

DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY


Motor Memories for Expectation and Uncertainty in the Environment

Citation

Hadjiosif, Alkis M. 2015. Motor Memories for Expectation and Uncertainty in the Environment. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Link

http://nrs.harvard.edu/urn-3:HUL.InstRepos:17467305

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material (LAA), as set forth at
https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFlOWI2M2EwYTg

Accessibility

https://accessibility.huit.harvard.edu/digital-accessibility-policy

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story

© 2015 –Alkis Hadjiosif
All rights reserved.

Thesis Advisor
Maurice A. Smith

Author
Alkis M. Hadjiosif

MOTOR MEMORIES FOR EXPECTATION AND UNCERTAINTY IN THE ENVIRONMENT

ABSTRACT

The motor system endows us with the ability to control movement – an ability that provides us, almost exclusively, with the means through which we interact with our environment. To accomplish this, the motor system continuously maintains, adapts, or forms anew memories that support the execution of a multitude of motor tasks, from the basic but critical like breathing or reaching, to complex and elaborate patterns of motion like those of an NBA player or a master cellist. Here, I use a series of behavioral paradigms to probe the mechanisms behind the formation and retention of motor memories related to motor adaptation. Specifically, I investigate the learning of both internal representations of the expected state of the environment but also representations of the uncertainty associated with that expectation.

First, I focus on the interplay between the formation of estimates for the expected environmental dynamics and the formation of estimates about the variability of these dynamics. I show how variability estimates are critical for the control of the safety margin associated with grip forces, and that grip force control is, in fact, three times more sensitive to changes in variability – specifically, the standard deviation, σ – than the expected value of environmental dynamics. This ratio of sensitivities effectively acts to provide a 3-σ confidence level (>99%) against slip. A consequence of heightened sensitivity of grip forces to uncertainty that we observe, is that grip forces strongly increase upon a change in environmental dynamics, regardless of whether that change increased or decreased the strength of the load. This in turn leads to bizarre, asymmetric grip force learning curves to step increases vs. step decreases in dynamics, explaining previous studies where grip force adaptation appeared to outpace manipulatory force adaptation. I proceed to show that this variability-driven safety margin control is based on a low-order variability metric that is robust to outliers in the distribution of environmental dynamics.

Second, I examine the temporal stability of motor adaptation. Using a visuomotor adaptation task, I investigate the speed and breadth of decay due to the passage of time, and find that adaptation consists of distinct temporally-labile and temporally-stable components. Whereas temporally-labile adaptation decays within the timecourse of only a few tens of seconds, the remaining temporally-stable adaptation persists well beyond that timeframe. Moreover, this temporally-stable adaptation predicts and acts as a gateway towards the long-term retention of the trained adaptation. Comparison of the temporal stability of motor adaptation formed from one training session to the next reveals wide-ranging differences that are not only driven by differences across individuals, but, remarkably, by considerable differences across different sessions for the same individual.

Based on a brain stimulation study and the reanalysis of previous studies on cerebellar patients, I proceed to identify the primary motor cortex and the posterior parietal cortex as likely parts of a network involved with temporally-labile adaptation, and the cerebellum as a part of a network involved with temporally-stable adaptation. A different study revealed that, in spite of the fact that temporally-stable adaptation is remarkably persistent even after prolonged active washout, savings – the phenomenon of faster relearning of a previously learned adaptation – is not driven by the reemergence of this strong, stable memory, but, instead, by an increased propensity of the temporally-labile component to relearn the task at hand.

Third, I examine how the rate of adaptation to changes in the environment is modulated by the statistics of that environment. Specifically, I show how the rate at which the motor system adapts to changes in the environment is primarily determined not by the degree to which environment change occurs – i.e. environmental variability – but by the degree to which the changes that do occur persist from one movement to the next – i.e. the trial-to-trial consistency of the environment. Finally, I find that repetition of the same learning stimulus can greatly potentiate the effect of consistency, although, unlike consistency, repetition alone does not increase adaptation rate.

Table of contents

Chapter 1: Introduction and Motivation 1
1.1 Adaptation paradigms used to study motor learning 3
1.2 Forward and inverse internal models for motor adaptation 5
1.3 Generalization of motor adaptation 6
1.4 Scope of this work 7
1.5 References 10
Chapter 2: Flexible control of safety margins for action based on internal estimates of environmental variability 13
2.1 Summary 14
2.2 Introduction 15
2.3 Materials and Methods 17
2.3.1 Participants and Ethics statement 17
2.3.2 General task 17
2.3.3 Data analysis and computational modeling 19
2.4 Results 26
2.4.1 Effects of increased environmental variability on grip and manipulatory forces 26
2.4.2 Increased environmental variability leads to increases in GF but not MF control 29
2.4.3 GF control is even more sensitive to changes in the LF variability than the mean LF 31
2.4.4 Grip force control maintains a safety margin based on statistical confidence 33
2.4.5 A computational model for understanding how GF control responds to both the mean and variability of load force dynamics. 36
2.4.6 The ACEVU model predicts non-monotonic adaptation 37
2.4.7 Grip forces adapt to an additive combination of uncertainty and the expected value 39
2.4.8 Manipulatory forces adapt chiefly to the expected value of environmental dynamics 41
2.4.9 Robustness of results using different metrics 41
2.4.10 Decomposition of grip forces into perturbation specific and nonspecific components 42
2.4.11 Learning curves for the specific and nonspecific components of GF and MF adaptation 44
2.5 Discussion 47
2.5.1 Previous ideas about safety margins for grip force control 48
2.5.2 Previous findings comparing the rates of grip force and manipulatory force adaptation 49
2.5.3 Uncertainty vs Variability 50
2.5.4 Uncertainty-driven control of grip forces and muscular co-contraction 50

2.5.5 Impact of neurologic disease on uncertainty-driven grip force control 51
2.5.6 Uncertainty estimation and metacognition 52
2.6 References 54
Chapter 3: Safety margin control is based on a metric of environmental variability robust to outliers 59
3.1 Summary 60
3.2 Introduction 61
3.3 Materials and Methods 64
3.3.1 Participants and Ethics statement 64
3.3.2 Task and analysis 64
3.4 Results 65
3.3.1 The control of the safety margin for grip force is based on sensitivity to a measure of environmental variability 65
3.3.2 Grip force control is based on a variability measure that is robust to outliers 67
3.3.3 Size of the memory window for estimating variability 71
3.5 Discussion 74
3.6 References 75
Chapter 4: Temporally-stable and temporally-labile memories of motor adaptation 76
4.1 Summary 77
4.2 Introduction 78
4.3 Materials and Methods 80
4.3.1 Apparatus 80
4.3.2 General task 80
4.3.3 Data analysis 81
4.3.4 Specific procedures 82
4.4 Results 87
4.4.1 Visuomotor adaptation consists of distinct temporally-stable and temporally-labile components 87
4.4.2 Intra- vs. inter- individual differences in time-driven decay: moods and personalities in the decay of motor adaptation 93
4.4.3 The neural circuits behind the temporally-labile process of adaptation 99
4.4.4 The neural circuits behind the temporally-stable component of adaptation 102
4.4.5 Temporally-stable adaptation predicts second-day retention 114
4.4.6 The temporal stability of motor adaptation is sensitive to the temporal correspondence between action and feedback 121
4.5 Discussion 127

