, a slow process that learns little from error but exhibits strong retention over trials, and a fast process that learns greatly from error but has poor ability to retain its memory over trials. To validate the choice of a single state over a two-state model, we fit a two-state model of learning to the *A* and *B* sequences of subject pointing angles, and compared the single state model and two-state model in their abilities to describe subject behavior using the Bayesian Information Criterion (BIC).

In a two-state model, the states evolve over trials according to Eq. (4.4).

$$\begin{aligned} \mathbf{x}_s^{(t+1)} &= \mathbf{a}_s \mathbf{x}_s^{(t)} + \mathbf{b}_s \mathbf{e}^{(t)} + \varepsilon_x^{(t)} \\ \mathbf{x}_f^{(t+1)} &= \mathbf{a}_f \mathbf{x}_f^{(t)} + \mathbf{b}_f \mathbf{e}^{(t)} + \varepsilon_x^{(t)} \end{aligned} \quad (\text{Eq. 4.4})$$

Here, the slow and the fast states are represented by the quantities \mathbf{X}_s and \mathbf{X}_f , respectively. As with the single state model (Eq. 4.2) each state changes due to forgetting (described by its retention factor a) and error-based learning (described by its error sensitivity b). These internal estimates of the perturbation are additively combined to determine motor behavior according to Eq. 4.5.

$$\mathbf{y}^{(t)} = \mathbf{x}_s^{(t)} + \mathbf{x}_f^{(t)} + \varepsilon_y^{(t)} \quad (\text{Eq. 4.5})$$

We fit this two-state model of learning to subject behavior during the *A* and *B* periods using the EM algorithm¹⁰¹. The algorithm identified the parameter set that maximized the likelihood of observing each sequence of subject pointing angles. We fit the model to the same cycles in *A* and *B* described for the single state model fits. To fit the model, we started the EM algorithm from 20 different initial parameter sets, performed 250 iterations of the algorithm, and selected the parameter set with the greatest likelihood. The model parameter set consisted of 9 variables: slow and fast retention factors \mathbf{a}_s and \mathbf{a}_f , slow and fast error sensitivities \mathbf{b}_s and \mathbf{b}_f , the variances of state evolution and motor execution, and three parameters for the initial state of the learner. We modeled the initial fast and slow states as normally distributed random variables with mean $\mathbf{x}_s^{(1)}$ and $\mathbf{x}_f^{(1)}$, and variance σ_1^2 . Each model was fit under the linear constraints $\mathbf{b}_f > \mathbf{b}_s$ and $\mathbf{a}_s > \mathbf{a}_f$. These constraints enforce that the slow state learns more slowly than the fast state, but also retains its memory better from one trial to the next⁸.

Then, we computed the Bayesian Information Criterion (BIC) for both models:

$$BIC = k \log(n) - 2 \log(L_{\max}) \quad (\text{Eq. 4.6})$$

Here k represents the number of model parameters (6 for the single state model, 9 for the two-state model), n represents the number of data points, and L_{\max} represents the maximum likelihood for the model fit obtained using the EM algorithm. To obtain a single estimate of BIC for each subject, we averaged the BIC over the A and B periods. To quantify the evidence for each model, we compared the BIC distributions for the single state and two-state models for all subjects in the experimental groups using a paired t -test.

4.2.7 Statistical assessment

Statistical differences were assessed at the 95% level of confidence. Prior to statistical testing, outlying parameter values were detected and removed based on a threshold of three median absolute deviations from the group median. For cases where our variables of interest did not fail tests for normality and equality of variance, we used a one-way ANOVA for our statistical testing. In cases where the statistical distributions failed tests for both equal variance across groups (Bartlett's test) and normality (Shapiro-Wilk test) we used the Kruskal-Wallis test to detect non-parametric differences across experimental groups. In cases where our statistical tests indicated a significant effect of group ($p < 0.05$), we used either Tukey's test (following one-way ANOVA) or Dunn's test (following Kruskal-Wallis) for post-hoc testing. For the latter (Dunn's test), pairwise tests of all experimental groups were conducted against the control group and Bonferroni corrected. In cases where one-way ANOVA was used for statistical testing, complementary figures depict the mean statistical quantity for each group as well as the standard error of the mean, calculated assuming a normal distribution. In cases where Kruskal-Wallis was used for statistical testing, complementary figures depict the median statistical quantity for each group as well as the standard error of the median (estimated with bootstrapping). When comparing mean values against zero, a one-sample t -test test was used followed by the Bonferroni correction for multiple comparisons.

4.3 Results

When people adapt to perturbation A , and then switch to the opposite perturbation B , performance in B appears impaired^{42,142,145–147}. Pinpointing the origin of this behavioral deficit is hard because performance in B may reflect two different processes: the level of retention of the memory of A , and the

ability to learn *B*. In addition, these factors may vary independently as a function of time. Our study aimed to dissociate between these two factors by varying the time interval elapsed between *A* and *B* as subjects adapted to conflicting visuomotor rotations.

4.3.1 Memory of *A* decays in time, but persists even after 24 hours

On each trial, subjects moved a joystick to displace a cursor to one of 8 targets. On average, movement time for correct trials was 125.5 ± 26.6 ms (mean \pm 1 std. dev.), providing little or no opportunity for within-movement corrections. All groups initially trained in a baseline period of null trials (no perturbation) followed by adaptation to perturbation *A* (Figure 4.1A). After completion of training in *A*, subjects in each group waited for a specific amount of time (5 min, 1 h, 6 h or 24 h), and then were exposed to perturbation *B*. Figure 4.1B shows the pointing angle during null trials (cycles 1 to 11), learning of *A* (cycles 12 to 77) and learning of *B* (cycles 78 to 143) for each of the experimental groups (gray curves) and the control group (black curve). Pointing angle refers to angle of motion of the joystick relative to the line segment connecting the start and target positions. As expected, the pointing angle during null trials was close to zero. During exposure to perturbation *A*, subjects shifted their pointing angle gradually, approaching -30° , but maintained small, sustained residual errors¹⁰. After adapting to *A* and waiting the assigned time, subjects returned and were exposed to perturbation *B*.

How did learning of *A* impact performance in *B*? We quantified the initial level of performance in *B* as the mean pointing angle during the first cycle of adaptation for each group (Fig. 4.2A). Given that little or no learning is expected to take place in one cycle (1 cycle = one trial per target), this measure allowed us to estimate the recall of *A*.

The initial level of performance in *B* was biased towards *A* and decayed as a function of time (Fig. 4.2A, one-way ANOVA, $F_{(88,4)} = 39.59$, $p < 0.001$; Dunnett's test, control different from all experimental groups with $p < 0.001$. Tukey's test, 5 min different from 24 h with $p = 0.045$, all other comparisons are non-significant). Notably, even at 24 h the memory of *A* remained strong, exhibiting nearly 50% retention (one-sample t-test against zero with Bonferroni correction: $p < 0.001$ for all experimental groups), while the control's group pointing angle was not different from zero ($p=0.28$). This observation is consistent with the presence of a lingering memory of *A*^{104,145}. This level of memory retention (48%) is comparable to that found for reaching under a visuomotor rotation of 30 degrees, and for velocity dependent force-field⁴², suggesting that our findings are generalizable to other experimental paradigms.

In summary, during initial performance in *B* the movements were strongly influenced by the presence of a memory of *A*. This memory decayed with time, but was still present at 24 h.




Figure 4.1. Paradigm to measure anterograde interference. **A.** Paradigm. Subjects held a joystick and made pointing movements towards one of eight visual targets shown on a display. The experiment began with 11 cycles of null trials (Null) after which a 30° counterclockwise rotation was applied to the cursor for 66 cycles (perturbation *A*). Next, each experimental group waited a period of time ranging from 5 min to 24 h. After this break, subjects were immediately exposed to a 30° clockwise rotation (perturbation *B*) for 66 cycles. **B.** Behavior. Pointing angles on each trial were collapsed into cycles by

identifying the median pointing angle across each cycle of 8 trials. The shaded region indicates ± 1 standard error of the median. Each group differs in the amount of time that elapsed between the exposure to the *A* and *B* periods (from top to bottom: 5 min, 1 h, 6 h, and 24 h). The behavior for each experimental group (gray) is compared with that of a control group (black) that was exposed to 11 cycles of null and then 66 cycles of the *B* perturbation.




Figure 4.2. Effect of prior learning on the initial level of performance and the ability to learn. A. The initial level of performance in *B*, estimated from the mean pointing angle on the first cycle is displayed here for all groups. Given that learning within one cycle is minimal, in the experimental groups this measure reflects the retention of the memory of *A*. Even at 24 h, there is roughly 50% retention of *A*. B. The rate of improvement (*i.e.*, the learning rate) in *B* for all experimental groups and the control group is shown. In A and B, error bars indicate ± 1 standard error of the mean and median, respectively. Asterisks indicate a level of significance of $p < 0.05$ (*) or $p < 0.001$ (***) .

4.3.2 Anterograde interference dissipates with increasing time separating *A* and *B*

In order to assess the rate of learning, we fit the motor output for each subject in *A* and *B* with an exponential function (Eq. 4.1). The exponential model accounted for $67.2 \pm 13.2\%$ (mean ± 1 std. dev.) of the variance in individual subject behavior. Note that this is almost double of the variance accounted for when fitting individual trials instead of cycles ($35.2 \pm 10.6\%$, mean ± 1 std. dev.). If the exponential curve was fit to the median behavior across the group, thus reducing cycle-by-cycle variability in reaching movements, the variance accounted for 92.1 to 95.9% depending on the group.

We found that during the *A* period there was no difference in the learning rates across the four experimental groups (Kruskal-Wallis, $X^2(62) = 4.75$, $p = 0.19$). That is, the various groups were indistinguishable during learning of *A*.

To examine if prior learning of *A* impaired the ability to learn *B*, we compared the rate of learning in *B* with that of the control group (Figure 4.2B). Non-parametric testing revealed a significant

effect of group on the ability to learn *B* (Fig. 4.2B; Kruskal-Wallis, $X^2(80) = 10.84$, $p = 0.029$). Post-hoc comparison between each experimental group and the control group identified a significant difference at 5 min and 1 h (Dunn's test with Bonferroni correction, 5 min different from control with $p = 0.044$, 1 h different from control with $p = 0.024$), that disappeared by 6 h (6 h and 24 h not different from control with $p > 0.952$). This temporal pattern in the impairment of motor learning is consistent with the theory of consolidation¹⁴¹.

Finally, to rule out the possibility that our results may be explained by an order effect (subjects always learned the CCW rotation before the CW rotation), we statistically compared the rate of learning of the control group with those of the experimental groups during learning in *A*. No differences were found between the learning rates of *A* and *B* control (Kruskal-Wallis, $X^2(78) = 5.53$, $p = 0.237$). Therefore, the control condition rules out the possibility that our results are explained by the order in which the perturbations were learned.

In summary, while the lingering memory of *A* caused the starting point of the adaptation to be strongly biased in all experimental groups, the learning process itself was significantly impaired only at 5 min and 1 h. As the time interval lengthened, the ability to learn recovered.

4.3.3 Anterograde interference is caused by a decrease in sensitivity to error that recovers with time

The exponential model we employed for our empirical analysis implicitly assumed that the rate of learning remained constant across trials. For the *B* period, this assumption is unlikely to be true because initially, learning from the errors induced by *B* is aided by forgetting of the memory of *A*. That is, as the *B* period starts, performance falls toward baseline, and the rate of this fall is due to two processes: forgetting of *A*, and learning from error in *B*. During this period, forgetting and learning act in the same direction. However, once the performance crosses baseline levels, the influence of memory decay on behavior is in the opposite direction to learning from error. State-space models of learning disentangle these processes of forgetting and learning. For this reason, we fit a state-space model to each subject separately during the *A* and *B* periods (Eq. 4.2 and 4.3).

To assess which model was more appropriate to explain our data, we fit both a one-state and a two-state model of learning separately to the *A* and *B* periods and compared the likelihood of each model using the Bayesian Information Criterion (BIC). At the level of individual subjects, we found that a two-state model of learning was justified in only 5 of the 73 subjects across the experimental groups

(Fig. 4.3A, black lines). Therefore, in our task, the measured behavior was better described by a single state (Fig. 4.3B, lower BIC for single state model, paired t-test, $t(92) = 16.133$, $p < 0.001$) than a two-state model of learning. This is in agreement with previous work on visuomotor adaptation in reaching⁴⁹. Therefore, we fit a single-state model to the data.

The state-space model assumes that learning is governed by two processes: a process that learns from error, and a process that retains a fraction of that memory from one trial to the next. The model closely tracked the observed behavior (Fig. 4.4A).




Figure 4.3. Validation of the single state-space model for interference dataset. A. We calculated the Bayesian Information Criterion (BIC) for single and two-state model fits to individual subject behavior. The endpoint of each line shows the average BIC for the A and B periods (left, single state model; right, two-state model). Each line depicts the result for a single subject. Black lines indicate subjects for which the two-state model was superior to the single state model. B. We calculated the difference in BIC for the single state and two-state models. Negative values indicate higher evidence for the single-state model. The bar depicts mean BIC, and error bars indicate ± 1 standard error of the mean.

To quantify the model's goodness-of-fit, we computed the fraction of each subject's behavioral variance accounted for by our model fit (R^2). To measure this coefficient of determination, we computed the expected value of the behavior predicted by our stochastic model (Eqs. 4.2 and 4.3) and compared this prediction with each individual subject's data. We found that for single subjects, our model accounted for approximately $81.4 \pm 8.4\%$ (mean ± 1 std. dev.) of the variance in subject behavior. We repeated this analysis at the group level, where noise in the process of learning (Eq. 4.2) and production of a movement (Eq. 4.3) is smoothed over subjects. For each group, we computed the median behavior (Fig. 4.4A, black curves for experimental groups, red curve for control), the median behavior predicted

by our model (Fig. 4.4A, blue curves for experimental groups, green for control), and then the coefficient of determination for these two time-courses. At the group level, the model accounted for 96.0 to 98.2% of the variance in median subject behavior.




Figure 4.4. State-space model fit to interference dataset. A. We fit individual behavior using a single state model that estimated cycle-by-cycle forgetting, error-based learning, and initial bias. Each plot depicts the median pointing angle for each experimental group (black lines) as well as the median pointing angle predicted from simulating the state-space model without noise (blue lines) using the maximum likelihood model parameter sets identified for each subject. Behavior (red) and state-space predictions (green) are provided for the control group in the top-left plot. The shaded error region indicates ± 1 standard error of the median. B. Initial state of the learner at the start of the *B* period. C. Error sensitivity during the *B* period. D. Retention factor during the *B* period. In panels E-G we repeated our primary analysis, excluding the initial cycles of large errors from the model fit. To this aim, based on our exponential model, we selected the cycle during which the pointing angle was nearest zero. E. New initial pointing angle for this analysis (black) and the true initial pointing angles at the start of *B* (gray). F.

Error sensitivity of a state-space model after removing the initial period of large errors. G. Error sensitivity of the state-space model fit after removing the initial period of large errors (black) and fit to the entire *B* period (gray). Each bar represents the mean (B and E) or median (C, D, F, and G) parameter value for each group. Error bars indicate ± 1 standard error of the mean or median. Asterisks indicate a level of significance of $p < 0.05$ (*), $p < 0.01$ (**), or $p < 0.001$ (***) .

Using the model parameters, we measured the impact of prior learning on: (1) biases in performance due to the memory of *A*, (2) the rate of memory decay in *B*, and (3) the capacity to learn from error in *B*. These three processes are represented separately by three specific model parameters: (1) initial state of the learner in *B*, (2) retention factor (3) and error-sensitivity.

Unsurprisingly, the initial state of the learner in *B* (Fig. 4.4B) closely followed our empirical estimate of the initial level of performance in *B* (Fig. 4.2A). As the interval between *A* and *B* increased, the initial state of the learner in *B*, *i.e.*, the amount of the *A* memory retained over time, decreased (one-way ANOVA, $F_{(88,4)} = 52.16$, $p < 0.001$; Dunnett's test, control different from all experimental groups with $p < 0.001$. Tukey's test, 5 min different than 24 h with $p = 0.023$, all other comparisons are non-significant). However, despite this temporal decay, all experimental groups retained at least 50% of the memory of *A* (one-sample t-test with Bonferroni correction, all experimental groups $p < 0.001$), while the control group was not different from zero ($p=0.11$). Therefore, impairment of performance in *B* was in part caused by a lingering memory of *A* that was present even at 24 h.

To what extent was the impairment in *B* driven by changes in the rate of error-based learning and the strength of memory retention? Similar to our empirical analysis, we first confirmed that the experimental groups did not differ in performance during the *A* period. That is, there was no difference in error sensitivity (Kruskal-Wallis, $X^2(65) = 1.16$, $p = 0.763$) or retention factor (Kruskal-Wallis, $X^2(66) = 0.53$, $p = 0.912$) across the experimental groups during adaptation to *A*. Yet, we found that error sensitivity was affected by prior learning (Fig. 4.4C; Kruskal-Wallis, $X^2(83) = 14.47$, $p = 0.006$). Post-hoc tests against the control group unveiled a significant reduction in error sensitivity at 5 min and 1 h but not at longer time intervals (Dunn's test with Bonferroni correction, 5 min different from control with $p = 0.008$, 1 h different from control with $p = 0.004$, 6 h and 24 h not different from control with $p > 0.132$). In contrast, we found no difference in the retention factor during learning in *B* for any of the experimental groups, including the control group (Fig. 4.4D; Kruskal-Wallis, $X^2(79) = 5.66$, $p = 0.226$).

In summary, our state-space model pointed to a similar conclusion drawn from our empirical findings. Prior exposure to *A* resulted in a bias in the initial state of *B* that persisted through 24 h. However, despite this lingering initial bias, prior exposure produced a short-lived impairment in error sensitivity: error sensitivity in *B* resembled control values when the time between *A* and *B* was 6 h or

more. Therefore, differences in performance in *B* for any timescale greater than 6 h were likely related to a prior memory of *A* and not to a deficit in learning.

4.3.4 Error sensitivity is independent of initial error size

We and others have shown that sensitivity to error declines as a function of error size^{51,52,148}. This raises the concern that the differences in error sensitivity reported in Fig. 4.4C may be driven by the differences in the magnitude of the initial error experienced at the start of perturbation *B*. To address this possibility, we re-analyzed behavior in *B*, this time controlling for initial error size. To this aim, we fit an exponential function (Eq. 4.1) and identified the cycle in which each participant exhibited a pointing angle near zero, and re-fit our state-space model to the behavior after this point. In this way, model parameters could no longer be impacted by differences in initial error size across groups (Fig. 4.4E). We found that the pattern of error sensitivity described when fitting the whole behavior (Fig. 4.4C; Kruskal-Wallis, $X^2(83)=14.47$, $p=0.006$; Dunn's test 5 min vs control, $p=0.008$; Dunn's test 1 hour vs control, $p=0.004$) persisted even when the initial error was near zero (Fig. 4.4F; Kruskal-Wallis, $X^2(81)=11.25$, $p=0.024$; Dunn's test 5 min vs control, $p=0.041$; Dunn's test 1 hour vs control, $p=0.007$). Furthermore, no significant differences were found when the two analyses were compared (Fig. 4.4G; Wilcoxon signed rank test for each experimental group with Bonferroni correction, $p>0.177$ for all groups).

In conclusion, these results indicate that the impact of prior learning on error sensitivity cannot be explained by the initial level of error experienced in *B*. Rather, the effect appears to be related to a deficit in the ability to learn.

4.3.5 Alternate hypotheses: a two-state model account of anterograde interference

The decrease in error sensitivity observed here is at odds with a prior account of anterograde interference in force field adaptation¹⁴³, in which an impairment of learning arises from differing initial biases in the underlying adaptive states of a two-state system, rather than a change in sensitivity to error. Briefly, in the two-state framework, learning in *B* is hindered when the slower state of learning is biased towards earlier adaptation. We tested this idea in a supplementary analysis using a two-state model, and found that it fails to account for our empirical data (Section 4.3.6). Therefore, the fact that the slow state is heavily biased towards *A* at the start of *B* cannot explain the deficit we observed in the speed of learning in *B*. In fact, fitting a two-state model in which parameters are free to vary from *A* to *B* (Section 4.3.7), yields a similar pattern of impairment in error sensitivity of the slow module to the one

obtained with a single-state model, with no impact of anterograde interference on the fast module. This analysis suggests that the recovery of behavior we observe over time was not caused by the interaction of the fast and slow states (i.e., the initial bias) but rather by the restoration of error sensitivity of the slow process. Altogether, our findings indicate that anterograde interference in visuomotor adaptation is caused by a genuine impairment in error sensitivity both for one-state and two-state models of learning.

4.3.6. Differences in the initial bias of two parallel adaptive states cannot explain interference

An earlier account of anterograde interference¹⁴³ investigated how the emergent properties of a learning system composed of two parallel adaptive states⁸, could demonstrate impaired learning in *B* after the experience of *A*. In a two-state model of learning, adaptation is achieved through the combined output of two parallel adaptive states, a fast learning state (Fig. 4.5A, green) and a slow learning state (Fig. 4.5A, red). Sing and Smith (2010), demonstrated that a two-state learning system exposed to two opposing perturbations *A* and *B*, would demonstrate slower learning in *B* (Fig. 4.5A, compare *B* curve with the naïve *A* curve shown in blue) because the slow state is heavily biased towards *A* at the start of *B*. That is, a two-state system can exhibit impairment during the *B* period, even though error sensitivity remains the same in *A* and *B*. It is critical to distinguish this earlier hypothesis from our conclusion that anterograde interference is caused by a true reduction in error sensitivity from *A* to *B*.

To test if the model by Sing and Smith (2010)¹⁴³ could explain the pattern of interference reported herein, we simulated a two-state model of learning using the same parameters for *A* and *B* (Eqs. 4.4 and 4.5). We provide an example of such a simulation in Fig. 4.5B. We obtained the parameters for this simulation by fitting the two-state model to the control group behavior in *B*, as described in our Methods. The initial fast and slow states in *B* were obtained from the final fast and slow states in *A*. Passage of time was accounted for by forcing the fast and slow states to decay between the *A* and *B* periods. To do this, we calculated the amount of forgetting that occurred from the last cycle in *A* to the first cycle in *B* in our visuomotor dataset, and scaled the fast and slow state down by the median level of forgetting observed in each experimental group. To determine if each simulation showed signs of anterograde interference, we compared the rate of learning in *B* to that of the control group using an exponential model (Eq. 4.1) fit to all cycles in the *B* period.

Finally, we fit the two-state model to the *A* period behavior of the 5 min experimental group that exhibited the largest reduction in error sensitivity. We used these parameters to simulate behavior

in the *B* period and then compared the rate of learning in this simulated *B* period to that of the actual *B* period (Fig. 4.5E and F).

Fig. 4.5C depicts the rate of learning obtained from the *actual* performance normalized to the control condition, to allow comparisons with the two-state model simulations. We found that the learning rate observed in our two-state model simulations of visuomotor adaptation was not affected in *B* at any of the time points (Fig. 4.5D). Fig. 4.5E shows the actual learning curve for the 5 min group compared to the prediction of the two-state model for the same condition. We found that the exponential rate of improvement in *B* exhibited by our two-state model simulation was a rather poor predictor of actual behavior (Fig. 4.5F, paired t-test, $t(15) = 4.235$, $p < 0.001$). That is, our subjects learned much slower in the *B* perturbation than predicted by a two-state model in which parameters remain constant in *A* and *B*. These results argue against the hypothesis that biases in the initial fast and slow states at the start of *B* could have led to the reduction in learning rate we observed in actual subject behavior.




Figure 4.5. Initial biases in a two-state learning system cannot account for anterograde interference. A. The two-state model (parameters obtained from Sing and Smith, 2010) posits that behavior (black) can be decomposed into parallel contributions from a slow (red) and fast (green) state. After learning the *A* perturbation (left half of the figure, +30° perturbation) both the fast and slow states are biased towards a memory of *A*. At the start of *B* (right half of the figure, -30° perturbation) the initial bias of the slow state towards *A* can slow the rate of learning of the *B* perturbation, as compared to initial learning of *A* (blue trace shows the pointing angle during *A* with opposite sign and shifted to the 0° pointing angle of *B*). B. Simulations using the two-state model parameters identified for the control group in our

visuomotor rotation task. Simulations were performed for individual participants, and included noise in motor execution and in the process of learning. Parameters were the same in *A* and *B*. We show one batch of simulations, for a hypothetical group of 15 participants. We show the simulated behavior (black), simulated slow state of learning (red) and the simulated fast state of learning (green). C. Rate of learning, measured from an exponential fit to our visuomotor dataset, normalized to the median rate in the control group. Values less than one indicate anterograde interference. D. Metrics for our two-state model simulation for our visuomotor parameter set. Learning rate in *B* is reported, normalized to control values. Bars show the median interference in each group. Lines denote one standard deviation, across 1000 batches of simulated groups with 15 participants each, with noise. An exponential was then fit to the median simulated group behavior in *A* and *B*. E. We simulated behavior of each subject during the *B* perturbation (red) using the two-state model parameters fit to the *A* period and compared it to the rate of learning of our 5 min experimental group (black). Solid lines indicate median prediction or behavior across subjects. Error bars indicate ± 1 standard error of the median. F. We fit our exponential model to the actual and simulated *B* behaviors. Bars indicate the median learning rate. Error bars indicate ± 1 standard error of the median. Asterisks indicate a level of significance of $p < 0.001$. In *B* and *C*, shaded values denote ± 1 standard error of the median.

4.3.7 Error sensitivity of the slow state is impaired in a two-state model of anterograde interference

Even though we found better evidence for a single state learning process than a two-state learning process, we still considered how a two-state model of learning would account for the pattern of anterograde interference observed in our task. To this aim, we fit our two-state model of learning (Eqs. 4.4 and 4.5) to the pointing angles of each participant in the *A* and *B* periods, and allowed the model parameters to freely vary between *A* and *B*.

We found that the pattern of interference differed in the slow and fast states of learning (Fig. 4.6). The slow state of learning showed clear signs of interference that mirrored our findings from a single state model (Fig 4.6A; one-way ANOVA, $F_{(79,4)} = 2.86$, $p = 0.0289$). That is, the slow state error sensitivity of the 5 min group significantly differed from that of the control group (Dunnett's test, 5 min different from control with $p = 0.015$, 1 h marginally different from control with $p = 0.062$, all other comparisons have $p > 0.574$), and gradually recovered over time. The fast state error sensitivity, on the other hand, was no different across each of the experimental time points and the control group (Fig. 4.6B; one-way ANOVA, $F_{(79,4)} = 0.34$, $p = 0.852$). We speculate that the faster state of learning was not as salient due to the large number of targets used in our experimental paradigm. Increasing the number of targets also increases the time between successive visits to the same target, which leads to more decay in the faster temporally-unstable learning process²⁹.

Therefore, we conclude that even for a two-state model of learning, participant behavior is best explained by a model in which error sensitivity is reduced when *B* closely follows *A* in time, but recovers

with longer separation between each period. The reduction in error sensitivity is specific to the slow state of learning, but spares faster components of learning.




Figure 4.6. Error sensitivity of the slow state of learning is impaired with anterograde interference. We fit a two-state model of learning to participant behavior in *A* and *B*, allowing the model parameters to freely vary in each rotation period. Here we report the mean slow state error sensitivity (*A*) and fast state error sensitivity (*B*) during the *B* period. We tested for differences in error sensitivity using a one-way ANOVA, with post-hoc testing of each experimental group against the control group. We observed differences in error sensitivity only for the slow state of learning. Asterisks indicate a level of significance of $p < 0.05$. Bars indicate the mean parameter values across participants in each group. Lines indicate ± 1 standard error of the mean.

4.3.8 Relative implicit/explicit contributions are unlikely to explain differences in error sensitivity

It is possible that the changes we observed in error sensitivity are due to differential contributions from implicit and explicit components of learning. One of the signatures of explicit, strategy-based, learning is a requirement for extra time before the onset of a movement^{32,62}. Therefore, we would expect that if a change in error sensitivity is caused by a change in the relative contribution of the explicit system, there would be a concomitant variation in Reaction Times (RT). To address this possibility, we analyzed RTs throughout adaptation in *B*. RTs were computed for each movement as the time elapsed from target presentation to movement onset. Median RT cycles were then calculated for each participant. We found no statistical differences in RT across groups when all behavior during adaptation to *B* was considered (Fig. 4.7A; repeated measures ANOVA=main effect of cycle $p < 0.001$; main effect of group: $p = 0.536$; cycle x group interaction: $p = 0.442$).

Further examination of the first block (first 11 cycles) to capture possible changes in RT during exposure to higher initial errors during the critical early cycles, did not reach significance either (Fig. 4.7B; one-way ANOVA, $F_{(88,4)} = 1.06$, $p = 0.382$). These results suggest that the relative implicit/explicit contribution

to visuomotor adaptation was similar between experimental and control groups and did not change with the amount of time separating A.

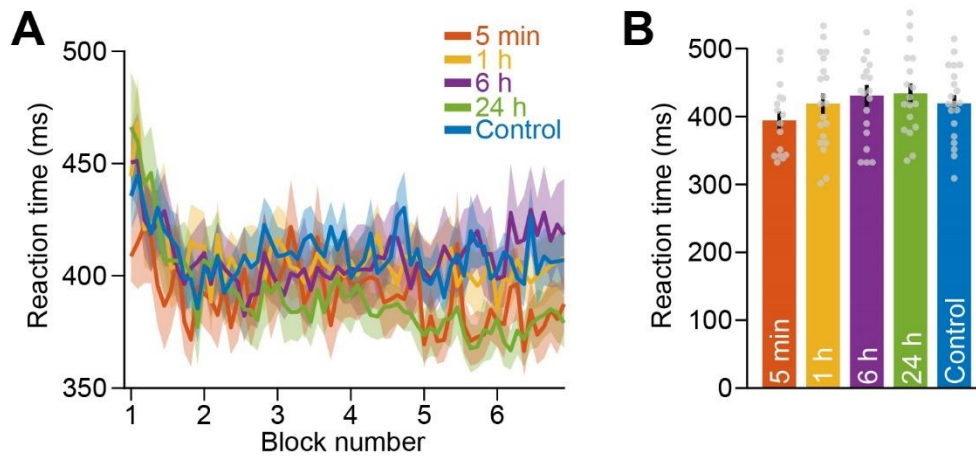


Figure 4.7. Reaction times during adaptation to the countermanding perturbation. A. Mean RT \pm standard error of the mean corresponding to each cycle as a function of time for all groups. B. Mean RT \pm standard error of the mean corresponding to the first block, when error size is large.

4.3.9 The implicit adaptive system is strongly limited by anterograde interference

Many authors have found that changes in the rate of learning can be attributed to explicit learning systems³¹. In these cases, implicit systems are generally thought to generate inflexible responses to error. Here, we sought to test if changes in implicit systems contribute to anterograde interference. To do this, we repeated the primary experiment, but this time limited reaction time to suppress explicit contributions to learning⁶². To do this, we instructed participants to begin their reaching movement as soon as possible, after the target location was revealed. To enforce this, we limited the amount of time available for the participants to start their movement after the target location was shown. This upper bound on reaction time was set to either 230-250 ms. We found that lower than these durations, participants were prone to executing random reaching movements or started their movement too late on most trials. If the reaction time of the participant exceeded the desired upper bound, the participant was punished with a 1 second timeout in the experiment after providing feedback of the endpoint. In addition, a low unpleasant tone (200 Hz) was played, and a message was provided on screen that read “React faster”.

Participants held the handle of a two-link robotic arm and made “shooting” movements with their arm in the x-y plane. Vision of the arm was occluded by an opaque white screen. Instead, feedback

of the hand position was provided through a small white cursor (3 mm diameter) that was placed on the screen by a projector.

At the start of each trial, the participant brought their hand to a center starting position (circle with 1 cm diameter). After maintaining the hand within the start circle for 750 ms, a target circle (1 cm diameter) appeared in 1 of 8 positions (0° , 45° , 90° , 135° , 180° , 225° , 270° , 315°) at a displacement of 8 cm from the starting circle. Participants then performed a “shooting” movement to move their hand briskly through the target.

Participants were provided audiovisual feedback about their movement speed and accuracy. If a movement was too fast (duration < 75 ms) the target turned red. If a movement was too slow (duration > 325 ms) the target turned blue. If the movement was the correct speed, but the cursor missed the target, the target turned white. Successful movements (correct speed and placement) were rewarded with a point that was displayed in the top left corner of the screen and also a pleasing tone (1000 Hz). If the movement was not successful, no point was awarded and a negative tone was played (200 Hz). Once the hand reached the target, visual feedback of the cursor was removed, and a yellow marker was frozen on screen to provide static feedback of the final hand position. At this point, participants were instructed to move their hand back to the starting position. During this time, the cursor continued to be hidden until the hand was moved within 2 cm of the starting circle. After 1 second, the robot motors turned on to assist the participant in finding the starting circle if it was not already reached. In the experiment, each target was visited once in a pseudorandom order in cycles of 8 trials.

Participants performed two blocks of movements, A and then B. The A period consisted of 5 cycles of normal movements followed by 80 cycles of movements where the cursor feedback was rotated by 30° in the CCW direction. The B period consisted of 80 cycles of movements where the cursor feedback was rotation by 30° in the CW direction. A and B were separated by a specified period of time. In the 5 min group ($n=9$), participants waited 5 minutes after finishing A before starting B. In the 24 hr group ($n=11$), participants waited 24 hours after finishing A before starting B.

Our results (Fig. 4.8) strongly suggested that implicit adaptation is altered by anterograde interference. First, the application of the upper bound on reaction time served to strongly limit the time participants used to prepare their movements (Figs. 4.8A and 4.8B), nearly halving that of the original experiment. This reduction in reaction time is known to isolate implicit learning, by preventing the expression of explicit memories that require additional time to prepare. If anterograde interference solely affected explicit adaptation and not implicit adaptation, we would expect to see no slowing of learning when reaction time is severely restricted. This was not the case (Fig. 4.8C). In fact, in both the 5

min group and 24 hr group, we noted a significant slowing of adaptation in B as compared with A. As before, the learning rate appeared to partially recover over time (Fig. 4.8D), but remained suppressed in B even after 24 hrs. Therefore, we can conclude that anterograde interference strongly interacts on the implicit learning process, but its magnitude diminishes in time. The strong reduction in learning rate in the 24 hr group relative to that of our original data set (Fig. 4.8D red) suggests that explicit learning partially rescued the deficit caused by implicit adaptation. Therefore, it may be that anterograde interference is specific to implicit adaptation.




Figure 4.8. Anterograde interference in implicit adaptation. A. The mean RT. Left (“free”) is the original data set. Right (“limit”) is the control data set where reaction time is limited. B. The mean RT. For “free” this is the first block average that you reported in the original data set. For “limit” this is the average over the first 11 cycles of the limit reaction time group. C. This is the behavior in B compared to the “flipped” behavior in A. All data is for the “limit” RT data set. Left is the 5-minute group (n=9). Right is the 24-hour group (n=11). To align the traces, we fit an exponential to the mean, and shifted along the time axis so that the exponential fit for each curve intersects zero at the same epoch number. D. Here is a quantification of slowing of learning. To account for potential differences in the speed of learning A, here this is the rate of learning in B normalized to the rate of learning in A. In other words, this is the rate from an exponential fit to the B period divided by the same quantity for the A period. You can see for both the 5 min and 24 hr groups, limiting RT increases the amount of interference.

4.3.10 Reduced error sensitivity is likely the cause of anterograde interference in force field adaptation

Is the reduction in error sensitivity we observed for visuomotor rotation a general property of sensorimotor adaptation? To answer this question, we revisited a force field adaptation experiment described in Sing and Smith (2010)¹⁴³. In this work, the authors argued that anterograde interference (Fig. 4.9A) is caused by the interaction of a fast and slow adaptive state. To quantify interference, the authors focused on behavior in the A period and compared it to the behavior in the B period, but only after behavior in B had crossed the point of zero-adaptation. Here we note a potential issue with this metric. Slowing of learning after the zero-point of adaptation could be caused either by (1) the interactions between a slow and fast state, or (2) the reduction in error sensitivity of the slow state, fast state, or both. On the contrary, in our work, we consider the rate of adaptation for all of the behavior in B, not solely after the zero-crossing. The benefit of this metric, is that it will only show a reduction if there is an adaptive process whose error sensitivity decreases. In other words, it will not show a reduction solely from the interactions of a fast and slow state, whose constitutive parameters do not change.

Therefore, we fit an exponential to the A period and the B period behavior of each group (Fig. 4.9B). To obtain the behavior of each group, we obtained the mean performance and standard error of each group from Fig. 4A of Sing and Smith (2010)¹⁴³. Then we generated theoretical subjects by sampling data on each trial assuming a normal distribution. We created theoretical sets of subjects, computed their mean, and then fit the exponential model. We repeated this bootstrapping approach, and then isolated the innermost 95% of the bootstrapped data. As a result, we found that repetition of A led to a precipitous drop in the learning rate for all groups in B. This result effectively means that one state of adaptation must have experienced a reduction in error sensitivity.

To further rule out the possibility that anterograde interference could not be caused solely by the interactions of a fast and slow state, we fit the two-state model to behavior in A and used these parameters to simulate performance in B, specifically for the 112 B trial group that performed roughly the same number of trials in A and B. We then fit our exponential model to the simulated behavior in A and the simulated B behavior (Fig. 4.9C, red) and compared the difference in adaptation rate to that observed in the actual subject data (Fig. 4.9C black). Whereas the data showed a decrease in learning rate from A to B, this reduction in learning rate was not observed when the B period behavior was simulated from two state model parameters obtained from the A period.

Now that we had confirmed that behavior in B had truly slowed due to a reduction in error sensitivity, we next attempted to identify whether this reduction occurred in the slow state, fast state, or both states of learning. To do this we fit a two-state model to behavior in A and B. When fitting the state-space model, we only possessed measurements on error-clamp trials randomly interspersed throughout the experiment. The missing intervening data points make use of the generalized EM algorithm in Chapter 3 impractical. Therefore, to fit the state-space model we used the technique adopted in Sing and Smith (2010)¹⁴³. We fit the two-state model to the data in the least-squares sense. However, we used a hypothesis-driven adjustment; we fit behavior in A and B simultaneously, but allowed error sensitivity of the fast and slow state to vary across A and B. In other words, we started with a noise-free two-state model:

$$\begin{aligned}
 y(n) &= x_s(n) + x_f(n) \\
 x_s(n) &= a_s x_s(n-1) + b_s e(n-1) \\
 x_f(n) &= a_f x_f(n-1) + b_f e(n-1)
 \end{aligned}
 \tag{Eq. 4.7}$$

Here, the slow and the fast states are represented by the quantities x_s and x_f . There error is given by $r(n) - y(n)$ where the perturbation r is zero during baseline, 1 during A, and -1 during B, and y is the adaptation index.

To fit behavior, we first collapsed data across the A period for each group. Then we fit the two-state model to the A period. We used the retention factors estimated from this two-state fit for further fitting. We next fit the two-state model to both A and B at the same time, requiring continuity in the fast and slow state between A and B. But we allowed there to be two separate slow state error sensitivities and two separate fast state error sensitivities, one for A and one for B. Therefore, A and B were allowed to have different error sensitivities, but shared the same retention factors. Using the retention factors for our model fit to A, we next fit the 4 free error sensitivity parameters for each group separately.

This modeling approach provided an excellent fit to behavior during both the A and B periods (dashed lines in Fig. 4.9A). Furthermore, the model suggested that anterograde interference was caused by a decrease in both the error sensitivity of the slow state (Fig. 4.9D) and fast state (Fig. 4.9E). Exposure to A prior to B appeared to completely impair the fast state of learning irrespective of the number of training trials in A. On the contrary, the impairment in error sensitivity in the slow state appeared to increase as the number of training trials in the A period increased. Altogether, our results from

visuomotor rotation and re-analysis of force field data suggested that time in A decreases the learning rate in B, but time after B allows this error sensitivity to recover.