4.5.1 The stable-labile dichotomy and state-space models of motor adaptation 128
4.5.2 Time-dependent vs. experience-dependent decay 129
4.5.3 Implications for rehabilitation 130
4.5.4 Possible reasons on why visual latencies reduce the temporal stability of motor adaptation 130
4.5.5 Reverse hierarchy theory and the robustness of the globally-generalizing component to increased visual latencies 131
4.6 References 132
Chapter 5: Savings arises from faster acquisition of short-term memories rather than the reemergence of stable long-term memories 136
5.1 Summary 137
5.2 Introduction 138
5.3 Materials and Methods 140
5.3.1 Apparatus 140
5.3.2 Experiment protocol 140
5.3.3 Data analysis 141
Statistical comparisons 141
Data inclusion criteria 141
Estimation of washout time constants 141
Normalization of adaptation data 142
5.4 Results 143
5.4.1 Stable adaptation washes out much slower than overall adaptation 143
5.4.2 Temporally-stable adaptation leads to a long-term memory 145
5.4.3 Overall adaptation displays savings even after complete washout of temporally-stable adaptation 148
5.4.4 Stable adaptation does not display savings 152
5.4.5 Savings is driven exclusively by temporally-labile adaptation 152
5.4.6 Savings does not transfer to the learning of the opposite perturbation 155
5.5 Discussion 158
5.5.1 Linear time-invariant models do not account for all savings in motor adaptation 158
5.5.2 What drives savings in the temporally-labile process? 159
5.5.3 Using savings as an indicator of memory consolidation can be problematic 160
5.6 References 161
Chapter 6: Environmental Consistency Determines the Rate of Motor Adaptation 163
6.1 Summary 164

6.2 Introduction 165
6.3 Experimental Procedures 168
6.3.1 Ethics statement 168
6.3.2 General task description 168
6.3.3 Learning environment experiments 170
6.3.4 Data inclusion criteria 174
6.3.5 Statistical analyses 174
6.3.6 Calculation of confidence intervals for the responses predicted by the combined-CR model 175
6.4 Results 177
6.4.1 Environmental consistency can upregulate and downregulate motor adaptation rates 178
6.4.2 Comparing next-trial adaptive responses with same-trial feedback responses 180
6.4.3 Next-trial adaptive responses are driven by environmental consistency, whereas same-trial feedback responses are driven by environmental variability 181
6.4.4 Increases in motor adaptation rate are not explained by the amount of environmental exposure 185
6.4.5 Specificity of adaptation rate increases for the experienced environment dynamics 188
6.4.6 Repetition contributes to, but is not required for consistency-driven adaptation rate upregulation 191
6.5 Discussion 194
6.5.1 Interactions between environmental consistency and repetition 195
6.5.2 Attempts to manipulate measurement noise and state noise to alter adaptation rates 197
6.5.3 Implications for retention of motor memories 198
6.5.4 Implications for rehabilitation 199
6.6 References 200

Acknowledgements

Completing a PhD dissertation can be a long and arduous task. Through the time I spent working on the research presented in the next pages, I received generous guidance and support from my mentors, academic peers, and personal relationships alike. For that, I am eternally grateful.

First, I would like to thank my advisor, Maurice Smith, for mentoring me through the past 6 years. His remarkable talent for framing intellectual challenges that lead to exciting and fruitful research projects is something I can only aspire to emulate. Moreover, the quantitative and analytical rigor which he brings to every situation has been one of his greatest gifts, and I join my old labmates in hoping that some of it has rubbed off onto me. It is a gift that helped me dive into the exciting field of motor control, which, I have to admit, had one asked me about a year before I met Maurice, my mind would think of a PID control loop. I would also like to thank the other members of my Committee, Bence Ölveczky, Daniel Press, and Rob Howe, for their continued help and support in my research. Particular thanks to Dan, for his help and guidance throughout a stimulating research project (part of chapter 4), and for always enjoying a pun.

Next, I would like to give thanks to my labmates, past and present (Gary Sing, Nick Gonzalez-Castro, Jordan Brayanov, Howard Wu, Yohsuke Miyamoto, Biljana Petreska and Andrew Brennan, in the order I met them), for their help, engaging conversations, and sharing some of the fun aspects of life. I would especially like to thank Nick, with whom I collaborated in the exciting research presented in chapter 6. His great sense of humor, impeccable sense of style, and insanely long variable names in Matlab are things to remember and reckon with. Special thanks to Biljana, with whom I worked with on some of the research presented in chapter 4. Special thanks to Andrew as well, for taking the courage to read through parts of what you are about to witness in the next pages, and make insightful comments. I would also like to extend my thanks to Katherine Mallett, with whom I had the chance to work together in another exciting project, presented in section 4.4.6. In addition, I would like to share my gratitude towards my collaborators who ran some of the numerous experiments presented in this dissertation, Jocelyn Breton,

Simon Orozco, Allison Randolph, Stella St. Hubert, Ben Glauser, Shreya Maheshwari, Jasmine Bailey, and Emerson Fang, as well as the legions of willing test subjects who took part in these experiments.

I would like to thank my friends, both here and back home, for their support through the last six years. My warmest thanks go to my partner, Samuel, for always being there for me; for his love, care, and extraordinary patience with my schedule and peculiar sleeping patterns these last few weeks. Last, but not least, I would like to thank my family for their continued love and support– my parents, Michael and Daphne, and my sister, Georgia. Especially dad, who, throughout our long skype conversations, kept asking “is that published yet?”

Chapter 1: Introduction and Motivation

The motor system endows us with the ability to control movement – an ability that provides us, almost exclusively, with the means through which we interact with our environment. To accomplish this, the motor system continuously maintains, adapts, or forms anew memories that support the execution of a multitude of motor tasks, from the basic but critical like reaching for a glass of water, to complex elaborate patterns of motion like those in an Olympic ice skating performance.

Successful and efficient motor control is critical in everyday life: we use motion to breathe, feed ourselves, travel from home to work, and so on. Also importantly, we command motion in order to communicate: from controlling the dynamics of our vocal apparatus to articulate speech, to controlling our fingers to write or type down our ideas, to sharing our emotions through dance or music. Because of the importance of motor control, the impact of neurological disease that affects the motor system can be devastating to both patients and the ones close to them. Consequences include the loss of ability to perform essential everyday functions, the psychological toll that can be associated with this loss, the need to rely on caregivers, and the economic burden of reduced productivity and increased medical costs. For example, Parkinson's disease alone was estimated, in 2002, to impose an annual burden of 23 billion dollars in the US economy (Huse et al., 2005). Even for Alzheimer's disease, an ailment usually associated with deficits in the cognitive domain, signs of motor dysfunction can predict both cognitive impairments and the need for patient institutionalization (Scarmeas et al., 2005).

At the same time, extraordinary skill in motor control can be source of great awe, appreciation, and reward from society. Consider a virtuoso cellist like Pablo Casals, rapidly jumping from one string to another with great precision across space and time, delivering a cascade of emotions through subtle variations in volume; an acclaimed soprano like Maria Callas, confidently and precisely jumping through a series of rapidly-changing, sometimes impossibly high vocal pitches; a master golfer like Brendon Todd, who aced a hole, 150 yards away, not once, but twice in a row. Understanding the brain mechanisms that can lead to such mastery is another inspiration for motor neuroscience.

In general, motor neuroscience aims to elucidate the determinants that drive the learning of motion and the creation of motor memories, from the chemistry that drives interactions in a neuronal synapse to the quantitative analysis of behavioral movement patterns. Here, I use a series of behavioral paradigms to focus on the mechanisms behind the formation and retention of motor memories related to motor adaptation, investigating the learning of both internal representations of the expected state of the environment and representations of the uncertainty associated with that expectation. In the rest of this chapter, I will briefly mention and discuss a few core concepts in motor learning that are relevant to this dissertation, before offering a brief overview of the other chapters.

1.1 Adaptation paradigms used to study motor learning

A powerful method for studying motor control at the behavioral level is to isolate specific aspects of skilled movement. This allows us to focus on motor tasks with a reduced number of learning parameters, consequently making the identification of the underlying learning elements and processes easier. Once these have been identified, the motor neuroscientist can then examine how the findings can generalize to more complex and natural motor tasks. Here I discuss a few common adaptation paradigms that are in line with this idea.

A motor learning paradigm that led to a wide range of insights about motor control is force-field adaptation, which simulates interactions with a novel object. Participants in this task perform movements (usually trying to move directly from one point to another) that, at some point during training, are perturbed by a force-field whose dynamics can be a function of the motion's kinematics such as position, velocity, acceleration, or linear combinations of them (Shadmehr and Mussa-Ivaldi, 1994; Flanagan et al., 2003; Smith and Shadmehr, 2005; Hwang et al., 2006; Joiner and Smith, 2008; Sing et al., 2009; Criscimagna-Hemminger et al., 2010; Shadmehr et al., 2010; Joiner et al., 2011; Gibo et al., 2013). Forces that depend on kinematics are appropriate for force-field adaptation tasks, because motion dynamics themselves are encoded in the brain as a function of kinematics rather than time (Ashe and Georgopoulos, 1994; Conditt et al., 1997; Hwang et al., 2003; Paninski et al., 2004; Sing et al., 2009; Sing et al., 2013).