Figure 4.9. Anterograde interference in force field adaptation. A. Here we show the data as reported in Sing and Smith (2010). The dashed black lines show the two-state model fit to the data. B. We combined data from all groups for the A period. We fit an exponential to the learning curve (in A). Then we fit exponentials to the B period for each group separately. Error bars show 95% CI intervals. C. At left, we show the exponential fit to the A period (as in B) as well as the exponential fit to the 112-trial group that experienced A for 112 trials. This group was chosen because it is the group that most nearly had just as many perturbation trials in A as in B, like the visuomotor rotation experiment. At right, we show simulations of the two-state model. Here, we fit the model to the A period, and then simulated behavior in B with constant parameters. Note that without a reduction in error sensitivity, there is no change in the learning rate. Therefore, anterograde interference in force field adaptation is not caused solely by

biases in the slow and fast states. Error bars show 95% CI intervals. D. Here we show the error sensitivity for the slow state. The A period here refers to the 112-trial group shown in C. All of the other bars refer to the B period. Error bars show 95% CI intervals. E. Here we show the error sensitivity for the fast state. The A period here refers to the 112-trial group shown in C. All of the other bars refer to the B period. Error bars show 95% CI intervals.

4.4 Discussion

How do motor memories influence one another? In this work, we studied the expression of anterograde interference in visuomotor adaptation by varying the time elapsed between learning opposing perturbations. We examined the impact of prior learning on the initial level of performance as well as the rate of learning over the time course of 5 min through 24 h. We found that these two parameters behaved very differently as a function of time. On one hand, adaptation in *A* biased the initial level of performance in *B*. Although the magnitude of this effect decreased with time, it remained strong at 24 h. On the other hand, prior adaptation impaired the ability to learn from error, resulting in reduced error sensitivity when perturbations were separated by 5 min and 1 h. Unlike the bias caused by prior learning, error-sensitivity completely recovered with the passage of time. To the best of our knowledge, these findings demonstrate for the first time that anterograde interference, a fundamental concept in memory research, is caused by a reduction in error sensitivity that recovers over time.

4.4.1 Anterograde interference differs from a lingering memory of a prior

There has been no general agreement in the sensorimotor literature regarding how to define and, therefore, quantify anterograde interference. With the exception of Sing and Smith (2010)¹⁴³, who measured the relative change in learning rate, most previous studies estimated anterograde interference based on the initial level of performance, by averaging across the first trials/cycles/blocks^{40,48,142,145,146}. For example, Tong and Flanagan (2003)¹⁴⁶ reported interference at 5 min based on the average of the second and third cycles. Likewise, Miall and collaborators (2004)³⁹ reported interference at 15 min based on the initial state obtained from fitting a power function, while noting that the rate of learning was not affected. Yet, it is likely that the initial level of performance, when averaged across trials, not only reflects the capacity for learning in *B* but also the bias of a lingering memory of *A*. Evidence supporting this possibility comes from force field studies showing that the preferred direction of the biceps and triceps during exposure to a second opposing force field is appropriate to solve the first force field (e.g. Thoroughman and Shadmehr 1999¹⁰⁴). Therefore, assessing anterograde interference based on the initial level of performance may overestimate its magnitude.

Here, we compared the bias imposed by the memory of *A* with the deficit observed in the rate of learning. We reasoned that if, as suggested by previous work, the initial level of performance reflects the level of anterograde interference then the two measures should behave similarly as a function of time. In contrast, we found that initial performance was profoundly hindered throughout the 24 h of testing, whereas the ability to learn resembled control levels starting at 6 h. Furthermore, the fact that the pattern of error sensitivity persists when controlling for initial error size suggests that anterograde interference is caused by a genuine down-regulation of error-sensitivity.

4.4.2 Anterograde interference and memory stabilization

Our work sheds light on a long-standing debate regarding the failure of retrograde protocols at unveiling the time course of memory consolidation. Over the past two decades, several laboratories have attempted to uncover the time course of memory stabilization using behavioral protocols based on retrograde interference^{40,42,142}. In these studies, subjects usually adapt to opposing perturbations *A* (A_1) and *B* separated by a time interval that varies between minutes to 24 h. Next, they wait for a further period of time (usually 24 h) and are again exposed to *A* (A_2) to assess the integrity of the motor memory. Consolidation of the memory of *A* should be reflected as the presence of savings (a faster rate of learning) in A_2 . Although this approach has proved successful in declarative^{41,150} and some kinds of motor skill learning tasks^{151,152}, it has led to inconclusive results in sensorimotor adaptation. In fact, with the exception of three force-field studies reporting release from interference at around 6 h^{142,145} or later¹⁵³, other experiments have shown complete lack of savings even if 24 h are interposed between A_1 and B ^{40,42,154,155}. These inconsistencies have also been reported for perceptual and motor sequence learning, prism adaptation^{155,156}, and declarative tasks including the paired associate task^{157–160}. Miall and collaborators (2004)³⁹ have claimed that naïve performance at recall (A_2) reported in retrograde protocols^{40,42,155} reflects a mixture of anterograde interference from *B* and the integrity of the memory of *A*, and not catastrophic retrograde interference. It is important to note, however, that these authors measured anterograde interference based on the initial level of performance. In light of our findings, the interpretation of these studies may need to be revisited. Our data indicates that because release from interference starts at around 6 h, anterograde interference is not likely to cause naïve performance in A_2 . Tracking the time course of recovery in learning rate reported here, provides a new path forward for understanding the process of memory stabilization.

The temporal dissociation we observed between the initial level of performance and the rate of learning likely reflects contributions of two distinct processes: (1) the persistence of a prior memory and

(2) competition for neural resources that support learning. The formation of memory involves learning-dependent synaptic plasticity as part of a process known as long-term potentiation and depression (LTP and LTD). Given that biological substrates underlying synaptic plasticity are limited by nature, cellular modifications induced by learning temporarily constrain the capacity for further LTP induction. This phenomenon is known as occlusion, and reflects competition for neural resources that support plasticity (Ling et al. 2002). Using this approach, it has been reported that motor skill learning in rats and humans is associated with LTP^{161–163}. Cantarero and collaborators showed that in fact, in humans, occlusion fades around 6 h after motor skill learning. In this light, we may speculate that adaptation in *A* may have partially occluded the capacity for further synaptic plasticity, thereby hindering adaptation in *B*.

In our study, all experimental groups received the same amount of training in *A*. However, others have demonstrated that increased repetition of the *A* perturbation enhances the amount of interference in *B*^{30,143}. Our results suggest that these changes in interference are caused not by differences in the amount of the *A* memory retained at the start of *B*, but potentially because of the amount of occlusion that occurred during the *A* period. That is, it may be that greater amounts of adaptation during the *A* period lead to further depletion of synaptic resources, thereby increasing the amount of interference in *B*. If this is correct, time in *A* appears to increase the amount of interference, and time after *A* improves the recovery from this occlusion.

The timing of recovery from interference we describe here (starting around 6 h) coincides with the peak in functional connectivity of a visuomotor adaptation network that includes the primary motor cortex (M1), the posterior parietal cortex (PPC) and the cerebellum¹⁶⁴. These regions have been linked to memory formation in this paradigm^{165–168}. Whether this timing reflects the process of motor memory consolidation is now a hypothesis amenable for testing.

4.4.3 Anterograde interference results from an impairment in error sensitivity that recovers with time

Using a state-space model allowed us to identify which aspect of learning was affected by anterograde interference, i.e., a deficit in the ability to learn from error or in the ability to retain information cycle-by-cycle. Here we found that anterograde interference could be attributed to a change in error sensitivity. It is well established that humans have the ability to change their error sensitivity depending on the errors they have experienced in the past¹²⁵. Current models of error sensitivity⁷ posit that sensitivity to a specific error increases gradually in an environment where that error is likely to occur

again on the next trial, and decreases gradually if the error is unlikely to occur again. In this way, sensitivity to an error is specific to one's prior history of error. However, the decrease in error sensitivity we report here, is of a different nature. Here, error sensitivity in *B* is reduced even though errors in *B* were never experienced during the *A* perturbation. This points to a different mechanism for error sensitivity modification, one in which learning of a perturbation actively suppresses learning of a different perturbation, potentially through a competition for neural resources that subsides at longer time intervals.

It is not currently clear if this interference affects all errors in *B*, or only a subset of error sizes. Because error sensitivity from the start of the *B* period (Fig. 4C) did not differ from the adaptation rate after the point of zero-crossing (Fig. 4F), we expect that interference generalizes broadly over errors opposite to those experienced in *A*. We would therefore expect the same temporal pattern of interference on the rate of de-adaptation during washout trials after *A*, with faster de-adaptation at long intervals (>6h) than short intervals, though to our knowledge, this prediction has not been clearly tested.

Substantial evidence indicates that visuomotor adaptation results from the interplay between explicit learning (driven by target error) and implicit learning (driven by prediction error)^{11,14,31,32}. Could the changes we observed in error sensitivity be caused by differential contributions of implicit and explicit components of learning? This possibility is unlikely for two reasons. The reaction time (a marker for explicit learning) of experimental and control groups were similar in *B* (Section 4.3.8), suggesting that the relative implicit/explicit contribution to adaptation did not change with the amount of time separating *A* and *B*. In addition, the pattern of error sensitivity persisted even after the critical early cycles, generally associated with the use of explicit strategies, were excluded (compare Figs. 4C and 4F). Together these results undermine the possibility that differences in the relative implicit/explicit contributions could account for our results.

In conclusion, we have examined the strength and duration of anterograde interference in visuomotor adaptation by tracking its impact on behavior when learning opposing perturbations was separated from 5 min through 24 h. We found that prior learning dramatically hindered the initial state at all time intervals. This was likely due to a bias imposed by a lingering memory associated with adapting to the initial perturbation. Prior learning also impaired the ability to learn from errors for at least 1 h, but a release from interference was detected starting as early as 6 h post training. This finding is consistent with a process of memory stabilization for this type of learning. Our work suggests that the poor performance observed when opposing rotations are learned consecutively, is driven by two distinct

phenomena operating on different time scales (days vs. hours): a long-lasting influence of a memory that acts as a prior which negatively influences the initial level of performance, and a shorter-lasting impairment of learning.

1 Chapter 5. Asymptotic limits of sensorimotor 2 adaptation

3
4 In many forms of motor adaptation, after extended practice performance reaches a saturation point in
5 which learning appears to stop, despite the fact that residual errors persist. Why is there a limit to
6 adaptation? Here we measured reaching in visuomotor and force field perturbation environments and
7 found that while reach adaptation exhibited an asymptotic limit, this limit changed with the second
8 order statistics of the perturbation: when the variance of the perturbation decreased, the adaptation
9 limit increased, leading to better performance. This relationship appeared to be causal, as the
10 adaptation limit could be changed in real-time by altering perturbation variance. In addition, the
11 relationship persisted at low reaction times, indicating that implicit processes played a role in changing
12 the adaptation limit. To better understand why variance impacted adaptation, we considered a
13 mathematical model in which the adaptation limit was due to a balance between two competing
14 biological forces: error sensitivity, which promoted learning, and forgetting, which promoted
15 performance decay. To test the model, we measured each of these forces during adaptation, and found
16 that the adaptation limit was controlled exclusively through error sensitivity modulation. By sorting
17 movements according to error size, we found that changes in error sensitivity were linked to the
18 consistency of the error experience: participants learned to adapt more from consistent error sequences
19 and less from inconsistent error sequences. These observations suggest that during adaptation, time-
20 varying error sensitivities interact with constant forgetting, setting the boundary conditions that
21 produce an apparent limit in the total extent of adaptation.

22

23 5.1 Introduction

24 During motor adaptation, humans and other animals experience perturbations that alter the sensory
25 consequences of motor commands, yielding sensory prediction errors. The brain responds to each error
26 by learning, thus adjusting its motor commands on the subsequent attempt. Over many trials, these
27 adjustments to the motor commands accumulate, but surprisingly, adaptation often remains
28 incomplete: even after an extended period of practice, residual errors persist. Residual errors are
29 observed in various behaviors including reaching^{14,123,142,169}, saccades^{9,170}, and walking¹⁷¹. Thus, an
30 unknown factor appears to limit the total amount of adaptation. Why does adaptation stop despite the
31 fact that residual errors remain?

32 An important clue is the observation that the mean of the residual errors is not fixed, but varies
33 with the variance of the past perturbations. For example, when people are exposed to perturbations
34 that vary in magnitude from one trial to the next, an increase in the perturbation variance also increases

35 the mean of the residual errors^{172–174}. Thus, the asymptote of adaptation is not a hard limit, but a
36 dynamic variable that depends on the second order statistics of the perturbation history.

37 To understand these observations, we begin with a well-established model in which the process
38 of adaptation is driven by two competing forces: sensitivity to error which produces learning, and trial-
39 by-trial forgetting which produces decay^{8,36,55,58}. When errors are large, learning dominates, yielding
40 changes in motor commands that improve performance. However, as errors become small, forces that
41 encourage forgetting reach an equilibrium with forces that induce learning. At this equilibrium,
42 performance stops changing, yielding persistent steady-state errors¹⁰.

43 While this theoretical framework predicts that residual errors should vary with the *mean* of the
44 perturbation, it fails to explain why the residual errors are altered by the *variance* of the perturbation.
45 One possibility is that perturbation variance affects error sensitivity^{55–60,125,175}, while another possibility is
46 that perturbation variance affects the forgetting rate. To compare these ideas, we trained volunteers to
47 make reaching movements in the presence of variable visual or force field perturbations. In some
48 experiments, we imposed a strict upper bound on reaction time, thus limiting the contributions of
49 explicit strategy^{32,62,132,176}. In all cases, we found that an increase in perturbation variance decreased
50 asymptotic performance, thereby increasing residual errors. Critically, these changes in residual errors
51 were caused by modulation of error sensitivity, not forgetting rates.

52 But why did increases in perturbation variance alter error sensitivity? To answer this question,
53 we sorted pairs of movements according to error size, and found that sensitivity to each error size
54 changed during adaptation: an error size that was consistently encountered was followed by an increase
55 in sensitivity to that error⁷. This simple rule, increase error sensitivity when errors are consistent,
56 correctly predicted not only changes to the rate of learning during adaptation, but also the magnitude of
57 the residual errors that remained at the end of adaptation.

58 Together, the results suggest that the brain does not simply learn from error and then discard
59 the error information. Rather, the effect of an error lasts beyond the trial in which it was experienced,
60 forming a memory of past errors that serves to regulate error sensitivity. Changes in error sensitivity
61 reach an equilibrium with forgetting, resulting in an asymptotic limit on the extent of adaptation.

62

63 5.2 Materials and methods

64 A total of 117 volunteers participated in our experiments. All experiments were approved by the
65 Institutional Review Board at the Johns Hopkins School of Medicine. In addition, we re-analyzed an
66 earlier study with 16 participants¹⁷². In Experiments 1-5, participants held the handle of a planar robotic

67 arm (Fig. 5.1A) and made reaching movements to different target locations in the horizontal plane. The
68 forearm was obscured from view by an opaque screen. An overhead projector displayed a small white
69 cursor (diameter = 3mm) on the screen that tracked the motion of the hand. Throughout testing we
70 recorded the position of the robot handle using a differential encoder with submillimeter precision. We
71 also recorded the forces produced on the handle by the subject using a 6-axis force transducer. Data
72 were recorded at 200 Hz.

73

74 5.2.1 Visuomotor rotation

75 Experiments 1, 3, 4, and 5 followed a similar protocol. At the start of each trial, the participant brought
76 their hand to a center starting position (circle with 1 cm diameter). After maintaining the hand within
77 the start circle, a target circle (1 cm diameter) appeared in 1 of 4 positions (0°, 90°, 180°, and 270°) at a
78 displacement of 8 cm from the starting circle. Participants then performed a “shooting” movement to
79 move their hand briskly through the target. Each experiment consisted of epochs of 4 trials where each
80 target was visited once in a pseudorandom order.

81 Participants were provided audiovisual feedback about their movement speed and accuracy. If a
82 movement was too fast (duration < 75 ms) the target turned red. If a movement was too slow (duration
83 > 325 ms) the target turned blue. If the movement was the correct speed, but the cursor missed the
84 target, the target turned white. Successful movements (correct speed and placement) were rewarded
85 with a point (total score displayed on-screen), an on-screen animation, and also a pleasing tone (1000
86 Hz). If the movement was unsuccessful, no point was awarded and a negative tone was played (200 Hz).
87 Participants were instructed to obtain as many points as possible throughout the experimental session.

88 Once the hand reached the target, visual feedback of the cursor was removed, and a yellow
89 marker was frozen on-screen to provide static feedback of the final hand position. At this point,
90 participants were instructed to move their hand back to the starting position. The cursor continued to
91 be hidden until the hand was moved within 2 cm of the starting circle. In most experiments, participants
92 actively moved their hand back to the start position. However, in Experiments 3 and 5, the robot
93 assisted the subject if their hand had not returned to the start position after 1 second.

94 Movements were performed in one of three conditions: null trials, rotation trials, and no
95 feedback trials. On null trials, veridical feedback of hand position was provided. On rotation trials, once
96 the target appeared on screen (the cue for the reaching movement), the on-screen cursor was rotated
97 relative to the start position (Fig. 5.1A). Each rotation experiment terminated with a period of no

98 feedback trials. On these trials, the subject cursor was hidden during the entire trial. In addition, the
99 subject was provided no feedback regarding movement endpoint, accuracy, or timing.

100 As a measure of adaptation, we analyzed the reach angle on each trial. The reaching angle was
101 measured as the angle between the line segment connecting the start and target positions, and the line
102 segment connecting the start and final hand position. The final hand position was taken to be the
103 location of the hand right after the hand exceeded 95% of the target displacement. The sign of the
104 reaching angle was determined by whether the final hand position was rotated CW or CCW relative to
105 the line connecting the start and target positions. For analysis of reaching errors, we computed the
106 same quantity, but for the final cursor position rather than the final hand position. On perturbation
107 trials, these two quantities are related by the rotation that occurred on that trial.

108

109 5.2.2 Force field adaptation

110 In Experiment 2, participants were perturbed by a velocity-dependent force field (Fig. 5.1A), as opposed
111 to a visuomotor rotation. At trial onset, a circular target (diameter = 1 cm) appeared in the workspace,
112 coincident with a tone that cued subject movement. Participants then reached from the starting
113 position to the target. The trial ended when the hand stopped within the target location. After stopping
114 the hand within the target, movement timing feedback was provided. If the preceding reach was too
115 slow, the target turned blue and a low tone was played. If the reach was too fast, the target turned red
116 and a low tone was played. If the reach fell within the desired movement interval (450-550 ms), the
117 subject was rewarded with a point (total score displayed on-screen), an on-screen animation, and also a
118 pleasing tone (1000 Hz). Participants were instructed to obtain as many points as possible throughout
119 the experimental session. After completing each outward reaching movement, participants were
120 instructed to then bring their hand back to the starting position. This return movement was not
121 rewarded and was always guided by a “channel” (see description below in this section).

122 As in the rotation experiments, the target appeared in 1 of 4 positions (0° , 90° , 180° , and 270°)
123 at a displacement of 10 cm from the starting circle. Each experiment consisted of epochs of 4 trials
124 where each target was visited once in a pseudorandom order. The experiment began with a set of null
125 field trials (no perturbations from the robot). After this period, participants were exposed to a force
126 field. The force field was a velocity-dependent curl field (Fig. 5.1A) in which the robot generated forces
127 proportional and perpendicular to the velocity of the hand according to:

128
$$\begin{bmatrix} f_x \\ f_y \end{bmatrix} = b \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix} \quad (5.3)$$

129 where v_x and v_y represent the x and y velocity of the hand, f_x and f_y represent the x and y force
130 generated by the robot on the handle, and b represents the magnitude of the force field. Note that the
131 sign of b determines the orientation of the field (the field is CW if $b > 0$ and CCW if $b < 0$).

132 Subject reaching forces were measured on designated “channel” trials²⁷ where the motion of
133 the handle was restricted to a linear path connecting the start and target locations (Fig. 5.1A). To restrict
134 hand motion to the straight-line channel trajectory, the robot applied perpendicular stiff spring-like
135 forces with damping (stiffness = 6000 N/m, viscosity = 250 N-s/m). Reaching forces were measured on
136 every 5th epoch of movements with a cycle of 4 channel trials, one for each target. Additionally, the
137 experiment terminated with a block of channel trials to measure how well the adapted state of each
138 participant was retained over time.

139 Offline we isolated forces produced on each trial against the channel wall, perpendicular to the
140 direction of the primary movement. To do this, we calculated the average force during baseline channel
141 trials. We then subtracted this baseline force timeseries from all of the force timeseries recorded during
142 channel trials throughout the experiment. After correcting for this baseline force, we then computed an
143 adaptation index. The adaptation index represents the scaling factor relating the force produced on a
144 given trial and the ideal force the subject would produce if they were fully adapted to the perturbation⁸.
145 To calculate this scaling factor, we linearly regressed the ideal force timecourse onto the actual force
146 timecourse. The ideal force was calculated by multiplying the velocity timecourse of the movement by
147 the perturbation magnitude. The adaptation index served as our behavioral measurement for model
148 fitting and parameter estimation.

149 In addition to analyzing the forces produced on channel trials, we also analyzed the trajectory of
150 the hand on perturbation trials. From each trajectory we isolated a signed movement error, which we
151 used to calculate the probability that an error switched sign from one trial to the next (Fig. 5.4C, Exp. 2).
152 To calculate the movement error, we isolated the portion of each reaching movement between 20% and
153 90% of target displacement. Within this region we detected the maximum absolute error and treated
154 this as the error magnitude. We signed this error according to whether the hand was to the left or right
155 (or top or bottom) of the line connecting the start position and target position. To prevent minor
156 overcompensations for the force field magnitude from being treated as movement errors, deviations

157 that fell within 3 mm of the line connecting the start and target locations were not treated as errors.
158 Using smaller thresholds of 1 or 2 mm did not qualitatively affect our results.

159

160 5.2.3 Statistics

161 In this work, we employed several types of statistical tests: repeated measures ANOVA, two-way
162 ANOVA, and mixed-ANOVA. These tests were carried out in IBM SPSS 25. In all cases we report the p-
163 value, F-value, and η_p^2 for each test. For post-hoc testing we employed t-tests with Bonferroni
164 corrections. For these tests, we report the p-value and Cohen's d as a measure of effect size. Our mixed-
165 ANOVA contained a between-subjects factor and a within-subjects repeated measure. For the within-
166 subjects repeated measure, data are binned within small windows defined by differences in error size. In
167 the event that a participant is missing data within a bin (data are missing in approximately 15% of all
168 bins), we replaced the missing data point with the mean of the appropriate distribution.

169

170 5.2.4 Experiment 1

171 We tested how variance in the perturbation affected the total extent of visuomotor adaptation. The
172 experiment started with 10 epochs (40 trials) of no perturbation. After this a perturbation period began
173 that consisted of 60 rotation epochs (240 trials total). At the end of the perturbation period, retention of
174 the visuomotor memory was tested in a series of 15 epochs (60 trials) of no feedback. To test the effect
175 of perturbation variance on behavior, participants were divided into 1 of 2 groups. In the zero-variance
176 group, participants (n=19) were exposed to a constant visuomotor rotation of 30°. In the high variance
177 group, participants (n=14) were exposed to a visuomotor rotation that changed on each trial. The
178 rotation was sampled from a normal distribution with a mean of 30° and a standard deviation of 12°.

179

180 5.2.5 Experiment 2

181 We found that perturbation variance reduced the total amount of adaptation in Experiment 1. To test if
182 this impairment was a general property of sensorimotor adaptation, we tested another group of
183 subjects with a force field. The experiment started with 10 epochs (40 trials) of no perturbation (2 of
184 these epochs were channel trials). After this a perturbation period began that consisted of 75 epochs
185 (300 trials, 20% were channel trials) of force field perturbations. At the end of the perturbation period,
186 retention of the adapted state was tested in a series of 10 epochs (40 trials) of channel trial movements.
187 To test the effect of perturbation variance on behavior, participants were divided into 1 of 2 groups. In

188 the zero-variance group, participants (n=12) were exposed to a constant force field magnitude of 14 N-
189 s/m. In the high variance group, participants (n=13) were exposed to a force field magnitude that
190 changed on each trial. The force field magnitude was sampled from a normal distribution with a mean of
191 14 N-s/m and a standard deviation of 6 N-s/m.

192

193 5.2.6 Experiment 3

194 Inspection of the learning curves in Experiment 1 indicated that performance may not have completely
195 saturated by the end of the perturbation period. Therefore, to confirm that perturbation variance
196 induces different performance saturation levels, we repeated Experiment 1, but this time more than
197 doubled the number of perturbation trials. The experiment started with 5 epochs (20 trials) of no
198 perturbation. The following perturbation period consisted of 160 rotation epochs (640 trials).

199 As in Experiment 1, participants were divided into a zero-variance group (n=10) and a high variance
200 group (n=10). Perturbation statistics remained identical to Experiment 1.

201

202 5.2.6 Experiment 4

203 To determine if perturbation variance causally altered the total extent of adaptation, we designed a
204 control experiment. In this experiment, participants started with a visuomotor rotation in the zero-
205 variance condition, and then after reaching asymptotic performance, were exposed to the high variance
206 condition. If variance causally determined the total amount of learning, we expected that asymptotic
207 performance would decrease after the addition of variability to the perturbation.

208 Participants (n=14) began the experiment with 5 epochs (20 trials) of null trials. After this, the
209 zero-variance period started. Participants were exposed to either a CW or CCW visuomotor rotation of
210 30° for a total of 80 epochs (320 trials). At the end of this period, participants switched to a high
211 variance condition where the rotation was sampled on each trial from a normal distribution with a mean
212 of 30° and a standard deviation of 12°. This period lasted for an additional 80 epochs (320 trials). Finally,
213 the experiment concluded with 15 epochs (60 trials) of no feedback.

214

215 5.2.7 Experiment 5

216 Sensorimotor adaptation is supported by both explicit strategy and implicit learning¹¹. To determine
217 which of these types of learning were impaired by perturbation variance we performed an experiment

218 where we limited the time participants had to prepare their movements. Limiting reaction time is known
219 to suppress explicit strategy⁶².

220 To limit reaction time, we instructed participants to begin their reaching movement as soon as
221 possible, after the target location was revealed. To enforce this, we limited the amount of time available
222 for the participants to start their movement after the target location was shown. This upper bound on
223 reaction time was set to either 225, 235, or 245 ms (taking into account screen delay). To enforce the
224 desired preparation time, if the reaction time of the participant exceeded the desired upper bound, the
225 participant was punished with a screen timeout after providing feedback of the movement endpoint. In
226 addition, a low unpleasant tone (200 Hz) was played, and a message was provided on screen that read
227 “React faster”. As in Experiment 1, participants were divided into a zero-variance perturbation group
228 (n=13) and a high variance group (n=12). All other details were identical to Experiment 1.

229

230 5.2.8 Re-analysis of prior work

231 In Fig. 5.1B, we reference earlier work from a study by Fernandes and colleagues¹⁷². The experiment
232 methodology is fully described in their original manuscript. Briefly, participants (n=16) made a center-
233 out reaching movement to a target. After the movement ended, participants were shown the endpoint
234 location of an otherwise hidden cursor that tracked the position of the right index finger. Participants
235 performed three experimental blocks. Each block had the same general structure. At the start of the
236 block, participants made 40 reaching movements to 8 different targets (5 for each target) with
237 continuous visual feedback of the cursor. Next, participants made an additional 80 reaching movements
238 to 8 different targets (10 for each target) using only endpoint feedback of the cursor position. After this
239 baseline period, a single target position was selected, and 240 reaching movements were performed
240 under the influence of a visuomotor rotation. The visuomotor rotation was sampled on each trial from a
241 normal distribution with a mean of 30° and a standard deviation of either 0°, 4°, or 12°. The block ended
242 in a set of 160 generalization trials that are not relevant to the current study. The experiment had a
243 within-subject design. Each participant was exposed to all three perturbation variances, but in a random
244 order. The orientation of the rotation (CW or CCW) was randomly chosen on each block. In addition, the
245 target selected during the adaptation period was randomly chosen from 1 of the 4 diagonal targets.

246

247 5.2.9 State-space model of learning

248 After the experience of a movement error, humans and other animals change their behavior on future
249 trials. In the absence of error, adapted behavior decays over time. Here we used a state-space model¹⁰¹

250 to capture this process of error-based learning. Here, the internal state of an individual x , changes from
251 trials n to $n+1$ due to learning and forgetting.

$$252 \quad x^{(n+1)} = ax^{(n)} + b^{(n)}e^{(n)} + \varepsilon_x^{(n)} \quad (5.4)$$

253 Forgetting is controlled by the retention factor a . The rate of learning is controlled by the error
254 sensitivity b . We describe modulation of error sensitivity in a later section. Learning and forgetting are
255 stochastic processes affected by internal state noise ε_x : a normal random variable with zero-mean and
256 standard deviation of σ_x . While we cannot directly measure the internal state of an individual, we can
257 measure their movements. The internal state x leads to a movement y according to:

$$258 \quad y^{(n)} = x^{(n)} + \varepsilon_y^{(n)} \quad (5.5)$$

259 The desired movement is affected by execution noise, represented by ε_y : a normal random variable
260 with zero-mean and standard deviation of σ_y .

261 To complete the state-space model in Eqs. 5.3 and 5.4, we must operationalize the value of an
262 error, e . In sensorimotor adaptation, movement errors are determined both by motor output of the
263 participant (y) and the size of the external perturbation (r):

$$264 \quad e^{(n)} = r^{(n)} - y^{(n)} \quad (5.6)$$

265 In our studies, the perturbation took the form of either a visuomotor rotation or a velocity-dependent
266 force field. Eq. 5.6 clearly demonstrates why perturbation variance influences error variance.

267

268 5.2.10 Asymptotic properties of learning

269 State-space models of learning predict that performance can saturate despite presence of residual
270 errors. This saturation is caused by a steady state condition where the amount of learning from error is
271 exactly counterbalanced by forgetting (Fig. 5.3A). Steady state can be derived from Eqs. 5.4-6:

$$272 \quad y_{ss} = \frac{b\bar{r}}{1-a+b} \quad (5.7)$$

273 The formula for steady-state adaptation (y_{ss}) shows that one's learning extent depends on 3 factors: (1)
274 error sensitivity b , (2) retention factor a , and (3) the mean of the perturbation \bar{r} . If there is no
275 forgetting ($a = 1$), an individual will adapt completely to the mean of the perturbation. However, if
276 retention is incomplete ($a < 1$), the steady state behavior (y_{ss}) will always fall short of the mean of the
277 perturbation, resulting in residual errors. Eq. 5.7 is important for three reasons. (1) It demonstrates why

278 the total extent of learning varies with a change in forgetting rate (Fig. 5.3B). (2) It demonstrates why
279 the total extent of learning varies with a change in error sensitivity (Fig. 5.3C). (3) It demonstrates that
280 the total amount of learning does not directly depend on variability in the perturbation, only the mean
281 of the perturbation (Fig. 5.3A).

282

283 5.2.11 Calculate of the retention factor

284 To determine if differences in learning extent were caused by a change in the rate of forgetting, we
285 estimated the retention factor (a) of each participant. To do this, we quantified how behavior decayed
286 during the error-free periods that terminated Experiments 1, 2, 4, and 5 (Figs. 5.1-3). During these error-
287 free periods, trial errors were either hidden (no feedback condition in visuomotor rotation experiments)
288 or fixed to zero (channel trials in the force field adaptation experiment). In the absence of error ($e=0$),
289 our state-space model simplifies to exponential decay (omitting noise terms):

$$290 \quad y^{(n)} = a^{n-m} y^{(m)} \quad (5.8)$$

291 Eq. 5.8 relates the motor output (y) on trial n of the error-free period to the initial motor behavior
292 measured at the end of the adaptation period, $y^{(m)}$. The term $n - m$ represents the number of trials
293 that elapsed from the start of the error-free period until the current trial n .

294 For visuomotor rotation experiments, we estimated the retention factor separately for each
295 target by fitting Eq. 5.8 to subject behavior in the least-squares sense. We report the mean retention
296 factor in Fig. 5.3E. For force field adaptation, we estimated a single retention factor, by first averaging
297 the adaptation index across the 4 targets in each epoch, and then fitting Eq. 5.8 to the epoch-by-epoch
298 behavior in the least-squares sense. In Fig. 5.3E, we converted this epoch-based retention factor to a
299 trial-based retention factor by raising the epoch-based retention factor to the power of 1/4 (an epoch of
300 4 trials has 4 trial-by-trial decay events).

301

302 5.2.12 Calculation of error sensitivity

303 Using Eq. 5.8, we found that changes in learning saturation were not caused by modulation of forgetting
304 rates. Next, we determined how variability impacted error sensitivity (b), using its empirical definition:

$$305 \quad b^{(n_1)} = \frac{y^{(n_2)} - a^{n_2-n_1} y^{(n_1)}}{e^{(n_1)}} \quad (5.9)$$

306 Eq. 5.9 determines the sensitivity to an error experienced on trial n_1 when the participant visited a
307 particular target T. This error sensitivity is equal to the change in behavior between two consecutive
308 visits to target T, on trials n_1 and n_2 (i.e., there are no intervening trials where target T was visited)
309 divided by the error that had been experienced on trial n_1 . In the numerator, we account for decay in
310 the behavior by multiplying the behavior on trial n_1 by a decay factor that accounted for the number of
311 intervening trials between trials n_1 and n_2 . For each target, we used the specific retention factor
312 estimated for that target with Eq. 5.8.

313 We used Eq. 5.9 to calculate error sensitivity for all of our visuomotor rotation experiments.
314 When reporting error sensitivity, we averaged across the four targets (Figs. 5.3F, 5.4A, 5.5C, 5.5D, and
315 5.5E). In some cases (Fig. 5.3F) we collapsed trial-by-trial measurements of error sensitivity across all
316 trials and all errors. In other cases, we calculated the change in error sensitivity over different periods of
317 training. For Fig. 5.5D, we measured the change in sensitivity from the beginning (epochs 1-15) to the
318 end (epochs 40-59) of the perturbation block in Exp. 1. For this, we calculated two error sensitivities,
319 one for errors less than 20° and the other for errors greater than 20° . To remove outliers, we identified
320 error sensitivity estimates that deviated from the population median by over two median absolute
321 deviations. We used a similar process for our analysis of Exp. 4 in Fig. 5.5E. Here we had three periods of
322 interest, the start of adaptation (epochs 1-3), the end of the zero-variance period (epochs 78-80), and
323 the end of the high variance period (epochs 157-159).

324 Our second to last method for analyzing error sensitivity is shown in Fig. 5.4A. Here, we
325 calculated sensitivity to errors of specific sizes. For this, we separated trial pairs into bins that depended
326 on error size: 5° - 10° , 10° - 15° , 15° - 20° , 20° - 25° , and 25° - 30° . For each subject-error bin pair, we required
327 there to be at least 10 measurements. We did not consider errors smaller than 5° because the empirical
328 estimator in Eq. 5.9 becomes unstable for small error sizes. Our final method of analyzing error
329 sensitivity is shown in Fig. 5.5C. Here we calculated error sensitivity both as a function of error size and
330 also trial number. These measurements are particularly noisy because on any given trial, only a subset of
331 participants experienced errors of a given size. Therefore, to reduce this noise, we calculated error
332 sensitivity within bins of consecutive trials. For the zero-variance perturbation we generally included 5
333 epochs (20 trials) in each bin. For the high variance perturbation (Fig. 5.5C middle) we included 10
334 epochs in each bin.

335 For force field adaptation, we could not empirically estimate error sensitivity, as this approach
336 requires the measurement of forces directly before and after the experience of an error. However, in
337 reality, forces are measured only on infrequent channel trials, making such an empirical calculation

338 impossible. For this reason, we used a model-based approach to measure error sensitivity (Fig. 5.3F, Exp.
339 2). We fit our state-space model Eqs. 5.4-6 to single subject data in the least-squares sense, over the last
340 5 channel trial epochs of the adaptation period. To do this, we needed to describe four states of learning
341 (one for each target). We describe multitarget state-space models in more detail in an earlier work¹⁰¹.
342 As a brief summary, we modeled our multitarget experiment by applying Eqs. 5.4-6 separately for each
343 target. On any given trial, the state corresponding to the relevant target learned from the error on that
344 trial. The other three states exhibited only decay on that trial. We described the perturbation r in terms
345 of the force field magnitude on that trial (14 N-s/m was considered a perturbation of unit 1 in the
346 model). Using this framework, we found the error sensitivity that minimized the squared difference
347 between our model simulation and participant behavior.

348

349 5.2.13 Memory of errors model

350 Using Eq. 5.9, we found that the capacity for learning changed in the different perturbation
351 environments due to a modulation in error sensitivity (Fig. 5.3F). Curiously, error sensitivity differed for
352 errors of certain sizes, not for all errors (Fig. 5.4A). To account for these findings, we used a memory of
353 errors model that was recently proposed by Herzfeld and colleagues⁷.

354 This model uses a simple normative framework. When the errors on trial n and trial $n+1$ have
355 the same sign (a consistent error), this signals that the brain under-corrected for the first error (Fig.
356 5.4B). Therefore, the brain should increase its sensitivity to the initial error. On the other hand, when
357 the errors on trials n and $n+1$ have opposite signs (an inconsistent error), this signals that the brain over-
358 corrected for the first error. Therefore, the brain should decrease its sensitivity to the initial error. These
359 rules are encapsulated by Eq. 5.2. Note that unlike the original memory of errors model, here, we allow
360 for decay in error sensitivity over time through the decay factor α .

361 In our simulations, we applied Eq. 5.2 to errors of different sizes. We divided up the error space
362 into 5° error bins. We assumed that all errors in that bin shared the same error sensitivity (Fig. 5.5B), but
363 sensitivity could differ across bins.

364

365 5.2.14 Simulation of the memory of errors model

366 In Fig. 5.5, we asked if our memory of errors model (Eq. 5.2) would accurately predict the behavior of
367 our subject population. To answer this question, we first fit the free parameters in Eq. 5.2 to our data
368 (Fig. 5.5A), and then used the fully specified model to simulate error sensitivity patterns (Fig. 5.5B).

369 These free parameters included two terms: α and β in Eq. 5.2.

370 We fit these two model parameters to the mean behavior in Experiments 3 and 4, in the least-
371 squares sense. We focused on Experiments 3 and 4, as these experiments included the longest training
372 periods, and therefore produced reliable steady-state adaptation limits. We fit the model to the
373 combined behavior of the zero-variance group in Experiment 3 (Fig. 5.5A, left), the high variance group
374 in Experiment 3 (Fig. 5.5A, middle), and the zero-to-high variance group in Experiment 4 (Fig. 5.5A,
375 right). In this way, we identified a single parameter set that minimized the sum of squared errors across
376 all three groups. The error we considered was the difference between mean subject reaching angles and
377 the mean reaching angle predicted through simulation of our model. By model, we are referring to the
378 reaching angle predicted by our state-space model (Eqs. 5.4 and 5.5) that used an error sensitivity that
379 started at 10% (chosen to match the initial error sensitivity empirically measured in subject behavior;
380 see curves in Fig. 5.5C) and then varied according to a memory of errors (Eq. 5.2). Our model assumed
381 that error sensitivity was common to errors within bins of 5° but could vary across bins (as in Fig. 5.4A).
382 We spaced these bins from -45 to 45° . For our simulations, we used a retention factor of $\alpha = 0.9736$ (the
383 mean retention factor measured in the error-free period in Experiment 1).

384 Due to variability in the process of learning, moving, and the high variance perturbation, the
385 output of our model was stochastic. Therefore, to obtain the model prediction for a given set of free
386 parameters, we calculated the mean output of the model over 10,000 simulations. For each simulation,
387 we resampled the trial-to-trial variations in the high variance perturbation. We also resampled trial-to-
388 trial variations in the state of the learner (ε_x in Eq. 5.4) and the motor output (ε_y in Eq. 5.5). For these
389 simulations, we set σ_x (Eq. 5.4) and σ_y (Eq. 5.5) both equal to 2° . We chose this variability level by
390 calculating the standard deviation of the last 100 reaching angles in the zero-variance group of
391 Experiment 3 (median standard deviation was 4.1° , and we divided this up evenly for state noise and
392 motor noise to arrive at our 2° estimate).

393 To identify the optimal parameter set we first attempted to use *fmincon* in MATLAB R2019a. We
394 found however, that *fmincon* yielded different parameter estimates with each change to the algorithm's
395 initial conditions. Therefore, to confirm that we identified a global, rather than a local minimum, in the
396 squared-error cost function, we performed a secondary grid search in the proximity of the parameter
397 sets identified by *fmincon*. For this grid search, we tested all 2,601 pairwise combinations of α and β ,
398 where α was varied from 0.95 to 1 in increments of 0.001, and β was varied from 0.03 to 0.08 in
399 increments of 0.001. This grid search identified an optimal parameter set of $\alpha = 0.987$ and $\beta = 0.042$.

400 We used this parameter set for our simulations in Figs. 5.5B-E. In Fig. 5.5A, we simulated our
401 model a total of 100,000. In Figs. 5.5C-5E we combined our model parameters with the actual error

402 sequences experienced by individual participants. In Fig. 5.5B, we tracked subject errors in each 5° bin,
403 and used Eq. 5.2 to predict if error sensitivity should increase or decrease from one trial to the next. For
404 the left and middle insets, we combined subjects across Experiments 1 and 3. However, Experiment 1
405 ended at epoch 60, so after this point, only subjects in Experiment 3 are represented (hence the change
406 in the size of the error bars after the vertical black line in the left and middle insets in Fig. 5.5B). To
407 transition between epoch 60 and 61, and maintain continuity in our predictions, we used a
408 bootstrapping approach. For each bootstrap, we sampled the initial error sensitivity on epoch 61 in each
409 error bin from a normal distribution with a mean and variance that were determined from the model
410 predictions on epoch 60. With that particular set of error sensitivities, we simulated the response from
411 epochs 61 to 160 for subjects in Experiment 3. We repeated this procedure a total of 1,000 times, each
412 time resampling the initial error sensitivity on epoch 61.

413 Finally, to obtain the model predictions in Figs. 5.5D and 5.5E, we used the error sensitivity
414 timecourses depicted in Fig. 5.5B. In Fig. 5.5D, we focused on participants in Experiment 1, where we
415 included an error-free period, thus permitting robust measurement of error sensitivity and forgetting
416 rates. In all cases, we measured the differences in error sensitivity for small (between 5° and 20°) and
417 large (between 20° and 30°) errors from the start (epoch 1) to the last epoch (epoch 60), predicted by
418 the memory of errors model (Fig. 5.5C). We used a similar method in Fig. 5.5E, only for Experiment 4.
419 Here we measured the change in error sensitivity for all errors (5° to 30°) over three separate periods.
420 From the beginning of the zero-variance period to the end of the zero-variance period (Fig. 5.5E, change
421 from 1 to 80), from the end of the zero-variance period to the end of the high variance period (Fig. 5.5E,
422 change from 80 to 160), and from the beginning to the end of the experiment (Fig. 5.5E, change from 1
423 to 160). In both Figs. 5.5D and 5.5E, we obtained single estimates of error sensitivity, by collapsing
424 across the appropriate 5° error sensitivity bins. When collapsing across bins, we weighted each bin by
425 the number of times errors in that bin were experienced by the subject population, over the
426 corresponding time window.

427

428 5.3 Results

429 To quantify the relationship between mean residual error and perturbation variance, we started by re-
430 analyzing data collected in an earlier study¹⁷² in which the authors perturbed reaching movements in
431 three groups of participants using visuomotor rotations (Fig. 5.1A, Rotation). All groups experienced
432 perturbation sequences that had the same mean (30°), but different levels of trial-to-trial variability: one


433 group experienced a constant perturbation of 30° (zero variance), while the other two groups
434 experienced perturbations with low or high variance (Fig. 5.1B, top). At the end of training, there were
435 residual errors in all groups. However, the residual errors increased with the perturbation variance (Fig.
436 5.1H, Fernandes, median residual error on last 10 trials; repeated measures ANOVA: $F(2,14)=17.8$,
437 $p<0.001$, $\eta_p^2=0.54$).

438 Here we designed a set of experiments to answer two questions: (1) why does adaptation suffer
439 from residual errors and (2) why does the size of the residual errors depend on perturbation variance?
440

441 5.3.1 Perturbation variance limits sensorimotor adaptation

442 We began by repeating the experiment performed by Fernandes and colleagues¹⁷², but with an
443 important adjustment. In the earlier work, all three perturbation conditions were experienced by the
444 same set of subjects, raising the possibility that prior exposure to the visuomotor rotation could have
445 altered subsequent learning in the other environments^{7,30,149}. To avoid this possibility, we recruited
446 different sets of participants for each perturbation condition.

447



448

449 **Figure 5.1.** Perturbation variance impairs sensorimotor adaptation. A. Schematic of our experiment
 450 setup. B. Fernandes and colleagues (8) measured the reach angle of participants (bottom, n=16) during
 451 adaptation to variable visuomotor rotations (top: SD = 0, 4, and 12° for zero, low, and high variance;
 452 mean is 30° for all). Participants demonstrated differing residual errors (reported in inset H, Fernandes;
 453 median error on the last 48 trials). C. In Experiment 1, we repeated the experiment of Fernandes et al.
 454 (2012) with a between-subjects design. Participants adapted to a zero (n=19) or high (n=14) variance
 455 perturbation (SD = 0 and 12° for zero and high variance; mean is 30° for both). The residual error is
 456 shown in H, Exp. 1 (median of the last 48 trials). D. In Experiment 2, we tested force field adaptation.
 457 Occasionally, we measured reaching forces on channel trials that restricted motion of the hand to a

458 straight path. Participants experienced a zero (n=12) or high (n=13) variance perturbation (top: SD = 0
459 and 6 N-s/m for zero and high variance; mean = 14 N-s/m for both). We computed an adaptation index
460 on each channel trial (bottom). Residual error (inset H, Exp. 2) is one minus mean adaptation index on
461 last 5 error clamp trials. E. In Experiment 3, we exposed participants to an extended period of
462 visuomotor rotations (160 epochs = 640 trials). The vertical dashed line indicates the total number of
463 rotation trials in Experiment 1. Participants adapted to a zero (n=19) or high (n=14) variance
464 perturbation (top: SD = 0 and 12° for zero and high variance; mean is 30° for both). Mean residual error
465 (inset H, Exp. 3) was computed over the last 50 epochs. To confirm that performance had reached a
466 plateau, we measured the slope of a line fit to the same period (inset G). For comparison, horizontal
467 dashed lines show the mean slope over the first 5 epochs of the perturbation. F. In Experiment 4, we
468 adapted participants (n=14) to a zero-variance perturbation, and then abruptly switched to a high
469 variance perturbation. Residual errors (inset H, Exp. 4) were computed over the last 10 epochs of each
470 period. Error bars are mean \pm SEM. Statistics denote the result of a repeated-measured ANOVA (H,
471 Fernandes) or two-sample t-tests (H, all other insets). Statistics: **p<0.01 and ***p<0.001.
472

473 In each of our experiments, participants held the handle of a robotic arm (Fig. 5.1A) and made
474 reaching movements in a two-dimensional workspace. In Experiment 1, we introduced a visual
475 perturbation and divided the participants into two groups: a zero-variance group (n=19) in which the
476 perturbation magnitude remained invariant at 30° (Fig. 5.1C, black), and a high variance group (n=14) in
477 which the perturbation was sampled on each trial from a normal distribution with a mean of 30° and
478 standard deviation of 12° (Fig. 5.1C, red). Our results confirmed the earlier observation: after extended
479 training, participants in the zero-variance group learned more than the high variance group (Fig. 5.1C,
480 bottom; Fig. 5.1H, Exp. 1, mean error on last 10 epochs, two-sample t-test, p=0.002; Cohen's d = 1.49).

481 In Experiment 2, we tested the generality of this observation by measuring how participants
482 responded to variability in force field perturbations (Fig. 5.1A, Force field). As before, we divided the
483 participants into two groups, a zero-variance group (n=12) in which the perturbation magnitude
484 remained constant at 14 N·sec/m (Fig. 5.1D, top, black), and a high variance group (n=13) in which the
485 perturbation magnitude was sampled on each trial from a normal distribution with mean 14 N·sec/m
486 and standard deviation of 6 N·sec/m (Fig. 5.1D, top, red). To track the learning process, we measured
487 the forces that the subjects produced via channel trials²⁷ (Fig. 5.1A, channel). As in visuomotor
488 adaptation, variance in the force field perturbation reduced the total amount of learning (Fig. 5.1D,
489 bottom; Fig. 5.1H, Exp. 2, mean error on the last 5 epochs; two-sample t-test, p=0.001; Cohen's d =
490 1.46). Thus, perturbation variability consistently affected the adaptation limit across various modalities
491 of reach adaptation.

492

493 5.3.2 Perturbation variance limits the total extent of adaptation

494 Closer examination of the late stage of training (Figs. 5.1B, 5.1C, and 5.1D, bottom) raised the possibility
495 that adaptation had not completely saturated; perhaps with additional exposure, adaptation might have
496 converged across variance conditions, even eliminating the residual errors. To examine this possibility,
497 we repeated Experiment 1, but this time more than doubled the training trials (Fig. 5.1E, top). Addition
498 of these trials allowed performance to saturate, as evidenced by the slope of the reach angles (Fig. 5.1G,
499 slope of the line fit to individual performance over the last 50 epochs was not different than zero;
500 $p=0.71$ and $p=0.83$ for the low and high variance groups). Notably, despite extended training, errors
501 persisted (Fig. 5.1H, Exp. 3, residual errors \pm SD on last 50 epochs; zero-variance: $1.7 \pm 0.9^\circ$; high
502 variance: $8.7 \pm 1.7^\circ$; t-test against zero; both groups, $p<0.001$). Furthermore, once again we found that
503 increased perturbation variance coincided with an increase in residual error: at the end of adaptation,
504 there was no overlap in the distributions of residual errors in the low and high variance groups (Fig.
505 5.1H, Exp. 3; two-sample t-test, $p<0.001$; Cohen's $d = 5.24$).

506 If perturbation variability causally modulated the asymptotic limit of adaptation, we reasoned
507 that we could switch between two different asymptotic states by changing perturbation variance mid-
508 experiment. To test this prediction, in Experiment 4 participants ($n=14$) first adapted to a zero-variance
509 30° visuomotor perturbation (Fig. 5.1F, black). With training, performance approached a plateau. We
510 next increased the perturbation variance (while keeping the mean constant) by sampling from a normal
511 distribution with a standard deviation of 12° (Fig. 5.1F, red). Coincident with the increase in perturbation
512 variance there was a reduction in reaching angle (Fig. 5.1H, Exp. 4, mean residual error on last 10
513 epochs; two-sample t-test, $p=0.005$; Cohen's $d = 1.16$). Thus, despite having already learned to
514 compensate for much of the perturbation, when the variance of the perturbation increased, the residual
515 error in every subject also increased (Fig. 5.1H, Exp. 4).

516 Together, Experiments 1-4 demonstrated that even after extended practice, motor adaptation
517 reached an asymptotic limit, resulting in small persistent errors. However, this asymptotic limit was
518 dynamic, responding to the second order statistics of the perturbation.

519

520 5.3.3 Perturbation variance impairs the implicit component of learning

521 While reach adaptation can occur despite severe damage to the explicit, conscious learning system of
522 the brain¹⁷⁷, under normal circumstances performance benefits from both implicit and explicit learning
523 systems^{11,12,65,178}. A plausible hypothesis is that when perturbation variance increases, there is
524 impairment in the explicit system, resulting in increased residual error. One way to reduce or eliminate

525 the influence of the explicit system during reach adaptation is to limit the preparation time that people
526 have to initiate their movements^{32,62,132,176}. If perturbation variance impairs explicit learning, we would
527 predict that when reaction time is strongly limited, thus reducing or eliminating the influence of the
528 explicit system, there will be increased residual errors. However, increased perturbation variance should
529 have little or no additional effects on residual errors.

530 Alternatively, perturbation variance may affect predominately the implicit system. In that case,
531 when reaction time is strongly limited, suppressing contributions of the explicit system, there will be
532 increased residual errors. Critically, increased perturbation variance should produce further reductions
533 in residual errors. In this scenario, the change in residual errors following increased perturbation
534 variance should be the same regardless of whether reaction time is limited or not.

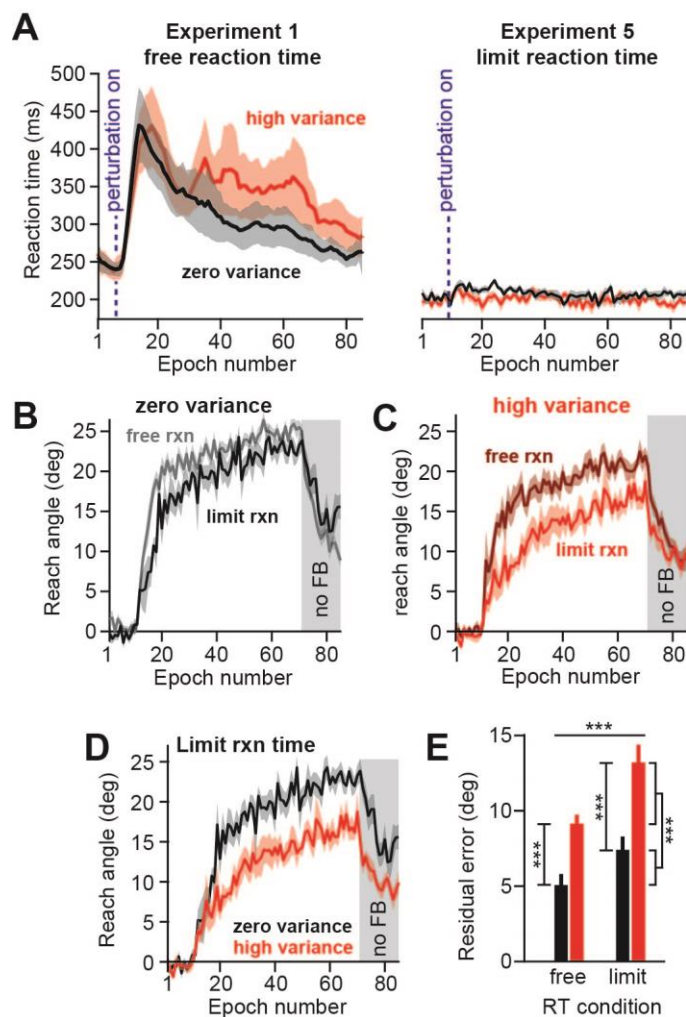
535 To suppress the contributions of explicit learning, we repeated Experiment 1, but forced the
536 subjects to respond to the target with very low reaction times (less than half the time that they typically
537 required to initiate a movement at the start of adaptation, Fig. 5.2A). As before, we divided participants
538 into two groups: a zero-variance group (n=13) in which the perturbation magnitude remained invariant
539 at 30°, and a high variance group (n=12) in which the perturbation was sampled on each trial from a
540 normal distribution with a mean of 30° and standard deviation of 12°.

541 Under normal conditions in which there was no constraint on reaction time (Fig. 5.2A,
542 Experiment 1), introduction of the perturbation led to a dramatic increase in reaction time: participants
543 nearly doubled their preparation time, potentially signaling the development and expression of explicit
544 strategies. In contrast, in the constrained reaction time group (Fig. 5.2A, Experiment 5), subjects
545 executed reaching movements at considerably lower latencies. In this group, the time required for
546 movement preparation remained roughly constant throughout the experiment, even after the
547 introduction of the perturbation.

548 As expected, limiting reaction time impaired adaptation rates. In the zero-variance (Fig. 5.2B)
549 and high variance (Fig. 5.2C) conditions, performance at short reaction times was worse than that of
550 long reaction times (two-sample t-test on last 10 epochs; $p=0.041$ and $p=0.007$ for zero and high
551 variance; Cohen's $d = 0.77$ and 1.17 for zero and high variance). Thus, reducing reaction time impaired
552 performance in both the zero and high variance perturbation groups.

553 Next, we compared the effects of perturbation variance on residual errors at low reaction times.
554 We found that even when reaction times were constrained to suppress the explicit component of
555 adaptation, increased perturbation variance produced a clear increase in the residual errors (Fig. 5.2D),
556 reducing the total extent of learning by approximately 5° (Fig. 5.2E; difference in residual errors during

557 the last 10 epochs; two-sample t-test, $p < 0.001$; Cohen's $d = 1.53$). Therefore, despite suppression of
 558 explicit learning, the deficit in performance caused by high variance perturbations persisted: high
 559 variability led to a reduction in the extent of learning.



560

561 **Figure 5.2.** Perturbation variance decreases the total amount of implicit adaptation. A. We measured
 562 reaction time (time to start the reach in response to presentation of a visual target) during adaptation to
 563 a visuomotor rotation. In Experiment 1 (left) we did not limit reaction time. Reaction time greatly
 564 increased after the introduction of the perturbation (blue dashed line). In Experiment 5 (right) we
 565 repeated the paradigm used in Experiment 1, but limited reaction time. B. Limiting reaction time
 566 reduced adaptation rates in the zero-variance perturbation group. C. Limiting reaction time reduced
 567 adaptation rates in the high variance perturbation group. D. We compared adaptation to the zero-
 568 variance and high variance perturbations when reaction time was limited. Even in the limit rxn
 569 condition, perturbation variance impaired adaptation. E. To measure differences in adaptation extent as
 570 a function of reaction time and perturbation variance, we used a two-way ANOVA. For this test, we
 571 measured the total amount of adaptation over the last 10 epochs of Experiments 1 and 5. The top-most
 572 and right-most statistical bars denote the significance of the main effects of reaction time and
 573 perturbation variance, respectively. The other two vertical statistical bars display the results of two-
 574 sample t-tests. Error bars are mean \pm SEM. Statistics: *** $p < 0.001$.

Critically, a 2-way ANOVA yielded a significant effect of both perturbation variance ($F=33.08$, $p<0.001$, $\eta_p^2=0.38$) and reaction time ($F=13.76$, $p<0.001$, $\eta_p^2=0.20$), but no interaction effect ($F=1.07$, $p=0.31$). The absence of an interaction effect indicated that the reduction in asymptotic performance induced by perturbation variability was no different when reaction time was high or low. If we assume that reduced reaction time impairs the ability of the explicit system to contribute to reach adaptation, then these results suggest that the reductions in performance observed in Experiments 1 and 5 were predominantly caused by impairments in implicit learning.

5.3.4 Perturbation variance reduces error sensitivity, but not forgetting rates

An examination of the data in Figs. 5.1 and 5.2 illustrates a frequently observed characteristic of motor adaptation: even at the late stages of training, performance continues to suffer from small, persistent errors^{10,11,14,116,123,172,179,180}. Such persistent, steady-state errors are one of the fundamental predictions of mathematical models of adaptation^{8,10,36,55,58} in which performance is driven by an interaction between two opposing forces, error-based learning, and trial-to-trial forgetting:

$$x^{(n+1)} = ax^{(n)} + be^{(n)} \quad (5.1)$$

Here the adapted state of the individual, x , changes from trial n to trial $n+1$ due to both forgetting, and learning. Forgetting is controlled by the retention factor a , which determines the fraction of memory retained from one trial to the next. Learning is controlled by sensitivity to error, denoted by b , which determines how much learning will occur from the experience of error e .

In this model, performance reaches a steady state (Fig. 5.3A) in which residual errors persist. This occurs because as training progresses, the errors which drive the learning process eventually become small enough that there is a balance between the forces that promote forgetting, and those that promote learning. At this stage learning appears to stop, despite the fact that errors remain. Thus, in principle, residual errors can increase because of a decrease in error sensitivity (Fig. 5.3C), or an increase in forgetting (Fig. 5.3B). Which factor changes due to increases in perturbation variance?

To answer this question, we estimated the forgetting rate (a in Eq. 5.1) of each participant by including an error-free movement period at the end of the experiments (gray region in Figs. 5.1C, 5.1D, 5.1F, and 5.2D). During these periods, behavior naturally decayed towards the baseline state (Fig. 5.3D), thus providing a means to isolate the rate of forgetting (i.e., the rate of decay of behavior). Interestingly, we found that in all experiments, the rate of forgetting was similar in the low and high variance groups

(Fig. 5.3E, two-sample t-test; Exp. 1, $p=0.72$; Exp. 2, $p=0.19$; Exp. 5, $p=0.79$). Thus, perturbation variance did not alter the rate of forgetting.




Figure 5.3. Perturbation variance decreases error sensitivity, not decay rates. A. State-space model of adaptation predicts that learning will reach an asymptote when the amount of learning from an error exactly counterbalances the amount of forgetting that occurs between trials. The plot demonstrates the behavior of such a model during adaptation to a perturbation of unit 1. According to the model, changes in asymptotic levels of performance can occur because of changes in forgetting (B, Possibility 1 schematic; $a = 0.98$ for low forgetting and 0.96 for high forgetting), or changes in error sensitivity (C, Possibility 2 schematic; $b = 0.05$ for low error sensitivity and 0.1 for high error sensitivity). D. To test Possibility 1, we measured the retention during error-free periods at the end of Experiments 1 (Exp. 1), 2 (Exp. 2), and 4 (Exp. 5). We normalized reach angle to the first trial in the no-feedback period. Decay rates did not differ between the low and high variance groups. Each point on the x-axis is a cycle of 4 trials. E. We measured the retention factor during error-free periods in each experiment. We found no difference in retention for the zero-variance and high variance groups. F. To test Possibility 2, we measured sensitivity to error in each experiment. Error sensitivity was greater for the zero-variance perturbation in every experiment. Error bars are mean \pm SEM. Statistics: * $p<0.05$, ** $p<0.01$, and n.s. indicates no statistical significance.

Next, we empirically estimated error sensitivity (b in Eq. 5.1). To do this, we calculated the difference between the reach angle in each pair of consecutive trials (adjusting for forgetting) and divided this by the error experienced on the first of the two trials. By definition, this quotient represents one's sensitivity to error, i.e., the fraction of the error that is compensated for on the next trial. In sharp contrast to forgetting rates, we found consistent differences in sensitivity to error (Fig. 5.3F) among the various groups: in all experiments, participants who experienced a zero-variance perturbation exhibited

an error sensitivity nearly twice that of individuals who were exposed to a high variance perturbation (two-sample t-test; Exp. 1, $p=0.002$, Cohen's $d = 1.18$; Exp. 2, $p=0.039$, Cohen's $d = 0.87$; Exp. 4, $p=0.006$, Cohen's $d = 1.12$; Exp. 5, $p=0.016$, Cohen's $d = 1.05$).

In summary, we asked whether perturbation variance affected forgetting rates, error sensitivity, or both. We found that in all experiments, the groups that had experienced high perturbation variance exhibited reduced error sensitivity, but no change in forgetting rates.

5.3.5 Perturbation variance reduced the ability to learn from small errors, not large errors

Our quantification of error sensitivity in Fig. 5.3F made the assumption that the brain is equally sensitive to errors of all sizes. However, it is well-documented that error sensitivity varies with magnitude of error: one tends to learn proportionally more from small errors^{51,52,68,148}. In other words, error sensitivity is not constant, but declines as error size increases. How did perturbation variance alter the functional relationship between error magnitude and sensitivity to error?

To answer this question, we re-estimated error sensitivity, but this time controlled for the magnitude of error. We placed pairs of consecutive movements into bins according to the error experienced on the first trial, and then calculated error sensitivity within each bin. As expected, in both perturbation variance conditions, as error size increased, error sensitivity decreased (Fig. 5.4A; mixed-ANOVA, within-subjects effect of error size, $F=24.69$, $p<0.001$, $\eta_p^2=0.30$). This confirmed that indeed, people tended to learn less from larger errors. However, for a given error size, the high variance perturbation group exhibited lower error sensitivity than the zero-variance group (Fig. 5.4A; mixed-ANOVA, between-subjects effect of perturbation variance, $F=10.9$, $p=0.002$, $\eta_p^2=0.16$). Notably, increased perturbation variance reduced the ability to learn from small errors ($<20^\circ$), but had no effect on the larger errors ($>20^\circ$) (Fig. 5.4A, post-hoc testing with t-test adjusted with Bonferroni correction, $p=0.012$ and Cohen's $d = 0.75$ for 5-10°, $p=0.012$ and Cohen's $d = 0.78$ for 10-15°, $p=0.005$ and Cohen's $d = 0.83$ for 15-20°, $p=1$ for other two bins). Why should increases in perturbation variance selectively affect learning from small errors, but not large errors?

5.3.6 The consistency of past errors modulated error sensitivity

A current model of sensorimotor adaptation⁷ posits that the brain adjusts its sensitivity to error in response to the consistency of past errors. In this memory of errors model, when the error on trial n has the same sign as the error on trial $n+1$, it signals that the brain has undercompensated for error on trial n , and so should increase sensitivity to that error (Fig. 5.4B, left). Conversely, when the errors in two

consecutive trials differ in sign, the brain has overcompensated for the first error, and so should decrease sensitivity to that error (Fig. 5.4B, right). Herzfeld and colleagues demonstrated that these changes in error sensitivity are local to specific errors, meaning that the brain can simultaneously increase sensitivity to one error size, while decreasing sensitivity to another⁷. Thus, in the context of a variable perturbation, the memory of errors model provides an interesting prediction: the perturbation affects the trial-to-trial consistency of errors, producing less consistency for some error sizes, particularly when perturbation variance increases.




Figure 5.4. The consistency of trial-to-trial errors coincides with changes in error sensitivity. **A.** To determine how error sensitivity varied as a function of error size, we sorted pairs of movements into different bins (5° width) according to the size of the error on the first movement. Next, we computed the mean error sensitivity across all trials within each error size bin. In the zero-variance environment, subjects exhibited greater error sensitivity to error sizes that were less than 20°. **B.** We considered the possibility that the trial-to-trial consistency of errors caused changes in error sensitivity. Consistent errors (left) are consecutive pairs of trials where the errors have the same sign. Inconsistent errors (right) are consecutive pairs of trials where the errors have opposite signs. The black and brown traces show example reach trajectories from a single participant. **C.** We measured the total fraction of inconsistent error trials. The high variance perturbation caused a higher probability of inconsistent errors in every experiment. **D.** The difference between the number of consistent and inconsistent errors during adaptation to the visuomotor rotation. We found that high variance caused a change in consistency of small errors (less than 20°) but not large errors. Error bars are mean \pm SEM. For **A** and **D**, we used a mixed-ANOVA followed by post-hoc two-sample t-tests with Bonferroni corrections. In **C**, two-sample t-tests were used for statistical testing. Statistics: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ and n.s. indicates no statistical significance.

We tested the predictions of this model by evaluating all consecutive trial pairs for each subject, and counting all instances of consistent and inconsistent errors. As expected, we found that in all five experiments, increased perturbation variance produced an increase in the probability of experiencing an inconsistent error (Figs. 5.4C; Exps. 1 & 3, $p=0.017$, Cohen's $d = 0.68$; Exp. 2, $p<0.001$, Cohen's $d = 2.84$; Exp. 4, $p<0.001$, Cohen's $d = 2.22$; Exp. 5, $p=0.048$, Cohen's $d = 0.84$). However, when we binned the data based on error size, the differences in the relative number of consistent and inconsistent errors exhibited a surprising pattern (Fig. 5.4D, mixed-ANOVA, between-subjects effect of perturbation variance, $F=26.04$, $p<0.001$, $\eta_p^2=0.31$; within-subjects effect of error size, $F=61.02$, $p<0.001$, $\eta_p^2=0.51$): for errors less than 20° , the zero-variance group had many more consistent error events and fewer inconsistent error events than the high variance group (Fig. 5.4D; post-hoc testing with t-test adjusted with Bonferroni correction, $p<0.001$ and Cohen's $d = 1.99$ for $5-10^\circ$, $p<0.001$ and Cohen's $d = 1.47$ for $10-15^\circ$, $p=0.003$ and Cohen's $d = 0.78$ for $15-20^\circ$). However, for errors larger than 20° , there was no difference in the relative number of consistent and inconsistent error pairs (Fig. 5.4D; post-hoc testing with t-test adjusted with Bonferroni correction, $p=1$ for $20-25^\circ$ and $p=0.161$ for $25-30^\circ$).

In summary, the high variance perturbation led to a reduction in the trial-to-trial consistency of small errors, but not large errors (Fig. 5.4D). The memory of errors model (Fig. 5.4B) predicts that this pattern of error history will lead to a reduction in sensitivity to small errors for the high variance condition, but no difference in sensitivity to large errors. Indeed, both of these patterns were present in the measured behavior (Fig. 5.4A).

5.3.7 Error sensitivity changes throughout training according to the consistency of error

To determine if error sensitivity could truly have been altered by a memory of past errors, we next simulated a memory of errors model. This model expresses how error sensitivity varies both as a function of time, and error size. That is, on trial n , the sensitivity to error e , is expressed as $b^{(n)}(e)$. Suppose that on trials $n-1$ and n , one experiences errors $e^{(n-1)}$ and $e^{(n)}$. The memory of errors model predicts that sensitivity to the first error will either increase or decrease, depending on the consistency of the two errors:

$$\Delta b^{(n+1)}(e^{(n-1)}) = \alpha \Delta b^{(n)}(e^{(n-1)}) + \beta \text{sign}(e^{(n)} e^{(n-1)}) \quad (5.2)$$

In the above equation, Δb represents the change in error sensitivity. Like Eq. 5.1, changes to error sensitivity are due to two forces: one that depends on consistency of consecutive errors (contained

within the sign function), and one that imposes a decay on error sensitivity (parameter α). Thus, the model has only two unknown parameters, α and β .




Figure 5.5. The memory of errors model predicts that error sensitivity changes over the course of training. We simulated a memory of errors model that increases/decreases error sensitivity from one trial to the next when a consistent/inconsistent error occurs. A. We simulated the response of this model to the zero-variance perturbation in Exp. 3 (left, Exp. 3, zero var.), the high variance perturbation in Exp. 3 (middle, Exp. 3, high var.), and the zero-to-high variance perturbation in Exp. 4 (right, Exp. 4). The same parameter set was used for each simulation. The model predictions represent the mean model prediction over 100,000 simulations. B. Next, we took the specific errors experienced by each

participant and used the memory of errors model to predict trial-by-trial changes in error sensitivity. At left, the plot shows the predictions for the zero-variance group in Experiments 1 and 3 (epochs 1-60 represent both experiments, but from 61 onwards only Experiment 3). The same is true for the middle plot, but for the high variance group predictions. At right, the plot shows predictions for the zero-to-high variance group in Experiment 4. C. To corroborate the model predictions, we measured error sensitivity on each trial as a function of error size. Then, we binned data across sequences of consecutive trials. Missing data points in each curve represent periods over which errors were scarce or entirely absent. D. The memory of errors model predicted that error sensitivity should increase for both the zero-variance and high variance perturbation, but less so for the high variance condition. Light gray bars show the predicted change in error sensitivity. Dark gray bars show the actual change in error sensitivity measured in the subject population (difference in error sensitivity from epochs 1-15 to epochs 40-59). Predictions are split for small (less than 20°) and large (greater than 20°) errors. E. In Experiment 4, the model predicts error sensitivity should first increase in the zero-variance perturbation, and then decrease after adding variance to the perturbation. Light gray bars show the predicted change in error sensitivity. Dark gray bars show the actual change in error sensitivity measured in the subject population over three separate periods (change from 1 to 80 shows the change from the beginning to the end of the zero-variance period; change from 80 to 160 shows the effect of introducing variance to the perturbation; change from 1 to 160 shows the change in error sensitivity from the very start of the experiment, to the very end of the experiment). All bars include errors in the range of 5° to 30°. Error bars are mean \pm SEM. Statistics: * $p < 0.05$, ** $p < 0.01$, and *** $p < 0.001$.

To find these two parameters, we simulated the response of Eqs. 5.1 and 5.2 to the zero and high variance perturbation patterns we used in the various experiments (see Methods for a more complete model description; all other model parameters apart from α and β were obtained empirically from the data as in Figs. 5.3E and 5.3F). We varied our parameter set to identify a single pair of parameters (α and β) that best accounted for the measured reach angles in Experiments 3 and 4, in a least squares sense. The resulting parameter set yielded predictions that closely tracked subject behavior (Fig. 5.5A).

Having estimated the two unknown parameters of the model, we next applied Eq. 5.2 to predict how error sensitivity should have changed if the subjects stored a memory of past errors (Fig. 5.5B). To do this, we applied Eq. 5.2 to the actual sequence of errors experienced by each subject ($e^{(1)}$, $e^{(2)}$, etc.). Thus, the model predicted how error sensitivity should develop for each error size. To corroborate the model predictions, we measured error sensitivity on each trial (Fig. 5.5C).

We focused on Exp. 1 and 4, as these experiments included error-free periods, permitting empirical measurement of both error sensitivity and forgetting (note the difference in the ranges of the x-axes of Figs. 5.5B and 5.5C in the left and middle columns). In addition, we omitted measurements of error sensitivity over periods of the experiment where the number of observed data points was low or absent entirely (note gaps in time-courses in Fig. 5.5C left and right).

The model made the following predictions: First, the model predicted that with the zero-variance perturbation, error sensitivity should increase generally for all error sizes, but the extent of this increase should be greater for smaller errors (Fig. 5.5B, left). To test this prediction, for each subject we measured the change in error sensitivity across the first 60 epochs of Experiment 1. In the zero-variance group, error sensitivity indeed increased over the training period as predicted (Fig. 5.5D, left, zero variance, small errors. We could only estimate the change in sensitivity to small errors, because there were virtually no large errors at the end of the learning period).

Second, the model predicted that if perturbation variance was initially low, but then increased in the middle of adaptation (as in Exp. 4), error sensitivity should initially increase but then decline causing an increase in the residual errors (Fig. 5.5B, right). To test these predictions, for Exp. 4 we measured error sensitivity at the start of adaptation, at the end of the zero-variance period, and at the end of the high variance period. In accordance with model predictions, error sensitivity increased substantially over the initial zero-variance period (Fig. 5.5E, change from 1 to 80; paired t-test, $p=0.003$, Cohen's $d = 0.96$) and then dropped precipitously after increasing perturbation variance (Fig. 5.5E, data, change from 80 to 160; paired t-test, $p=0.044$, Cohen's $d = 0.60$).

Third, and most surprisingly, the model predicted that even if participants were exposed to a high variance perturbation from the onset, error sensitivity should increase for all errors, but only by a small amount (Fig. 5.5B, middle). This final prediction was interesting because it implied that the residual errors in the high variance perturbation were larger not because of a decrease in error sensitivity, but instead because trial-by-trial increases in error sensitivity were stunted relative to the zero-variance group (compare Figs. 5.5B left and middle). Indeed, we found that error sensitivity also increased for the subjects in the high variance group (Fig. 5.5D, high var.), but the magnitude of this increase was smaller relative to the zero-variance group (Fig. 5.5D, two-sample t-test, $p<0.001$ and Cohen's $d = 1.65$ for small-small error comparison, $p<0.001$ and Cohen's $d = 1.91$ for small-large error comparison). The same was true for Exp. 4, when comparing error sensitivity at the end of the high variance period, to the initial error sensitivity at the start of the zero-variance period (Fig. 5.5E, change from 1 to 160; paired t-test, $p=0.042$, Cohen's $d = 0.60$).

In summary, the memory of errors model made the surprising prediction that error sensitivity should generally increase in the early phase of training, more so for small than large errors, and then saturate or even decline as the training trials continued. It also predicted that introducing variance into the perturbation should not decrease error sensitivity, but rather stunt its growth. Our measurements confirmed these predictions.

Overall, variance in the perturbation altered the sequence of trial-to-trial errors. A memory of these errors modulated error sensitivity, which in turn, resulted in different asymptotic limits to adaptation.

5.3.8 Asymptotic differences are caused by changes in implicit, not explicit adaptation

In Section 5.3.3, we investigated if differences in asymptotic performance were caused by the implicit or explicit component of adaptation. By limiting reaction time, we found that differences in asymptote persisted, consistent with modulation in implicit adaptation. However, we cannot be completely certain that the implicit process also differed in the primary experiments in which reaction time was free. In addition, our analysis did not allow us to make direct conclusions about the contributions of explicit learning to the total extent of adaptation. To address both of these concerns, we performed a control experiment. Here, we did not manipulate reaction time and exposed participants to a noise-free ($n=9$) and a high-noise ($n=9$) perturbation as in our primary experiments. As in Fig. 5.1, we observed that adding variability to the perturbation reduced the total extent of adaptation (Fig. 5.6A), resulting in an increase in the residual error that persisted at the end of the learning period (Fig. 5.6B, two-sample t-test, $p<0.001$).

Next, unlike our primary experiments, we probed the amount of implicit and explicit adaptation. To do this, we instructed participants to reach to each of the 4 targets, but this time there would be no cursor, no feedback, and no perturbation, and that their goal was to move their physical hand straight through the target on the screen (Fig. 5.6A, middle schematic). The reach angle we measured on these probe trials would serve as the final level of implicit adaptation. To obtain an estimate of the explicit aiming angle, we subtracted this implicit measure from the reach angle measured on the standard trials at the end of adaptation.

As a result of this new instruction, the reach angle immediately dropped precipitously from its pre-probe level (note the reach angle shown in the gray probe region in Fig. 5.6A). Critically, the difference between the no noise and high noise conditions persisted on these probe trials (Fig. 5.6C, left; two-sample t-test, $p=0.023$), indicating that differences in implicit adaptation actively contributed to the change in asymptotic performance. On the contrary, the change in reach angle corresponding to the explicit component of adaptation (Fig. 5.6C, right) did not differ across the no noise and high noise conditions (two-sample t-test, $p=0.688$).

To confirm these results, we performed another probe after the implicit reach angle instruction. Here we went through each of the 4 targets again (Fig. 5A, right-most schematic), and asked the

participant to indicate the location they believed they were aiming towards in order to move the cursor through the target at the end of the adaptation period. To assist this measurement, we displayed small circles radially around the workspace at a spacing of 3.333° , along with a colored character. Remarkably, we again observed that participants in the no noise and high noise reported no differences in their perceived explicit reach angle (Fig. 5.6D, two-sample t-test, $p=0.453$), confirming our findings.




Figure 5.6. Perturbation variability reduces the total amount of implicit, but not explicit adaptation. A. Two sets of participants adapted to a rotation with the same mean, but different levels of trial-to-trial variance. At the end of the perturbation period, we probed the level of implicit and explicit adaptation (gray probe period). The first probe (middle schematic) consisted of instructing participants to aim their hand and not the cursor straight through the target. The second probe consisted of asking participants to report the location they aimed their hand to score a point (explicit probe, right-most schematic). B. The difference between the perturbation mean and the reach angle on the last 10 epochs of the perturbation period. C. At left we show the reach angle that persisted after participants were instructed to move their hand through the target. At right we show the drop in reach angle from the end of the adaptation period to the implicit probe trials. D. Here we show the reported reach angle, averaged across the 4 targets. Statistics: n.s. indicates $p>0.5$, $*p<0.05$, $*p<0.01$, and $***p<0.001$.

Altogether, these results indicate that perturbation variance altered asymptotic performance through changes in implicit adaptation alone. To understand why we observed a change in implicit adaptation, when so many others have found that this process has an invariant response to error, we considered the idea of competition. When a learner experiences a perturbation that causes an error,

both implicit adaptation and explicit adaptation compete with one another to learn from error. Therefore, it is not always possible to know if one measures a change in one process, if that change is caused by a shift in the underlying dynamics of that process, or the parallel competitive process.