For example, a popular type of force-field used in motor studies (including experiments in this dissertation) involves a viscous curl component which produces forces lateral to the direction of motion that scale with the tangential velocity of motion. This force field presents a motor task that is novel and engages the natural motion-dependent elements of motor learning.

Another class of motor adaptation paradigms involves learning transformations that alter the relationship between motion and the resulting visual feedback, engaging the neural maps between motor commands and their visual effect. With direct vision of their hand occluded, subjects in these experiments instead receive computer-controlled visual feedback about their hand position. This allows the experimenter to alter the mapping between hand motion and visual feedback about it (which is usually in the form of a cursor) through transformations that can impose a rotation about the starting position of movement (Krakauer et al., 2000; Krakauer et al., 2005; Braun et al., 2009; Huang et al., 2011; Brayanov et al., 2012; Shmuelof et al., 2012; Kitago et al., 2013), change the gain between the extent of hand and cursor motion (Pine et al., 1996; Vindras and Viviani, 2002; Pearson et al., 2010; Brayanov, 2013), or displace the cursor laterally to the movement direction in proportion to the tangential distance traveled (shear transformations) (Braun et al., 2009).

Other commonly used adaptation paradigms include saccade adaptation, whereby the saccade gain is changed by rapidly moving the target during the saccade (Kojima et al., 2004; Ethier et al., 2008; Jenkinson and Miall, 2010), and split treadmill adaptation, whereby subjects walk on a treadmill that imposes different step sizes for each leg (Morton and Bastian, 2006; Reisman et al., 2007; Vasudevan and Bastian, 2010). In this work, we focus on studying point-to-point motion under two types of motor adaptation paradigms: viscous force-field adaptation (Chapters 2, 3, 4, and 6) and visuomotor rotation adaptation (Chapters 4 and 5).

1.2 Forward and inverse internal models for motor adaptation

Motor memories are thought to primarily reflect internal models of the environment. Such internal models allow for successful motor control, and they can adapt during motor learning. Of these, *inverse models* map a desired state of the environment (for example, a specific motion that we want to execute) into a set of motor commands that can achieve it, whereas *forward models* predict the consequences of a given motor command on the environment (Wolpert et al., 1995; Wolpert and Miall, 1996; Wolpert et al., 1998). This prediction is available before the arrival of sensory feedback about these consequences, thus allowing the motor system to potentially alleviate the effects of delays in sensorimotor loops.

When a novel perturbation changes the environment – for example, when a force-field alters the dynamics of motion and thus the mapping between motor commands and their consequences – these models become insufficient, resulting in erroneous motion. The consequent sensory error – the difference between the predicted vs. the observed results of motor commands – then feeds into these models, acting as a training signal that gradually adapts the models to the new situation. Previous work has offered some insights behind the mechanisms of this adaptation. The sensitivity of the internal models to sensory errors (i.e. the learning rate) is hypothesized to scale with the certainty associated with sensory feedback vs. the certainty associated with the prior state of the internal model (Kording et al., 2004; Kording and Wolpert, 2004; Franklin et al., 2012), implementing a Kalman filter (Kalman, 1960) to optimally combine these two sources of information in a manner that minimizes variance. This is analogous to the Bayesian integration between different sources of sensory input (Ernst and Banks, 2002), where the information from each one of these sources is weighed inversely to its variance. Moreover, the learning rate has been found to inversely correlate with the size of the error (Marko et al., 2012), whereas we recently showed that the learning rate is also dependent on the consistency of trial-to-trial environmental dynamics (Gonzalez Castro et al., 2014) – also see Chapter 6.

1.3 Generalization of motor adaptation

Generalization refers to the degree to which motor skill or adaptation applies to conditions different than the one it was trained on. Consider the common skill of tying one's shoelaces, on which we are trained from as early as few years of age, and we utilize throughout our life. Compare this skill to the rather exotic one of tying a bow-tie – an uncommon skill (unless, maybe, you go to Harvard), and one that most people consider hard and will not even bother attempting. These two skills are more intimately tied than one might think: they employ the same knot. Why, then, is it hard to generalize the motor skill of tying a shoelace to the skill of tying a bow tie?

Behind this question lie various parameters that determine how skill or adaptation transfers from the trained set of conditions to the other. Transfer can depend on the similarity between one set of dynamics and the other (Sing et al., 2013); here, a shoelace might be characterized by different dynamics than a silk bowtie. It can also depend on the extrinsic location of the trained movement within the workspace (Krakauer et al., 2000; Brayanov et al., 2012); in our example, close to the floor for the shoelace vs. on shoulder height for the bowtie. Also important can be the similarity with trained motion trajectory in terms of its intrinsic coordinates – the joint displacements associated with it – (Shadmehr and Mussa-Ivaldi, 1994; Shadmehr and Moussavi, 2000); here, the baseline extension of the shoulders to reach down for one's laces vs. to reach for a bow tie around the neck is different, meaning that the two skills are thus performed in considerably different joint coordinates. Transfer can also depend on how the workspace is viewed, like through a different angle on the retina (Chou and Lisberger, 2002), or, through a mirror transformation (Milner, 1970), which is sometimes the case when tying a bow tie. Other factors that can determine transfer / generalization include whether the same or the opposite arm is used to perform the task (Hamilton, 1964; Criscimagna-Hemminger et al., 2003; Taylor et al., 2011; Brayanov, 2013; Joiner et al., 2013), differences in the speed (Joiner et al., 2011) or the curvature (Brayanov, 2013) of motion, et cetera.

The study of motor skill generalization can be invaluable to the motor neuroscientist, for two main reasons. One is that the value of training and studying simple adaptation paradigms, like the ones discussed in paragraph 1.1, depends on whether both the learning improvements and the insights about the mechanisms of motor learning that we obtain from these paradigms generalizes to more complex, natural motor tasks. Along these lines, simple point-to-point movement tasks have been successfully utilized to help patients recovering deficits in upper-limb function after a stroke (Lum et al., 2002) while visuomotor rotations have been utilized to alleviate the symptoms of patients suffering from hemispatial neglect (Rode et al., 2003). The other main reason is that generalization patterns might offer insight as to how the memories underlying motor learning are encoded in the brain. The ability to generalize training from one condition to another can tell us the degree to which there is a shared neural representation between these two conditions. For example, previous work suggested that elements of motor learning can each be tuned to a different preferred movement direction or velocity amplitude (Hwang et al., 2003; Shadmehr, 2004); the width of generalization across different movement directions or velocities, correspondingly, can inform us about the width of these tuning functions of the underlying neuronal populations.

1.4 Scope of this work

This dissertation investigates the formation and retention of memories of both the expected state of the environment and the uncertainty associated with it.

In chapter 2, we study the interplay between the formation of estimates for the expected environmental dynamics and the formation of estimates about the variability of these dynamics. Specifically, we study how variability estimates are critical for the control of the safety margin associated with grip forces. The idea here is that we can afford to keep the grip force safety margin small without risking slip when environmental variability is low, but, we would benefit from maintaining a large safety margin to prevent frequent slips when variability is high. In line with the above, we find that grip force control is, in fact, three times more sensitive to changes in variability than the expected value of environmental dynamics.