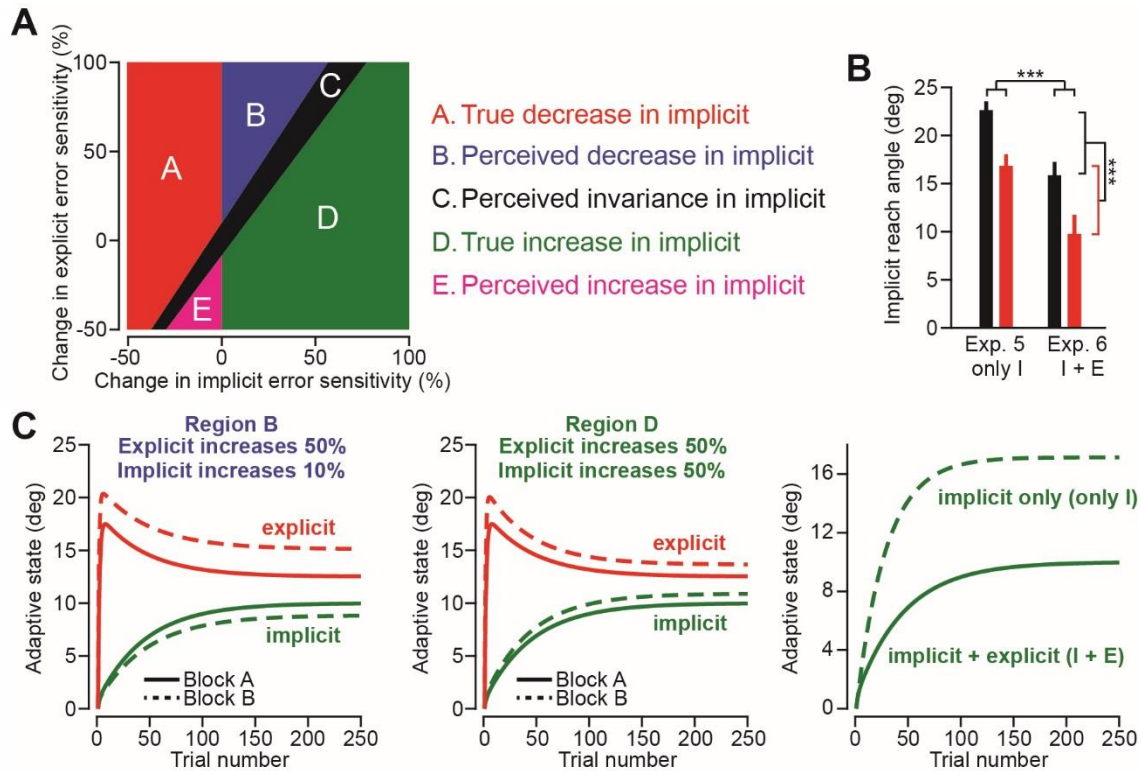


Figure 5.7. Implicit and explicit adaptation compete for error. A. Here we imagine that implicit and explicit learning are measured in two different blocks. Across the blocks there is a change to the error sensitivity of the implicit and explicit adaptation. On the y-axis we show the change in explicit error sensitivity. On the x-axis we show the change in implicit error sensitivity. The colors indicate how the implicit state would change at asymptotic performance. The red region A shows instances where the implicit process becomes less sensitive to errors and also decreases its asymptotic contributions. The blue B region shows instances where the implicit process actually become faster, but still has a smaller asymptotic contribution. The black C region represents where the amount of change in the implicit asymptotic contribution has changed by an absolute magnitude of less than 5%. The green D region shows instances where the implicit process is faster and also contributes more to asymptotic performance. Finally, the magenta E region shows instances where the implicit process slows down, yet still increases its contributions to asymptotic performance. B. Here we show the two-way ANOVA for a perturbation variance by learning type comparison. Experiment 5 shows the amount of implicit learning measured when reaction time is limited and there is no explicit adaptation. Experiment 6 shows the amount of implicit learning probed at the end of adaptation where both explicit and implicit processes learn from error. C. The left-most schematic shows an example of the blue B region in inset A. The middle schematic shows an example of the green D region in inset A. For both the left and middle insets, the change in the rate of explicit adaptation is the same. For the right-most inset, we show the effect of removing explicit learning on asymptotic performance of implicit adaptation.

To illustrate this point, consider Fig. 5.7. Here we simulated a two-state model of learning⁷² assuming that implicit adaptation can be represented by a slowly adapting system (retention is 0.97, and error sensitivity is 2%) and explicit adaptation can be represented by a quickly adapting system (retention is 0.85 and error sensitivity is 30%). Now imagine we measure implicit learning and explicit learning at asymptote in two different blocks (as occurs in anterograde interference, savings, and the experiments reported herein). To compare implicit and explicit learning across these two blocks, we created an error sensitivity map (Fig. 5.7A). Along the y-axis we show the percent change in the error sensitivity of the explicit system and along the x-axis we show the percent change in the error sensitivity of the implicit system. The colors divide the map into different regions. Regions A and D are the simplest to understand. Here the implicit process has a decrease in error sensitivity (Region A) or an increase in error sensitivity (Region D) and consequently would be observed to have a smaller or larger magnitude at asymptote respectively. However, Regions B and E could be misleading to the observer. In Region B, the implicit process becomes more sensitive to error, yet has a *decrease* in its asymptotic performance. This decrease is caused by competition with explicit learning, which has an increase in error sensitivity that dwarfs that of implicit learning. In Region E, the opposite occurs. The implicit process lowers its error sensitivity, yet has a larger asymptotic value. This is because the explicit system experiences a more drastic slowing, allowing the implicit system to experience larger errors.

All-in-all this error sensitivity map demonstrates that for two processes that compete with one another, measuring a change in one process, does not necessarily reveal to you how that process has changed. To reiterate once further, consider the simulations in Fig. 5.7C. At left, we show changes in explicit learning and implicit learning, when explicit error sensitivity has increased by 50%, and implicit adaptation increases by 10%. Even though implicit adaptation becomes more sensitive to error, it decreases its asymptotic magnitude, corresponding to region B in Fig. 5.7A. In the middle section, explicit learning changes by the same magnitude, but implicit learning changes by 50%. Here, we find that the change in error sensitivity of implicit learning is reflected by its increase in contributions to asymptotic performance (Region D in Fig. 5.7A). Therefore, if explicit adaptation becomes faster by 50%, we cannot know by measuring the asymptotic amount of implicit adaptation, whether the implicit system has also become more sensitive to errors or not.

The best that can be done to determine how an intervention affects a particular learning process is to isolate it. If this learning process competes with others, eliminating its competitors will cause it to increase its overall dynamics (Fig. 5.7C). In support of this idea consider Fig. 5.7B; here we show implicit learning measured under limited reaction time (Experiment 5) versus in the implicit probe

trials at the end of learning in Exp. 6, where both explicit and implicit systems compete for error. As predicted, elimination of explicit learning, through limiting reaction time on all trials, causes the implicit system to increase its total contributions (two-way ANOVA, effect of experiment, $p < 0.001$). The same phenomenon is shown in Fig. 5.7C at right.

This framework might provide a way to unify our results with previous observations. In earlier work, the implicit process has been found to be inflexible, and invariant in its response to error^{31,32,52,64,66–69}, whereas here we find that the response of the implicit system to be quite responsive to the perturbation environment. We speculate that in these earlier studies, the implicit system did in fact change its response to error, but its modulation was dwarfed with respect to that of explicit adaptation, which is quite fluid and flexible in the context of visuomotor adaptation. This is consistent with Region C of Fig. 5.7A, where implicit adaptation changes its response to error, but such a change will not be discovered if one were to measure its magnitude alongside that of the explicit system. Therefore, we predict that limiting reaction time to isolate the implicit response will indeed demonstrate that implicit, like explicit systems, also modulate their response to error, as we have demonstrated in our experiments.

5.4 Discussion

Adaptation exhibits a curious property: even after prolonged training, learning appears to stop, leaving behind small persistent residual errors^{9,123,142,171}. Here we demonstrated that this asymptotic limit of adaptation is consistent with a learning system in which sensitivity to error is not constant, but rather changes as a function of the history of past errors. When the past errors are temporally consistent, which occurs when perturbation variance is low, error sensitivity rises quickly. However, as performance improves and errors become smaller, error sensitivity peaks and then declines. The history-dependent change in error sensitivity, which in turn is driven by the variance of the perturbation, appears to play a causal role in setting an upper bound on adaptation. Thus, errors are not simply used for learning and then discarded, but rather they appear to become a form of memory that describes how much the nervous system should learn from that error.

5.4.1 Controlling the total extent of adaptation through modulation of error sensitivity

Earlier studies that documented modulation of error sensitivity focused on understanding the rate of motor adaptation^{124,125} in the context of savings^{7,30,35,36,49,61,171}, meta-learning^{7,49}, and anterograde interference⁵⁰. Here we describe yet another feature of adaptation that appears to be controlled

through error sensitivity: the total extent of motor adaptation. In Experiments 3 and 4 (Fig. 5.1) we demonstrated that differences in residual errors were coupled to differences in error sensitivity, but not the rate of forgetting. The balance between error sensitivity, which changed with perturbation variance, and forgetting, which remained constant, created different levels of asymptotic performance.

This observation raised an interesting question: why didn't the brain continue to increase its error sensitivity until residual errors were completely eliminated? Here we proposed a potential solution: increases in error sensitivity are bounded by decay. Our model (Eq. 5.2) accounts for saturation in error sensitivity in the same way the state-space model (Eq. 5.1) accounts for saturation in adaptation: through forgetting. The counterbalancing of a memory of consistent errors (which increases error sensitivity) with decay (which decreases error sensitivity) leads to steady-state saturation in error sensitivity.

The addition of forgetting to the original memory of errors model⁷ offers a potential explanation for at least three separate behavioral observations. First, when a participant is exposed to the same perturbation twice, they exhibit faster adaptation during the second exposure⁸. To account for this savings, the memory of errors model predicts that error sensitivity is up-regulated during the first exposure to the perturbation, leading to faster learning in the future. However, in some cases, long periods of washout or error-free trials appear to prevent the occurrence of savings^{34,61}. This observation is consistent with Eq. 5.2: during sufficiently long periods of washout, error sensitivity would decay towards baseline due to forgetting, thereby preventing savings upon the next exposure to the perturbation.

Another peculiar feature of savings is the observation that the mean residual error is often indistinguishable during the first and second blocks of exposure to a perturbation, even though the initial rate of learning is faster during the second block^{34,36,54}. Why doesn't the faster rate of learning during the second exposure also lead to a higher extent of learning? Our model provides an explanation: because error sensitivity tends to saturate during the initial exposure to the perturbation, behavior will reach the same asymptotic level during the second exposure despite presence of savings and higher initial rates of learning.

Finally, experiments that employ a constant error-clamp condition also demonstrate saturation in error sensitivity. In these experiments^{52,53,68,170,181} subjects are exposed to the same error time and time again, irrespective of their motor output. Without decay, a memory of errors model would predict that error sensitivity, and thus the extent of learning, would increase without bound. However, in reality, learning reaches a steady-state in this constant error condition. In the context of visuomotor

adaptation, an earlier report speculated that this saturation of learning implies a limitation in the amount of adaptation that can be stored within implicit motor learning systems⁵². However, more generally, this saturation in adaptation is also predicted by Eq. 5.2 due to the eventual balancing of learning and forgetting.

5.4.2 Plasticity in implicit learning systems

Motor adaptation is known to be supported by both implicit (subconscious) and explicit (cognitive) corrective systems^{11,12,65,177,178}. Recent studies have interrogated the flexibility of these learning systems, with some^{31,32,64,69,182} suggesting that implicit processes have a response to error that does not change with multiple exposures to a perturbation.

While our primary purpose was not to carefully tease apart implicit and explicit contributions to adaptation, we did explore this issue by limiting reaction time, which is thought to substantially reduce the contributions of explicit learning^{32,62,132,176}. When reaction time was limited, the residual errors were not only present, but increased by the same amount if perturbation variance increased. This observation suggested that perturbation variance acted on the implicit learning system. These results argue against the idea that implicit learning systems generate inflexible and unchanging responses to sensorimotor errors.

This finding does not discount the possibility that explicit systems also contribute to residual errors. For example, recent work has demonstrated that declines in explicit learning associated with aging^{63,183} also lead to differences in the total extent of adaptation¹⁸⁴. It is tempting to suggest that these differences in asymptotic performance may also reflect a change in error sensitivity, like the implicit mechanism described herein. While this possibility remains untested, such an interpretation is consistent with recent work demonstrating that environment variability modulates explicit error sensitivity during adaptation to random walk perturbations¹⁸².

In summary, our data support the more inclusive view that with experience, both implicit and explicit processes change the way they respond to error, and together determine the total extent of sensorimotor adaptation.

5.4.3. Alternate models

Perturbation variance can affect uncertainty of the learner. Over the past few decades, numerous studies^{55,125,175} have used a Kalman filter⁶⁰ to study the relationship between uncertainty and learning rate. The Kalman filter describes the optimal way in which an observer should adjust their rate of

learning in response to different sources of variability. This Bayesian framework has proved useful in understanding the slowing of adaptation in response to reductions in the reliability of sensory feedback^{56,57,59}, acceleration of adaptation in response to uncertainty in the state of the individual or environment^{55,56,175}, and even the optimal tuning of adaptation rates in individual subjects⁵⁸.

Could this Bayesian framework also account for our results? Indeed, it is possible that the inconsistent errors generated by the high variance perturbation are interpreted as an observation noise, thus decreasing the brain's confidence in its sensory feedback. In the Kalman framework, this would correctly predict that adaptation to the high variance perturbation would proceed more slowly than adaptation to the low variance perturbation (Figs. 5.1-2).

With that said, two of our behavioral observations do not easily fit into the Kalman framework. First, as expected, we observed that in both the zero-variance and high variance groups, error sensitivity declined as a function of error size^{51,52,68,148}. However, across the variance conditions, differences in error sensitivity were specific to small errors, not large errors (Fig. 5.4A). The memory of errors model⁷ provided a way to understand this pattern: differences in perturbation variability led to changes in the consistency of small errors, but not large errors (Fig. 5.4D) matching the pattern we observed in error sensitivity. It is unclear how to account for this phenomenon with a Kalman filter which responds to all error sizes with the same learning rate.

The second observation that does not easily comply with the Kalman view is the fact that sensitivity to error increased over its baseline level with repeated exposure to the high variance perturbation (Fig. 5.5D, high var.; Fig. 5.5E, change from 1 to 160). For a Kalman filter, the Kalman gain (i.e., rate of learning) would necessarily drop from its baseline level (where there is no variance) when the uncertainty in its observations increases (during the perturbation block where variance is added). This prediction of the Kalman framework is opposite to our data. In contrast, the memory of errors model (Figs. 5.5B-E) correctly predicts an increase in error sensitivity, because the experience of consistent errors remains more probable than inconsistent errors even in the high variance condition (Figs. 5.4C and 5.4D).

5.4.4 Neural basis of error sensitivity

Motor adaptation depends critically on the cerebellum^{2,4,47,116,139}, where Purkinje cells learn to associate efference copies of motor commands with sensory consequences²⁵. This learning is guided by sensory prediction errors, which are transmitted to the Purkinje cells via the inferior olive, resulting in complex spikes. Notably, plasticity in Purkinje cells exhibits both sensitivity to error, and forgetting. Sensitivity to

error is encoded via probability of complex spikes: in each Purkinje cell, the probability of complex spikes is greatest for a particular error vector^{25,185}, and this tuning modulates learning from error. Forgetting is present in the time-dependent retention of the plasticity caused by the complex spikes^{22,118}, resulting in decay of plasticity with passage of time. Thus, the presence of error sensitivity and forgetting in the plasticity of Purkinje cells provides one mechanism by which cerebellar-dependent adaptation can exhibit asymptotic performance with non-zero residual errors.

Our behavioral experiments demonstrated that during adaptation, error sensitivity increased with training, but this increase was suppressed if the perturbations were variable (Figs. 5.5D and 5.5E). Our model (Figs. 5.5B and 5.5D) suggested that these changes in error sensitivity may have arisen from the temporal consistency of errors such that the presence of two errors in the same direction would result in increased sensitivity. Thus, the theory makes the interesting prediction that in the framework of cerebellar learning, the temporal proximity of complex spikes might modulate error sensitivity. Specifically, when two consecutive errors of similar direction and magnitude occur, Purkinje cells that prefer that error are more likely to experience complex spikes in close temporal proximity. The model predicts that the temporal proximity of these complex spikes would result in up-regulation of error sensitivity in that Purkinje cell. This idea remains to be tested.

In this framework, Eq. 5.2 serves as a simple abstraction of a more complicated reality. That is, in contrast to the binary output of the sign function in Eq. 5.2, the consistency of error should vary along a spatial continuum of preferred and non-preferred errors, as well as a temporal continuum that separates error events. These predictions provide a platform for future experiments, both in the context of behavior and neurophysiology.

In summary, our work suggests that over the course of training, the brain relies on a memory of past errors to adjust error sensitivity. These variations in error sensitivity produce an asymptotic limit that prevents further improvements in performance, thus resulting in persistent residual errors.

Chapter 6. Postural control of arm and fingers through integration of movement commands

Every movement ends in a period of stillness. Current models assume that commands that hold the limb at a target location do not depend on the commands that moved the limb to that location. Here, we report a surprising relationship between movement and posture in primates: on a within-trial basis, the commands that hold the arm and finger at a target location depend on the mathematical integration of the commands that moved the limb to that location. Following damage to the corticospinal tract, both the move and hold period commands become more variable. However, the hold period commands retain their dependence on the integral of the move period commands. Thus, our data suggest that the postural controller possesses a feedforward module that uses move commands to calculate a component of hold commands. This computation may arise within an unknown subcortical system that integrates cortical commands to stabilize limb posture.

6.1 Introduction

To hold the limb still, the muscles are not quiet. Rather, they are actively engaged with coordinated inputs that maintain postural stability. Current models assume that these inputs are produced by an impedance controller that translates the sensory representation of a desired location to patterns of muscle activity^{186–188}. To move and then hold, one feedback controller generates the commands that move the limb^{189,190}, and a separate controller generates commands that hold the limb still following movement¹⁸⁶. This architecture (Fig. 6.1A, left) in which movement and postural controllers are independent is implicit in optimal control formulations of reaching^{186,187,189,191} and forms the basis for interpreting how neurons in the motor cortex encode reach kinematics¹⁹². While many predictions of this theory have been confirmed for control of movement¹⁸⁹, here we provide evidence that challenges the assumption that posture and movement are controlled independently.

Our idea starts with consideration of a simpler control system: the eye and the head. In order to hold the eyes at a target, the oculomotor system uses a hold controller whose output directly depends on the move controller^{1,13,86,88,89,91,94,193}. The move controller produces a set of commands that displace the eyes¹. Simultaneously, these commands are integrated in real-time by a distinct brainstem structure, yielding sustained commands that hold the eyes and the head still when the movement ends^{86,193}. Thus, the architecture assumed for move and hold controllers of the arm (Fig. 6.1A, left) is not consistent with that of the eye^{1,13} and the head^{8,11,14,15}.

While we do not know if the output of the reach controller serves as an input to the hold controller, there is evidence that moving and holding are controlled by separate neural structures; total

inhibition of the mouse motor cortex during reaching causes the arm to stop moving, but the muscles continue to receive commands that hold the arm steady against gravity⁹⁷. It is difficult to reconcile this observation with the idea that the cortex drives both moving and holding^{194,195}.

Yet, there are also reasons to doubt that the neural control of the arm shares a design principle with control of the eyes and the head. The physical dynamics of the arm are much more complicated than the eye, casting doubt that any straightforward relationship might exist between commands that move the arm to a location, and the commands that subsequently hold the arm there. Furthermore, whereas damage to the brainstem structure that holds the eyes produces nystagmus^{90,196}, we know of no condition that resembles nystagmus in the context of reaching.

We began by asking a simple question: are the commands that hold the limb at the target solely determined by the target position (Fig. 6.1A, left), or are they dependent in part on the preceding move commands (Fig. 6.1A, right)? We began by measuring activity across arm muscles during point-to-point reaching. When monkeys reached to a single target from various directions, we found that the integral of reach activity predicted hold activity after the movement ended. Furthermore, as the target location varied, the same integration function accounted for hold period activity at the various endpoints. Thus, across a range of directions and endpoints, the hold period activity was related, through integration, with the preceding move period activity.

To ask whether this pattern held across other types of movements, we considered goal-directed finger movements in which the start and end target locations were kept constant. As monkeys flexed their finger, there was natural variability in muscle activity both during the movement and during the ensuing hold period. However, this variability had structure: on a within-trial basis, changes in hold period commands in all recorded muscles were correlated with the integral of the preceding commands that had brought the finger to its current location.

These patterns revealed a correlation, not causation. To test whether there might be a causal link between movement and holding, we imposed a change to the commands that moved the arm to a given target location, and quantified whether on a within-trial basis, the change in move period commands influenced the subsequent hold period commands. To do this, we altered the reach commands of humans through adaptation⁸¹. To measure the properties of the hold controller, we designed a procedure in which we slowly displaced the hand during the hold period while subjects were engaged in a working memory task. We recorded the forces produced by the hand in response to the involuntary displacements during the hold period, thus measuring the postural field that held the arm still. We found that as the reach period commands changed, the entire postural field shifted, indicating

that changes to reach commands altered the null point of the hold controller. Notably, the function that related the hold period controller to the preceding reach period was the same integration function that we had observed in point-to-point reaching and finger movements.

Finally, we probed the neural circuits that might support this putative link between movement and holding by examining reaching in patients who had suffered damage to their corticospinal tract (CST) above the level of the brainstem. As expected, these stroke survivors exhibited large trial-to-trial variability in the commands that they produced during both the reach and the subsequent hold periods. Remarkably, the link between move and hold period commands appeared intact: on a within-trial basis, the hold period forces were related via a form of integration to the immediately preceding, but now imperfect, reach period forces. Thus, in monkeys, healthy humans, and stroke survivors, across arm movements and finger movements, the hold period commands depended on the preceding commands that had moved the limb to its current location. These results raise the possibility that the postural controller possesses a subcortical feedforward module that calculates hold period commands through real-time integration of the move period commands. This feedforward computation then combines with visual and proprioceptive feedback to produce the sustained commands that result in postural stability.

6.2 Methods

Here we divide our methodology into three sections. Section A pertains to a theoretical derivation of the integration equation. Section B pertains to non-human primate data. Section C pertains to human data.

6.2.A The reach integrator hypothesis

Muscles are engaged during the period of moving as well as the subsequent period of holding still. We hypothesized that for arm and finger muscles, activity during the movement $u(t)$ could be decomposed into contributions from a move controller $m(t)$ and a hold controller $h(t)$:

$$u(t) = m(t) + h(t) \quad (3)$$

If the move and hold controllers are connected in series, like the control system present for the eye and the head, the hold controller produces its output by integrating in real-time the output of the move controller. An example of this is shown in Fig. 6.1B, in which EMG of ant. deltoid is decomposed into move and hold commands. Here, we explain this decomposition and derive some of its predictions.

For a movement from position h_1 to position h_2 the hold controller first produces commands for holding at h_1 and then transitions to holding commands at h_2 through integration of $m(t)$. Here is an integration function that could accomplish this task:

$$h(t) = u(h_1) + k \int_{t_0}^t m(\tau) d\tau \quad (4)$$

In the above equation, $u(h_1)$ represents hold activity at position 1. Combination of Eqs. 3 and 4 yields:

$$u(t) = m(t) + u(h_1) + k \int_{t_0}^t m(\tau) d\tau \quad (5)$$

Given the measurement $u(t)$, e.g., EMG from a given muscle, we can decompose the measurement into its constituents $m(t)$ and $h(t)$ by solving the above equation iteratively for $m(t)$. This is what we did to plot the traces in Fig. 6.1B (we assumed $k=1$).

To evaluate the predictions of this equation, we measured movements of duration T , between two positions h_1 and h_2 . At the end of the movement the move commands go to zero yielding:

$$u(h_2) = u(h_1) + k \int_{t_0}^T m(\tau) d\tau \quad (6)$$

Unfortunately, we do not have an *a priori* estimate of the integration gain k . Thus, during movement we cannot uniquely estimate $m(t)$ and $h(t)$. To proceed, we made a simplifying assumption: the move commands could be approximated by taking the overall EMG signal $u(t)$ and subtracting off the hold commands measured at the start of movement:

$$u(h_2) \approx u(h_1) + k \int_{t_0}^T [u(\tau) - u(h_1)] d\tau + a \quad (7)$$

Here we added a bias term, a , to account for systematic error introduced by our approximation. Rearranging the terms in Eq. (7) yields the prediction of the hypothesis:

$$u(h_2) - u(h_1) \approx k \int_{t_0}^T [u(\tau) - u(h_1)] d\tau + a \quad (8)$$

This states that the change in hold period activity should be roughly a linear function of the integral of muscle activity during the preceding movement with a gain k . Thus, from this approximation we arrive at Eqs. (1) and (2) in the main text.

6.2.B Monkey experiments

We performed four sets of experiments: (1) reaching movements in non-human primates, (2) finger movements in non-human primates, (3) reaching movements in healthy humans, and (4) reaching movements in stroke patients.

6.2.B.1 Reaching movements

Two monkeys participated in a reaching task described elsewhere^{197,198}. Briefly, at the start of each trial the monkey held its hand at a central home location. After a “GO” cue, the monkey reached to one of eight peripheral targets displayed in the vertical plane on a monitor (Fig. 6.1B), and then held its arm at the target for at least 0.5 sec. After this holding period, the monkey returned its arm to the central home location and held it there for at least 0.5 sec, until the start of the next trial.

On each trial, hand position was recorded through infrared optical tracking of a bead fixed to the third and fourth fingers. The activity of several muscles was also recorded using intramuscular electrodes. In Monkey A, these muscles included the anterior, medial, and posterior deltoid, the medial and lateral bicep, and the upper and lower trapezius. From Monkey B, these muscles included the anterior, medial, and posterior deltoid, the pectoralis, the brachialis, the medial and outer bicep, and the upper and lower trapezius. EMG signals were filtered (10-500 Hz), digitized at 1 kHz, rectified, and smoothed with a Gaussian kernel (standard deviation of 20 ms). Muscle activity was then averaged across movements towards each target, separately. We normalized the data by setting to 0 the average muscle activity at the center hold location, and setting to 1 its maximum activity in the task.

We asked whether activity in a given muscle during the hold period could be related to its activity during the preceding reach period (Eq. 1). For our analysis, the pre-movement hold activity $u(h_1)$ was quantified as the mean activity in the [-700, -350] ms period relative to movement onset for outwards reaches, and [-300, -200] ms period relative to movement onset for return reaches. The post-movement holding activity $u(h_2)$ was quantified as the mean activity over [+300, +450] ms period after movement termination for both outwards and return movements. We started this interval 300 ms after movement offset to allow time for muscle dynamics to settle. And finally, the bounds for integration (i.e., t_0 and T in Eq. 1) were set as 140 ms before movement onset up until movement offset. We started integration prior to movement onset to capture changes in muscle activity that preceded change in kinematics.

Of the 320 different movement types (20 muscles x 2 reaches per trial x 8 targets), 6 movements (1.9% of trials) had reach durations that were too slow to gain an accurate measurement of holding activity prior to the start of the next movement. To identify these trials, we used a cutoff for movement duration of 850 ms. The 6 trials with movement durations that exceeded 850 ms were not included in our analysis.

We fit the integration parameters, k and α , in Eq. (1) in the least-squares sense. To determine if the integration gain differed for outward and return movements (Fig. 6.1G), we fit the outwards and

return movements separately, and then tested for a difference in integration gain with a paired t-test. To determine if integration gain differed for movements of different durations, for each movement type (16 possible movements, 8 directions for outwards and return movements), we sorted movements into fast and slow, based on the median movement duration. Because muscles were recorded on different sets of sessions, the average movement durations for each muscle differed. For the slow and fast movement groups, we fit the integration parameters in Eq. (1), and tested for a difference in integration gain with a paired t-test.

To determine how gravity impacted the gain of integration, we separated movements into two groups: (1) the horizontal movements to and from Targets 2 and 7 in Fig. 6.1B, and (2) the other six movements (all contained a vertical component). Because the horizontal group groups consisted of only 4 movements (2 out and 2 back) for each muscle, we used a different technique to test for differences in integration gain. We collapsed data across all muscles and movements in the horizontal group, and then separately in the vertical group. We then fit integration parameters in Eq. (1). To test if there was a difference in the integration gain (i.e., the slope of the linear regression) we tested for a group by move period interaction effect on the hold period activity, using an ANCOVA.

We also tested if the maximum activity during moving (as opposed to the integral) was an equally good predictor of holding activity. For this, we quantified either the max activity (if the muscle increased activity) or min activity (if the muscle decreased activity) over the entire movement. We then regressed holding activity onto the max or min activity, for each muscle separately. We compared the R^2 of this fit to that of the integral fit using a paired t-test (Fig. 6.1F, right).

6.2.B.2 Finger movements

To examine the within-trial covariance between the commands generated during the movement period and the subsequent holding period, we considered finger movements¹⁹⁹. Two monkeys (R and D) used their index finger to track a visual target (Fig. 6.2A). The index finger was splinted within a narrow plastic tube, constraining movements to the metacarpophalangeal joint. The hand on the recording arm and all other digits were placed in a padded pocket which prevented movement. The recording arm was placed in a sleeve to prevent movement. The contralateral arm was not restrained.

On each trial the target moved between two positions at a speed of 12 deg/sec. On flexion trials the target moved from 12° to 24°. On extension trials, this order was reversed. The finger flexion movements were resisted with a spring load that measured 0.026 N·m at 12° and 0.048 N·m at 24°. Each trial started with a rapid movement to the start position. The monkey then held its finger at the start

location (hold 1) for 1 sec, and then made a slow tracking movement to the target over a 1 sec interval (e.g., Fig. 6.2B). The trial ended with a 1 sec hold at the target (hold 2).

We recorded activity of 9 muscles of the finger and forearm (1DI, AbPL, EDC, ECR, ECU, FDS, FDP, FCR, and FCU) on each trial and on each monkey (e.g., Fig. 6.2C). Muscle activity was measured using subcutaneous electrodes. In order to analyze single trial activity, we first rectified the data, and then smoothed it with a Gaussian kernel (standard deviation of 200 ms). After this, we normalized the data by setting to zero the average muscle activity at the 12° hold period, and setting to one its maximum activity during movement.

We used Eq. (1) to test whether there was a within-trial relationship between move period and the subsequent hold period activity of each muscle. We focused on flexion movements in which the finger moved against the external load. For this analysis, pre-movement holding activity $u(h_1)$ was the mean activity over a 400 ms interval starting 1 sec prior to movement onset. The post-movement holding activity $u(h_2)$ was the mean activity over a 200 ms period starting 700 ms after movement termination. And finally, the bounds for integration (i.e., t_0 and T in Eq. (1)) were set as 150 ms before movement onset up until movement offset.

We considered data across all muscles and trials ($n=6070$ for Monkey R, $n=2796$ for Monkey D). We then fit the integration parameters, k and a , in Eq. (1) in the least-squares sense, for each monkey separately. We quantified the variance in holding activity accounted for by the integral of moving activity. For visualization, we also computed the 95% confidence ellipses that describe the joint distribution of holding force and the integral of moving force across all trials (Fig. 6.2E).

We compared integration to an alternate hypothesis: fluctuation in the final holding activity was caused by fluctuation in the initial holding activity (Fig. 6.2D). This is described by the linear equation $u(h_2) \approx ku(h_1) + a$. To test the alternate hypothesis, we collapsed data across all muscles and trials and regressed hold 2 activity onto hold 1 activity, and computed the variance accounted for (i.e., R^2). For visualization we also computed the confidence ellipses that describe the joint distribution of hold 2 activity and hold 1 activity.

We also considered the possibility that the observed correlations between move and hold period activity could be spuriously generated by the physical constraints required to move and hold against the spring load. That is, larger movements would require more work against the spring, leading to larger displacements that in turn would require greater hold force. This alternate hypothesis hinges on two relationships. One, trial-by-trial changes in hold position must correlate with trial-by-trial

changes in hold EMG activity. Two, trial-by-trial changes in the finger displacement (i.e., movement size) must correlate with the integral of moving EMG activity. To test these possibilities, we used linear regression. First, we regressed muscle activity during the target hold period onto the position of the finger during the hold period. In addition, we also regressed the change in hold activity from the start location to the target location, onto the change in finger position across these two periods (i.e., the total displacement of the spring). Finally, we also regressed the integral of move period activity onto the change in finger position. In all cases, we performed separate regressions for each muscle, and then averaged the resulting R^2 values across muscles.

Note that there is a critical difference between regressing hold activity on hold 2 position alone versus regressing hold activity onto the combination of the hold 1 and hold 2 conditions. Regression within, but not across conditions, is the appropriate way to measure the trial-by-trial correlation between EMG and position without the influence of spurious across-condition correlations²⁰⁰. If hold 1 and hold 2 are combined into the same regression (thus doubling the number of data points), these combined measurements account for $42 \pm 4\%$ of the variance in EMG activity in individual muscles, as compared to approximately 2% for hold 2 activity alone. The inflated variance accounted for is caused by the large separation between the associated distributions for hold 1 and hold 2 (see Fig. 6.3C at right for an example), not by a strong coupling between EMG and position *within* any of these distributions. Therefore, we were careful not to collapse across both hold 1 and hold 2 when quantifying trial-by-trial correlation between EMG and finger position, to avoid these spurious correlations²⁰⁰. With that said, we mention the combination of hold 1 and hold 2, to confirm that the spring load does indeed require a substantial modification in the EMG activity of individual muscles, even though trial-by-trial variations in EMG activity at the hold 2 position are only very weakly related to position (see the main text).

Finally, we considered the possibility that the observed correlations between move and hold period activity could have been caused by trial-to-trial fluctuations in co-contraction. If on some trials, the finger was stiff and co-contracted its muscles, and on other trials less so, we may observe correlations between move and hold activity, without any relation to hold position. To change finger stiffness, agonist and antagonist pairs of muscles would exhibit correlated increases or decreases in activity. In other words, we should be able to predict the hold period activity in one muscle not solely based on its move period activity (as in the integration hypothesis) but also on the activity of simultaneously recorded muscles. To test this idea, we regressed hold period activity of each muscle onto the integral of move period EMG in other muscles. For a given muscle, we performed this regression separately for all of its possible pairs, calculated the R^2 value for each pair, and then averaged across pairs to obtain a

single R^2 value per muscle. We report the average R^2 value across all muscles. For context, we compare this R^2 value with that obtained by regressing hold period activity in a muscle onto the integral of move period activity in that same muscle.

6.2.C Human experiments

All human subject experiments were approved by the Institutional Review Board at the Johns Hopkins School of Medicine. Our healthy human cohort consisted of $n=223$ individuals. Healthy participants ranged from 18-61 years of age (mean \pm SD, 25.2 ± 7.9) and included 128 males and 95 females.

ID	Age	Sex	Time since stroke	Handedness	Paretic arm	FMA (/66)	ARAT (/57)	CST lesion location
S001	80	M	2y	Right	Left	57.5	57	Right internal capsule
S002	51	M	6y	Right	Left	40	47.5	Right fronto-parietal white matter
S003	68	F	7y	Right	Right	34.5	19	Left corona radiata
S004	30	F	5y	Right	Left	55.5	43.5	Right precentral and postcentral gyri
S005	78	M	13mo	Right	Right	43.5	34	Left corona radiata
S007	54	F	2mo	Left	Right	63	57	Left corona radiata
S008	53	F	14mo	Right	Left	41	25	Right centrum semiovale
S010	70	M	5y	Right	Left	20	12	Posterior limb of right internal capsule
S011	43	F	20mo	Right	Right	64	57	Left corona radiata
S012	48	M	6y	Right	Left	18.5	6.5	Right corona radiata
S013	68	M	9y	Right	Left	14	8	Right internal capsule
S014	45	F	16mo	Right	Left	40	39.5	Right precentral and postcentral gyri
S015	64	F	10y	Right	Left	22	4.5	Right corona radiata
S016	38	F	21mo	Both	Right	62.5	57	Left corona radiata

Table 6.1. Measures of impairment in stroke patients. Patients completed the Fugl-Meyer Assessment (FMA) and the Action Research Arm Test (ARAT), as well as strength testing in the shoulder and elbow. Strength measurements were repeated twice and the maximal force was recorded on each effort and then averaged across repetitions. Two separate raters scored the FMA and ARAT assessments, and scores were averaged across raters. Missing entries in table indicate that the patient was unable to perform the desired action. Patients were selected based on MRI or CT scans, and/or available radiologic reports. Scans and/or reports were corroborated to determine the level at which the white matter of the corticospinal tract (CST) was lesioned. Here we provide the level of the brain at which the white matter was lesioned.

Our stroke patient cohort consisted of $n=14$ adults that had suffered damage to the corticospinal tract (CST). The stroke patients ranged from 30-80 years of age (mean \pm SD, 56.4 ± 15.2) and included 6 males and 8 females. For comparison, we recruited a cohort of healthy age-matched controls who ranged from 28-81 years of age (mean \pm SD, 60.6 ± 16.3) and included 5 males and 5 females. There was no significant difference in age between the patient and older healthy control populations (2-sample t-test, $p=0.53$).

The stroke patients we recruited had survived a stroke affecting cortical or subcortical white matter associated with the CST. Patients were selected based on MRI or CT scans, and/or available radiologic reports. Scans and/or reports were corroborated to determine the level at which the white matter of the corticospinal tract (CST) was lesioned. Table S1 provides the level of the brain at which the white matter was damaged.

We measured the degree of motor impairment in the patients using the Fugl-Meyer Assessment (FMA) and the Action Research Arm Task (ARAT). Two separate raters scored these assessments in each patient. In each limb, we measured the strength of elbow flexion and extension and shoulder horizontal adduction and abduction using a dynamometer (microFet 2). During measurements, participants rested their arm on a side table so the arm rested slightly below shoulder level. Strength measurements were repeated twice, the maximal force was recorded on each effort, and forces were averaged over repetitions. FMA scores, ARAT scores, strength, and other patient characteristics are reported in Table S1. Missing entries in table indicate that the patient was unable to perform the desired action.

6.2.C.1 Overview of human reaching experiments

In all our human experiments (healthy participants and stroke patients), participants held the handle of a planar robotic arm (Fig. 6.3A) and made point-to-point reaching movements between targets in the horizontal plane. For stroke patients and age-matched controls, the arm was supported by a frictionless air-sled. In addition, both the paretic (contralateral to lesion) and non-paretic (ipsilateral to lesion) arms were tested. For all other participants, the subject supported the weight of their own arm, and only the dominant arm was tested.

As the subject held the robot handle the forearm was obscured from view by an opaque screen. An overhead projector displayed a small white cursor (diameter = 3mm) on the screen that tracked the motion of the hand. Visual feedback of the cursor was provided continuously throughout the entirety of the testing period, except where otherwise noted. Throughout testing we recorded the position of the robot handle using a differential encoder with submillimeter precision. We also recorded the forces

produced on the handle by the subject using a 6-axis force transducer. Data were recorded at 200 Hz. Except where otherwise noted, kinematic time series were aligned to the onset of movement at the time point where hand velocity crossed a threshold of 1 cm/s. At trial onset, a circular target (diameter = 1 cm) appeared in the workspace, coincident with a tone that cued reach onset. After stopping the hand within the target, a holding period of various durations (1.8 to 6.5 seconds) ensued where subjects were instructed to continue holding the handle within the target. After this holding period, a random inter-trial-interval sampled uniformly between 0.3 and 0.4 seconds elapsed prior to the start of the next trial.

At the end of each reach, coincident with the start of the holding period, movement timing feedback was provided. If the reach was too fast (or too slow), the target turned red (or blue) and a low tone was played. If the reach fell within the desired movement interval the target “exploded” in rings of concentric circles, a pleasing tone was played, and a point was added to a score displayed in the upper-left-hand corner of the workspace. For stroke patients and age-matched controls, the desired movement duration was 600-800 ms. For all other participants this interval was 450-550 ms. Participants were instructed to obtain as many points as possible throughout the experimental session.

In all human reaching experiments, trials were ordered in pairs of outwards and backwards movements. In other words, each pair started with a reach from a start position to a target (outward reach). The subject then held the arm still at the target position, and then reached back to the start position (the backward reach). Only outwards movements were analyzed here. All backwards movements were performed in a channel, or a partial channel condition (described below).

6.2.C.2 Measurement of moving and holding forces in human subjects

At regular intervals throughout each experiment (generally every 5th outwards trial) we measured forces in a channel trial²⁷. On these trials, the motion of the handle was restricted to a linear path connecting the start and target locations (Fig. 6.3A). To restrict hand motion to the straight-line channel trajectory, the robot applied perpendicular stiff spring-like forces with damping (stiffness = 6000 N/m, viscosity = 250 N·s/m). This condition maintained the hand in equilibrium along the axis perpendicular to movement. Therefore, the force applied by the robot was equal and opposite to the lateral force applied by the subject, thus serving as a precise measurement of lateral reaching forces. Before analyzing these forces offline, we first subtracted the baseline force from all force time series. We obtained the baseline force by averaging the forces recorded on the channel trials within the null field period at the start of each experiment.

One of the primary objectives of our study was the comparison of moving and holding forces on channel trials. In each of our experiments, the hand remained in the channel during both the moving period and the holding period (a period of time of at least 1.8 sec after the reach ended) to allow us to measure both moving and holding forces.

Our primary hypothesis was that the holding forces could be described as an integral of moving forces according to Eq. (2). In this equation, $F(t)$ refers to the channel forces exerted by the subject during movement. The quantities $F(h_1)$ and $F(h_2)$ refer to the forces applied by the hand while holding still at the start position prior to the reach (1) and the target after the reach (2). The parameters k and a refer to the integration gain and offset.

We calculated the initial holding force, $F(h_1)$, as the mean force over a 100 ms period starting 500 ms prior to reach onset. The final holding force, $F(h_2)$, was calculated as the mean holding force over a 100 ms interval starting 900 ms after reach termination. The termination of the reach was determined based on a velocity threshold of 3.5 cm/s. The integral of reach forces was computed over the entire movement duration (from movement onset to movement offset). Movement onset was determined based on a velocity threshold of 1 cm/s.

We conclude this section with a critical point. Holding forces would tend to move the hand off the target during periods of holding still. We hypothesized that these departures might result in adaptation of postural control. For this reason, we wanted to prevent this unwanted motion. On channel trials, we prevented this by keeping the hand in the channel, as described above. However, on all other trials, we also wanted to prevent motion of the hand during the holding process. Therefore, for all outwards movements not performed in the channel (this does not apply to the return movements), we applied a two-dimensional clamp to the hand at the end of the reaching movement (Fig. 6.3A). This clamp prevented motion of the hand during the hold period, despite any forces the participant might have applied to the handle, and was programmed as a “well” within the target location that attracted the hand in two dimensions, with stiff spring-like mechanics (stiffness = 4000 N/m, viscosity = 75 N·s/m). The target-hold well was applied when the hand entered the target location and the hand velocity fell below a threshold value of 3.5 cm/s.

And finally, to make sure that holding forces on outwards movements did not affect the initial motion of the hand on the following backwards movement trial, all return movements were performed in a partial channel. The channel was removed after the hand had traveled 40% of the desired movement amplitude. Therefore, the hand terminated at the start position without any external forces.

6.2.C.3 Working memory task

In some of our experiments, we employed a cognitive task to distract participants during measurement of holding forces. The working memory task consisted of a modified 2-back task where subjects were randomly shown an integer between 1 and 4. The integers appeared one at a time so that the next integer replaced the previous integer on the screen (Fig. 6.5A). Participants were told to determine if the integer on the display matched the integer shown two numbers in the past. If the integers matched, participants verbally responded with the keyword “same”. If the integers did not match, participants verbally responded with the keyword “different”. If the response was correct a pleasant tone was played and a point was added to the experiment score. If the response was incorrect a low pitch tone was played and no point was awarded. To confirm that participants were engaged in the cognitive task, we recorded each correct and incorrect response. Participants were clearly engaged in the cognitive task and responded correctly to $91.8 \pm 0.6\%$ of items correctly, at rates of approximately 0.77 ± 0.6 items per second.

6.2.C.4 Reaching movements in the null field

A total of 220 healthy subjects participated in these experiments. The general structure of the task is described in the previous two sections. Almost all movements were performed in a null field, i.e., the subjects freely reached between the start and target positions. At the end of these null field movements, the robot applied a target-hold well during the holding period, as described in a previous section. On some trials we applied a channel and measured moving and holding forces as described in a previous section. All return movements were performed in a channel, or partial channel as described in a previous section.

This process was the same for healthy subjects and stroke patients. The only differences were that we tested both arms of the stroke patients (paretic then non-paretic), supported the weight of the arm of the stroke patients and their age-matched controls with a frictionless air sled, and allowed for differences in movement timing are described above.

We measured the relationship between the move period forces and the change in holding forces as described in C2. Specifically, we measured the within-trial relationship between these two quantities (Fig. 6.3B right and Fig. 6.3C for healthy controls; Fig. 6.7F, left panel and Fig. 6.7H for stroke patients). At various points throughout the manuscript, we also report the integration gain, or in other words, the slope of the linear relationship between hold forces and the integral of reach forces. We

compared the gain of integration for the control group in Fig. 6.7 to the experiment groups in Fig. 6.3, to test if adding weight support altered the integration gain (Fig. 6.3B, bottom-right).

It is important to note that we combined several datasets for our analysis in Fig. 6.3C. While the structure of each dataset was the same, they differed in the kinematics of the reaching movement. Across these tasks, we varied several parameters of the movement, including the target location (center of the body, to the left of the body, and to the right of the body), the reaching direction (towards the body, away from the body, or at an oblique angle), and the reach amplitude (10 cm or 20 cm). Most subjects reached between the same two locations, but for some subjects, there were two potential targets for each trial. The number of reaching trials performed varied across tasks. They ranged from 40 trials to 288 trials (half were outwards movements; the other half were backwards movements).

In Fig. 6.3D, we tested to see if the integration gain differed for fast and slow movements. For each participant, we sorted their channel trial reaching movements in the null field period according to their duration. Then, we selected the two fastest and two slowest movements for each subject (Fig. 6.3D, left, individual points) and combined these movements across all subjects. Then we fit Eq. (2) to both the slow movement and fast movement distributions. We tested for differences in integration gain by reporting the move duration by move force integral interaction effect on hold force within an ANCOVA.

6.2.C.5 Reaching movements in a velocity-dependent force field

The serial architecture between moving and holding (Fig. 6.1A) makes a strong prediction: external adaptation of moving activity will lead to changes in holding still. To test this idea, we used a force field paradigm. We gradually adapted reaching movements of a set of participants ($n=32$) to a force field that exerted forces on the hand that were perpendicular and proportional to its velocity according to:

$$\begin{bmatrix} f_x \\ f_y \end{bmatrix} = b \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix} \quad (9)$$

Here, v_x and v_y represent the x and y velocity of the hand, f_x and f_y represent the x and y force applied to the hand by the robot, and b represents the magnitude of the force field (in units of N·s/m). When $b > 0$, this corresponds to a clockwise (CW) field, and when $b < 0$, this corresponds to a counterclockwise (CCW) field.

In the experiment shown in Fig. 6.3, participants were exposed to both CW and CCW force fields while making 10 cm movements. The experiment was structured so that the force field magnitude would start at zero (a null field trial) and then gradually increase to its maximum strength over many

trials. After this, the magnitude would then be reduced back to zero over many trials. The exact perturbation schedule for healthy subjects is shown in Fig. 6.3D. The experiment started with 40 trials (20 outwards movements and 20 backwards movements) of null field trials. Next subjects were adapted to either a CW or CCW force field. The magnitude of CW/CCW perturbation was increased/decreased from 0 to 15/-15 N·s/m over the course of 100 outwards reaching trials (200 trials in total). The perturbation magnitude was then maintained at a constant level of 15/-15 N·s/m over the course of 50 outwards reaches (100 trials total) and then brought back to zero gradually in a de-adaptation period of 100 outwards reaching trials (200 trials total). After this de-adaptation period, participants continued to reach in a washout period of 20 outwards reaches (40 trials total) where no force field was applied. Participants were then given a short break and this structure was repeated (either for the same force field, but a different target position (n=17), or the opposite force field and the same target position (n=15)).

The experimental protocol was nearly identical (Fig. 6.3D) for our stroke patient experiments, with two small differences. First, trial counts differed. Adaptation and de-adaptation periods were 160 trials, as opposed to 200 trials. The period of maximal perturbation magnitude was reduced from 100 trials to 80 trials. The second difference is that we increased the maximal force field magnitude to 18 N·s/m. We increased this magnitude to compensate for slower movements (we required 600-800 ms movement duration for stroke patients and their age-matched controls, but 450-550 ms in other experiments). Each arm was tested in the stroke experiments in four blocks: paretic, non-paretic, paretic, non-paretic. The perturbation magnitude went in an A-B-B-A order (where A and B refer to either a CW or CCW field). The arm was supported by a frictionless air sled for both patients and age-matched controls.

For both healthy subjects and stroke patients, we measured the relationship between the move period forces and the change in holding forces at regular intervals throughout the process of adaptation and de-adaptation. To do this, every 5th outwards reaching movement was performed in a channel. We measured the within-trial relationship between moving and holding forces (Eq. 2).

We performed a control experiment (n=7 subjects, Fig. 6.9) to determine if holding forces were stable over longer periods of time. In this experiment, the length of the hold period was increased from approximately 1.8 to 6.5 sec. To keep subjects engaged over this period of time, subjects were exposed to a working memory task during the holding period of channel trials as described in C3. The experiment was otherwise similar to the other tasks described in this section.

6.2.C.6 Reaching movements in a zero integral force field

We considered an alternate hypothesis: holding forces could be a trivial continuation of moving forces (as opposed to an integral of moving forces). To test this idea, we designed two position-dependent force fields, A and B, with the latter integrating to zero (Fig. 6.4).

Subjects reached between two targets separated by 20 cm. To form a zero integral force field, we designed a perturbation with two components, FF_1 and FF_2 . Each perturbation produced force along the x-axis, while movements were made along the y-axis. The first component FF_1 was applied during the first 10 cm of the reach, and the second component FF_2 was applied during the last 10 cm of the reach. Each component was programmed as a quadratic function of position. For FF_1 zero force was applied at the start position and at 10 cm. The maximal force was reached at 5 cm (the vertex of the first parabola). For FF_2 zero force was applied at 10 cm and at the target position. The maximal force was reached at 15 cm (the vertex of the second parabola). Here we refer to the magnitude of FF_1 and FF_2 as the maximal force of each perturbation. To obtain a zero integral force field, FF_1 and FF_2 produced forces in opposite directions. The experiment started with 25 outwards (50 trials total) null field trials (FF_1 and FF_2 were both equal to zero). Then we gradually increased the magnitude of FF_2 while FF_1 remained at zero (Fig. 6.4B, Phase 1). FF_2 was increased from 0 N to 3.5 N in even increments, over the course of 100 outwards reaching movements (200 trials total). Then Phase 2 of the experiment started. In this phase, FF_2 was maintained at 3.5 N on all trials, while FF_1 was gradually changed. The magnitude of FF_1 was decreased from 0 N to -3.5 N over the course of 200 outwards reaching trials (400 trials total). In this way, at the end of the experiment, participants were exposed to two force fields within the same reaching movement that perturbed the hand in opposite directions but with equal magnitude. Throughout this paradigm, we measured forces on every 5th outwards reach in a channel.

We found that holding forces gradually decreased in Phase 2 of the experiment, consistent with our hypothesis of integration. To make sure that the introduction of FF_1 caused this decrease in holding force, as opposed to repetition or fatigue, we performed a control experiment ($n=11$, Fig. 6.4A). In the control experiment, Phase 1 was identical to the main experiment described above. However, during Phase 2, the magnitude of FF_1 was maintained at 0 and the magnitude of FF_2 was maintained at 3.5 N. All other details were identical.

6.2.C.7 Measurement of the postural field

In order to hold the limb still, the postural controller generates a converging field of position-dependent forces that counters unwanted displacement of the limb^{201,202}. In some of our experiments (Fig. 6.5) we

set out to measure this field. To do this, we designed a postural probe (Fig. 6.5A). On a postural probe trial, the robot moved the hand slowly in a random direction after the hand stopped within the target. As the hand was moved, visual feedback of hand position was prevented: the display cursor was frozen at the holding location.

To quantify the output of the holding controller, we measured the forces the subject applied to the handle while the robot moved the hand. To prevent participants from voluntarily opposing the imposed hand displacement, we distracted each participant with the working memory task described in an earlier section. We did not inform participants as to the nature or presence of the postural perturbation. Instead, we instructed participants to solely concentrate on the working memory task and obtain as many points as possible by answering memory questions correctly. Points for correct responses were combined with the points awarded for successful reaching movements.

The postural probe consisted of a straight-line displacement designed to make the probe as imperceptible as possible. To move the hand, we placed the hand in a two-dimensional clamp with stiff spring-like mechanics (stiffness = 4000 N/m, viscosity = 75 N·s/m) and moved the equilibrium position of the clamp through the workspace a total of either 2.5 cm, 4 cm or 5 cm, depending on the trial. The imposed motion consisted of three phases.

In Phase 1, the hand was moved a short distance (0.15 cm for 2.5 cm probes, 0.15 cm for 4 cm probes, and 0.3 cm for 5 cm probe) along a minimum jerk trajectory, over a short duration (0.75 seconds for 2.5 cm probes, 0.75 seconds for 4 cm probes, and 1.5 seconds for 5 cm probes). At the end of this displacement the velocity of the hand was equal to 0.375 cm/s. In Phase 2, the hand was then moved at this constant velocity for a specified displacement (2.2 cm for 2.5 cm probes, 3.7 cm for 4 cm probes, and 4.4 cm for 5 cm probes). This constant velocity displacement lasted for 5.87 seconds for 2.5 cm probes, 9.87 seconds for 4 cm probes, and 11.73 seconds for 5 cm probes. In Phase 3, the hand was slowed to rest over a short distance (0.15 cm for 2.5 cm probes, 0.15 cm for 4 cm probes, and 0.3 cm for 5 cm probe) along a minimum jerk trajectory, over a short duration (0.75 seconds for 2.5 cm probes, 0.75 seconds for 4 cm probes, and 1.5 seconds for 5 cm probes). Finally, an additional buffer period of 0.3 seconds was added after reaching the final displaced position, prior to the end of the probe trial. The total duration of the probe was therefore 7.67 seconds for 2.5 cm probes, 11.67 seconds for 4 cm probes, and 15.03 seconds for 5 cm probes.

Critically, as stated earlier, the participant was not provided position feedback during the postural probe. Instead, cursor feedback of hand position was frozen at the target. Therefore, at the end of the postural probe, there was a discrepancy between the location of the hand and the location of

cursor feedback. To seamlessly reunite the hand with its cursor feedback without drawing the attention of the participant, we manipulated visual feedback during the following reaching trial; as the next reach was executed, we projected the cursor position onto the line connecting the frozen cursor position and the position of the next target. In this way, it appeared to the participant as if they were reaching perfectly straight between the start and target position. At the same time, we confined the motion of the hand to a straight line connecting its displaced position with that of the next target. When the hand entered the target, a small and brief force pulse was applied to move the hand to the center of the target at which point x and y feedback was reunited with the true hand position.

6.2.C.8 Quantifying the null point and shape of the postural field

As the hand of the participant was moved by the robot during postural probe trials, the displacement of the hand was opposed by postural restoring forces^{201,202} (Fig. 6.5B). To mathematically characterize the two-dimensional field of restoring forces (i.e., the postural field), we fit a simple mathematical model²⁰¹ that treated the arm as a linear two-dimensional spring with a single equilibrium point:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = K \begin{bmatrix} x - x_{null} \\ y - y_{null} \end{bmatrix} \quad (10)$$

where F_x and F_y are the forces applied to the handle due to displacement of the hand from the null point of the system (x_{null}, y_{null}) to some position (x, y) . The stiffness matrix K ,

$$K = \begin{bmatrix} k_{xx} & k_{xy} \\ k_{yx} & k_{yy} \end{bmatrix} \quad (11)$$

describes the magnitude and orientation of the stiffness field. We constrained K to be a symmetric matrix (i.e., $k_{xy} = k_{yx}$). We fit this linear spring model to the postural restoring forces by identifying the parameter set (5 free parameters, x_{null} , y_{null} , k_{xx} , k_{xy} , k_{yy}) that minimized the sum of squared error between the hand forces (collapsed across the x and y axes) predicted by Eq. (10) and the hand forces measured during all of the postural probe trials. For this fit, we used the forces measured within the ellipse bounded by -2.25 to 2.25 cm along the x -axis and -1.5 to 1.5 cm along the y -axis, relative to the final hand position. To locate the optimal parameter set, we used the genetic algorithm in MATLAB R2018a. We repeated the genetic algorithm search 8 times and selected the one that minimized the squared error cost function. The optimal parameter set provided a good fit to the data, accounting for approximately 70% of the variance in the observed postural field (R^2 during baseline period, mean \pm SEM: 0.70 ± 0.03 ; R^2 after adaptation: 0.69 ± 0.02).

To summarize the shape of the postural field, we considered three properties: (1) its null point, (2) its orientation, and (3) its stiffness. The null point was equivalent to x_{null} and y_{null} . To calculate the orientation of the field, we considered the eigenvector of the stiffness matrix K corresponding to the largest eigenvalue of K . We calculated the angle of this eigenvector in the x - y plane. To compute the stiffness of the field, we calculated the Frobenius norm of the stiffness matrix K .

6.2.C.9 Measuring how changes in movement forces alter the postural field

To determine if changes in reaching forces altered the postural field, we measured the postural field before and after adaptation to a velocity-dependent force field. Subjects ($n=27$) were adapted to a CCW velocity-dependent force field. To measure the postural field, postural probes moved the hand in 1 of 12 directions (0° , 15° , 30° , 45° , 90° , 135° , 180° , 225° , 270° , 315° , 330° , and 345° with respect to the x -axis) while participants were distracted with a working memory task.

We measured the postural field before and after adaptation. Before adaptation, participants completed 3 blocks of trials, each separated by a short break. In each block, all 12 postural probe directions were visited a single time. The probe displacement was 2.5 cm for all probe directions. Postural probes were given on every 4th outward reaching movement. Therefore, participants completed a total of 288 baseline trials (3 blocks x 12 probes/block x 4 reaching movement pairs/probe x 2 movements/reaching movement pair). Outwards reaching movements of 10 cm were all performed directly away from the body.

Participants were then gradually adapted to a velocity-dependent force field. The field magnitude was decreased from 0 to -10 N·s/m in constant increments over the course of 65 outwards reaching trials (130 trials total). After this adaptation period, the postural field was re-measured (Fig. 6.5B, right). As before, all 12 probe directions were probed in a random order, 3 times. No breaks were provided in between blocks. We anticipated that the postural field would shift after adaptation to the force field. Therefore, we extended the probe displacement to 4 cm for probe angles of 0° , 15° , 30° , 45° , 315° , 330° , and 345° . Postural probes occurred at the same frequency as before adaptation for a total of 288 trials. To maintain participants in an adapted state, on all outwards reaching trials other than postural probe trials, a velocity-dependent perturbation was maintained at -10 N·s/m. Target-hold wells were applied to the final hand position during all outwards trials (with the exception of postural probe trials) during the holding period (1.5 second duration, with an addition 0.3-0.4 inter-trial-interval).

Our analysis focused on changes in the postural field due to adaptation. First, we looked for within-subject changes to the location of the null point of the field. Second, we looked for within-subject

changes to the orientation and stiffness magnitude of the field. For visualization purposes we constructed a two-dimensional postural field from the forces measured during probe trials using linear interpolation. To do this, along each probe direction we resampled forces in x and y spatially in intervals of approximately 0.1 cm. For each of the resampled restoring forces we calculated the corresponding polar coordinates (*i.e.*, the radius and angle). In polar coordinates, all x and y forces lied along a rectangular grid. We used bilinear interpolation along these polar coordinates to estimate the postural field within the space between the 12 probe angles.

6.2.C.10 Measuring the relationship between holding forces and the null point of the postural field

We reasoned that the holding forces measured in our other experiments (Figs. 6.3, 6.4, 6.6, and 6.7) were potentially caused by a misalignment between the hand (fixed in the channel at the target) and the postural null point (somewhere displaced from the target). If this is true, we could gradually eliminate these forces if we displaced the hand towards its null point.

We recruited a set of subjects ($n=19$) to test these predictions at several points during adaptation to a velocity-dependent force field. At regular intervals during adaptation, we inserted postural probe trials along 0° , with respect to the x axis. This corresponded to the direction of the holding force. During the probe, participants were distracted with the working memory task. For the first ten participants, we used 5 cm postural probe trials. For the last nine participants, we shortened the probe length to 4 cm. Here we analyze only the first 4 cm of displacement.

Before adaptation, we measured the postural forces a total of 10 times. Postural probes were inserted regularly on every 5th outward reach, within a baseline period of 100 trials (50 outwards and 50 backwards movements). Outwards reaching movements of 10 cm were all performed directly away from the body. Next, we adapted participants gradually to a CCW velocity-dependent force field, where we decreased the field magnitude from 0 to -10.5 N·s/m in constant increments over 175 outwards reaching trials (350 trials total). We measured moving forces, holding forces, and the response to the postural probe on every 5th outwards reaching movement.

We found that hand forces during the postural probe passed through a null point as the hand was displaced from its terminal position (Fig. 6.5E). To determine the location of the null point of the arm on a trial-to-trial basis we fit a three-parameter exponential function to the hand forces as a function of hand displacement in the probe, and recorded its x-intercept. To do this, we first resampled subject forces spatially in increments of 0.05 cm. Next, to reduce noise inherent in the single trial force measurements, we used a bootstrapping approach. On each trial, we randomly sampled subjects with

replacement, calculated the mean postural force as a function of distance across the sample, and fit the exponential to this mean behavior. We repeated this process 2000 times, and used this distribution to estimate 95% confidence intervals around the mean (Fig. 6.5F).

6.2.C.11 Measuring adaptation in the gain of integration

To maintain stability at the endpoint, the integrator must adapt as the reach controller adapts. For the oculomotor system, adaptation of the integrator occurs when there are consistent errors between the terminal position of the eye, and its target²⁰³. We speculated that such a mechanism might also adapt the gain of the reach integrator. To detect these errors, we considered the reach trajectories of participants as they adapted to velocity-dependent force fields (Fig. 6.6B). We spatially aligned reach trajectories by subtracting off the terminal hand position.

We observed that participants often exhibited deviations in their hand position as they attempted to stop their hand at the end of movement (Fig. 6.6B, endpoint correction). To quantify the size of these errors, we measured the largest “positive” deviation of the hand from its terminal position, after the hand exceeded 80% of its reach trajectory. On movements in which no such error occurred, we instead calculated the largest “negative” deviation of the hand from its terminal position. In Fig. 6.6B, we highlight these errors for two example participants by temporally aligning reach trajectories to the point at which the hand exhibited the largest “positive” deviation from its terminal position.

In Fig. 6.6D, we compare the magnitude and sign of the late trajectory deviations, to the gain of integration at the end of adaptation. To calculate the integration gain, we divided the median hold force by the median reach force integral, measured over 10 channel trials after the perturbation magnitude had plateaued (see horizontal line after CCW+ and CW+ in Fig. 6.3E, bottom). Similarly, in Fig. 6.6E, the integration gain was calculated by dividing the hold force by the reach force integral, either over the first one-third of CCW+/CW+ trials (for the FF early period) or the final 50% of CCW+/CW+ trials (for the FF late period). For the null period, we could not use this technique, as the mean reach force integral and the hold force were both nearly zero. Therefore, to calculate integration gain, we linearly regressed hold period force onto the integral of reach period force, and reported the slope of the regression line.

6.3 Results

We performed experiments in which monkeys and humans made goal-directed movements toward a target and then held their arm or finger at the target location. In each case we asked whether the commands during the move period influenced the subsequent commands during the hold period.

6.3.1 Muscle activity during the hold period following reaching

When reaching movements are made from various starting points to the same target, we know of no model that predicts a consistent relationship between the reach period commands, which depend on reach direction, and the subsequent hold period commands, which depend on target location. Yet, if control of holding depends on the movement period (Fig. 6.1A, right), a single function must exist that transforms the commands that were generated during the movement period to the subsequent commands that are generated to hold the arm. To explore the plausibility of such a relationship, we used intramuscular electrodes to measure activity of 20 shoulder and elbow muscles during point-to-point reaching in the vertical plane (9 muscles in monkey A, 11 muscles in monkey B). On each trial, monkeys first moved their hand from a central location to one of eight targets, held their hand at that target for at least 0.5 sec, then reached back to the central location and again held their hand for at least 0.5 sec (Fig. 6.1B). We normalized muscle activity by setting to zero its average activity at the central hold location, and setting to one its maximum activity in the task.

Consider the activity of the anterior deltoid (Fig. 6.1B) as the arm reached from the central position to a target. For some directions, this muscle exhibited a burst of activity during the reach, and then sustained activity during holding (Fig. 6.1B, targets 3 and 5). For other targets, the muscle exhibited a smaller burst of activity during movement, and little or no activity during holding (target 1). Let us imagine that the measured EMG, denoted as $u(t)$, is actually the sum of two signals: a “hold” command $h(t)$ (brown traces in Fig. 6.1B) and a “move” command $m(t)$ (red traces in Fig. 6.1B). The hold command is computed by adding the real-time integral of the move command to the initial hold activity that precedes movement (see Methods Section 6.2.A). If the hold commands are computed from the move commands in this way, the two commands should be related by a common function across different types of reaching movements: the change in muscle activity from before the movement to after the movement, should be related to the intervening move period muscle activity (see Methods A for derivation). We measured muscle activity with respect to its pre-movement hold period, $u(h_1)$, and then integrated that change with respect to time until the end of the reach, $t = T$.

$$u(h_2) - u(h_1) \approx k \int_{t_0}^T [u(t) - u(h_1)] dt + a \quad (1)$$

This equation predicted that change in the hold period activity of a given muscle from before reach onset h_1 to the target h_2 should be approximately proportional to the integral of its activity during the reach period.




Figure 6.1. Integral of muscle activity during the reach correlates with subsequent activity during the hold period. **A.** In current models (left), a feedback controller generates commands that move the arm, and then upon reach end, a postural controller holds the limb still. For this model, hold commands depend only on the target position. In the model considered here (right), the move commands are integrated in real-time by a postural controller. Thus, the hold commands depend on the preceding reach commands, not solely the target position. **B.** Monkeys reached out to one of eight targets, waited, and then reached back to the home position. EMG from ant. deltoid is shown for three targets, and decomposed into $m(t)$ and $h(t)$ using Eq. (5), with $k = 1$. **C.** Normalized activity of anterior deltoid in Monkey B starting from the center location. Colors correspond to targets in part B. **D.** Change in ant. deltoid EMG from the initial hold period for reach out and reach back components of the task. The bars for hold 1 and hold 2 indicate periods where hold activity was calculated. **E** and **F.** Change in hold period activity for reach out (green) and reach back (red) components of the task as a function of the integral of the preceding reach period in two representative muscles for each monkey, and all muscles. **G.** The integration gain (slope of the line in E) across various conditions: outward vs. return, fast vs. slow, horizontal target positions vs. other targets. **H.** Comparison of two hypotheses: hold period activity

relates linearly to integral of previous reach period, or hold period activity relates linearly to the maximum activity of the muscle in the previous reach period. Each point is a single muscle. Error bars are SEM. Statistics: ** $p < 0.01$ and n.s. $p > 0.05$.

To test the validity of Eq. (1), we defined the hold period at the target h_2 starting at 300 ms after reach end (Fig. 6.1D, hold 2), thus allowing time for muscle activation dynamics to settle. Indeed, for most muscles (14/20), change in the hold period activity was well predicted by Eq. (1). For example, when the target location was constant (Fig. 6.1D, reach back) across various movement directions, the changes in the hold period activities of many muscles were proportional to the integral of their respective activity during the preceding reach (red lines, Fig. 6.1E). When the target position varied (Fig. 6.1D, reach out), Eq. (1) was again a good predictor of the change in hold period activity, despite the fact that both direction and endpoint of the reach changed (green lines, Fig. 6.1E).

Fig. 6.1F presents the data across all muscles, directions, and endpoints. Remarkably, we found that integration of the reach period activity was a reasonable predictor of the change in hold period activity across all conditions ($R^2 = 0.58$ for Monkey A and $R^2 = 0.53$ for Monkey B). Within each muscle the integration gain k was no different for outward reaches and return reaches (Fig. 6.1G, paired t-test on single muscle regression slopes, $p = 0.943$). In addition, the same integration gain predicted hold activity for fast and slow movements (Fig. 6.1G, fast vs. slow, two-sample t-test, $p = 0.30$) which differed modestly in terms of movement duration (two-sample t-test, $p < 0.001$, fast movement duration of 350.3 ± 11.1 ms and slow movement duration of 453.6 ± 4.9 ms, mean \pm SEM). This indicated that a single function (Eq. (1)) could account for various movements and speeds, despite the differing dynamics of these reaches.

Notably, despite these general trends, some muscles (6/20) did not exhibit the pattern described by Eq. (1). These muscles shared a specific property: they had little to no activity during the hold periods (Figure 6.8). Thus, Eq. (1) seemed to apply primarily to those muscles that modulated their activity during the hold period, contributing to maintenance of arm posture in this task. However, to not bias our results, we included all muscles in our regressions in Figs. 6.1F, 6.1G, and 6.1H.

We considered an alternate hypothesis: a muscle that is more active to lift the arm will be also be more active to hold the arm. Perhaps the correlations are driven mostly by biomechanical constraints of the reaching movements. Thus, we separated movements based on directions that were not affected by a change in gravitational forces (horizontal reaches), vs. other directions (Fig. 6.1B). If the relationship between move period and hold period was solely due to the gravitational field, we would expect that the integration function would differ for horizontal versus vertical movements. However, the gain of

integration was similar for the two groups of movements (Fig. 6.1G, ANCOVA, movement type by moving EMG integral interaction effect on holding activity, $F=0.37$, $p=0.54$). To broaden the scope of this alternate hypothesis, we considered the possibility that the change in hold period activity in each muscle depended on the maximum or minimum activity of that muscle during the reach period, not the integral of the entire reach period activity. This alternate hypothesis also proved to be significantly less accurate than Eq. (1) (within muscle comparison, paired t-test, $p = 0.005$, Fig. 6.1H). In 17/20 muscles, integration of the reach period activity was a better predictor of the hold period activity than either the maximum or minimum muscle activity.

In summary, for reaching across various directions and endpoints, the change in a muscle's activity from the pre-movement hold period to the post-movement hold period was partially predicted by the integral of that muscle's activity during the intervening reach period.

6.3.2 Hold period activity for finger movements

According to Eq. (1), trial-to-trial variation in the move period commands should lead to consistent trial-by-trial changes in the subsequent hold period commands, even if the target location remains constant. That is, if the integral of a muscle's activity is greater on some trials, then that muscle should also be more active during the hold period that immediately follows.

It is difficult to test this prediction for reaching because there are numerous configurations of the wrist, elbow, and shoulder joints that would maintain the hand at a target location. Therefore, to more precisely examine within-trial covariance between move and hold periods, we simplified the problem to a single degree of freedom: finger flexion¹⁹⁹.

Monkeys were trained to use their index finger to track a visual target that moved at 12 deg/sec over a 1 second period between a start (12°) and an end location (24°) against a spring load that resisted flexion (Fig. 6.2A). At the start location the load was 0.026 N·m. As the finger flexed, the load increased linearly, reaching 0.048 N·m at the target. We measured muscle activity using subcutaneous electrodes implanted over 17 muscles (8 muscles in monkey R, 9 in monkey D). For each session, we normalized each muscle's activity by setting the average activity at the start location to 0 (period h_1 , 400 ms in duration, began 1 sec before movement onset, Fig. 6.2C), and setting its maximum activity in the task to 1. The hold period h_2 at the target was 200 ms in duration and began 700 ms after movement end, thus allowing time for muscle dynamics to settle.




Figure 6.2. The integral of muscle activity during finger flexion correlates with subsequent activity during the hold period. A and B. Monkeys were trained to move their index finger from an initial position to a target against a load. The traces show representative movements. Positive displacements correspond to flexion. C. Activity of flexor digitorum profundus (FDP) in monkey R. Left panel shows FDP activity for the trials shown in part B. The bars for hold 1 and hold 2 indicate periods when hold activity was calculated. Right panel shows FDP activity during the hold periods. Activity increased with flexion of the finger, but was variable from one trial to the next. D. Evaluation of the hypothesis that variability in muscle activity at the hold 2 position could be explained by variability in the preceding hold 1 activity. Left panel is for the FDP muscle during a single session in monkey R. Center and right panels present data across all muscles recorded in each monkey. Each ellipse is the 95% confidence interval for a single muscle. R^2 value refers to a linear fit across all trials and muscles. E. Same as for D, except here we test the hypothesis that variability in hold period activity is related to the integral of the preceding moving activity.

The finger accurately tracked the target on each trial, moving along very similar trajectories (Fig. 6.2B). To both flex the finger and support the mechanical load, muscle activity during each hold period was strongly modulated by finger location (Fig. 6.2C). For example, the flexor digitorum profundus (FDP) muscle, one of the prime movers in this task, was more active when the finger was at the target as

compared to the start location (Fig. 6.2C, right panel). However, from one trial to the next, hold period EMG exhibited marked variability (Fig. 6.2C, note the vertical spread of the trial distribution at hold 2).

One possibility is that the variability in hold period EMG is due to position-related variability in the spring force applied to the finger (larger displacements lead to larger spring forces on the finger). To assess this possibility, for each muscle we regressed its activity onto its position during the target hold period. We found that variation in hold position accounted for less than 2% ($R^2 = 0.0107 \pm 0.0041$, mean \pm SEM across all muscles in both monkeys, individual regressions for each muscle) of the trial-to-trial variability in hold period EMG (we do not mean for this to give the impression that EMG was not strongly modulated by position, only that trial-by-trial differences in EMG were poorly explained by differences in position; see Methods Section B2 for more information).

If variability in position could not explain the trial-to-trial changes in hold period EMG, what was the source of hold period variability? Inspection of muscle activity in Fig. 6.2C suggested two possibilities: trial-by-trial changes in the hold period activity at the target (the hold 2 period) could be explained by (1) modulation in the initial holding activity prior to movement onset (the hold 1 period), or (2) modulation in move period activity. We first investigated the former possibility that when a muscle was more active during the initial hold period, it was also more active during the final hold period. This hypothesis stated that $u(h_2) \approx ku(h_1) + a$. To evaluate this hypothesis, we regressed the activity of each muscle during hold 2 onto its activity during hold 1. The left panel in Fig. 6.2D shows the strength of this correlation in the FDP muscle for a single session. Each ellipse in the middle panel of Fig. 6.2D represents the 95% confidence boundary for the trial-by-trial joint distribution between hold period activity at the target and hold period activity at the starting position, for each muscle, across all trials and sessions. Overall, hold activity at the start appeared to be a rather poor predictor of hold activity at the target, accounting for about 8% (monkey R) and 7% (monkey D) of the variance in the data (Fig. 6.2D).

We next considered the hypothesis that variation in hold period activity may be due to variation in preceding move period activity. Using Eq. (1), we integrated the move period activity in the muscle, $u(t)$ with respect to its pre-movement activity, $u(h_1)$, and asked if this integral could predict the change in hold period activity from the start location, $u(h_1)$, to the target location, $u(h_2)$. This comparison could be confounded by trial-to-trial differences in movement displacement. Under spring forces, a larger displacement might lead to greater move period as well as greater hold period activity. This did not appear consistent with the data: trial-to-trial displacement explained less than 1% of the variance in both the integral of move period activity ($R^2 = 0.006 \pm 0.0018$, mean \pm SEM), and the change in hold

activity ($R^2 = 0.005 \pm 0.001$, mean \pm SEM). On the other hand, the integral of move period activity exhibited robust correlation with the change in hold period activity, as shown for an example recording session in the left panel of Fig. 6.2E. For this session, about 56% of the trial-by-trial variance in FDP hold period activity $u(h_2) - u(h_1)$ was accounted for by Eq. (2) (linear fit for this session, $p < 0.001$).

To determine if all muscles exhibited similar within-trial correlations between moving and holding, we considered the data across all muscles and sessions (Fig. 6.2E, middle and right panels). Each ellipse in Fig. 6.2E represents the 95% confidence boundary for the within-trial joint distribution between holding activity and the integral of moving activity, for each muscle. Strikingly, the orientations of various muscle distributions (the angle of the major axis) were roughly parallel with each other. Thus, the gain of the integration function was similar across muscles, and a single function accounted for approximately 50% and 40% (monkeys R and D) of the trial-by-trial variability in holding activity.

Finally, we considered another potential source of correlation between move and hold periods: co-contraction. If finger stiffness varied from one trial to another, we would observe correlations between move and hold periods. To change finger stiffness, agonist and antagonist pairs of muscles would exhibit coordinated increases or decreases in their activities. In other words, we should be able to predict the change in hold period activity in one muscle based on the activity of other muscles. To test this idea, we regressed the change in hold period activity in each muscle onto the integral of move period EMG in other muscles. Roughly 10% of the variability in the change in hold period EMG could be explained by the integral of move period activity of other muscles ($R^2 = 0.10 \pm 0.02$, mean \pm SEM). Therefore, while some of the variance in hold period activity could be explained by the move period activity in other muscles, move period activity in a given muscle remained a much better predictor of the change in hold period activity in that same muscle ($R^2 = 0.42 \pm 0.031$, mean \pm SEM). In summary, for a constant target location, on trials where a muscle moved the finger with greater activity, it also produced greater activity during the hold period.

6.3.3 Change in reach period commands alters hold period commands

These EMG patterns illustrated a correlation between move and hold period commands, but did not test whether there was a causal link between the two. That is, trial-to-trial coupling between move and hold period commands arose from variability that was internally generated by the animal. To rigorously test if hold period commands directly followed from move period commands, we next imposed external changes on move period commands and measured if hold period commands changed in a manner consistent with integration. To do this, we instructed participants to reach to a target while holding the

handle of a robotic arm (Fig. 6.3A), and adapted their move period commands by imposing a velocity-dependent force field¹⁰⁰.

With the force field our goal was to gradually bias the move period commands through the process of adaptation. If hold period commands partly depended on the preceding move period, then biases in move commands should lead to biases in hold commands, even though (1) the hold period commands were never perturbed, and (2) the hold location remained constant. To ensure that hold period commands were never perturbed, on all force field trials the hand was placed in a “well” that held the hand still at the end of the movement for at least 1.5 sec (Fig. 6.3A, target-hold well). Next, we used a channel to prevent the hand from suddenly moving off the target while participants transitioned from the target-hold well to the next reaching movement (not shown in Fig. 6.3A; see “partial channel” in Methods).

Unlike the monkey experiments, we did not record EMG during these experiments. Instead, we measured changes in the forces that participants exerted against the handle. These forces served as a low-dimensional proxy for the motor commands sent to the arm muscles. Critically, we drove the adaptation process with forces that acted perpendicular to the trajectory of the hand. Because the learning axis was orthogonal to the primary movement, we could cleanly isolate the component of the motor commands that varied in response to the field, from the motor commands responsible for the primary movement. To measure the forces perpendicular to the primary movement, on some trials the robot produced a stiff channel that connected the start position to the target via a straight line (Fig. 6.3A, channel trial). To test the integration hypothesis, we recorded forces perpendicular to the direction of the target during the reach and hold periods, and asked if they were related through the following integration function (i.e., the force analogue of Eq. 1):

$$F(h_2) - F(h_1) = k \int_0^T [F(t) - F(h_1)] dt + a \quad (2)$$

At the start of each experiment, participants (n=220 in total) reached to the target in a neutral (i.e., null field) condition in which the robot did not produce any forces on the hand (Fig. 6.3A, null field trial). Even in the null field period we observed significant trial-by-trial variability in the perpendicular reach period forces, as shown for a representative subject in Fig. 6.3B. During the reach period the hand pushed slightly to the right on some trials, left on others, or exhibited a bimodal profile. However, because the movements were guided within a channel, the hand followed a straight line that ended at the center of the target (Fig. 6.3A, bottom left). Notably, following conclusion of the reach, we observed that the arm generated forces during the hold period (period h_2 , Fig. 6.3B) that were often different

than baseline (period h_1). Indeed, the change in the hold period forces $F(h_2) - F(h_1)$ was well predicted by Eq. (2), as illustrated by data for the representative subject in Fig. 6.3B (regression at right), and the entire population in Fig. 6.3C. On a within-trial basis, the integral of move period force accounted for $39 \pm 2\%$ (mean \pm SEM) of the variance in hold period force. Thus, just as EMG patterns exhibited a within-trial relationship between the reach and hold periods, so did the force patterns.

Like the EMG patterns, the force patterns did not appear to be trivially related to biomechanical constraints imposed on the arm due to gravity: the gain of the integration function was the same whether or not the weight of the arm was supported by a frictionless air-sled (Fig. 6.3B, two-sample t-test, $p=0.90$). Remarkably, the relationship between move and hold period forces remained unchanged when we divided the null period reaches of each subject into fast and slow movements (Fig. 6.3D, ANCOVA, movement type by move force integral interaction effect on hold force, $F=0.007$, $p=0.935$). Thus, in the null field, forces during the reach and subsequent hold periods showed natural variability. However, on a within-trial basis, the integral of the move period forces appeared to influence the subsequent hold period forces.