This ratio of sensitivities effectively acts to provide a 3-σ confidence level (>99%) against slip. A consequence of heightened sensitivity of grip forces to uncertainty that we observe is that it leads to strong increases in grip forces upon a change in environmental dynamics, regardless of whether that change increased or decreased the strength of the load. We verify this experimentally, showing bizarre, asymmetric grip force learning curves to step increases vs. step decreases in dynamics. Specifically, the variability-driven component will make grip force learning appear faster for grip force increases, but slower for grip force decreases. Our results thus negate the findings of a previous influential study (Flanagan et al., 2003) which observed that, under a load-increasing perturbation, grip force adaptation outpaces manipulatory force adaptation, taking this as general evidence for a grip force rate advantage and, critically, using to support the idea that a forward model for predicting the sensory consequences of environmental dynamics adapts faster than an inverse model

In chapter 3, we extend our investigation to cases where environmental dynamics are not normally distributed, in order to examine the mechanism behind how the apparent confidence-based safety margin control is calculated in the motor system. Specifically, we find that the motor system does not directly estimate confidence by maintaining a memory of worst-case conditions of the environment, but instead indirectly estimates confidence by scaling grip force responses by a measure of variability. We proceed to examine how this variability measure is calculated, and find that it utilizes a low-order metric that is robust to outliers in the distribution of environmental dynamics.

In chapter 4, we study the temporal stability of motor adaptation. Using a visuomotor adaptation task, we investigate the speed and breadth of decay due to the passage of time, and find that adaptation consists of a temporally-labile component, which decays within the course of tens of seconds, and a temporally-stable component, which persists at least over tens of minutes… and can predict the long-term retention of the trained adaptation. Based on a brain stimulation study and the reanalysis of previous studies on cerebellar patients, we identify the primary motor cortex and the posterior parietal cortex as a likely parts of a network involved with temporally-labile adaptation, and the cerebellum as involved with temporally-

stable adaptation. In addition, we examine differences in decay from one training session to the next, and find that these differences are not only driven by differences across individuals, but, remarkably, by considerable differences across different sessions for the same individual.

In chapter 5, we examine how the temporally-stable and temporally-labile components are involved in savings – the phenomenon whereby retraining is faster compared to initial adaptation. We find that, in spite of the fact that temporally-stable adaptation is remarkably persistent even after prolonged active washout (reversal of the training stimulus), savings is not driven by the reemergence of this strong, stable memory, but it is instead driven by an increased propensity of the temporally-labile component to relearn the task at hand.

In chapter 6, we examine how the rate of adaptation to changes in the environment is modulated by the statistics of that environment. We show that the rate at which the motor system adapts to changes in the environment is primarily determined not by the degree to which environment change occurs – i.e. environmental variability – but by the degree to which the changes that do occur persist from one movement to the next – i.e. the consistency of the environment. Specifically, we demonstrate a striking double dissociation whereby feedback response strength is predicted by environmental variability rather than consistency, whereas adaptation rate is predicted by environmental consistency rather than variability. Finally, we elucidate the role of stimulus repetition in speeding up adaptation, finding that repetition can greatly potentiate the effect of consistency, although, unlike consistency, repetition alone does not increase adaptation rate. We conclude that controlling for environmental consistency and repetition during training may improve regimens for rehabilitation and motor training, leading to faster and more robust motor learning.

1.5 References

Ashe J, Georgopoulos AP (1994) Movement parameters and neural activity in motor cortex and area 5. Cerebral cortex 4:590-600.

Braun DA, Aertsen A, Wolpert DM, Mehring C (2009) Motor task variation induces structural learning. Current biology : CB 19:352-357.

Brayanov JB (2013) Internal Representations for the Generalization of Motor Memories.

Brayanov JB, Press DZ, Smith MA (2012) Motor memory is encoded as a gain-field combination of intrinsic and extrinsic action representations. The Journal of neuroscience : the official journal of the Society for Neuroscience 32:14951-14965.

Chou IH, Lisberger SG (2002) Spatial generalization of learning in smooth pursuit eye movements: implications for the coordinate frame and sites of learning. The Journal of neuroscience : the official journal of the Society for Neuroscience 22:4728-4739.

Conditt MA, Gandolfo F, MussaIvaldi FA (1997) The motor system does not learn the dynamics of the arm by rote memorization of past experience. Journal of neurophysiology 78:554-560.

Criscimagna-Hemminger SE, Bastian AJ, Shadmehr R (2010) Size of error affects cerebellar contributions to motor learning. Journal of neurophysiology 103:2275-2284.

Criscimagna-Hemminger SE, Donchin O, Gazzaniga MS, Shadmehr R (2003) Learned dynamics of reaching movements generalize from dominant to nondominant arm. Journal of neurophysiology 89:168-176.

Ernst MO, Banks MS (2002) Humans integrate visual and haptic information in a statistically optimal fashion. Nature 415:429-433.

Ethier V, Zee DS, Shadmehr R (2008) Spontaneous recovery of motor memory during saccade adaptation. Journal of neurophysiology 99:2577-2583.

Flanagan JR, Vetter P, Johansson RS, Wolpert DM (2003) Prediction precedes control in motor learning. Current biology : CB 13:146-150.

Franklin S, Wolpert DM, Franklin DW (2012) Visuomotor feedback gains upregulate during the learning of novel dynamics. Journal of neurophysiology 108:467-478.

Gibo TL, Criscimagna-Hemminger SE, Okamura AM, Bastian AJ (2013) Cerebellar motor learning: are environment dynamics more important than error size? Journal of neurophysiology 110:322-333.

Gonzalez Castro LN, Hadjiosif AM, Hemphill MA, Smith MA (2014) Environmental consistency determines the rate of motor adaptation. Current biology : CB 24:1050-1061.

Hamilton CR (1964) Intermanual Transfer of Adaptation to Prisms. The American journal of psychology 77:457-462.

Huang VS, Haith A, Mazzoni P, Krakauer JW (2011) Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. Neuron 70:787-801.

Huse DM, Schulman K, Orsini L, Castelli-Haley J, Kennedy S, Lenhart G (2005) Burden of illness in Parkinson's disease. Movement disorders : official journal of the Movement Disorder Society 20:1449-1454.

Hwang EJ, Smith MA, Shadmehr R (2006) Adaptation and generalization in acceleration-dependent force fields. Experimental brain research 169:496-506.

Hwang EJ, Donchin O, Smith MA, Shadmehr R (2003) A gain-field encoding of limb position and velocity in the internal model of arm dynamics. PLoS biology 1:E25.

Jenkinson N, Miall RC (2010) Disruption of saccadic adaptation with repetitive transcranial magnetic stimulation of the posterior cerebellum in humans. Cerebellum 9:548-555.

Joiner WM, Smith MA (2008) Long-term retention explained by a model of short-term learning in the adaptive control of reaching. Journal of neurophysiology 100:2948-2955.

Joiner WM, Brayanov JB, Smith MA (2013) The training schedule affects the stability, not the magnitude, of the interlimb transfer of learned dynamics. Journal of neurophysiology 110:984-998.

Joiner WM, Ajayi O, Sing GC, Smith MA (2011) Linear hypergeneralization of learned dynamics across movement speeds reveals anisotropic, gain-encoding primitives for motor adaptation. Journal of neurophysiology 105:45-59.

Kalman RE (1960) A new approach to linear filtering and prediction problems. Journal of Fluids Engineering 82:35-45.

Kitago T, Ryan SL, Mazzoni P, Krakauer JW, Haith AM (2013) Unlearning versus savings in visuomotor adaptation: comparing effects of washout, passage of time, and removal of errors on motor memory. Frontiers in human neuroscience 7:307.

Kojima Y, Iwamoto Y, Yoshida K (2004) Memory of learning facilitates saccadic adaptation in the monkey. The Journal of neuroscience : the official journal of the Society for Neuroscience 24:7531-7539.

Kording KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. Nature 427:244-247.

Kording KP, Ku SP, Wolpert DM (2004) Bayesian integration in force estimation. Journal of neurophysiology 92:3161-3165.

Krakauer JW, Ghez C, Ghilardi MF (2005) Adaptation to visuomotor transformations: Consolidation, interference, and forgetting. Journal of Neuroscience 25:473-478.

Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. Journal of Neuroscience 20:8916-8924.

Lum PS, Burgar CG, Shor PC, Majmundar M, Van der Loos M (2002) Robot-assisted movement training compared with conventional therapy techniques for the rehabilitation of upper-limb motor function after stroke. Archives of physical medicine and rehabilitation 83:952-959.