At the conclusion of the null field period, we gradually imposed a velocity-dependent force field during the reach (Fig. 6.3E). The velocity-dependent forces perturbed the hand perpendicular to the reach trajectory, thus leading to adaptation of reach period forces (Fig. 6.3F). The gradual onset of the perturbation produced gradual changes in reach period forces (while also minimizing positional errors throughout the trajectory). Remarkably, as the reach period force changed from one trial to the next, so did the hold period force (Figs. 6.3F and 6.3G). Notably, the hold period forces were not transient, but sustained, persisting up to 6 seconds during the entire hold period interval (Figure 6.9). The relationship between the change in hold period force and the (now externally-driven) reach period forces was again consistent with integration: over the course of adaptation, Eq. (2) accounted for $48 \pm 3\%$ (mean \pm SEM across 32 participants) of the variance in the change in hold period forces.

As an alternate hypothesis we considered the possibility that hold period forces may be a trivial continuation of the forces produced near the end of the preceding reach, not an integration of the entire history of the reach period. To test this idea, we conducted a pair of experiments. In the first experiment (Fig. 6.4A), participants ($n=11$) reached in a force field that was active only during the second half of the reach. In Phase 1 of the experiment, trial after trial we gradually increased the magnitude of the force field (Fig. 6.4A, Phase 1). As expected, participants produced hold period forces that increased with the integral of the reach period forces. In Phase 2, we maintained the force field at a

constant magnitude for hundreds of additional trials. The change in hold period forces and the integral of reach period forces remained correlated during all trials (Fig. 6.4A, Phase 2).




Figure 6.3. The integral of reaching forces correlates with forces produced during the subsequent hold period. A. Human participants held the handle of a robotic arm and made point-to-point reaching movements (top). On most trials, participants reached freely to the target. After the reach ended, a target-hold well held the hand in place (null field trial). On some trials, hand trajectory was constrained to a straight line (channel trial, trajectories shown at left). B. Example lateral force traces on channel trials during the null field period (in a single subject). The bars hold 1 and hold 2 indicate the periods in which holding force levels were quantified. At top-right, we show the correlation between the integral

of reach period forces and the change in holding force for this representative subject. At bottom-right we compare the gain of integration (slope of the line at top-right) between conditions with and without weight support for the arm. C. We calculated the correlation coefficient between the time-integral of reach forces and holding forces across all channel trials in the null field, within each individual ($n=220$). The vertical blue line denotes the mean of the distribution. D. For each subject, we selected their two fastest and two slowest movements in the null field, resulting in two distributions, with each subject represented equally in each distribution. We then performed linear regression on each distribution separately. Error bars for the integration gain are 95% CI at right. At middle, we show the distribution of reach duration. E. After the null field period, we gradually introduced a velocity-dependent force field (top). We measured moving and holding forces as subjects adapted and de-adapted to counterclockwise (CCW) and clockwise (CW) force fields (bottom). F. Each trace represents the force on one channel trial, averaged across participants during the CCW force field adaptation period. The vertical gray bars denote the start and end of the reaching movement. The color of each traces indicates the force field magnitude at each point in the experiment. The hold 1 and hold 2 bars indicate periods over which holding forces were quantified. G. On each trial, we calculated the time-integral of forces during reaching (green) and compared these to the change in holding force (red). Values are mean \pm SEM across all participants. Statistics: n.s. $p>0.05$.

In the next experiment (Fig. 6.4B) we again exposed participants to a force field that was active only during the second half of the movement. But then, in Phase 2, we gradually added a second force field that was active during the first half of the movement, but in the opposite direction (Fig. 6.4B, Phase 2). In this way, the reach period forces should integrate to approximately zero. If holding forces were a simple continuation of the reach period forces, then the addition of the force field during the first half of the movement should not alter the hold forces at the end of the movement. However, this is not what we observed: as the integral of reach period forces approached zero, hold period forces gradually vanished (Fig. 6.4B, Phase 2). That is, even though reach period forces were matched just before the end of movement (Fig. 6.4C), the ensuing hold period force depended on the entire history of the reach. In addition, we observed no difference in the integration function between each phase of the experiment (paired t-test on slope, $p=0.22$, paired t-test on intercept, $p=0.09$).

Together, these observations demonstrated that on a within-trial basis, as the reach period forces changed, so did the ensuing hold period forces. The change in hold period force was partially described via a function that integrated in time the temporal history of the preceding reach period forces.

6.3.4 Move and hold correlations are only weakly influenced by initial hold activity

Eqs. (1) and (2) use initial hold activity ($u(h_1)$ or $F(h_1)$) both to calculate the change in hold activity, and to estimate move period activity (Methods Section A). We were concerned that some or all of the observed correlations may be due to this common factor that appears on both sides of the equation, and not the integral of the move period activity.




Figure 6.4. Holding forces are an integration of moving forces, not a continuation. We designed a set of experiments to test the possibility that hold period forces are a continuation, not an integral, of moving forces. A. Participants ($n=11$) reached in a force field that was active only during the second half of the reach. In Phase 1 (left), we gradually increased the magnitude of the force field. In Phase 2 (right) we maintained this force field for several hundred trials. We measured the change in holding force (right, red) and the integral of moving force (left, green) throughout adaptation. B. A new set of participants ($n=14$) repeated Phase 1 (left), but during Phase 2 (right) an opposite force field was gradually applied to the first half of the movement. As the integral of move period force approached zero in Phase 2, so did holding force. C. The mean force profile over trials sampled from Phases 1 and 2 for the experiment in part B. Values are mean \pm SEM across all participants.

To consider this problem, we noted that for outwards reaching movements in Fig. 6.1, this was not a concern because the trial-averaged EMGs were nearly identical before movement onset at the center location, but differed greatly during the movement to various directions. To address this concern for our data in Figs. 6.2 and 6.3, we re-analyzed the correlation between move and hold activity, but only on trials in which hold 1 activity fell within one standard deviation of the mean. This criterion reduced the variance in hold 1 activity by 75% for our finger movement dataset, and 88.6% for our reach force dataset. Despite this dramatic reduction in hold 1 variability, we found little effect on the measured correlations: the correlation coefficient between integral of move activity and change in hold

activity remained 58.4% for muscles in the finger (compared to 64.2% if all trials are included) and 48.5% for reach forces on channel trials in the null field (compared to 52.2% if all trials are included).

In summary, the observed trial-by-trial correlations between move activity and change in hold activity were almost entirely driven by the integration of the move commands, with little dependence on variability in the initial hold activity.

6.3.5 The postural field during the hold period

Thus far we have described the state of the limb during the hold period in terms of muscle activity or force generation. However, in order to hold the limb still, the postural controller does not simply produce a force, rather it generates a converging field of position-dependent forces^{201,202}. We next asked if this postural field also relied on the commands generated during the preceding reach.

We designed a new experiment in which we measured the postural field following completion of a reach. Participants ($n=27$) reached to a target as before, but now, during the hold period they were engaged in a short-term memory task (2-back, Fig. 6.5A). As they performed the memory task, the robot slowly displaced their arm in a random direction. In response to the displacement, the postural controller produced restoring forces against the handle, thus allowing us to measure the postural field (Fig. 6.5B).

As expected, the postural field's null position was near the target (Fig. 6.5B, null point of postural field). However, after participants were exposed to a force field, the postural field changed: the null position was no longer aligned with the target (Fig. 6.5B, right). Rather, it shifted by approximately 1 cm (Fig. 6.5D; paired t-test, $p < 10^{-4}$) in the direction of the force produced during the reach. In contrast, the orientation (Fig. 6.5D, paired t-test, $p=0.84$) and stiffness (Fig. 6.5D, paired t-test, $p=0.62$) of the postural field remained unchanged.

After the reach had ended and the cursor was at the target, we slowly displaced the hand toward the postural null position. We observed that the hold period forces gradually approached zero, and then switched direction and grew larger as the hand was displaced beyond the null position (Fig. 6.5E). The holding force at the hand scaled linearly with the distance between hand position and the postural null position (Fig. 6.5F). This implied that the hold period forces we had measured in previous experiments (Figs. 6.3-4) were a proxy for the location of the null position of the postural controller: the larger the hold period force at the target, the farther the null position of the postural field.

Thus, as the reach period forces changed, so did postural control: the null position of the postural field shifted in the same direction as the change in the preceding move period forces.




Figure 6.5. The null point of the postural field is set by the integral of reaching forces. A. To measure the arm's postural field, we slowly displaced the hand during the holding period, while participants were distracted with a working memory task. B. We measured the forces applied to the handle (left). We re-measured these forces after participants were exposed to a velocity-dependent curl field (at right). Forces were measured by displacing the hand outwards along 12 different lines. Interior estimates for the force were made using two-dimensional linear interpolation. The magnitude and direction of these interpolated forces are indicated by the black arrows. Color reiterates the restoring force magnitude. The holding position at reach end is located at the intersection of the two dashed white lines. C. We measured lateral forces applied to the channel walls during reaching movements (null field period, black, and curl field period, red). D. We used a two-dimensional spring model to quantify postural field properties: null point, orientation, and stiffness (null field and curl field in black and red). E. To test if holding forces were related to the null point of the postural field, participants ($n=19$) were exposed to a curl field that gradually increased over trials. During holding, we recorded hand forces (right inset) as the arm was displaced in the direction of holding forces. Arrows show the location of the null point (zero-crossing) on selected trials. F. We calculated the holding force before displacement of the hand, and the corresponding postural null point on each trial. Values are trial means and 95% CIs for distributions bootstrapped across participants. Linear regression was performed on the bootstrapped estimates (black line). Error bars denote mean \pm SEM in panels C and D. Statistics: *** $p < 10^{-3}$ and n.s. $p > 0.05$.

6.3.6 Adaptation of the integration gain

These results create a puzzling scenario. In the presence of a velocity-dependent force field, the reach controller readily adapts and changes the move period forces. However, changes to the move period forces are integrated and cause the hold system to program an entirely different null position. This implies that postural stability will be compromised in the face of an adapting reach controller. To solve this problem, the integrator must also be adaptive: the integration function must change when there is an error between its current output and the desired movement endpoint (Fig. 6.6A). Presumably this

error-based adaptation would be driven by unexpected deviations from the hand’s desired trajectory as the reach period ends and the hand arrives at the target location.

To test this idea, we examined reach trajectories of individual subjects as they gradually adapted to velocity-dependent force fields. These trajectories exhibited two primary types of errors. The first error happened midway through the movement, and was caused by incomplete compensation for the velocity-dependent perturbation (Fig. 6.6B, the large negative mid-movement error). The second error happened near the end of the movement, and was oriented in the opposite direction (Fig. 6.6B, “endpoint correction”). This near endpoint error possessed two properties that were well-suited for integrator adaptation: (1) they occurred late in the movement as the participant attempted to stop their hand within the target, and (2) they were oriented in the direction opposite the shift in postural null point reported in Fig. 6.5. To quantify the size of these endpoint errors, we measured the largest “positive” deviation (or “negative” if the lateral deviation was in the opposite direction) from the terminal hand position, after the hand exceeded 80% of its reach displacement.




Figure 6.6. Adaptation of the integration gain. A. To maintain endpoint stability after adaptation of the reach controller, the postural controller must also adapt. B. We hypothesized that integrator adaptation would be driven by errors in hand trajectory that occur near the end of the reach. To detect these errors, we looked for deviations in the reach trajectory after the reach exceeded 80% of its

displacement. We spatially aligned reaches by subtracting off the terminal hand position and then temporally aligned the reach trajectories to the point in time at which the hand had the largest endpoint error. These errors are marked as “endpoint correction”. Here we show the average reach trajectory during adaptation to a gradual force field for two subjects, one who exhibited integration (S5), and one who did not (S11). C. The average forces produced at the end of adaptation for the same two subjects. Note that the subject with large errors near reach endpoint no longer generated holding forces at the end of adaptation. D. We measured the gain of integration at the end of the adaptation period, for subjects that adapted to gradual force fields. We also measured the magnitude and sign of the late reach errors. Larger errors led to a reduction in gain. Each data point is one subject. E. To confirm that the integration gain changed over the course of adaptation, and not immediately upon introduction of the force field, we compared the gain during the null period, with gains measured during early and late parts of adaptation using a repeated measures ANOVA. Values are mean \pm SEM across participants. Statistics: *** $p < 10^{-3}$ and n.s. $p > 0.05$.

To determine if endpoint errors caused integrator adaptation, we compared the size of these errors during the adaptation process, to the gain of integration observed at the end of adaptation. The size of endpoint error was heterogeneous across our subjects; some participants exhibited large endpoint errors (Fig. 6.6B, S11) while others exhibited small endpoint errors (Fig. 6.6B, S5). Critically, we found that participants with larger endpoint errors ultimately produced smaller holding forces (Fig. 6.6C). In other words, these errors appeared to reduce the gain of integration. In fact, about 40% of the variability in integration gain could be explained by the magnitude of the endpoint errors (Fig. 6.6D).

This adaptation of integration gain progressed over time (Fig. 6.6E, repeated measures ANOVA, $F=12.60$, $p < 0.001$). In the early part of training, the integration gain was no different than that of the pre-adaptation null trials (Fig. 6.6E, post-hoc comparison, $p=0.88$). However, late in training, the integration gain had decreased substantially (Fig. 6.6E, post-hoc tests, $p < 0.001$ for comparison of late adaptation with both null field trials as well as early adaptation trials).

These data suggest that Eq. (2) alone cannot predict the change in holding forces. In certain conditions, for example velocity-dependent force fields, the reach period commands change, but if one is to hold the hand at the target, then the hold period commands cannot simply integrate the preceding move period commands. Rather, as the move period commands adapt, so too must the integration function. Like the adaptation of movement commands, the adaptation of hold commands does not occur instantaneously, but appears to emerge gradually as errors near the end of the movement are experienced repeatedly. This process may also explain why holding forces gradually diminish during adaptation to an abrupt force field, where endpoint errors are large⁸².

6.3.7 Differential contributions of the corticospinal tract to reaching and holding

The CST conveys the cortically generated reach commands to downstream motor structures. Does this same pathway also convey postural signals, or does a separate, downstream structure receive and then integrate the reach commands? If both reaching and holding commands are conveyed in the CST, then damage to the CST above the level of the brainstem should disrupt both the generation of forces during reaching, and its integration during the hold period. However, if the integrator is downstream to this level, then damage to the CST might result in deficient reach commands, but spare the process of integration, resulting in hold commands that reflect the within-trial integration of the now deficient reach commands.




Figure 6.7. Cortical reaching commands are integrated in a subcortical area. A. Stroke survivors (n=14) participated in a set of clinical exams to measure functional impairment. Shown are isolated images for an extension-based task, for the non-paretic (top) and paretic (bottom) arms of an example participant (S015). The instruction is to place a rectangular block on the elevated surface. Images show the moment of maximal extension for the paretic (right) and nonparetic (left) arms. B. To improve the range of motion of the arm, patients and healthy controls performed reaching movements holding the robotic handle, with the arm supported by an air sled. Shown are example trajectories during an initial null field period for the representative patient (black is nonparetic, red is paretic) in A, and a control participant (blue). C. Example force traces during null field block in channel trials. The solid line denotes forces during moving. The dashed line denotes forces during holding still. D and E. We measured the integral of moving forces (D) and holding forces (E) on each channel trial. We measured the trial-by-trial variability (standard deviation) of these quantities across all movements in the null field. F, H, I, and J. We compared trial-by-trial fluctuations in moving and holding forces during the null field period (F, left panel). Next, we gradually adapted subjects to a velocity-dependent force field and compared within-trial integral of moving forces with subsequent holding force (F, right panel). Data are shown for representative stroke patient and healthy control. We calculated the correlation coefficient between reaching and holding forces during the initial null field period (H) and force field period (I). We measured the slope of the integration function (i.e., the integration gain) across all trials within individual subjects (J). G. Our conjecture that the cortex generates reaching commands which are then integrated in a subcortical area spared by cortical stroke. Values are mean \pm SEM across participants. Points represent individual trials in F. Points represent individual subjects in D, E, H, and I. Statistics: * $p < 0.05$, and n.s. $p > 0.05$.

To examine these possibilities, we recruited stroke patients (n=14) who had suffered lesions affecting the CST pathway from the cortex through the internal capsule (Table 1). The patients exhibited profound impairments, as demonstrated by difficulty with extension of their arm during unsupported reaching^{204,205} (Fig. 6.7A, patient S015). To improve their reach capacity, we supported the weight of their arm in the horizontal plane (frictionless air sled), which allowed them to better extend their arm at the elbow, enabling them to make planar, point-to-point reaching movements while holding the handle of the robot arm (as in Fig. 6.3A).

As has been noted before^{206,207}, movements of the paretic arm exhibited erratic trajectories (Fig. 6.7B), increased movement duration (paretic vs. control) of approximately 41% (paretic vs. non-paretic, paired t-test, $p < 0.01$; paretic vs. control, two-sample t-test, $p < 10^{-4}$), and reach endpoints that terminated nearly 89% (paretic vs. control) further away from the target location (paretic vs. non-paretic, paired t-test, $p < 0.01$; paretic vs. control, two-sample t-test, $p < 0.001$). The reaching impairment coincided with a marked increase in the trial-to-trial variability of move period forces (traces in Fig. 6.7C solid lines; Fig. 6.7D paretic vs. control, Wilcoxon rank-sum, $p < 0.001$; Fig. 6.7D paretic vs. non-paretic, Wilcoxon signed-rank, $p = 0.058$). However, like healthy subjects, these move period forces, no matter how variable, terminated with stable holding forces (traces in Fig. 6.7C dashed lines). The trial-by-trial

variability of the hold period forces, like the move period forces that preceded them, was significantly greater in the paretic arm of the patients (Fig. 6.7E; paretic vs. control, two-sample t-test, $p=0.01$; paretic vs. non-paretic, paired t-test, $p=0.06$).

If move period commands were integrated into hold period commands, the increased variability in holding forces (Fig. 6.7E) could arise indirectly from the normal integration of the highly variable moving commands. If this were true, the variability in the moving and holding commands would be similarly structured in both healthy subjects and stroke patients. To test this idea, we quantified the within-trial correlation between change in hold period forces and the integral of the preceding move period forces (Fig. 6.7F, left column, representative subjects). Remarkably, in the null field trials, the coupling between move and hold periods was intact in stroke patients (Fig. 6.7H, paretic vs. control, two-sample t-test, $p=0.14$; paretic vs. non-paretic, paired t-test, $p=0.63$).

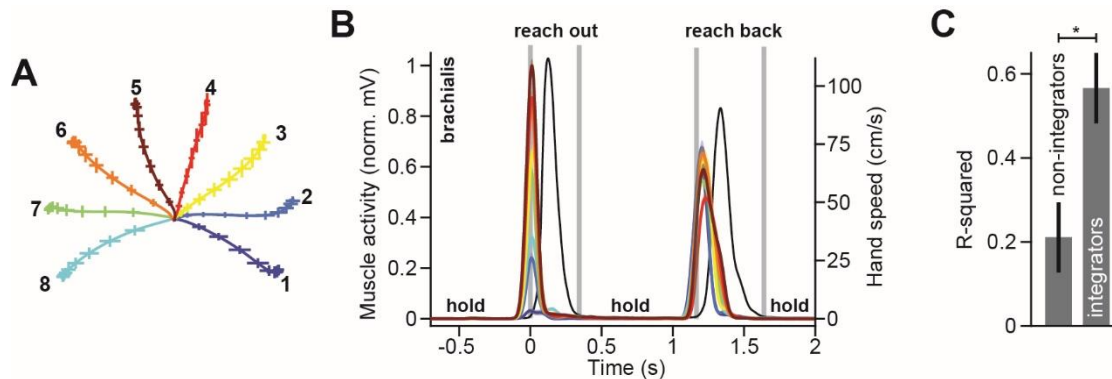


Figure 6.8. Muscles that do not integrate also do not contribute to posture. A. Two monkeys performed a reaching task in the vertical plane. The monkeys reached out to one of eight targets, waited, and then reached back to the home position. Trajectories for the outward reach are shown. B. Normalized activity of brachialis in Monkey B. Colors correspond to targets in A. We selected this representative muscle to demonstrate that some muscles show little to no holding activity at most, if not all, targets. These muscles all tend to show early moving activity, well before peak speed, and then cease activity around halfway through the movement. The hand speed is shown in black. C. We selected muscles that have little to no holding activity, i.e., those whose activity during the hold period has a magnitude less than 10% of the peak moving activity. There are 6 out of 20 muscles who had such a property for most, if not all holding locations. We quantified the change in holding activity (from before to after the movement) and the integral of moving activity. We asked how well the integral of moving activity predicted the change in holding activity (the variance accounted for, or R^2) in muscles with little holding activity (non-integrators) and others that were more active during holding still (integrators). Error bars are SEM. Statistics: $*p < 0.05$.

We next used adaptation to systematically manipulate move period forces. Because force field adaptation is largely a cerebellar-dependent process¹⁶, despite damage to the CST the patients learned

to alter their move period forces. As the move period force changed in the paretic arm, so did the change in hold period force (Fig. 6.7F, right column, example subjects). Once again Eq. (2) provided a reasonable account of the within-trial relationship between the move period and the change in hold period forces for the paretic arm, non-paretic arm, and the dominant arm of age-matched control subjects (paretic vs. control, two-sample t-test, $p=0.08$; Fig. 6.7I: paretic vs. non-paretic, paired t-test, $p=0.24$). Notably, the integration gain was not significantly different across the paretic and non-paretic limbs of the patients, nor across the patients and age-matched controls (Fig. 6.7J; paretic vs. control, two-sample t-test, $p=0.86$; paretic vs. non-paretic, paired t-test, $p=0.91$). In other words, the integration function was similar in healthy participants and stroke patients.

In summary, damage to the CST severely affected the reach period commands, resulting in high trial-to-trial variability. However, the change in hold period commands remained coupled to the integral of the preceding reach commands in both null field and force field trials. The gain of this integration in the stroke patients was not different than that of healthy controls, suggesting that CST damage impaired the reach commands, but not the process of integration that may have generated the hold commands.




Figure 6.9. Holding forces are sustained across long time intervals. Participants ($n=7$) were exposed to a CCW velocity-dependent curl force that increased in magnitude across many trials. Here we show the lateral force (top) and hand velocity (bottom) after exposure to maximal field strength. The vertical line denotes the end of the movement, hence, the start of the holding period. On the trials shown, participants held the hand at the target and waited for a long inter-trial-interval to start the next reaching movement. During this period of time, participants engaged in a working memory task.

6.3.8 Holding forces are not specific to a particular holding location

Earlier models of holding forces posited that they arose from simultaneous adaptation of velocity and position primitives⁸². If this is the case, it might follow that adapting individuals to both CW and CCW velocity-dependent force fields might lead to the elimination of holding forces, because of a process of cancellation. That is, if on half of the trials the force field is oriented to the left, and the other half to the right, interference may lead to cancellation of holding forces at the common reach endpoint.

To test this possibility, we used a dual adaptation paradigm²⁰⁸ to simultaneously adapt individuals ($n=14$) to a CW and CCW force field. To accomplish this, a rectangular bar was shown centered on the subject's hand. On some trials, a control point was shown on the left of the bar and a target was positioned vertically above the leftwards control point (Fig. 6.10A). On every other trial, the control point and target were moved to the right side of the bar. When the control point was on the left, the hand was perturbed with a CW force field. When the control point was on the right, the hand was perturbed with a CCW force field. We tested for adaptation on channel trials interspersed through the experiment. The alternating cue successfully allowed for simultaneous adaptation to both force fields, though did not completely eliminate interference (Fig. 6.10B); adaptation to the CCW exceeded that of the CW field, likely due to interference.

Critically, when reach forces were produced to the right, holding forces to the right followed. When reach forces were produced to the left, holding forces to the left followed (Fig. 6.10C). These holding forces persisted throughout the duration of the experiment. The integral of reach forces accurately predicted hold forces as before (Fig. 6.10D). Altogether, these results suggested that holding forces were not specific to a point in space, but instead were dependent on the direction and magnitude of the preceding reach forces.

6.3.9 The change in holding force, not final holding force, is predicted by the integral of reach forces

Our derivation of the integration equations suggested that the integral of reach forces should predict the change in reach force, rather than the terminal reach force at the target (Eq. 6.2). To investigate the validity of this prediction, we adapted to a set of participants to a CW force field experienced at two separate reach targets (Fig. 6.11). On some trials the participants made a reaching movement directly away from the body (Target 1) and experienced forces to the right (Fig. 6.11A, column one). On other trials, participants reached directly towards their body (Target 2) and experienced forces to the left.

Critically, when participants reached back from Target 1 and Target 2 to the start position, they always moved in a channel condition (Fig. 6.11A, columns 2 and 4).




Figure 6.10. Moving and holding forces during exposure to dual force fields. A. Participants ($n=14$, Experiment 7) reached between two locations in space (the filled blue circle and the dashed central green circle). Rather than controlling a single cursor, participants were shown a rectangular tool that was centered on the hand. On some trials (context 1, left schematic), a control point appeared on the left of the tool and participants were told to move that point into a target that appeared to the left in the workspace. On other trials (context 2, right schematic), a control point appeared on the right of the tool and participants were told to move that point into a target that appeared to the right in the workspace. Critically, the physical starting position and target position were the same for all of the trials, simply the context changed. Participants were exposed to CW and CCW force fields. The orientation of the field was paired to the context (context 1 paired to CW field and context 2 paired to CCW field). B. Forces were measured on occasional channel probe trials. Here we show the forces produced on probes in context 1 (dashed lines) and context 2 (solid lines). Color indicates the trial number (blue is early in the experiment, red is late). C. We calculate the moving force time-integral (green) and holding force (red) for context 1 trials (top) and context 2 trials (bottom). Values are mean \pm SEM across participants. D. Next we examined the scatterplot showing the holding force one each trial as a function of the moving force time-integral on that trial. Black points correspond to context 1 trials and red points correspond to context 2 trials. To reiterate, all of the trials start and end at the same location in space.

When participants reached to Targets 1 and 2 on channel trials (Fig. 6.11B), they successfully expressed reach forces consistent with the CW force field that increased over time due to the gradual onset of the perturbation (Fig. 6.11C). As before, the integral of reach force predicted the holding forces at the endpoint (Fig. 6.11D). Given the geometric arrangement of the targets, reaching movements from Target 1 back to center were oriented in the same direction as reaching movements from the center to Target 2. Similarly, reaching movements from Target 2 to center were oriented in the same direction as reaching movements from the center to Target 1. Because of these geometric similarities, participants generalized their adaptation to the return channel trials (Figs. 6.11E and 6.11F). Even for these generalized forces, the integral of reach forces accurately predicted hold forces at the endpoint, providing further evidence that adaptation was not necessary to express hold forces. That is, even though errors were never experienced when reaching to the start position, hold forces still followed the generalized reach forces.




Figure 6.11. Holding forces at the same target position are specific to reaches initiated from different start positions. A. Participants ($n=13$) reached in cycles of four trials. On the first trial (first column), participants reached away from the body from a center position (red) to Target 1 (green). At the end of the movement, the hand was held in place for a long holding period. On the next trial (third column), the participant returned the hand from Target 1 to center. On the third trial of the cycle (second column), participants reached from the center position to Target 2. And finally, on the fourth trial (fourth column), participants returned the hand back to center again. After an initial familiarization period, participants were exposed to force fields when reaching from the center to Target 1 (rightwards field) and from the center to Target 2 (leftwards field). Reaching movements that terminated at the center position were always executed in a channel that prevented lateral errors. B. Occasionally, channel probe trials were inserted to measure hand forces when reaching to Targets 1 and 2. C. Shown are the average reach forces (sign-flipped) for reaching movements from the start position to Targets 1 and 2. D. Shown is a scatterplot of the holding force as a function of the moving force time-integral for the probe trials depicted in C. Each point represents one trial. Color indicates the strength of the force field at that point in the experiment (blue means small, and red means large). Unfilled circles show reaches to Target 1 and filled circles show reaches to Target 2. The solid line shows the linear regression across trials. Values are mean \pm SEM across participants. E. All reaching movements made towards the center position were executed in a channel. F. Same as in C but for probe trials depicted in D. Color indicates the force field magnitude at a specific point in the experiment. G. Shown is a scatterplot of the hold forces and integral of reach forces in F. Values are mean \pm SEM across participants.

The generalization of forces presented another important scenario. Because reaching movements from the start position to the target were preceded by generalized reach forces in the same direction from the preceding channel trial, the initial forces were non-zero during the hold period at the start of the reach (Fig. 6.11C). That is, initial forces for the outwards movements were determined by the hold forces from the previous reach. Similarly, for reaching movements from the peripheral targets back to the center, hold forces biased the initial reach forces in the opposite direction to the upcoming generalized reach forces (Fig. 6.11F). This arrangement permitted us to ask if the final hold force was predicted by the integral of reach force, or if according to Eq. (6.2), the change in hold force from start to end was better predicted by the integral of reach force. Therefore, we considered channel trials to all of the adaptation targets (Fig. 6.12A, left) as well as the generalization trials back to center (Fig. 6.12A, right). First, we regressed the change in hold forces onto the integral of reach forces (Fig. 6.12B) and calculated the gain that related these two quantities, i.e., the scaling factor that related the change in hold force to the integral of reach force. Similarly, we regressed the terminal hold force onto the integral of reach force (Fig. 6.12C) and calculated the gain that related these two quantities.

We compared these gains to that recorded on trials in other experiments where there were no initial hold forces (Figs. 6.12B and 6.12C, control, no bias). Critically, we found that the same integration gain described the relationship between the change in hold forces and the integral of reach forces whether the initial hold forces were zero or not (Fig. 6.12B, two-sample t-test). On the contrary, the

relationship between the integral of reach force and the final hold force possessed a different integration gain when reaching movements began with an initial bias in hold forces (Fig. 6.12C, initial bias vs. control, no bias; two-sample t-test, $p < 0.001$). In other words, we found that the integral of reach force accurately described the change in hold force across each of our datasets consistent with Eq. 6.2, but did not predict the terminal hold force the same across each experiment. Therefore, the terminal hold force is not simply a continuation of reach force; rather, the change in hold force is determined by the integral of reach force.




Figure 6.12. The integral of reach force predicts the change in reach force, not the terminal reach force. A. In our two-target experiment in Fig. 6.11, participants were adapted during outwards movements, and tested for generalization on return movements. We calculated the integral of reach force and the hold force for channel trials in each of the four directions. Because of the geometry of this task, adapted reach forces or generalized reach forces all led to initial biases in hold forces at the start of the reach. B. Here we tested our primary hypothesis in Eq. (6.2): does the integral of reach force predict the change in hold force? For the four trial types (Start to Target 1, Target 1 to Start, Start to Target 2, and Target 2 to Start) we calculated the integration gain relating the change in hold force with the integral of reach force. The integration gain is shown for each movement in black. Then we averaged across each of these integration gains (gray bar at right) and compared this gain to the integration measured in our primary experiments where there was no initial bias in hold force at the start of the reach (blue bar). C. We performed a similar comparison, but this time, for integration gains relating the terminal hold force with the integral of reach force. Statistics: n.s. means $p > 0.05$ and *** means $p < 0.001$.

6.3.10 Reach forces predict hold forces independent of spatial orientation of movement

To test the robustness of Eq. (6.2) we next repeated our primary experiments (Fig. 6.3) but for reaching movements oriented at different points in space. Whereas the majority of our tasks tested movements straight away or towards the body, we tested oblique movements oriented on a 45° angle (Fig. 6.13A) in a gradual force field adaptation paradigm (Fig. 6.13B). As in our primary experiment, hold forces increased or decreased as reach force increased or decreased in the various phases of the experiment

(Fig. 6.13C). The integral of reach forces accurately predicted the change in hold force (Figs. 6.13D and 6.13E). Therefore, our primary results (Fig. 6.3) generalized to reaching movements made at oblique angles to the body.




Figure 6.13. Moving and holding forces for reach movements of different angles. A. Participants ($n=11$, Experiment 2) reached between a start location and a target located at an oblique angle (135° with respect to the positive x-axis) from the midline of the body (represented by the dashed line). First, participants reached in a null field (left), but then were exposed to CW or CCW velocity-dependent force fields (right). At the end of null field and force field trials, the hand was held in place with an endpoint clamp. B. The force field schedule. Shown is the force field magnitude on each trial. The experiment was composed of cycles consisting of one outwards reach and backwards reach. On outwards trials were perturbed. The force field gradually increased to a constant value and then decreased. The + and - indicators represent periods of increasing and decreasing force field strength, respectively. CCW and CW force fields were applied in different blocks, separated by a short break. C. Occasionally, probe trials were inserted to measure forces during moving and holding. Shown are forces during different periods of the experiment (upper-left is increasing CCW field, upper-right is decreasing CCW field, lower-left is increasing CW field, and lower-right is decreasing CW field). Each trace is one trial. The color indicates the field magnitude in each respective period (blue is early and red is late). D. We calculated the time-integral of the moving forces (green) and the static holding force (red) on each trial. Values are mean \pm SEM across participants. E. Shown is a scatterplot of the holding forces as a function of the moving force

time-integrals. Colors are consistent with trials in C. Marker shape depicts the period of the experiment (legend bottom-right). The solid line depicts the linear regression. Values are mean \pm SEM across participants.

Next, we considered movements that started and ended at different points in space. We tested reaching movements where the arm started to the right of the midline of the body (Fig. 6.14A) and to the left of the midline of the body (Fig. 6.14B) in different blocks. Again, we used a gradual adaptation paradigm where subjects were exposed to either a CW or CCW velocity-dependent force field (Fig. 6.14E). As in our primary results (Fig. 6.3) reach forces and hold forces covaried with one another throughout each phase of the experiment (Figs. 6.14C and 6.14D). Again, Eq. (6.2) provided an accurate account of the relationship between the integral of reach force and the change in hold force (Fig. 6.14F). Altogether, these results demonstrated that Eq. (6.2) generalized to different spatial locations in the two-dimensional workspace.




Figure 6.14. The relationship between reaching and holding forces is consistent across different arm postures. A. The movement was centered about 15 cm to the left of the midline of the body (dashed line). First, participants reached in a null field (top) but were then exposed to a CCW (Experiment 3, $n=9$) or CW (Experiment 4, $n=8$) force field that gradually increased in strength over many trials. At the end of each null field or perturbation trial, the hand was held in place during an extended holding period within an endpoint clamp. B. Participants also reached in a separate block in the same manner as in A, but with the arm centered 5 cm to the right of the midline of the body. C. Occasional probe trials were inserted to measure hand forces in the left workspace. The time-integral of the reach force is shown on each trial in green. The holding force is shown in red. The + and - indicators represent periods of time in which the force field was increasing and decreasing, respectively. Values are mean \pm SEM across participants. D. Same as in C but for the right workspace. E. Participants were gradually adapted and then deadadapted to either a CCW (blue group, bottom) or CW (black group, top) force field. Participants were tested in two blocks separated by a short break. In one block, participants reached with the arm centered to the left of midline (shown to the left of dashed line). In the other block, participants reached with the arm

centered to the right of midline (shown to the right of dashed line). The order was counterbalanced across participants. F. Shown is a scatterplot depicting the holding forces as a function of the moving force time-integrals. Points represent individual trials. The black line shows the linear regression. Values are mean \pm SEM across participants.

6.3.11 Cortically-driven feedback responses to error are also integrated

In this work, we used channel trials to isolate feedforward components of motor plans in the absence of error. That is, all of the evidence to this point that reach commands are integrated into hold commands applies solely to movements executed in the absence of error. Does the feedback response to error also generate reach commands that undergo integration? To answer this question, we designed an experiment where participants reached from a start position to a target (Fig. 6.15A) as before. But every so often, participants received a visual perturbation to the cursor while the hand was restrained within the channel. On some trials, the cursor was fictitiously moved to the left, and on others to the right, to differing visual extents. On these channel probe trials we measured the reach forces exerted laterally against the channel wall (Fig. 6.15B). Consistent with a visual feedback response to error, corrective reach forces were produced starting 200 ms after reach offset.




Figure 6.15. The visual feedback response to error is integrated. A. Participants reached from a start position to a target position. On some trials the hand was restrained within a channel, but the cursor was fictitiously displaced to the left or right along a parabolic trajectory, thus simulating a reach error. B. Participants responded to this visual prediction error, generating a feedback response consistent with the direction of the visual perturbation that started approximately 200 ms after reach onset. Hold forces appeared to follow the different feedback responses. C. On each trial we calculated the change in hold force and plotted this as a function of the integral of reach force on that trial. D. The integral of reach forces accounted for trial-by-trial changes in hold forces. To quantify the strength of their relationship, we linearly regressed the change in hold force onto the integral of reach force, and reported the corresponding R^2 value.

These feedback driven reach forces showed clear signs of integration (hold period in Fig. 6.15B). In fact, the level of integration appeared strong on a trial-by-trial basis (Fig. 6.15C). Ultimately, the integral of reach force accounted for approximately 30% of the variation in the change in hold force (Fig. 6.15D) consistent with our primary results in Fig. 6.3. Therefore, cortical commands generated online in response to error are also integrated as are feedforward commands generated in the absence of error.

6.4 Discussion

Current computational models of reaching view move and hold periods as events that take place in sequence: the sensory representation of the target engages a feedback controller that moves the arm, and then once the desired endpoint has been achieved, a separate controller is engaged that produces the sustained commands that hold the arm^{186,187,189,191}. These models have usually assumed that the motor cortex is responsible for generating the move period as well as the ensuing hold period commands^{194,195}. Here, our experiments suggest the possibility of a different architecture (Fig. 6.7G), one in which movement commands are integrated in real-time (Fig. 6.1B) by a separate network of neurons of possibly subcortical origin, resulting in holding commands. If true, this would imply that control of the arm shares a design principle present in control of the eye^{1,13,86,88,89,91,94,193} and the head⁹⁴⁻⁹⁶.

We measured muscle activity during point-to-point reaching in the vertical plane and found that across directions and durations, a form of mathematical integration related muscle activity during the hold period with the preceding reach period activity (Fig. 6.1). When the start and end positions of finger movements were kept nearly constant, there was still large variability in the hold period EMG of many finger muscles. On a within-trial basis, for all muscles recorded the integral of the move period EMG partly accounted for the change in hold period EMG. In contrast, activity before the start of the movement, co-contraction, or even finger position itself were poor predictors of the final hold period EMG. Thus, fluctuations in the integral of the move period commands influenced the change in hold period commands.

Next, we altered the move period commands and asked whether that change altered the commands that were generated during the hold period. We approached this question in healthy participants, as well as patients who had survived a stroke affecting their CST above the brainstem. In both populations, as reach period forces changed during force field adaptation, so did the subsequent hold period forces, as predicted by integration (Figs. 6.3 and 6.7). Integration was also observed during the null period prior to the introduction of the force field, effectively ruling out the possibility that moving and holding correlations arose due to reach adaptation. Critically, the same integration function

was observed for both healthy participants and stroke patients, suggesting that the putative integrator might reside within a subcortical structure.

As an alternative to integration we considered the possibility that the hold period forces may be a continuation of the reach period forces, not an integration of the entire period. To test this idea, we changed the reach period forces via adaptation to a bidirectional force field (Fig. 6.4), one in which the integral of the move period forces was zero, but not the forces near the start or end of the move period. We found that as the integral of the move period force approached zero, so did the change in hold period force.

Finally, we considered the fact that in order to hold the arm at a specific location, the nervous system must produce a postural field^{99,201,202}. We measured this field by engaging subjects with a working memory task, while slowly moving their hand away from the target location. We found that as the reach period commands changed, so did the null position of the postural field. The magnitude of this shift was proportional to the integral of the preceding reach forces. These results suggested that the hold forces we measured in our adaptation experiments were generated by a disparity between hand position and the null point established by the integration of moving commands. This mismatch may help explain the paradoxical illusions in perception of arm position^{209,210} that accompany force field adaptation.

The idea that move period activity is integrated into a specific null point might explain why we observed poor trial-by-trial correlation between hold activity and finger position in Fig. 6.2. Note that the tube housing the finger constrained the motion of the finger to an arc, and thus may have prevented it from moving to its true “null position” in three-dimensional space. Therefore, if move commands integrated to a position outside of the tube, for a fixed tube rotation, there could be many null positions lateral to the tube, thus decreasing the observed correlation between the rotational position we measured and the associated EMG. In other words, trial-to-trial variability in move period activity should integrate to different magnitudes of hold activity, thus leading to hold positions that do not necessarily lie within the tube.

The notion that holding commands control a null position is reminiscent of earlier theories in which the principal role of the motor system was to specify equilibrium positions for the arm^{211,212}. These theories posited that movement arose from the transition between holding locations. In sharp contrast, our results imply that the hold commands are generated in real-time via integration of the ongoing move period commands.

6.4.1 A hypothetical architecture for control of arm posture

The model presented in Fig. 6.7G represents our conjecture regarding architecture of the reach and hold controllers. In this conjecture, movements are encoded by the cortex, resulting in the move signal, termed $m(t)$, which is integrated in real-time by an unknown subcortical area. The integral of move commands is then added to initial hold activity, yielding a feedforward estimate of the commands required for holding still, termed $h(t)$. During a movement, the commands that arrive at the motoneurons are a sum of commands for moving and holding, $u(t)$.

There are a number of predictions that arise from this model. Motor commands required to move the arm to a target position are not fixed, but vary because of interactions with external objects, interaction torques that arise when the body is in motion, and over time as our bodies change. We know that the move system continuously adapts to these novel dynamics. A similar process of adaptation would also be required of the proposed reach integrator.

Indeed, in the oculomotor system, move period and hold period commands can both undergo adaptation, and this adaptation depends on the cerebellum¹⁶. However, different regions of the cerebellum are required for adaptation of the move and hold periods¹. Optican and Miles²⁰³ demonstrated that the oculomotor integrator could be adapted by translating the target position as the eye transitioned from a saccade to gaze holding. In our reaching experiments, we found evidence for a similar adaptation mechanism. Errors near the end of the reach that were consistently encountered as the arm attempted to stop within the target location reduced the gain of the integration function (Fig. 6.6). Through adaptation of the hold system, the arm would be able to cope with novel dynamics during the reach while also maintaining the ability to hold the arm at the target.

While this adaptation mechanism would achieve endpoint stability over the sequence of many movements, immediate corrections in the integrator output would be handled through the parallel operation of sensory feedback (Fig. 6.7G, sensory feedback). This scheme would resemble a neural integrator for control of the head in the interstitial nucleus of Cajal⁹⁴⁻⁹⁶, which relies on proprioceptive and visual feedback⁹⁶. The importance of these feedback pathways is further illustrated by conditions in which the arm is passively moved to a new position.

We think that this new model of reaching might shed light on a number of interesting puzzles. For example, transient inhibition of the motor cortex during a reach results in “freezing of the arm” at its current posture, and not loss of muscle tone⁹⁷. That is, despite near complete removal of output from the motor cortex during a reach, commands of unknown origin continue to sustain arm position against

the force of gravity. Our model suggests that a distinct structure, possibly located in a subcortical area, integrates the cortically-generated reach commands up until the moment of cortical inhibition, and thus maintains posture despite removal of reach commands. This model is also consistent with the observation that cortical stimulation not only moves the limb, but also produces specific postures in primates²¹³ and rodents²¹⁴. Moreover, the total displacement of the limb appears to scale with the duration of stimulation, consistent with the idea that displacement is produced due to integration of motor commands over time.

The idea that integration occurs outside of the motor cortex potentially explains why successful decoders of cortical activity are designed to control cursor velocity, as opposed to cursor position^{192,215,216}. To hold the cursor still, these decoders assume that the cortex encodes a zero-velocity command rather than a position-based command. This observation is consistent with the possibility that the cortex is primarily concerned with dynamic quantities, i.e., changes in the state of the limb, rather than the maintenance of a specific limb state over time. This idea would account for the observation that neurons in the motor cortex modulate their activity when there is a change in force production, but less so when the constant force is maintained over time^{217,218}. To maintain a constant force over time, our conjecture states that a subcortical area integrates phasic activity from the cortex, and maintains its output over time. This idea is consistent with evidence that transient stimulation of the brainstem in decerebrate cats produces sustained (timescale of minutes) changes in extensor muscle force⁹⁸.

6.4.2 Limitations

Eqs. (1) and (2) describe how the integration of move activity may relate to changes in hold commands, but does not specify the hold command at the target. This reflects the reality that move period commands alone will not determine the terminal position of the arm: the initial arm position must also be taken into account. In other words, to hold the limb at the target, the integrator must not only integrate move commands, it must add this integral to the hold period activity that preceded movement (see. Eq. (6) in Methods; illustrated in Fig. 6.7G). We do not know if the integrator internally performs this addition, or if a downstream structure handles this adjustment for initial limb position.

Without a biomechanical model of the arm, it is not obvious why the linear functions in Eqs. (1) and (2) robustly predicted the relationship between change in hold activity and the integral of move activity. It may be that as we test movements that are more complex than simple point-to-point reaches, the observed linearity will break down. For example, how would the reach integrator maintain endpoint accuracy when the arm grasps an object, thus altering its mass? Such a scenario would require

an increase in muscle activity to move the larger mass, and then hold it still against the forces of gravity, so it may be the case that integration would naturally account for these positive correlations. Furthermore, over successive movements, inaccuracies in the output of the integrator could be reduced through adaptation of the integration gain as in Fig. 6.6. Finally, as is the case for the move controller^{208,219,220}, the reach integrator may store object-specific or environment-specific memories and implicitly alter its integration properties when an often-encountered dynamical context is re-experienced. These mechanisms, along with the parallel output of visual and proprioceptive controllers would allow the arm to be stabilized under different inertial conditions.

These possibilities could be tested in the future, by altering the mass of the arm and recording EMG from various arm muscles. With regards to EMG recording, it should be re-emphasized that the reach period forces measured in our human subject experiments were perpendicular to the direction of movement, and thus represented only part of the complete motor command. Because perturbations were applied perpendicular to the reach trajectory, our measure captured the primary dimension of adaptation. In future studies, a more complete picture of the integration process would be provided by recording EMG during adaptation.

It is unlikely that the motor cortex has no role in postural control. At the very least, the monosynaptic projections from corticomotoneurons²²¹ to alpha-motoneurons in the spinal cord are likely to be active during periods of holding still. We do not know if these hold period signals in the cortex arise from recurrent connections with a subcortical integration system, or from a separate position controller.

Understanding the differential contributions of the cortical systems and the putative subcortical integrator is essential to improve our understanding of neurological disorders such as stroke and dystonia²²². Identifying the locus of the reach integration may help explain why these patients exhibit abnormal postures at rest. These abnormalities could arise from a lesion to the integrator circuit, or perhaps more provocatively, from the normal integration of chronically abnormal moving commands.

Chapter 7. Conclusions and future directions

When we experience a movement error, the brain learns from the error and changes the way we perform that movement on the next attempt. Here we set out to address (1) how the brain learns the correct motor response to prevent future errors, (2) how this learning process changes over time, and (3) how the process of movement adaptation changes the way we maintain stillness. Here we revisit the primary conclusions from Chapters 2-6, describe possible ways to continue these lines of inquiry, and also speculate on the neural correlates of the behavioral mechanisms we highlighted in this work.

7.1 Mechanisms of error-based learning

When we make a saccade towards an object of interest, our movements often end with a visual error between the object and our fovea. When these errors occur, our brain quickly programs a corrective saccade to better bring the object into focus. The cerebellum contains a region called the oculomotor vermis which appears to care a lot about these situations. Purkinje cells of the oculomotor vermis experience a complex spike with higher probability when these errors occur in their preferred direction²⁵. This complex spike is thought to cause LTD at synapses that are active at that moment in time, thus tuning the synaptic plasticity of that neuron to improve future behavior²². But how should cerebellar output modify future behavior in order to improve behavioral performance?

One potential answer is described by the feedback error learning hypothesis²³. When our movements lead to sensory consequences that differ from our predictions, we either modify our movement online through a feedback controller, or perform a subsequent movement to correct the discrepancies left by the first (e.g., a corrective saccade). Therefore, the feedback corrections we make to correct for error serve as a model for the predictive motor commands the brain might produce in the future to eliminate error. In other words, feedback response to error may serve as a teaching signal for the motor learning system^{24,105}.

In Chapter 2, we investigated this possibility by perturbing the arm in an unpredictable manner, and then using EMG to extract the way muscles responded to error, and how they changed their activity on the next trial in the absence of error. In each muscle, we found evidence that the change in feedforward muscle activity (i.e., the learning response) mirrored the feedback response to error of that muscle. Specifically, we found that the learning response resembled a copy of the feedback response, but scaled down in magnitude, and shifted earlier in time by approximately 100-150 ms. This rule

accounted for each period of the learning response, from the start of the movement, up until the reach ended at the target.

These correlations do not of course imply causation. There are at least two models that might explain these findings. The first is that correlation between feedback and learning may arise because both are driven by the same error. In contrast, the feedback error learning hypothesis holds that the error drives the feedback response, and the feedback response then drives the learning response. To tease apart these possibilities, we divided up our trials into instances where the feedback response was low and others where the feedback response was high. Critically, the kinematic error did not differ across these trials, simply the way the muscle responded. In support of the feedback error learning hypothesis, we found that trials in which the feedback response was high in a muscle were followed by a learning response that was also high in that muscle.

How might this mechanism occur in the brain? Towards the end of Chapter 2, we showed that the single trial response was bi-directionally modified by cerebellar stimulation. Therefore, cerebellar learning may play a role in the feedback error learning hypothesis. We could imagine a design where the combination of visual and proprioceptive evokes complex spikes in the cerebellum via the inferior olive. These complex spikes occur at some visual and proprioceptive delay. While we do not know this latency, for saccades in the primate, this appears to be approximately 120 ms. The feedback response to the visual and proprioceptive errors is also delayed, requiring time to travel to the sensorimotor cortex, and then back down the CST. An oft-quoted latency for this response in the spinal cord is approximately 150 ms. If an efference copy of this feedback response to error arrives at the cerebellum a bit before this, then it may be the case that complex spikes in the cerebellum are occurring at precisely the same time as the arrival of the feedback response to error (note that it has been shown recently that complex spike latency is tuned to align with certain behavioral functions, adding credence to this possibility²²³). Therefore, if the parallel fibers in the cerebellar cortex also encode the feedback response to error, the synapses active during the feedback response to error are also those that will undergo plasticity due to the occurrence of the complex spike, forming an engram of the feedback response to error in the cerebellum. This engram may be the key to why the learning response of the next trial resembles the feedback response to error.

Therefore, the muscle-based learning rules we described in Chapter 2, could be caused by the representation of muscle commands in the cerebellar cortex. The stochasticity inherent in the feedback response to error makes it such that on some trials, certain muscles are updated more than others. If the trial-by-trial usage of a muscle in the feedback response follows an unbiased normal distribution, we

would not expect that one muscle would necessarily become dominant for the feedback response to learning. However, it is interesting to wonder if across different exposures in the same participant, if the feedback response could be encouraged somehow to include some muscles and not others, and then to demonstrate that the endpoint of adaptation can be altered in muscle space, depending on the feedback responses to error on the preceding trials. This idea could be tested in a future experiment.

We do not imagine the learning from the feedback response to error is the sole contributor to implicit learning processes. More than likely, feedback error learning is one of many processes by which the brain learns from error, and in many cases, a sensory prediction error alone is sufficient to drive adaptation^{14,115}. This latter idea comes from the observation that many movements like saccades and fast reaching movements terminate too quickly for feedback responses. How then could a feedback response serve as a teacher if it does not occur? While we think this is likely the case, we should also note that simply because a feedback response is not expressed does not mean that it was never computed implicitly. It may be the case that a feedback response to error is in fact sent to the cerebellum, but the cerebellum does not act on its presence because the reach is over. While this may seem an outlandish possibility, it has been demonstrated that cerebellar and cortical output is extremely context dependent, that is, whether or not it produces an output that affects movement depends on the state of the individual²²⁴. And so, it could be the case that feedback responses are computed but not executed when a movement ends due to a shift in context.

In Chapters 3-5 we explore how learning responses, more specifically error sensitivity, changes over time. It is also interesting to wonder how the feedback-instructed learning might also change over time. For example, in the context of savings, is it possible in some cases that the amount learned from the feedback response to error is increased by past experience (and vice versa)? If so, is it possible that part of the memory of past errors⁷ we observe is actually a memory of how we corrected for those errors in the past? Future experiments should record both the feedback response to a perturbation as well as the learning response over time and exposures, to understand how each is modified over time.

Finally, with the advent of Neuropixels and optogenetics, it may be finally possible to determine what complex spikes in the cerebellar cortex encode, in the context of arm movements. For example, do complex spikes encode low-level muscle properties, or errors in higher dimensional visual space, or joint proprioceptive space? When do complex spikes occur during a reaching movement: at discrete times during the movement, or continuously throughout a reaching movement? What is the consequence of complex spikes in the cerebellar cortex in the context of reaching? For example, in saccades, a complex spike in a Purkinje cell microcircuit that prefers errors to the right results in learning to the right²². Might

a similar mechanism be true for reaching movements? To answer these questions, we will need to record from the cerebellar cortex, locate the locus of reach adaptation, record during normal reaching movements, and other sensory manipulations of the limb. These experiments have great promise in helping link the theories of motor control and motor learning with actual neurophysiology of the brain.

Section 7.2 Models of error-based learning

Our understanding of the brain is accelerated by our attempt to mathematically describe observed behavior. In the field of motor learning, the state-space model of error-based adaptation has had a significant impact on the direction of the field, the questions we ask, and the way we interpret our findings. As an example of this, all of Chapters 3-5 was based on interrogating predictions of these equations, and updating the underlying models based on observation.

In the context of these models, most attention is given to understanding our sensitivity to error and also the decay of memory over time. However, in Chapter 3, we demonstrate that another critical component of these models is variability: the variability present in the production of movement, the observation of our errors, our learning processes, and the environment around us. The way our models account for these sources of variability influences how we select the most likely model that describes our data. We demonstrated that incorrectly selecting the wrong model, changes your belief about the properties of adaptation, and also effects your ability to do hypothesis testing, increasing variability in the estimation of model parameters and thereby the number of participants that are needed for statistical testing.

We showed that EM is a better tool for understanding human motor behavior than the standard least squares technique, and also provided a framework for its use in cases where model parameters are time varying. By no means however should this technique be a stopping point. First, of all EM is a tool that maximizes a likelihood function iteratively. However, the likelihood function could be searched directly, as demonstrated in the appendix of Chapter 3. Very preliminary evidence that we did not provide here suggested that searching the likelihood function directly tended to yield extreme estimates for parameters, often due to the convergence of certain noise variance parameters to zero. Therefore, it may be the case that there are parameters that correspond to a large likelihood, but nevertheless are less reasonable than a set with a lower likelihood.

Part of this issue might arise from the fact that EM is a frequentist approach that has no prior assumption on what sets of parameters are more likely than others (and that the parameter sets are point estimates for which there is no measure of uncertainty). Therefore, Bayesian techniques provide

an exciting opportunity to improve model fitting procedures in the human psychophysics literature⁵⁸. While Bayesian approaches to estimation are commonplace in other fields, the sensorimotor literature has not yet caught on to this trend. Future work should focus on making these techniques approachable for the user irrespective of their computational skills.

In Chapters 3-4 we used the EM algorithm to estimate learning parameters in various contexts: savings, spontaneous recovery, and anterograde interference. In Chapter 5, we considered the point at which motor learning saturates in short term paradigms. In all cases, we found that modulation of error sensitivity played a key role. That is, savings was caused by an increase in error sensitivity in fast learning processes. Anterograde interference was caused by a decrease in error sensitivity that recovered over time. And the point of saturation could be modulated not by changing the rate of decay, but by the sensitivity to error in environments with low variability and high variability. For this latter observation, we saw that much of the control of asymptote was through a memory of errors in the implicit learning system. Therefore, we speculate that the cerebellum may play a critical role in each phenomenon and could potentially house a memory of past errors.

At the end of Chapter 5 we speculated as to the form that such a memory of past errors might take. Given that Purkinje cells of the cerebellum have a preference for error, and that different microcircuits of the cerebellar cortex therefore encode different parts of the error space, it may be possible that olivary projections to the cortex have organized the cerebellum in a manner conducive to a memory of errors. We imagine a scenario in which the temporal frequency of complex spikes may modify the magnitude of LTD that occurs at the parallel fiber-Purkinje cell synapse (or perhaps even the synapses with the climbing fiber itself). In other words, consistent errors invoke complex spikes in the same set of neurons, whereas inconsistent errors invoke complex spikes in different neurons on each movement attempt. A high frequency of complex spikes in a neuron may induce synaptic changes that make that neuron more sensitive to the occurrence of that complex spike in the future, and vice versa for low frequency of complex spikes. This mechanism could be tested in a savings paradigm, where individual Purkinje cells could be held for two exposures to determine on the second exposure how LTD in that cell is affected by the occurrence of the first exposure.

There are other potential aspects of the memory of errors model that will be important to explore both in behavioral and neurophysiologic experiments. Our results in Chapter 5 require that the memory of errors, or error sensitivity modulation more generally, must have an upper bound. That is, at some point, error sensitivity must stop increasing even if errors are consistent on each trial. To account for this possibility, we speculated that there may be a decay in error sensitivity that maintains its

magnitude within a homeostatic range. This idea would be consistent with findings that savings is lost after long periods of washout, but preserved when washout is limited⁶¹. In the context of motor learning we know that decay of memory takes at least two forms: time-based and movement-based²⁹. That is, memories are forgotten by waiting, or by moving in the absence of error. And more interestingly, different components of memory may be differentially sensitive to each possibility. Future experiments should explore if this is also true of error sensitivity. That is, does decay in error sensitivity depend on time, the presence of other movements, or the presence of the errors for which sensitivity was upregulated in the first place? Many of these possibilities could be tested in savings paradigms in which washout, time, and errors between exposures are manipulated. In addition, this memory of errors could be coupled with neurophysiologic recordings from the cerebellar cortex, but of course, after the neural correlate of error sensitivity is revealed.

Finally, I should end by noting that when it comes to these models, one size does not fit all. For example, a two-state model of learning can explain spontaneous recovery, but not savings or interference. A memory of errors model with a single process explains savings, but not anterograde interference. And we currently do not understand why the experience of two dissimilar perturbations causes a reduction in error sensitivity of the latter. For this reason, the field is still in need of a model that could account for each of these phenomena, or for neuronal recordings that describe why the mechanism for each (and therefore the mathematical model) differs at the level of neural circuits. Continuing along these lines, while the memory of errors model has received much attention, we have a much poorer understanding of the processes that cause interference. While we describe how error sensitivity changes in the anterograde interference paradigm in Chapter 4, we do not know what elements of the initial learning period inhibit adaptation to the second learning period. This represents another area for future investigation.

Section 7.3 Cortical and subcortical control of holding still

One of the great promises of neuroengineering is the brain-computer interface^{192,225}. The idea is that clinicians can place an electrode over the motor cortex, decode the neurons' intentions for movement, and use this signal to power a robotic prosthesis for individuals that have lost use of their limb(s). While the feasibility of this idea has been partially explored for robots with multiple degrees of freedom, most have used this technology by placing arrays over M1 to control the movement of a cursor on a screen. However, when this done, a peculiar problem emerges. Whereas these devices can be trained to control accurate movements of a cursor, they lose their stability during periods of holding still. That is, stopping

the cursor on a target is very hard to do from the signals present in M1. For this reason, successful decoding algorithms exclusively control cursor velocity, and then apply zero velocity to the cursor once it drops below some threshold speed^{192,215,216}. In other words, cortical decoders hold the cursor still using a zero-velocity command, rather than a position-based command. Why should this be the case?

In Chapter 6, we demonstrated that damage to the corticospinal tract selectively impairs movement, but not the integration of movement into commands for holding still. Therefore, perhaps the difficulty with position-control of the cursor is because static position-related commands are simply not present within the cortex, but are computed elsewhere. Perhaps the cortex's role in movement is the production of dynamics for movement of the limb or cursor, but not the maintenance of static commands over time. This possibility is consistent with recordings from rostral parts of M1, where neurons demonstrate phasic activity during the production of a movement, but no sensitivity to the state of the arm when it is holding still, or applying a constant force over time^{217,226}. Remarkably, single neurons in the motor cortex do not modulate their activity when the arm is producing a constant force, irrespective of the orientation of that force in space²¹⁸.

If the encoding of static forces and positions is not present in the motor cortex, where might it be? Transient stimulation of the brainstem in decerebrate cats leads extensor muscles in the limb to change its tonic level of contractile activity⁹⁸. So perhaps tonic activity is produced from the integration of phasic activity in the brainstem (or spinal cord). This possibility is consistent with studies of cortical inactivation in mice. When mice reach for a food pellet, inactivation of the motor cortex halts the progression of the reaching movement, but does not prevent the hand from being held still in space. Therefore, there must be some subcortical region that produces persistent activity that maintains motoneuronal output, much like the prepositus (oculomotor integrator) does for the eye.

Therefore, we propose that the arm, like the eye, may in part depend on a neural integrator to hold the arm steady in a desired posture. Such a possibility may have important implications for our understanding of various neurological disorders, and the design of devices like neural prosthetics that interact with our brain to control external objects. How might this hypothetical integrator function? In Chapter 6, we demonstrated that the commands that move and hold the arm against gravity are, to a good approximation, related through linear integration. The change in hold commands from the start of the movement to the end of the movement are linearly proportional to the time-integral of the muscle's activity during the movement. The same is true of point to point movements of the finger, and for forces produced during periods of reaching and holding. Therefore, we may be able to record neural activity in M1, construct a low dimensional representation of the population activity, relate this to the velocity of

the effector, and simulate the output of the integrator by integrating this signal in time to hold a prosthetic device still at the end of movement. Or perhaps we can create a neural network whose role is to transform phasic movement activity from M1 into a static EMG output during periods of holding still.

How might this device program the correct posture in space? Here we demonstrated that to hold the hand at the target, the arm created a postural force field with a single stable null point²⁰² that was centered at the terminal position of the hand. However, after the adaptation of the adaptation of reach forces, this null point moved to a new position in space. Therefore, it appears that movement commands are integrated into a null point in space, about which a stabilizing postural field is applied. How can we maintain the accuracy of the device over time? Here we found evidence that the putative reach integrator adapted to errors near the endpoint of the movement. Therefore, when the relationship between the movement commands and the terminal position in space are altered, this leads to endpoint errors that then would teach the integrator to change its gain.

While the possibility of constructing a silicon-based integrator seems far away, there are a number of steps that we could follow to get there. The first is acquiring a better understanding of the basic behavioral and neurophysiologic properties of cortical and subcortical systems for moving and holding still. We will need to understand how the reach integrator interacts with sensory feedback controllers based on vision and proprioception. In addition, we must also understand the space in the movement commands are integrated. Here we operate in the space of muscles, but the more likely possibility is that integration occurs in a kinematic space that represents accelerations and velocities. Perhaps we will discover that there are two integrators connected in series that integrate acceleration of the shoulder and elbow joints into a change in position. In addition, we must perturb the system further to produce real life scenarios in which the integrator must operate. The first of these will be applying masses to the arm that alter its inertial properties and then measuring how this may or may not change the underlying integration equation that relates moving and holding commands.

Even if it turns out that the reach integrator is quite different than we suggest here, this framework provides a way forward to understanding cortical activity. First, what does the cortex do during periods of holding still. Now is the time to return to the cortex and record its activity during periods of holding still to understand how its neurons may or may not contribute to this process. In these experiments, we will need to be careful not to confuse holding activity with preparatory activity for future movements¹⁹⁸. In addition, we will need to consider other possible sources of position signals in the cortex. For instance, position-related modulation of M1 neurons may reflect efferent activity from S1 projections that encode the proprioceptive sense of the limb.

The only way to tease apart some of these possibilities is through the design of carefully constructed behavioral tasks that isolate the appropriate variables of interest. In addition, tools for fast timescale perturbations of the system will be critical, such as optogenetics. For this reason, the mouse represents an ideal system for working out the details of integration. However, principles gained from the mouse should be taken with a grain of salt, as they may not apply to the more cortically-developed primate. Nevertheless, the advent of the Neuropixel probe makes possible the ability to record from large populations of neurons in the brainstem during cortical inactivation, in an effort to determine where persistent holding activity originates in the mouse. Once this nucleus is identified, the next step will be to trace and understand its inputs, hopefully for the purpose of understanding the mathematical relationships that achieve integration.

Finally, it is reasonable to expect that at the end of the day, modern machine learning tools will be critical for the understanding of the integrator and how it might achieve its function. Now we have poor mathematical understanding of the integration functions that could transduce movement-related activity of the limb into static commands for holding still. The design of a recurrent neural network may offer a way forward to understanding these computational principles, in tandem with approaches from synthetic neuroscience.

Reference List

1. Shadmehr, R. Distinct neural circuits for control of movement vs. holding still. *J. Neurophysiol.* **117**, 1431–1460 (2017).
2. Xu-Wilson, M., Chen-Harris, H., Zee, D. S. & Shadmehr, R. Cerebellar Contributions to Adaptive Control of Saccades in Humans. *J. Neurosci.* **29**, 12930–12939 (2009).
3. Shadmehr, R. Learning to Predict and Control the Physics of Our Movements. *J. Neurosci.* **37**, 1663–1671 (2017).
4. Herzfeld, D. J., Pastor, D., Haith, A. M., Rossetti, Y., Shadmehr, R. & O'Shea, J. Contributions of the cerebellum and the motor cortex to acquisition and retention of motor memories. *Neuroimage* **98**, 147–158 (2014).
5. Albert, S. T. & Shadmehr, R. The cerebellum learns to predict the physics of our movements. in *Think Tank: Forty neuroscientists explore the biological roots of human experience* (ed. Linden, D. J.) (Yale University Press, 2018).
6. Shadmehr, R., Smith, M. A. & Krakauer, J. W. Error correction, sensory prediction, and adaptation in motor control. *Annu Rev Neurosci* **33**, 89–108 (2010).
7. Herzfeld, D. J., Vaswani, P. A., Marko, M. K. & Shadmehr, R. A memory of errors in sensorimotor learning. *Science* **345**, 1349–1353 (2014).
8. Smith, M. A., Ghazizadeh, A. & Shadmehr, R. Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol.* **4**, e179 (2006).
9. Ethier, V., Zee, D. S. & Shadmehr, R. Spontaneous Recovery of Motor Memory During Saccade Adaptation. *J. Neurophysiol.* **99**, 2577–2583 (2008).
10. Vaswani, P. A., Shmuelof, L., Haith, A. M., Delnicki, R. J., Huang, V. S., Mazzoni, P., Shadmehr, R. & Krakauer, J. W. Persistent Residual Errors in Motor Adaptation Tasks: Reversion to Baseline and Exploratory Escape. *J. Neurosci.* **35**, 6969–6977 (2015).
11. Taylor, J. A., Krakauer, J. W. & Ivry, R. B. Explicit and Implicit Contributions to Learning in a Sensorimotor Adaptation Task. *J. Neurosci.* **34**, 3023–3032 (2014).
12. Mazzoni, P. & Krakauer, J. W. An implicit plan overrides an explicit strategy during visuomotor adaptation. *J. Neurosci.* **26**, 3642–3645 (2006).
13. Cannon, S. C. & Robinson, D. A. Loss of the neural integrator of the oculomotor system from brain stem lesions in monkey. *J. Neurophysiol.* **57**, 1383–1409 (1987).

14. Tseng, Y. W., Diedrichsen, J., Krakauer, J. W., Shadmehr, R. & Bastian, A. J. Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J. Neurophysiol.* **98**, 54–62 (2007).
15. Morton, S. M. & Bastian, A. J. Cerebellar contributions to locomotor adaptations during splitbelt treadmill walking. *J. Neurosci.* **26**, 9107–9116 (2006).
16. Smith, M. A. & Shadmehr, R. Intact ability to learn internal models of arm dynamics in Huntington’s disease but not cerebellar degeneration. *J. Neurophysiol.* **93**, 2809–2821 (2005).
17. Bastian, A. J., Martin, T. A., Keating, J. G. & Thach, W. T. Cerebellar ataxia: abnormal control of interaction torques across multiple joints. *J. Neurophysiol.* **76**, 492–509 (1996).
18. Izawa, J., Criscimagna-Hemminger, S. E. & Shadmehr, R. Cerebellar Contributions to Reach Adaptation and Learning Sensory Consequences of Action. *J. Neurosci.* **32**, 4230–4239 (2012).
19. Wolpert, D. M. & Kawato, M. Multiple paired forward and inverse models for motor control. *Neural Netw.* **11**, 1317–1329 (1998).
20. Honda, T. *et al.* Tandem internal models execute motor learning in the cerebellum. *Proc. Natl. Acad. Sci. U. S. A.* **115**, 7428–7433 (2018).
21. Ohmae, S. & Medina, J. F. Climbing fibers encode a temporal-difference prediction error during cerebellar learning in mice. *Nat. Neurosci.* **18**, 1798–1803 (2015).
22. Herzfeld, D. J., Kojima, Y., Soetedjo, R. & Shadmehr, R. Encoding of error and learning to correct that error by the Purkinje cells of the cerebellum. *Nat. Neurosci.* **21**, 736–743 (2018).
23. Kawato, M., Furukawa, K. & Suzuki, R. A hierarchical neural-network model for control and learning of voluntary movement. *Biol. Cybern.* **57**, 169–185 (1987).
24. Franklin, D. W., Burdet, E., Tee, K. P., Osu, R., Chew, C. M., Milner, T. E. & Kawato, M. CNS learns stable, accurate, and efficient movements using a simple algorithm. *J. Neurosci.* **28**, 11165–11173 (2008).
25. Herzfeld, D. J., Kojima, Y., Soetedjo, R. & Shadmehr, R. Encoding of action by the Purkinje cells of the cerebellum. *Nature* **526**, 439–442 (2015).
26. Thoroughman, K. & Shadmehr, R. Learning of action through adaptive combination of motor primitives. *Nature* **407**, 742–7 (2000).
27. Scheidt, R. A., Reinkensmeyer, D. J., Conditt, M., Rymer, W. Z. & Mussa-ivaldi, F. A. Persistence of motor adaptation during constrained, multi-joint, arm movements. *J. Neurophysiol.* **84**, 853–862 (2000).
28. van der Kooij, K., Brenner, E., van Beers, R. J. & Smeets, J. B. J. Visuomotor adaptation: how forgetting keeps us conservative. *PLoS One* **10**, e0117901 (2015).

29. Hadjiosif, A. M. & Smith, M. A. Savings is restricted to the temporally labile component of motor adaptation. in *Translational and Computational Motor Control* 1–2 (2015).
30. Leow, L. A., de Rugy, A., Marinovic, W., Riek, S. & Carroll, T. J. Savings for visuomotor adaptation require prior history of error, not prior repetition of successful actions. *J. Neurophysiol.* **116**, 1603–1614 (2016).
31. Morehead, J. R., Qasim, S. E., Crossley, M. J. & Ivry, R. Savings upon Re-Aiming in Visuomotor Adaptation. *J. Neurosci.* **35**, 14386–14396 (2015).
32. Haith, A. M., Huberdeau, D. M. & Krakauer, J. W. The Influence of Movement Preparation Time on the Expression of Visuomotor Learning and Savings. *J. Neurosci.* **35**, 5109–5117 (2015).
33. Berniker, M. & Kording, K. P. Estimating the relevance of world disturbances to explain savings, interference and long-term motor adaptation effects. *PLoS Comput. Biol.* **7**, e1002210 (2011).
34. Kitago, T., Ryan, S. L., Mazzoni, P., Krakauer, J. W. & Haith, A. M. Unlearning versus savings in visuomotor adaptation: comparing effects of washout, passage of time, and removal of errors on motor memory. *Front. Hum. Neurosci.* **7**, 307 (2013).
35. Mawase, F., Shmuelof, L., Bar-Haim, S. & Karniel, A. Savings in locomotor adaptation explained by changes in learning parameters following initial adaptation. *J. Neurophysiol.* **111**, 1444–1454 (2014).
36. Coltman, S. K., Cashaback, J. G. A. & Gribble, P. L. Both fast and slow learning processes contribute to savings following sensorimotor adaptation. *J. Neurophysiol.* **121**, 1575–1583 (2019).
37. Day, K. A., Leech, K. A., Roemmich, R. T. & Bastian, A. J. Accelerating locomotor savings in learning: compressing four training days to one. *J. Neurophysiol.* **119**, 2100–2113 (2018).
38. Wigmore, V., Tong, C. & Flanagan, J. R. Visuomotor rotations of varying size and direction compete for a single internal model in motor working memory. *J. Exp. Psychol. Hum. Percept. Perform.* **28**, 447–457 (2002).
39. Miall, R. C., Jenkinson, N. & Kulkarni, K. Adaptation to rotated visual feedback: a re-examination of motor interference. *Exp. Brain Res.* **154**, 201–210 (2004).
40. Krakauer, J. W., Ghez, C. & Ghilardi, M. F. Adaptation to visuomotor transformations: consolidation, interference, and forgetting. *J. Neurosci.* **25**, 473–478 (2005).
41. Tulving, E. Retrograde amnesia in free recall. *Science* **164**, 88–90 (1969).
42. Caithness, G., Osu, R., Bays, P., Chase, H., Klassen, J., Kawato, M., Wolpert, D. M. & Flanagan, J. R. Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. *J. Neurosci.* **24**, 8662–8671 (2004).

43. Criscimagna-Hemminger, S. E. & Shadmehr, R. Consolidation Patterns of Human Motor Memory. *J. Neurosci.* **28**, 9610–9618 (2008).
44. Pekny, S. E., Criscimagna-Hemminger, S. E. & Shadmehr, R. Protection and Expression of Human Motor Memories. *J. Neurosci.* **31**, 13829–13839 (2011).
45. Colagiorgio, P., Bertolini, G., Bockisch, C. J., Straumann, D. & Ramat, S. Multiple timescales in the adaptation of the rotational VOR. *J. Neurophysiol.* **113**, 3130–3142 (2015).
46. Stollhoff, N. Spontaneous Recovery from Extinction Depends on the Reconsolidation of the Acquisition Memory in an Appetitive Learning Paradigm in the Honeybee (*Apis mellifera*). *J. Neurosci.* **25**, 4485–4492 (2005).
47. Kim, S., Ogawa, K., Lv, J., Schweighofer, N. & Imamizu, H. Neural Substrates Related to Motor Memory with Multiple Timescales in Sensorimotor Adaptation. *PLoS Biol.* **13**, e1002312 (2015).
48. Lee, J. Y. & Schweighofer, N. Dual adaptation supports a parallel architecture of motor memory. *J. Neurosci.* **29**, 10396–10404 (2009).
49. Zarah, E., Weston, G. D., Liang, J., Mazzoni, P. & Krakauer, J. W. Explaining Savings for Visuomotor Adaptation: Linear Time-Invariant State-Space Models Are Not Sufficient. *J. Neurophysiol.* **100**, 2537–2548 (2008).
50. Lerner, G., Albert, S., T., Caffaro, P. A., Villalta, J. I., Jacobacci, F., Shadmehr, R. & Della-Maggiore V. The Origins of Anterograde Interference in Visuomotor Adaptation. *Cerebral Cortex* (2020). doi:10.1093/cercor/bhaa016
51. Marko, M. K., Haith, A. M., Harran, M. D. & Shadmehr, R. Sensitivity to prediction error in reach adaptation. *J. Neurophysiol.* **108**, 1752–1763 (2012).
52. Kim, H. E., Morehead, J. R., Parvin, D. E., Moazzezi, R. & Ivry, R. B. Invariant errors reveal limitations in motor correction rather than constraints on error sensitivity. *Commun. Biol.* **1**, 19 (2018).
53. Kojima, Y. & Soetedjo, R. Change in sensitivity to visual error in superior colliculus during saccade adaptation. *Sci. Rep.* **7**, 9566 (2017).
54. Huang, V. S., Haith, A., Mazzoni, P. & Krakauer, J. W. Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. *Neuron* **70**, 787–801 (2011).
55. Kording, K. P., Tenenbaum, J. B. & Shadmehr, R. The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat. Neurosci.* **10**, 779–786 (2007).
56. Wei, K. & Körding, K. Uncertainty of feedback and state estimation determines the speed of

- motor adaptation. *Front. Comput. Neurosci.* **4**, 11 (2010).
57. Burge, J., Ernst, M. O. & Banks, M. S. The statistical determinants of adaptation rate in human reaching. *J. Vis.* **8**, 1–19 (2008).
 58. van der Vliet, R., Frens, M. A., de Vreede, L., Jonker, Z. D., Ribbers, G. M., Selles, R. W., van der Geest, J. N. & Donchin, O. Individual Differences in Motor Noise and Adaptation Rate Are Optimally Related. *eNeuro* **5**, (2018).
 59. van Beers, R. J. How does our motor system determine its learning rate? *PLoS One* **7**, e49373–e49373 (2012).
 60. Kalman, R. A New Approach to Linear Filtering and Prediction Problems. *ASME J. Basic Eng.* 34–45 (1960).
 61. Kojima, Y. Memory of Learning Facilitates Saccadic Adaptation in the Monkey. *J. Neurosci.* **24**, 7531–7539 (2004).
 62. Fernandez-Ruiz, J., Wong, W., Armstrong, I. T. & Flanagan, J. R. Relation between reaction time and reach errors during visuomotor adaptation. *Behav. Brain Res.* **219**, 8–14 (2011).
 63. Hegele, M. & Heuer, H. The impact of augmented information on visuo-motor adaptation in younger and older adults. *PLoS One* **5**, e12071–e12071 (2010).
 64. Wilterson, S. A. & Taylor, J. A. Implicit visuomotor adaptation remains limited after several days of training. *bioRxiv* (2019). doi:10.1101/711598
 65. Taylor, J. A. & Ivry, R. B. Flexible cognitive strategies during motor learning. *PLoS Comput. Biol.* **7**, e1001096 (2011).
 66. Wilterson, S. A. & Taylor, J. A. Implicit adaptation processes appear unable to account for learning in sensorimotor adaptation tasks. *bioRxiv* (2019). doi:10.1101/711598
 67. Leow, L. A., Marinovic, W., de Rugy, A. & Carroll, T. J. Task errors drive memories that improve sensorimotor adaptation. *J. Neurosci.* (2020). doi:10.1523/JNEUROSCI.1506-19.2020
 68. Morehead, J. R., Taylor, J. A., Parvin, D. E. & Ivry, R. B. Characteristics of Implicit Sensorimotor Adaptation Revealed by Task-irrelevant Clamped Feedback. *J. Cogn. Neurosci.* **29**, 1061–1074 (2017).
 69. Bond, K. M. & Taylor, J. A. Flexible explicit but rigid implicit learning in a visuomotor adaptation task. *J. Neurophysiol.* **113**, 3836–3849 (2015).
 70. Trewartha, K. M., Garcia, A., Wolpert, D. M. & Flanagan, J. R. Fast But Fleeting: Adaptive Motor Learning Processes Associated with Aging and Cognitive Decline. *J. Neurosci.* **34**, 13411–13421 (2014).

71. Galea, J. M., Mallia, E., Rothwell, J. & Diedrichsen, J. The dissociable effects of punishment and reward on motor learning. *Nat. Neurosci.* **18**, 597–602 (2015).
72. McDougle, S. D., Bond, K. M. & Taylor, J. A. Explicit and Implicit Processes Constitute the Fast and Slow Processes of Sensorimotor Learning. *J. Neurosci.* **35**, 9568–9579 (2015).
73. Inoue, M., Uchimura, M., Karibe, A., O'Shea, J., Rossetti, Y. & Kitazawa, S. Three timescales in prism adaptation. *J. Neurophysiol.* **113**, 328–338 (2015).
74. Cheng, S. & Sabes, P. N. Modeling sensorimotor learning with linear dynamical systems. *Neuronal Comput.* **18**, 760–793 (2006).
75. Wu, H. G., Miyamoto, Y. R., Castro, L. N. G., Ölveczky, B. P. & Smith, M. A. Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat. Neurosci.* **17**, 312–321 (2014).
76. Tanaka, H., Krakauer, J. W. & Sejnowski, T. J. Generalization and multirate models of motor adaptation. *Neural Comput.* **24**, 939–66 (2012).
77. Dempster, A. P., Laird, N. M. & Rubin, D. B. Maximum Likelihood from Incomplete Data via the EM Algorithm. *J. R. St* **39**, 1–38 (1977).
78. Wu, C. F. J. On the convergence properties of the EM algorithm. *Ann. Stat.* **11**, 95–103 (1983).
79. Ghahramani, Z. & Hinton, G. Parameter estimation for linear dynamical systems. *Tech. Rep. CRG-TR-96-2, Dept. Comp. Sci., Univ. Toronto* (1996).
80. Nettleton, D. Convergence properties of the EM algorithm in constrained parameter spaces. *Can. J. Stat.* **27**, 639–648 (1999).
81. Shadmehr, R. & Mussa-Ivaldi, F. A. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* **14**, 3208–3224 (1994).
82. Sing, G. C., Joiner, W. M., Nanayakkara, T., Brayanov, J. B. & Smith, M. A. Primitives for motor adaptation reflect correlated neural tuning to position and velocity. *Neuron* **64**, 575–589 (2009).
83. Strassman, A., Highstein, S. M. & McCrea, R. A. Anatomy and physiology of saccadic burst neurons in the alert squirrel monkey. I. Excitatory burst neurons. *J. Comp. Neurol.* **249**, 337–357 (1986).
84. Robinson, D. A. Oculomotor unit behavior in the monkey. *J. Neurophysiol.* **33**, 393–403 (1970).
85. Robinson, D. A. Models of the saccadic eye movement control system. *Kybernetik* **14**, 71–83 (1973).
86. Cohen, B. & Komatsuzaki, A. Eye movements induced by stimulation of the pontine reticular formation: evidence for integration in oculomotor pathways. *Exp. Neurol.* **36**, 101–117 (1972).