Marko MK, Haith AM, Harran MD, Shadmehr R (2012) Sensitivity to prediction error in reach adaptation. Journal of neurophysiology 108:1752-1763.

Milner B (1970) Memory and the medial temporal regions of the brain. Biology of memory 23:31-59.

Morton SM, Bastian AJ (2006) Cerebellar contributions to locomotor adaptations during splitbelt treadmill walking. The Journal of neuroscience : the official journal of the Society for Neuroscience 26:9107-9116.

Paninski L, Fellows MR, Hatsopoulos NG, Donoghue JP (2004) Spatiotemporal tuning of motor cortical neurons for hand position and velocity. Journal of neurophysiology 91:515-532.

Pearson TS, Krakauer JW, Mazzoni P (2010) Learning not to generalize: modular adaptation of visuomotor gain. Journal of neurophysiology 103:2938-2952.

Pine ZM, Krakauer JW, Gordon J, Ghez C (1996) Learning of scaling factors and reference axes for reaching movements. Neuroreport 7:2357-2361.

Reisman DS, Wityk R, Silver K, Bastian AJ (2007) Locomotor adaptation on a split-belt treadmill can improve walking symmetry post-stroke. Brain : a journal of neurology 130:1861-1872.

Rode G, Pisella L, Rossetti Y, Farne A, Boisson D (2003) Bottom-up transfer of sensory-motor plasticity to recovery of spatial cognition: visuomotor adaptation and spatial neglect. Progress in brain research 142:273-287.

Scarmeas N, Albert M, Brandt J, Blacker D, Hadjigeorgiou G, Papadimitriou A, Dubois B, Sarazin M, Wegesin D, Marder K, Bell K, Honig L, Stern Y (2005) Motor signs predict poor outcomes in Alzheimer disease. Neurology 64:1696-1703.

Shadmehr R (2004) Generalization as a behavioral window to the neural mechanisms of learning internal models. Human movement science 23:543-568.

Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. The Journal of neuroscience : the official journal of the Society for Neuroscience 14:3208-3224.

Shadmehr R, Moussavi ZM (2000) Spatial generalization from learning dynamics of reaching movements. The Journal of neuroscience : the official journal of the Society for Neuroscience 20:7807-7815.

Shadmehr R, Smith MA, Krakauer JW (2010) Error correction, sensory prediction, and adaptation in motor control. Annual review of neuroscience 33:89-108.

Shmuelof L, Huang VS, Haith AM, Delnicki RJ, Mazzoni P, Krakauer JW (2012) Overcoming motor "forgetting" through reinforcement of learned actions. The Journal of neuroscience : the official journal of the Society for Neuroscience 32:14617-14621.

Sing GC, Orozco SP, Smith MA (2013) Limb motion dictates how motor learning arises from arbitrary environmental dynamics. Journal of neurophysiology 109:2466-2482.

Sing GC, Joiner WM, Nanayakkara T, Brayanov JB, Smith MA (2009) Primitives for motor adaptation reflect correlated neural tuning to position and velocity. Neuron 64:575-589.

Smith MA, Shadmehr R (2005) Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. Journal of neurophysiology 93:2809-2821.

Taylor JA, Wojaczynski GJ, Ivry RB (2011) Trial-by-trial analysis of intermanual transfer during visuomotor adaptation. Journal of neurophysiology 106:3157-3172.

Vasudevan EV, Bastian AJ (2010) Split-belt treadmill adaptation shows different functional networks for fast and slow human walking. Journal of neurophysiology 103:183-191.

Vindras P, Viviani P (2002) Altering the visuomotor gain. Evidence that motor plans deal with vector quantities. Experimental brain research 147:280-295.

Wolpert DM, Miall RC (1996) Forward Models for Physiological Motor Control. Neural networks : the official journal of the International Neural Network Society 9:1265-1279.

Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensorimotor integration. Science 269:1880-1882.

Wolpert DM, Miall RC, Kawato M (1998) Internal models in the cerebellum. Trends in cognitive sciences 2:338-347.

Chapter 2: Flexible control of safety margins for action based on internal estimates of environmental variability

2.1 Summary[1]

To reduce the risk of slip, grip force (GF) control includes a safety margin above the force level ordinarily sufficient for the expected load force (LF) dynamics. The current view is that this safety margin is based on expected LF dynamics, like a static safety factor used in engineering design. More efficient control could be achieved, however, if the motor system reduces the safety margin when LF variability is low and increases it when this variability is high. Here we show that this is indeed the case by demonstrating that the human motor system sizes the GF safety margin in proportion to an internal estimate of LF variability, in order to maintain a fixed statistical confidence against slip. In contrast to current models of GF control that neglect the variability of LF dynamics, we demonstrate that GF is 3-fold more sensitive to the standard deviation than the expected value of LF dynamics, in line with the maintenance of a 3-sigma confidence level. We then show that a computational model of GF control that includes a variability-driven safety margin predicts highly asymmetric GF adaptation between increases vs decreases in load. We find clear experimental evidence for this asymmetry and show that it explains previously reported differences in how rapidly GFs and manipulatory forces adapt. This model further predicts bizarre non-monotonic shapes for GF learning curves, which are faithfully borne out in our experimental data. Our findings establish a new role for internal estimates of environmental variability in the control of action.

[1] Most of this chapter is under review in the *Journal of Neuroscience* as a paper with the same name.

2.2 Introduction

A skilled cellist masterfully controls the dynamics of his bow via a finely tuned pattern of motor output that ensures that it slides at the right speed, touching the right string, at the right time, with the right amount of force. Yet as Pablo Casals, the preeminent cellist of the early 20th century, learned all too well, this remarkable control fails completely if one loses grip on the bow, even momentarily. In his highly anticipated 1907 Vienna debut, a stage-frightened Casals lost his grip, sending his bow flying into the audience. The accommodating Viennese concertgoers quickly returned the bow, and Casals went on to deliver a landmark performance. Yet, amongst thousands of flawlessly-executed movements, that single bow slip was what he remembered 50 years later(Goldsmith, 1956).

Casal's story invites us to consider the importance of robust grip force (GF) control that generally prevents accidental slips. Indeed, how does the brain avoid making them all the time? To prevent such accidents, GFs generally include a safety margin above the force level usually sufficient to maintain grasp, reducing the chance that unexpected perturbations, misestimation of load forces (LFs), or errors in GF production result in slippage(Westling and Johansson, 1984).

Previous work has suggested that the predictive control of GF acts to maintain a constant GF/LF ratio which, in line with the maintenance of a safety margin, is consistently higher than the so-called slip ratio, the minimum GF/LF ratio that avoids slip(Johansson and Westling, 1984, 1988; Cole and Johansson, 1993). Moreover, previous work has generally maintained that both the overall GF/LF ratio and the ratio margin (GF/LF – slip ratio) are proportional to the slip ratio(Westling and Johansson, 1984; Cole and Johansson, 1993), making the GF safety margin proportional to the expected slip force. Thus the current theory corresponds to a margin described by a fixed *safety factor*, like that widely used in engineering design(ACI Committee 318, 2005; NASA, 2008). For example, with a 20% safety factor, a bridge built for a 1000-ton load would be designed to withstand 1200 tons, or when holding a 5kg object, one would apply a GF sufficient for 6.

We hypothesize that the CNS does not maintain a fixed safety factor for GF control, but instead adaptively controls the safety margin based on online estimates of the uncertainty about environmental dynamics. This allows the motor system to efficiently balance effort against the risk of slip: We can afford to keep the safety margin small without risking slip when environmental variability is low, but, would benefit from maintaining a large safety margin to prevent frequent slips when variability is high – analogous to the notion that, when driving, one should keep a larger distance away from erratically behaving vehicles. Quantitatively, a variability-based safety margin could [allow for GF control that] provides a *fixed* statistical confidence level against slip, rather than a confidence level that changes with environmental variability.