87. Gandhi, N. J., Barton, E. J. & Sparks, D. L. Coordination of eye and head components of movements evoked by stimulation of the paramedian pontine reticular formation. *Exp. brain Res.* **189**, 35–47 (2008).
88. Godaux, E., Mettens, P. & Cheron, G. Differential effect of injections of kainic acid into the prepositus and the vestibular nuclei of the cat. *J. Physiol.* **472**, 459–482 (1993).
89. Cheron, G., Godaux, E., Laune, J. M. & Vanderkelen, B. Lesions in the cat prepositus complex: effects on the vestibulo-ocular reflex and saccades. *J. Physiol.* **372**, 75–94 (1986).
90. Kaneko, C. R. Eye movement deficits following ibotenic acid lesions of the nucleus prepositus hypoglossi in monkeys II. Pursuit, vestibular, and optokinetic responses. *J. Neurophysiol.* **81**, 668–681 (1999).
91. McFarland, J. L. & Fuchs, A. F. Discharge patterns in nucleus prepositus hypoglossi and adjacent medial vestibular nucleus during horizontal eye movement in behaving macaques. *J. Neurophysiol.* **68**, 319–332 (1992).
92. Shemesh, A. A. & Zee, D. S. Eye Movement Disorders and the Cerebellum. *J. Clin. Neurophysiol.* **36**, 405–414 (2019).
93. Zee, D. S., Leigh, R. J. & Mathieu-Millaire, F. Cerebellar control of ocular gaze stability. *Ann. Neurol.* **7**, 37–40 (1980).
94. Crawford, J. D., Cadera, W. & Vilis, T. Generation of torsional and vertical eye position signals by the interstitial nucleus of Cajal. *Science* **252**, 1551–1553 (1991).
95. Klier, E. M., Wang, H., Constantin, A. G. & Crawford, J. D. Midbrain control of three-dimensional head orientation. *Science* **295**, 1314–1316 (2002).
96. Shaikh, A. G., Wong, A. L., Zee, D. S. & Jinnah, H. A. Keeping your head on target. *J. Neurosci.* **33**, 11281–11295 (2013).
97. Guo, J. Z., Graves, A. R., Guo, W. W., Zheng, J., Lee, A., Rodriguez-Gonzalez, J., Li, N., Macklin, J. J., Phillips, J. W., Mensh, B. D., Branson, K. & Hantman, A. W. Cortex commands the performance of skilled movement. *Elife* **4**, e10774 (2015).
98. Mori, S., Kawahara, K., Sakamoto, T., Aoki, M. & Tomiyama, T. Setting and resetting of level of postural muscle tone in decerebrate cat by stimulation of brain stem. *J. Neurophysiol.* **48**, 737–748 (1982).
99. Giszter, S. F., Mussa-Ivaldi, F. A. & Bizzi, E. Convergent force fields organized in the frog's spinal cord. *J. Neurosci.* **13**, 467–491 (1993).
100. Albert, S. T. & Shadmehr, R. The Neural Feedback Response to Error As a Teaching Signal for the

- Motor Learning System. *J. Neurosci.* **36**, 4832–4845 (2016).
101. Albert, S. T. & Shadmehr, R. Estimating properties of the fast and slow adaptive processes during sensorimotor adaptation. *J. Neurophysiol.* **119**, 1367–1393 (2018).
 102. Albert, S. T., Jang, J., Sheahan, H., Teunissen, L., Vandevoorde, K. & Shadmehr, R. Asymptotic limits of sensorimotor adaptation. *bioRxiv* (2019). doi:10.1101/868406.
 103. Albert, S. T., Hadjiosif, A. M., Jang, J., Zimnik, A. J., Soteropoulos, D. S., Baker, S. N., Churchland, M. M., Krakauer, J. W. & Shadmehr, R. Postural control of arm and fingers through integration of movement commands. *Elife* **9**, (2020).
 104. Thoroughman, K. A. & Shadmehr, R. Electromyographic correlates of learning an internal model of reaching movements. *J. Neurosci.* **19**, 8573–8588 (1999).
 105. Franklin, D. W., Osu, R., Burdet, E., Kawato, M. & Milner, T. E. Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. *J. Neurophysiol.* **90**, 3270–3282 (2003).
 106. Milner, T. E. & Franklin, D. W. Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. *J. Physiol.* **567**, 651–664 (2005).
 107. Ahmadi-Pajouh, M. A., Towhidkhan, F. & Shadmehr, R. Preparing to reach: selecting an adaptive long-latency feedback controller. *J. Neurosci.* **32**, 9537–9545 (2012).
 108. Strick, P. L. Cerebellar involvement in volitional muscle responses to load change. in *Cerebral motor control in man: long loop mechanisms* (ed. Desmedt, J.) 85–93 (Basel: Karger, 1978).
 109. Fine, M. S. & Thoroughman, K. A. Motor adaptation to single force pulses: sensitive to direction but insensitive to within-movement pulse placement and magnitude. *J. Neurophysiol.* **96**, 710–720 (2006).
 110. Wei, K., Wert, D. & Kording, K. The nervous system uses nonspecific motor learning in response to random perturbations of varying nature. *J. Neurophysiol.* **104**, 3053–3063 (2010).
 111. Osu, R., Franklin, D. W., Kato, H., Gomi, H., Domen, K., Yoshioka, T. & Kawato, M. Short- and long-term changes in joint co-contraction associated with motor learning as revealed from surface EMG. *J. Neurophysiol.* **88**, 991–1004 (2002).
 112. Haith, A., Pekny, S. E., Shadmehr, R. & Krakauer, J. W. Evidence for model-free learning during force field adaptation. in *Advances in Computational Motor Control* (2011).
 113. Burdet, E., Osu, R., Franklin, D. W., Milner, T. E. & Kawato, M. The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* **414**, 446–449 (2001).
 114. Franklin, D. W., Liaw, G., Milner, T. E., Osu, R., Burdet, E. & Kawato, M. Endpoint stiffness of the

- arm is directionally tuned to instability in the environment. *J. Neurosci.* **27**, 7705–7716 (2007).
115. Wallman, J. & Fuchs, A. F. Saccadic gain modification: visual error drives motor adaptation. *J. Neurophysiol.* **80**, 2405–2416 (1998).
 116. Galea, J. M., Vazquez, A., Pasricha, N., Orban De Xivry, J. J. & Celnik, P. Dissociating the roles of the cerebellum and motor cortex during adaptive learning: The motor cortex retains what the cerebellum learns. *Cereb. Cortex* **21**, 1761–1770 (2011).
 117. Kojima, Y., Soetedjo, R. & Fuchs, A. F. Changes in Simple Spike Activity of Some Purkinje Cells in the Oculomotor Vermis during Saccade Adaptation Are Appropriate to Participate in Motor Learning. *J. Neurosci.* **30**, 3715–3727 (2010).
 118. Yang, Y. & Lisberger, S. G. Role of plasticity at different sites across the time course of cerebellar motor learning. *J. Neurosci.* **34**, 7077–90 (2014).
 119. Vaswani, P. A. & Shadmehr, R. Decay of motor memories in the absence of error. *J. Neurosci.* **33**, 7700–7709 (2013).
 120. Brennan, A. E. & Smith, M. A. The Decay of Motor Memories Is Independent of Context Change Detection. *PLoS Comput. Biol.* **11**, e1004278 (2015).
 121. Pekny, S. E. & Shadmehr, R. Optimizing effort: increased efficiency of motor memory with time away from practice. *J. Neurophysiol.* **113**, 445–454 (2015).
 122. Joiner, W. M. & Smith, M. A. Long-term retention explained by a model of short-term learning in the adaptive control of reaching. *J. Neurophysiol.* **100**, 2948–2955 (2008).
 123. Krakauer, J. W., Pine, Z. M., Ghilardi, M. F. & Ghez, C. Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J. Neurosci.* **20**, 8916–8924 (2000).
 124. Smith, M. A. & Shadmehr, R. Modulation of the rate of error-dependent learning by statistical properties of the task. in *Advances in Computational Motor Control* (2004).
 125. Gonzalez Castro, L. N., Hadjiosif, A. M., Hemphill, M. A. & Smith, M. A. Environmental Consistency Determines the Rate of Motor Adaptation. *Curr. Biol.* **24**, 1050–1061 (2014).
 126. Saint-Cyr, J. A. & Courville, J. Descending projections to the inferior olive from the mesencephalon and superior colliculus in the cat. An autoradiographic study. *Exp. Brain Res.* **45**, 333–348 (1982).
 127. Kojima, Y. & Soetedjo, R. Selective reward affects the rate of saccade adaptation. *Neuroscience* **355**, 113–125 (2017).
 128. Leow, L. A., Loftus, A. M. & Hammond, G. R. Impaired savings despite intact initial learning of motor adaptation in Parkinson’s disease. *Exp. Brain Res.* **218**, 295–304 (2012).

129. Leow, L. A., de Rugy, A., Loftus, A. M. & Hammond, G. Different mechanisms contributing to savings and anterograde interference are impaired in Parkinson's disease. *Front. Hum. Neurosci.* **7**, 55 (2013).
130. Bedard, P. & Sanes, J. N. Basal ganglia-dependent processes in recalling learned visual-motor adaptations. *Exp. Brain Res.* **209**, 385–393 (2011).
131. Marinelli, L., Crupi, D., Di Rocco, A., Bove, M., Eidelberg, D., Abbruzzese, G. & Ghilardi, M. F. Learning and consolidation of visuo-motor adaptation in Parkinson's disease. *Parkinsonism Relat. Disord.* **15**, 6–11 (2009).
132. Leow, L. A., Gunn, R., Marinovic, W. & Carroll, T. J. Estimating the implicit component of visuomotor rotation learning by constraining movement preparation time. *J. Neurophysiol.* **118**, 666–676 (2017).
133. Takahashi, C. D., Scheidt, R. A. & Reinkensmeyer, D. J. Impedance control and internal model formation when reaching in a randomly varying dynamical environment. *J. Neurophysiol.* **86**, 1047–1051 (2001).
134. Srimal, R., Diedrichsen, J., Ryklin, E. B. & Curtis, C. E. Obligatory Adaptation of Saccade Gains. *J. Neurophysiol.* **99**, 1554–1558 (2008).
135. Donchin, O., Francis, J. T. & Shadmehr, R. 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**, 9032–45 (2003).
136. van Beers, R. J., Brenner, E. & Smeets, J. B. J. Random walk of motor planning in task-irrelevant dimensions. *J. Neurophysiol.* **109**, 969–977 (2013).
137. Brennan, A. E. & Smith, M. A. Motor memories are confined to distinct channels with differing stability across time and experience. in *Society for Neuroscience* 57.14 (2016).
138. Chen-Harris, H., Joiner, W. M., Ethier, V., Zee, D. S. & Shadmehr, R. Adaptive Control of Saccades via Internal Feedback. *J. Neurosci.* **28**, 2804–2813 (2008).
139. Hanajima, R. *et al.* Modulation of error-sensitivity during a prism adaptation task in people with cerebellar degeneration. *J. Neurophysiol.* **114**, 2460–2471 (2015).
140. Fujita, H. & Sugihara, I. Branching patterns of olivocerebellar axons in relation to the compartmental organization of the cerebellum. *Front. Neural Circuits* **7**, 3 (2013).
141. Krakauer, J. W. Motor learning and consolidation: the case of visuomotor rotation. *Adv. Exp. Med. Biol.* **629**, 405–421 (2009).
142. Brashers-Krug, T., Shadmehr, R. & Bizzi, E. Consolidation in human motor memory. *Nature* **382**,

- 252–255 (1996).
143. Sing, G. C. & Smith, M. A. Reduction in learning rates associated with anterograde interference results from interactions between different timescales in motor adaptation. *PLoS Comput. Biol.* **6**, e1000893 (2010).
 144. Oldfield, R. C. The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia* **9**, 97–113 (1971).
 145. Shadmehr, R. & Brashers-Krug, T. Functional stages in the formation of human long-term motor memory. *J. Neurosci.* **17**, 409–419 (1997).
 146. Tong, C. & Flanagan, J. R. Task-Specific Internal Models for Kinematic Transformations. *J. Neurophysiol.* **90**, 578–585 (2003).
 147. Braun, D. A., Aertsen, A., Wolpert, D. M. & Mehring, C. Motor task variation induces structural learning. *Curr. Biol.* **19**, 352–357 (2009).
 148. Wei, K. & Kording, K. Relevance of error: what drives motor adaptation? *J. Neurophysiol.* **101**, 655–664 (2009).
 149. Sing, G. C. & Smith, M. A. Reduction in learning rates associated with anterograde interference results from interactions between different timescales in motor adaptation. *PLoS Comput. Biol.* **6**, e1000893 (2010).
 150. Lechner, H. A., Squire, L. R. & Byrne, J. H. 100 years of consolidation--remembering Muller and Pilzecker. *Learn. Mem.* **6**, 77–87 (1999).
 151. Korman, M., Doyon, J., Doljansky, J., Carrier, J., Dagan, Y. & Karni, A. Daytime sleep condenses the time course of motor memory consolidation. *Nat. Neurosci.* **10**, 1206–1213 (2007).
 152. Walker, M. P., Brakefield, T., Hobson, J. A. & Stickgold, R. Dissociable stages of human memory consolidation and reconsolidation. *Nature* **425**, 616–620 (2003).
 153. Overduin, S. A., Richardson, A. G., Lane, C. E., Bizzi, E. & Press, D. Z. Intermittent practice facilitates stable motor memories. *J. Neurosci.* **26**, 11888–11892 (2006).
 154. Bock, O., Schneider, S. & Bloomberg, J. Conditions for interference versus facilitation during sequential sensorimotor adaptation. *Exp. Brain Res.* **138**, 359–365 (2001).
 155. Goedert, K. M. & Willingham, D. B. Patterns of interference in sequence learning and prism adaptation inconsistent with the consolidation hypothesis. *Learn. Mem.* **9**, 279–292 (2002).
 156. Aberg, K. C. & Herzog, M. H. Does perceptual learning suffer from retrograde interference? *PLoS One* **5**, e14161 (2010).
 157. McGeoch, J. A. The formal criteria of a systematic psychology. *Psychol. Rev.* **40**, 1–12 (1933).

158. Houston, J. P. Stimulus selection as influenced by degrees of learning, attention, prior associations, and experience with the stimulus components. *J. Exp. Psychol.* **73**, 509–516 (1967).
159. Howe, T. Effects of delayed interference on List 1 recall. *J. Exp. Psychol.* **80**, 120–124 (1969).
160. Wixted, J. T. The Psychology and Neuroscience of Forgetting. *Annu. Rev. Psychol.* **55**, 235–269 (2004).
161. Cantarero, G., Tang, B., O'Malley, R., Salas, R. & Celnik, P. Motor learning interference is proportional to occlusion of LTP-like plasticity. *J. Neurosci.* **33**, 4634–4641 (2013).
162. Rioult-Pedotti, M. S., Friedman, D., Hess, G. & Donoghue, J. P. Strengthening of horizontal cortical connections following skill learning. *Nat. Neurosci.* **1**, 230–234 (1998).
163. Rioult-Pedotti, M. S., Friedman, D. & Donoghue, J. P. Learning-induced LTP in neocortex. *Science* **290**, 533–536 (2000).
164. Della-Maggiore, V., Villalta, J. I., Kovacevic, N. & McIntosh, A. R. Functional Evidence for Memory Stabilization in Sensorimotor Adaptation: A 24-h Resting-State fMRI Study. *Cereb. Cortex* **27**, 1748–1757 (2017).
165. Della-Maggiore, V., Malfait, N., Ostry, D. J. & Paus, T. Stimulation of the posterior parietal cortex interferes with arm trajectory adjustments during the learning of new dynamics. *J. Neurosci.* **24**, 9971–9976 (2004).
166. Richardson, A. G., Overduin, S. A., Valero-Cabre, A., Padoa-Schioppa, C., Pascual-Leone, A., Bizzi, E. & Press, D. Z. Disruption of primary motor cortex before learning impairs memory of movement dynamics. *J. Neurosci.* **26**, 12466–12470 (2006).
167. Hadipour-Niktarash, A., Lee, C. K., Desmond, J. E. & Shadmehr R. Impairment of retention but not acquisition of a visuomotor skill through time-dependent disruption of primary motor cortex. *J. Neurosci.* **27**, 13413–13419 (2007).
168. Landi, S. M., Baguear, F. & Della-Maggiore, V. One week of motor adaptation induces structural changes in primary motor cortex that predict long-term memory one year later. *J. Neurosci.* **31**, 11808–11813 (2011).
169. Donchin, O., Rabe, K., Diedrichsen, J., Lally, N., Schooch, B., Gizewski, E. R. & Timmann, D. Cerebellar regions involved in adaptation to force field and visuomotor perturbation. *J. Neurophysiol.* **107**, 134–147 (2012).
170. Robinson, F. R., Noto, C. T. & Bevans, S. E. Effect of Visual Error Size on Saccade Adaptation in Monkey. *J. Neurophysiol.* **90**, 1235–1244 (2003).
171. Malone, L. A., Vasudevan, E. V. L. & Bastian, A. J. Motor Adaptation Training for Faster

- Relearning. *J. Neurosci.* **31**, 15136–15143 (2011).
172. Fernandes, H. L., Stevenson, I. H. & Kording, K. P. Generalization of stochastic visuomotor rotations. *PLoS One* **7**, e43016 (2012).
 173. Therrien, A. S., Wolpert, D. M. & Bastian, A. J. Increasing Motor Noise Impairs Reinforcement Learning in Healthy Individuals. *eNeuro* **5**, (2018).
 174. Havermann, K. & Lappe, M. The Influence of the Consistency of Postsaccadic Visual Errors on Saccadic Adaptation. *J. Neurophysiol.* **103**, 3302–3310 (2010).
 175. Baddeley, R. J., Ingram, H. A. & Miall, R. C. System identification applied to a visuomotor task: near-optimal human performance in a noisy changing task. *J. Neurosci.* **23**, 3066–3075 (2003).
 176. McDougle, S. D. & Taylor, J. A. Dissociable cognitive strategies for sensorimotor learning. *Nat. Commun.* **10**, 40 (2019).
 177. Shadmehr, R., Brandt, J. & Corkin, S. Time-dependent motor memory processes in amnesic subjects. *J. Neurophysiol.* **80**, 1590–1597 (1998).
 178. Hwang, E. J., Smith, M. A. & Shadmehr, R. Dissociable effects of the implicit and explicit memory systems on learning control of reaching. *Exp. brain Res.* **173**, 425–437 (2006).
 179. Krakauer, J. W., Mazzoni, P., Ghazizadeh, A., Ravindran, R. & Shadmehr, R. Generalization of motor learning depends on the history of prior action. *PLoS Biol.* **4**, e316 (2006).
 180. Rabe, K., Livne, O., Gizewski, E. R., Aurich, V., Beck, A., Timmann, D. & Donchin, O. Adaptation to Visuomotor Rotation and Force Field Perturbation Is Correlated to Different Brain Areas in Patients With Cerebellar Degeneration. *J. Neurophysiol.* **101**, 1961–1971 (2009).
 181. Kojima, Y. & Soetedjo, R. Elimination of the error signal in the superior colliculus impairs saccade motor learning. *Proc. Natl. Acad. Sci.* **115**, E8987–E8995 (2018).
 182. Avraham, G., Keizman, M. & Shmuelof, L. Environmental Consistency Modulation of Error Sensitivity During Motor Adaptation is Explicitly Controlled. *J. Neurophysiol.* (2019).
doi:10.1152/jn.00080.2019
 183. Heuer, H. & Hegele, M. Adaptation to visuomotor rotations in younger and older adults. *Psychol. Aging* **23**, 190–202 (2008).
 184. Vandevoorde, K. & Orban de Xivry, J.-J. Internal model recalibration does not deteriorate with age while motor adaptation does. *Neurobiol. Aging* **80**, 138–153 (2019).
 185. Soetedjo, R., Kojima, Y. & Fuchs, A. F. Complex spike activity in the oculomotor vermis of the cerebellum: a vectorial error signal for saccade motor learning? *J. Neurophysiol.* **100**, 1949–1966 (2008).

186. Yadav, V. & Sainburg, R. L. Motor lateralization is characterized by a serial hybrid control scheme. *Neuroscience* **196**, 153–167 (2011).
187. Lametti, D. R., Houle, G. & Ostry, D. J. Control of movement variability and the regulation of limb impedance. *J. Neurophysiol.* **98**, 3516–3524 (2007).
188. Schabowsky, C. N., Hidler, J. M. & Lum, P. S. Greater reliance on impedance control in the nondominant arm compared with the dominant arm when adapting to a novel dynamic environment. *Exp. brain Res.* **182**, 567–577 (2007).
189. Todorov, E. & Jordan, M. I. Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* **5**, 1226–1235 (2002).
190. Liu, D. & Todorov, E. Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J. Neurosci.* **27**, 9354–9368 (2007).
191. Ghez, C., Scheidt, R. & Heijink, H. Different learned coordinate frames for planning trajectories and final positions in reaching. *J. Neurophysiol.* **98**, 3614–3626 (2007).
192. Sachs, N. A., Ruiz-Torres, R., Perreault, E. J. & Miller, L. E. Brain-state classification and a dual-state decoder dramatically improve the control of cursor movement through a brain-machine interface. *J. Neural Eng.* **13**, 16009 (2016).
193. Miri, A., Daie, K., Arrenberg, A. B., Baier, H., Aksay, E. & Tank, D. W. Spatial gradients and multidimensional dynamics in a neural integrator circuit. *Nat. Neurosci.* **14**, 1150–1159 (2011).
194. Humphrey, D. R. & Reed, D. J. Separate cortical systems for control of joint movement and joint stiffness: reciprocal activation and coactivation of antagonist muscles. *Adv. Neurol.* **39**, 347–372 (1983).
195. Kurtzer, I., Herter, T. M. & Scott, S. H. Random change in cortical load representation suggests distinct control of posture and movement. *Nat. Neurosci.* **8**, 498–504 (2005).
196. Kaneko, C. R. Eye movement deficits after ibotenic acid lesions of the nucleus prepositus hypoglossi in monkeys. I. Saccades and fixation. *J. Neurophysiol.* **78**, 1753–1768 (1997).
197. Lara, A. H., Cunningham, J. P. & Churchland, M. M. Different population dynamics in the supplementary motor area and motor cortex during reaching. *Nat. Commun.* **9**, 2754 (2018).
198. Lara, A. H., Elsayed, G. F., Zimnik, A. J., Cunningham, J. P. & Churchland, M. M. Conservation of preparatory neural events in monkey motor cortex regardless of how movement is initiated. *Elife* **7**, (2018).
199. Soteropoulos, D. S., Williams, E. R. & Baker, S. N. Cells in the monkey ponto-medullary reticular formation modulate their activity with slow finger movements. *J. Physiol.* **590**, 4011–4027 (2012).

200. Makin, T. R. & Orban de Xivry, J.-J. Ten common statistical mistakes to watch out for when writing or reviewing a manuscript. *Elife* **8**, e48175 (2019).
201. Mussa-Ivaldi, F. A., Hogan, N. & Bizzi, E. Neural, mechanical, and geometric factors subserving arm posture in humans. *J. Neurosci.* **5**, 2732–2743 (1985).
202. Shadmehr, R., Mussa-Ivaldi, F. A. & Bizzi, E. Postural force fields of the human arm and their role in generating multijoint movements. *J. Neurosci.* **13**, 45–62 (1993).
203. Optican, L. M. & Miles, F. A. Visually induced adaptive changes in primate saccadic oculomotor control signals. *J. Neurophysiol.* **54**, 940–958 (1985).
204. Zackowski, K. M., Dromerick, A. W., Sahrman, S. A., Thach, W. T. & Bastian, A. J. How do strength, sensation, spasticity and joint individuation relate to the reaching deficits of people with chronic hemiparesis? *Brain* **127**, 1035–1046 (2004).
205. Roh, J., Rymer, W. Z. & Beer, R. F. Evidence for altered upper extremity muscle synergies in chronic stroke survivors with mild and moderate impairment. *Front. Hum. Neurosci.* **9**, 6 (2015).
206. Scheidt, R. A. & Stoekmann, T. Reach adaptation and final position control amid environmental uncertainty after stroke. *J. Neurophysiol.* **97**, 2824–2836 (2007).
207. Coderre, A. M., Zeid, A. A., Dukelow, S. P., Demmer, M. J., Moore, K. D., Demers, M. J., Bretzke, H., Herter, T. M., Glasgow, J. I., Norman, K. E., Bagg, S. D. & Scott, S. H. Assessment of upper-limb sensorimotor function of subacute stroke patients using visually guided reaching. *Neurorehabil. Neural Repair* **24**, 528–541 (2010).
208. Heald, J. B., Ingram, J. N., Flanagan, J. R. & Wolpert, D. M. Multiple motor memories are learned to control different points on a tool. *Nat. Hum. Behav.* **2**, 300–311 (2018).
209. Mattar, A. A. G., Nasir, S. M., Darainy, M. & Ostry, D. J. Sensory change following motor learning. *Prog. Brain Res.* **191**, 31–44 (2011).
210. Darainy, M., Vahdat, S. & Ostry, D. J. Perceptual learning in sensorimotor adaptation. *J. Neurophysiol.* **110**, 2152–2162 (2013).
211. Feldman, A. G. On the functional tuning of the nervous system in movement control or preservation of stationary pose. II. Adjustable parameters in muscles. *Biofizika* **11**, 498–508 (1966).
212. Feldman, A. G. Once more on the equilibrium-point hypothesis (lambda model) for motor control. *J. Mot. Behav.* **18**, 17–54 (1986).
213. Graziano, M. S. A., Taylor, C. S. R. & Moore, T. Complex movements evoked by microstimulation of precentral cortex. *Neuron* **34**, 841–851 (2002).

214. Harrison, T. C., Ayling, O. G. S. & Murphy, T. H. Distinct cortical circuit mechanisms for complex forelimb movement and motor map topography. *Neuron* **74**, 397–409 (2012).
215. Kim, S. P., Simeral, J. D., Hochberg, L. R., Donoghue, J. P., Friehs, G. M. & Black, M. J. Point-and-click cursor control with an intracortical neural interface system by humans with tetraplegia. *IEEE Trans. Neural Syst. Rehabil. Eng.* **19**, 193–203 (2011).
216. Gilja, V., Nuyujukian, P., Chestek, C. A., Cunningham, J. P., Yu, B. M., Fan, J. W., Churchland, M. M., Kaufman, M. T., Kao, J. C., Ryu, S. I. & Shenoy, K. V. A high-performance neural prosthesis enabled by control algorithm design. *Nat. Neurosci.* **15**, 1752–1757 (2012).
217. Shalit, U., Zinger, N., Joshua, M. & Prut, Y. Descending systems translate transient cortical commands into a sustained muscle activation signal. *Cereb. Cortex* **22**, 1904–1914 (2012).
218. Georgopoulos, A. P., Ashe, J., Smyrnis, N. & Taira, M. The motor cortex and the coding of force. *Science* **256**, 1692–1695 (1992).
219. Sheahan, H. R., Franklin, D. W. & Wolpert, D. M. Motor Planning, Not Execution, Separates Motor Memories. *Neuron* **92**, 773–779 (2016).
220. Howard, I. S., Wolpert, D. M. & Franklin, D. W. The effect of contextual cues on the encoding of motor memories. *J. Neurophysiol.* **109**, 2632–2644 (2013).
221. Griffin, D. M., Hoffman, D. S. & Strick, P. L. Corticomotoneuronal cells are ‘functionally tuned’. *Science* **350**, 667–670 (2015).
222. Sadnicka, A., Stevenson, A., Bhatia, K. P., Rothwell, J. C., Edwards, M. J. & Galea, J. M. High motor variability in DYT1 dystonia is associated with impaired visuomotor adaptation. *Sci. Rep.* **8**, 3653 (2018).
223. Suvrathan, A., Payne, H. L. & Raymond, J. L. Timing Rules for Synaptic Plasticity Matched to Behavioral Function. *Neuron* **92**, 959–967 (2016).
224. Becker, M. I. & Person, A. L. Cerebellar Control of Reach Kinematics for Endpoint Precision. *Neuron* **103**, 335–348.e5 (2019).
225. Abbasi, J. Advanced Brain-Computer Interface for People With Paralysis. *JAMA* **321**, 537 (2019).
226. Crammond, D. J. & Kalaska, J. F. Differential relation of discharge in primary motor cortex and premotor cortex to movements versus actively maintained postures during a reaching task. *Exp. Brain Res.* **108**, 45–61 (1996).

Curriculum Vitae: Scott T. Albert

Department of Biomedical Engineering, Johns Hopkins University
416 Traylor Building, 720 Rutland Ave., Baltimore, MD 21205
Ph: 908-670-7999
Email: salbert8@jhmi.edu

EDUCATION

1. Johns Hopkins University, Baltimore, MD (expected: May, 2020)
Ph.D. Candidate, Biomedical Engineering
Advisor: Reza Shadmehr
2. Lafayette College, Easton, PA (2013)
Major: Chemical Engineering, B.S.
Minor: Mathematics
GPA: 4.0

HONORS & AWARDS

1. **Mette Strand Award** (2020), awarded for outstanding research contributions as a trainee at Johns Hopkins
2. **NCM Scholarship** (2020), scholarship awarded to attend the Society for the Neural Control of Movement meeting
3. **Siebel Scholar Award** (2019), one of 94 graduate students world-wide, recognized for research and leadership
4. **Outstanding Teaching Assistant Award** (2017), awarded by Johns Hopkins Biomedical Engineering
5. **Best Independent Project at The Computational Sensory-Motor Neuroscience School** (2016, Minneapolis)
6. **Ruth L. Kirchstein National Research Service Award** (2015) F31 award from the NIH (NINDS)
7. **Summa Cum Laude with Honors** (2013), graduated with a 4.0 GPA and completed an honors thesis
8. **George Wharton Pepper Prize Finalist** (2013), highest non-academic award offered to a Lafayette senior
9. **Marquis Scholar** (2009-2013), highest academic scholarship offered by Lafayette College
10. **AICHE Donald F. Othmer Award** (2012), highest grade-point average as a sophomore student
11. **William G. McLean Tau Beta Pi Prize** (2011), awarded for campus citizenship, academics, and professionalism
12. **Eugene P. Chase Phi Beta Kappa Prize** (2011), awarded for scholarship as a first-year student
13. **Honor Societies**, Tau Beta Pi (2011), Phi Beta Kappa (2012), Pi Mu Epsilon (2012), Sigma Xi (2013)

PUBLICATIONS

1. **Albert SA**, Hadjosif AM, Jang J, Zimnik A, Churchland MM, Krakauer JW, Shadmehr R (2020) Holding the arm still through integration of cortical commands. *eLife*. <https://doi.org/10.7554/eLife.52507>.
2. **Albert SA**, Jang J, Sheahan H, Teunissen L, Vandevoorde K, Shadmehr R (2020) Asymptotic limits of sensorimotor adaptation. *In submission*. *bioRxiv*: <https://doi.org/10.1101/868406>.
3. Lerner G*, **Albert SA***, Caffaro PA, Villalta JI, Jacobacci F, Shadmehr R, Della-Maggiore V (2020) The origins of anterograde interference in visuomotor adaptation. *Cerebral Cortex*. *bioRxiv*: <https://doi.org/10.1093/cercor/bhaa016>.
*Co-first author.
4. Orozco S, **Albert SA**, Shadmehr R (2020) Spontaneous recovery and the multiple timescales of motor memory. *In preparation*.
5. **Albert SA**, Shadmehr R (2018) Estimating properties of the fast and slow adaptive processes during sensorimotor adaptation. *Journal of Neurophysiology* 119:1367-1393.
6. **Albert SA**, Shadmehr R (2018) The cerebellum learns to predict the physics of our movements. *Think Tank: Forty Neuroscientists Explore the Biological Roots of Human Experience*. Ed. David J. Linden. New Haven and London: Yale University Press.
7. **Albert SA**, Shadmehr R (2016) The neural feedback response to error as a teaching signal for the motor learning system. *Journal of Neuroscience* 36:4832-4845.
8. **Albert SA**, Balaban RS, Neufeld EB, Rossmann JS (2014) Influence of the renal artery ostium flow diverter on hemodynamics and atherogenesis. *Journal of Biomechanics* 47:1594-1602.

SOFTWARE

1. **Albert, SA.** A tool to fit state-space models of adaptation to sensorimotor data (2018). Available at <http://shadmehrlab.org/tools>. This openly distributed toolkit is used to fit state-space models of error-based learning to behavioral data.

CONFERENCE PROCEEDINGS – TALKS

1. **Albert SA,** Hadjosif A, Jang J, Zimnik A, Churchland MM, Krakauer JW, Shadmehr R (2020) Holding the arm still through integration of cortical commands. *Neural Control of Movement*, Dubrovnik, Croatia.
2. **Albert SA,** Hadjosif A, Jang J, Krakauer JW, Shadmehr R (2018) Neural integration in the serial control of reaching and holding still. *Advances in Motor Learning & Motor Control*, San Diego, CA.
3. **Albert SA,** Hadjosif A, Jang J, Krakauer JW, Shadmehr R (2018) Neural integration in the serial control of reaching and holding still. *Sensorimotor Research Day*, Johns Hopkins, Baltimore, MD.
4. **Albert SA,** Shadmehr R (2017) Estimating fast and slow processes of motor learning. *University of Southern California Motor Control Day*, Los Angeles, CA.
5. **Albert SA,** Sheahan, H, Teunissen, L, Vandevoorde, K (2016) A memory of errors modulates the extent of sensorimotor learning. *Advances in Motor Learning & Motor Control*, San Diego, CA.
6. **Albert SA,** Shadmehr R (2015) Identifying the fast and slow adaptive processes in sensorimotor learning. *Advances in Motor Learning & Motor Control*, Chicago, IL.
7. **Albert SA,** Shadmehr R (2014) Feedback-error: a potential teaching signal for motor learning. *Sensorimotor Research Day*, Johns Hopkins, Baltimore, MD.

CONFERENCE PROCEEDINGS – POSTER PRESENTATIONS

1. **Albert SA,** Hadjosif A, Jang J, Zimnik A, Churchland MM, Krakauer JW, Shadmehr R (2019) Holding the arm still through integration of cortical commands. *Society for Neuroscience*, Chicago, IL.
2. **Albert SA,** Lerner G, Caffaro P, Villalta J, Jacobacci F, Shadmehr R, Della-Maggiore V (2019) Anterograde interference develops from a transient reduction in error sensitivity. *Society for Neuroscience*, Chicago, IL.
3. **Albert SA,** Hadjosif A, Jang J, Krakauer JW, Shadmehr R (2018) Force production during holding suggests the presence of a neural integrator for reaching. *Society for Neuroscience*, San Diego, CA.
4. **Albert SA,** Jang J, Shadmehr R (2017) Force production during holding suggests the presence of a neural integrator in reach adaptation. *Studying Complex Behavior*, Weizmann Institute, Israel.
5. **Albert SA,** Jang J, Shadmehr R (2017) Force production during holding suggests the presence of a neural integrator in reach adaptation. *Society for Neuroscience*, Washington, DC.
6. **Albert SA,** Shadmehr R (2016) Estimating fast and slow adaptive processes during sensorimotor learning. *Neural Control of Movement*, Dublin, Ireland.
7. **Albert SA,** Shadmehr R (2016) Estimating properties of the fast and slow adaptive processes during sensorimotor learning. *Society for Neuroscience*, San Diego, CA.
8. **Albert SA,** Shadmehr R (2015) The feedback response to error is a teaching signal during motor adaptation. *Society for Neuroscience*, Chicago, IL.
9. **Albert SA,** Shadmehr R (2015) The feedback response to error is a teaching signal during motor adaptation. *Neural Control of Movement*, Charleston, SC.
10. **Albert SA,** Rossmann JS, Balaban RS (2013) Numerical simulation of blood flow in the renal arteries: influence of the ostium flow diverter. *American Society of Mechanical Engineering. Summer Bioengineering Conference*, Sunriver, OR.

DEPARTMENTAL TALKS

1. **Albert SA**, Hadjiosif A, Jang J, Zimnik A, Churchland MM, Krakauer JW, Shadmehr R (2019) Holding the arm still through integration of cortical commands. *Janelia Farm*, Ashburn, VA.
2. **Albert SA**, Hadjiosif A, Jang J, Krakauer JW, Shadmehr R (2018) Neural integration in the serial control of reaching and holding still. *Columbia University*, New York, NY.
3. **Albert SA**, Jang J, Shadmehr R (2018) Behavioral evidence for neural integration in the control of reaching and holding. *University College London*, London UK.
4. **Albert SA**, Jang J, Shadmehr R (2018) Behavioral evidence for neural integration in the control of reaching and holding. *University of Birmingham*, Birmingham, UK.
5. **Albert SA**, Shadmehr R (2017) Estimating fast and slow processes of motor learning. *University of Southern California*, Los Angeles, CA.
6. **Albert SA**, Shadmehr R (2017) Estimating properties of the fast and slow adaptive learning processes during sensorimotor adaptation. *BME Student Seminar Series*, Johns Hopkins, Baltimore MD.
7. **Albert SA**, Shadmehr R (2015) Feedback response to error as a teaching signal for motor learning. *University of California, Berkeley*, Berkeley, CA.

SERVICE

1. **Thread – The New Social Fabric**
GrandParent at Academy for College and Career Exploration (2018-present)
I lead 8 different families, that support 8 different high school students in Baltimore City with their academic pursuits, employment opportunities, and other socioeconomic barriers they face at home and in their community.
2. **Music and the Arts**
Charles Street Brass Quintet (Co-founder, 2016-2018), Johns Hopkins Jazz Ensemble (lead trumpet player, 2013-present, Johns Hopkins Wind Ensemble (co-lead trumpet player, 2013-2018), Johns Hopkins Theater
Lead actor in 2017 and lead trumpet player in 2019).
3. **Small Foods Party 2020**
Co-organizer of a food competition (300+ attendees) in Baltimore that raises money for Moveable Feast.
4. **Johns Hopkins Graduate Representative Organization**
Soccer team captain in GRO Summer League (2018), the Cooley Center League (2017), and the Charm City Soccer League (2015 and 2017).
5. **Johns Hopkins Biomedical Engineering**
Prospective Student Liaison (2013-2019)
I assisted with student recruitment, hosting, Q&A panels, poster presentations, tours, and peer advising.
6. **Contributing scientific reviewer**
eNeuro (2019), PNAS (2016,1028), Journal of Neuroscience (2017, 2018), Journal of Neurophysiology (2016-2018), Nature Neuroscience (2017), NeuroImage (2017), Current Biology (2016), Science (2016)
7. **Scientific reviewer**
PLoS Computational Biology (2020)

TEACHING & MENTORSHIP

1. **Course Advisor**
Senior Design Project (EN:580.580, Johns Hopkins University, 2019), research advisor to Jihoon Jang
2. **Teaching Assistant**
Learning Theory (EN:580.691, Johns Hopkins University, 2016)
Awarded Outstanding Teaching Assistant Award in the Department of Biomedical Engineering.
3. **Teaching Assistant**
BME Modeling and Design (EN.580.111.01, Johns Hopkins University, 2015).