Altogether, we propose that the control of GF is based on an additive combination of what is directly known about the environment (its expected dynamics) to estimate the ordinarily sufficient action, and the known unknown (the environment's variability) to size a safety margin based on statistical confidence. Because a safety margin of 2-3 standard deviations would be required to maintain 95-99% confidence bounds(Wasserman, 2004), this hypothesis would predict GF to be *more* sensitive to the standard deviation of the environment than to its mean, a radical departure from current thinking in which environmental variability has no role(Johansson and Westling, 1984; Westling and Johansson, 1984; Johansson and Westling, 1988; Cole and Johansson, 1993; Flanagan and Wing, 1993, 1997; Flanagan et al., 2003).

In contrast to GF control, safety margins do not readily apply to manipulatory force (MF) control, because both excessive and inadequate forces result in motor errors. Accordingly, MF control is largely based on estimates of the expected value of dynamics(Scheidt et al., 2001; Takahashi et al., 2001a; Mawase and Karniel, 2012) without a safety margin. We thus designed a series of experiments to compare how the control of GFs and MFs are affected by changes in environmental variability.

2.3 Materials and Methods

2.3.1 Participants and Ethics statement

A total of 70 healthy individuals without known neurological disease (21 male, age 22.1 ± 4.4 yrs) participated in the two experiments. All subjects were naïve to the purpose of the experiments and provided informed consent. Two subjects were (self-reportedly) left-handed, but used their right arm for the experiment, as the task did not require using the dominant hand. All experiment protocols were approved by the Harvard University Committee on the Use of Human Subjects in Research.

2.3.2 General task

Participants grasped a custom-designed grip force measurement device using a pinch grip (Fig. 2.1A). This object consisted of two load cells (FUTEK LRF 300) connecting two rubber-coated parallel grip surfaces 28.3 mm apart. We restricted subjects' choice of grasp position in order to avoid grasp-configuration-dependent differences in grip force adaptation (Fu et al., 2010) that might obscure effects related to differences in experienced dynamics. This device was attached to a robotic manipulandum (Interactive Motion Technologies, Inc.) in such a way that it could freely rotate about the vertical axis allowing an orientation lateral to movement direction to be maintained throughout each movement (Fig. 2.1A). While grasping the object, subjects performed 10cm reaching arm movements toward or away from the chest aimed at 1cm circular targets displayed on a vertically-oriented LCD screen in front of them. During each movement, the subject's hand position was represented by a 3mm diameter cursor. Subjects were trained to strive for a movement duration of 500±50ms, based on visual feedback about movement speed following each trial. Subjects were instructed to move directly to each target, and to maintain a fixed grasp orientation as illustrated in Fig. 2.1A. After a short familiarization session, a leftward load force (LF) consisting of static and dynamic components was applied to the object:

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} F_0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & b \\ -b & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix} \qquad \text{(Equation 2.1)}$$

where F_0 is the static offset, and b is the strength of a viscous force field FF, which could change from one trial to the next. We made the static force offset F_0 large enough to ensure that the overall LF would always remain leftward, even for negative values of b (Fig. 2.2C). Experiments consisted of baseline and training sessions, with breaks every 6-7 minutes. Each trial consisted of two movements, one toward (270° direction) and one away from the chest (90° direction). Note that, throughout this manuscript, we refer to "positive b" as one that increases the magnitude of the exerted LF regardless of the movement direction (i.e. a "positive b" would refer to $b > 0$ for 270° movements and to $b < 0$ for 90° movements). Thus, the sign of b in Equation 1 is correct when 270° movements are considered, whereas for 90° movements this sign has to be inverted.

Experiments 2.1, 2.1b and 2.1c

Experiments 2.1, 2.1b and 2.1c were designed to investigate the effect of environmental variability on GF and MF control. In Experiment 2.1 (N=11), the force offset F_0 was held constant at 4N, while the viscous force-field coefficient b was varied from trial to trial; b was drawn from a Gaussian distribution with mean zero and standard deviation σ_b which depended on the type of experiment block: zero, low, medium and high variability blocks corresponding to σ_b = 0, 1.8, 3.6 or 5.4 Ns/m (Fig. 2.2A,B). During the baseline period, b was set to zero (i.e. only the offset force was experienced). After the baseline period, participants were exposed to 18 50-trial blocks. The first and last were zero-variability blocks and the middle 16 randomly alternated between the zero, low, medium and high variability conditions such that each was presented four times. Brief ~1 minute rest breaks were instituted every 50 trials and occurred 10 trials prior to the onset of each block. Experiment 2.1b (N=23) was similar to Experiment 2.1 except that F_0=3N, σ_b = 0, 1.2, 2.4 or 3.6 Ns/m, blocks were 100 trials in duration, and there were a total of 8 training blocks, all randomly ordered, such that each variability condition was presented twice. Because the blocks were longer, breaks were given both 10 trials before and 40 trials after the onset of each block, to maintain 50-trial intervals between breaks. Experiment 2.1c (N=17) was similar to Experiment 2.1 except that b was drawn from a Gaussian distribution with a mean that was +3.75 Ns/m, and a 240-trial

constant +3.75 Ns/m FF epoch (Experiment 2.2c – see below) was inserted between the 18-block experiment period and baseline.

Experiments 2.2a, 2.2b and 2.2c

Experiments 2.2a, 2.2b and 2.2c were designed to investigate how a sudden change in the environment would affect GF and MF control. For Experiments 2.2a and 2.2b, the force offset F_0 was held constant at 3N. Following the 80-trial baseline period, b was switched from zero to a value of ± 7.5 Nm/s at the onset of the 240-trial training period. In Experiment 2.2a (N=18) participants trained with leftward FFs that we deem positive (+FF) because they increase the LF relative to F_0 (b = +7.5 Ns/m and -7.5 Ns/m for 270° and 90° movement directions, respectively), whereas in Experiment 2.2b (N=18) participants trained with rightward FFs that we deem negative (-FF) because they decrease the LF relative to F_0. We inserted short rest breaks every 64 trials, and the training period began 16 trials after a break to dissociate any effect of the break from the onset of the force-field. Experiment 2.2c (N = 17) was similar to Experiment 2.2a except that it was immediately followed by another experiment (2.1c), the force offset F_0 was held constant at 4N, and the FF strength was b=+3.75 Ns/m instead of +7.5 Ns/m.

2.3.3 Data analysis and computational modeling

Data inclusion criteria

Trials in which the peak velocity was less than 0.2 m/s or more than 0.5 m/s (less than 0.4% of trials in total) were excluded. In Experiment 2.2, we found strong GF increases in the first few trials following each rest break. To remove this post-break artifact we omitted the first 4 trials after every break from our main analyses for all experiments.

Estimation of adaptation levels for manipulatory forces and grip forces

We measured manipulatory forces during error-clamp trials(Scheidt et al., 2000; Smith et al., 2006) interspersed with a frequency of 1 in 5 in a pattern that was random but which avoided consecutive error-clamp trials. These trials minimize lateral error via the implementation of a virtual force channel with a

high lateral stiffness of 6000N/m in order to isolate the feed-forward component of adaptation. We estimated the manipulatory force adaptation level on each error-clamp trial by regressing(Smith et al., 2006; Joiner and Smith, 2008; Yousif and Diedrichsen, 2012) the baseline-subtracted (baseline defined as 250ms to 100ms before movement onset) force profile onto the velocity profile for that trial (on which the force-field would have been based). The regression coefficient obtained is an estimate of the value of *b* that the measured MF would have compensated for during each trial, because the viscous FFs we employ scale the velocity profile by *b* (Equation 2.1).

We measured grip forces during all trials using the measurement device described in the **General task** paragraph. The device consisted of two load cells sandwiched between the two grip surfaces (Figure 2.1A). Grip forces were measured as the sum of the readings from these two load cells. We assessed the level of grip force adaptation on each trial using the same method as for MFs: we regressed each baseline-subtracted GF profile onto the corresponding velocity profile for that trial and used the resulting regression coefficient as a measure of GF adaptation. In order to calculate the force adaptation level associated with each variability block, in experiments 2.1, 2.1b, and 2.1c, we averaged the regression coefficients obtained from trial 16 to the end of the block. We omitted the first 15 trials to exclude potential aftereffects from the variability level during the preceding block. However, as we note in the Results section, changing the size of this selection window did not yield meaningful differences in results.

Estimation of the sensitivity of manipulatory and grip forces to the standard deviation and mean of load force dynamics

We estimated the sensitivity of GF and MF to the standard deviation of environmental dynamics (i.e. the GF/σ_{LF} and MF/σ_{LF} sensitivities, respectively) in Experiment 2.1, where we measured GF and MF responses to different variability levels. Sensitivity to standard deviation corresponded to the slope of Figure 2.3A for GF and Figure 2.3D for MF, as it was defined as the change in force level divided by the change in standard deviation from Experiment 2.1:

$$Sensitivity\ to\ standard\ deviation = \frac{\Delta F}{\Delta \sigma_{LF}}$$

Correspondingly, we estimated the sensitivity of GF and MF to the mean of environmental dynamics (i.e. the GF/μ_{LF} and MF/μ_{LF} sensitivities, respectively) in Experiment 2.2, where we measured asymptotic GF and MF adaptation to one of three different step changes in the mean FF: 7.5 Ns/m (Experiment 2.2a), -7.5Ns/m (Experiment 2.2b), and 3.75Ns/m (Experiment 2.2c). Sensitivity to the mean corresponded to the slope of Figure 2.3B for GF and Figure 2.3E for MF, as it was defined as change in GF between baseline and asymptote adaptation, divided by the corresponding change in the mean:

$$Sensitivity\ to\ the\ mean = \frac{\Delta F}{\Delta \mu_{LF}}$$

Computational model of combined mean-driven and variability-driven adaptation

Based on the results of Experiments 2.1, 2.1b, 2.1c, 2.2a, 2.2b, and 2.2c presented in Figs. 2.2 and 2.3 which show that GFs are sensitive to both the mean and variability of environmental dynamics, we developed a model to explain how GFs respond to a step change in environmental dynamics. This model posits GF control to be the sum of two components, one driven by an internal estimate of the expected LF and the other driven by an estimate of LF uncertainty. We refer to this as the Additive Combination of Expected Value and Uncertainty (ACEVU) model. The model posits that the mean, μ, and standard deviation, σ, of the environmental dynamics (i.e. the value of b) at each trial k are both estimated based on an exponential memory window with time constant τ (equation 2.2). This is in line with recent work suggesting that the motor system registers a memory of squared errors that decreases exponentially from trial-to-trial(Landy et al., 2012); however, we note that any window type that weighs recent memories more heavily than older ones would yield qualitatively similar results. The computations required to generate the exponentially-windowed estimates of μ and σ required for the ACEVU model are shown below:

$$\hat{\mu}(k+1) = \frac{\sum_{i=0}^{k} b(i)\cdot e^{-(k-i)/\tau}}{\sum_{i=0}^{k} e^{-(k-i)/\tau}} = \frac{\sum_{i=0}^{k} b(i)\cdot e^{i/\tau}}{\sum_{i=0}^{k} e^{i/\tau}} = \left(\frac{1-e^{1/\tau}}{1-e^{k/\tau}}\right)\sum_{i=0}^{k} b(i)\cdot e^{i/\tau}$$

$$\hat{\sigma}^2(k+1) = \frac{\sum_{i=0}^{k} (b(i)-\hat{\mu}(i))^2\cdot e^{-(k-i)/\tau}}{\sum_{i=0}^{k} e^{-(k-i)/\tau}} = \frac{\sum_{i=0}^{k} (b(i)-\hat{\mu}(i))^2\cdot e^{i/\tau}}{\sum_{i=0}^{k} e^{i/\tau}} = \left(\frac{1-e^{1/\tau}}{1-e^{k/\tau}}\right)\sum_{i=0}^{k} (b(i)-\hat{\mu}(i))^2\cdot e^{i/\tau}$$

(Equation 2.2)

Note that Equation 2.2 can be rewritten recursively as:

$$\hat{\mu}(k+1) = A\cdot\hat{\mu}(k) + B\cdot b(k)$$

$$\hat{\sigma}^2(k+1) = A\cdot\hat{\sigma}^2(k) + B\cdot(b(k)-\hat{\mu}(k))^2$$

$$A = \frac{1-e^{k/\tau}}{1-e^{(k+1)/\tau}}, \qquad B = (1-e^{1/\tau})\frac{e^{k/\tau}}{1-e^{(k+1)/\tau}}$$

(Equation 2.3)

The recursive formulation presented in Equation 2.3 illustrates an iterative implementation of mean and variability estimates that computes the effect of an exponential memory window without requiring the maintenance of separate memories for all of the individual movements in the memory window. Note that for k sufficiently large (i.e. when the mean and variance have been estimated for a large number of trials compared to τ), the coefficients A and B can be simplified and no longer depend on k:

$$A = e^{-1/\tau}, \qquad B = 1-e^{-1/\tau} \qquad \text{(Equation 2.4)}$$

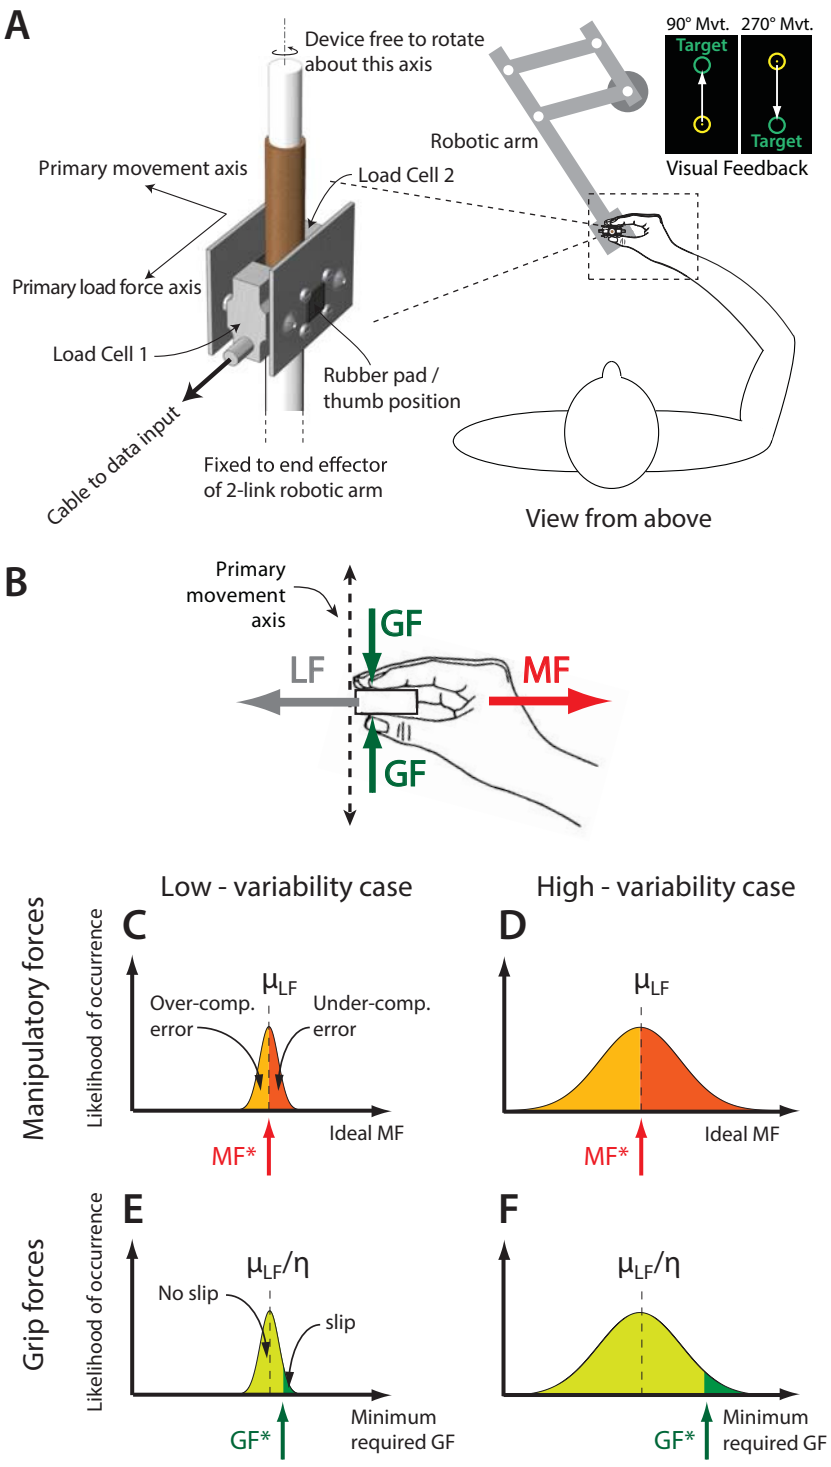
Statistical Comparisons

We used paired t-tests for all statistical comparisons between adaptation levels and sensitivity estimates across conditions within groups, and unpaired t-tests for comparisons between adaptation levels and sensitivity estimates across groups. We used bootstrapping for analyses where we expected the distribution of the data to be highly non-normal. Specifically, we used bootstrapping for ratio estimation, as the ratio between random variables is known to produce a distribution which can be far from normal, even when the random variables themselves are normally distributed(Marsaglia, 1965). We also used bootstrapping for estimating the time constants of task-specific adaptation for grip and manipulatory forces using exponential fits. We chose this method instead of performing a t-test between parameters derived from fits to individual subject data, because the high noise in these individual data leads to low confidence about the corresponding individual parameters.

Figure 2.1. Grip and Manipulatory forces in variable environments.

(A): Diagram of the grasping device (left) and top view of the experiment setup (B): Top view of the grasped object illustrating the forces exerted on and by the subject during movement. Rightward manipulatory forces (MFs) can compensate leftward load forces (LF) and compressive grip forces (GF) maintain grasp. (C-D): MF compensation in a low variability (C) or a high variability (D) environment, both with the same expected value μ_{LF}. The optimal MF level is equal to the expected value of the perturbation, which minimizes the amount of over- and under-compensation in both cases. (E-F): GF control in the same environments. The GF required to avoid slip has a mean value of the mean load divided by the coefficient of friction (μ_{LF} /η) in both cases. However, note that using this expected value (vertical dashed line) would result in slip whenever the required GF exceeds this value, in 50% of the cases. In a low variability environment (E), using a GF (GF*) that incorporates a small safety margin above this expected value would ensure the chance of slippage is low (dark green area under the bell curve). However, in a high variability environment (F), a larger safety margin would be required to maintain a low slip probability. This leads to the hypothesis that GF control includes a safety margin that scales with environmental variability.

Figure 2.1 (continued)



2.4 Results

We begin by studying changes in the control of grip force (GF) and manipulatory force (MF) when the mean and variability of load force (LF) dynamics are systematically manipulated to examine the hypothesis that GF control is more sensitive to the variability than the expected value of LF dynamics. We then examine how variability-driven responses affect the changes in GF and MF control observed after sudden changes in LF dynamics. With this understanding, we investigate whether variability-driven GF adaptation can explain previous results suggesting that GF control adapts faster than MF control(Flanagan et al., 2003).

2.4.1 Effects of increased environmental variability on grip and manipulatory forces

While previous work has focused on the adaptive responses of both GFs and MFs to changes in the expected value of environmental dynamics(Flanagan and Wing, 1993; Shadmehr and Mussaivaldi, 1994; Flanagan and Wing, 1997; Wolpert and Kawato, 1998; Wolpert et al., 1998; Kawato, 1999; Krakauer et al., 1999; Smith et al., 2006; Danion and Sarlegna, 2007; Sing et al., 2009; Yokoi et al., 2011; Yousif and Diedrichsen, 2012), here we hypothesize that safety margins, which apply to GF but not MF control, are specified based on the variability rather than the expected value of environmental dynamics, leading to fundamental differences between the control of MFs and GFs. In particular, this would predict the control of GFs, but not MFs, to be strongly modulated by environmental variability.

In order to hold steady against a load, the motor system must apply a MF equal and opposite to the load force (LF): a weaker MF would result in acceleration in the direction of the LF, whereas a stronger MF would accelerate the object against the load. Accordingly, a strategy that yields efficient MF control, minimizing the risk of accelerating toward either direction, would be to exert a MF that continually approximates the expected value of the LF during the course of each movement. This idea is in line with previous work showing that under variable dynamic environments, MFs adapt to the expected value of the perturbation(Scheidt et al., 2001; Takahashi et al., 2001a) and that under variable visuomotor

transformations, movements adapt to the expected value of the transformation(Burge et al., 2008; Saijo and Gomi, 2012), independent of the amount of variability in both cases.

To apply this strategy to GF control, however, would mean that the planned GF would often be insufficient to maintain grasp because, producing a GF equal to the expected value of the required GF (i.e. the expected value of the orthogonal LF divided by the coefficient of friction, η), would result in compromised grasp whenever the required GF exceeds its mean. This exposes a key dichotomy between the control of GFs and MFs: whereas under- and over-compensating MFs both result in motion errors, under-compensating GFs result in grasp errors while over-compensating GFs do not (see Fig. 2.1). A strategy for improved GF control would thus reduce the risk of slip by programming GFs to be greater than the expected value of the required GF (Johansson and Westling, 1984, 1988; Cole and Johansson, 1993). The specification of the size of this GF safety margin can be viewed as a tradeoff between the increased effort associated a higher margin and the reduced risk of slip that it affords. Here we hypothesize that when grip forces are not limited by strength or fatigue, the safety margin for GF control is specified based on a fixed level of statistical confidence. Fig. 2.1 illustrates that this hypothesis predicts higher safety margins as uncertainty increases, because the risk of slip at a fixed GF level would increase with higher uncertainty (Fig. 2.1E,F).

To investigate the fixed statistical confidence hypothesis, we tested the prediction that the safety margin for GF control would be modulated by environmental variability. We designed a task in which subjects were exposed to environments with different amounts of variability. Subjects performed rapid point-to-point reaching movements while grasping a small object with physical dynamics that changed randomly from one trial to the next, with the amplitude of these random changes systematically varying from one block to the next, as illustrated in Fig. 2.2A-C. In this task, the direction of the load force (LF) was orthogonal to the grasp axis so that grip force (GF) levels, rather than finger configuration, determined whether the object slipped (see Fig. 2.1A,B). In addition to specifying the orientation of grasp relative to the LF, we restricted the possible grasp locations, as freedom to vary the finger configuration can

influence the applied GF (Fu et al., 2010). Thus, subjects could modulate grip force to prevent slip and independently modulate manipulatory force (MF) to control the motion of the object. After a practice period of 150 trials in which participants were familiarized with the basic task, they were exposed to 16 50-trial blocks with different levels of variability. During each block, the strength (*b*) of a viscous curl force-field environment (see equation 1 in the methods) was varied from one trial to the next by drawing *b* from a Gaussian distribution with mean zero and standard deviation σ. σ was fixed within each block at one of four possible values: 0, 1.8, 3.6 and 5.4 Ns/m, and four blocks were presented for each σ value during the 16-block session, in a randomized order. In order to give subjects time to adapt to each variability level, we focused our analysis on the last 70% of the data from each block (see Materials and Methods). It should be noted, however, that the adaptation was fairly rapid and so other choices for the window size yield similar results.

Critically, we applied a leftward static force offset of 4N (F_0 in equation 2.1) throughout the experiment on top of the time-varying FF perturbation in order to ensure that the net direction of the force would always remain leftward even for negative values of *b* as shown in Fig. 2.2C. *This ensured that negative values of b required smaller GFs than baseline in order to prevent slip whereas positive values of b required higher GFs.* Without this offset, negative values of *b* would result in rightward forces whereas positive values of *b* would generate leftward forces, with increased GF required to maintain grasp in both cases. Moreover, since load-increasing and load-decreasing values of *b* were *equally likely*, any symmetric positive vs. negative trial-to-trial learning (Scheidt et al., 2001; Takahashi et al., 2001b; Quaney et al., 2003; Braun et al., 2009) would balance out. As we estimated variability-driven responses based on the average responses during each variability block, the net effect of trial-to-trial learning predicted by current models of motor adaptation (Thoroughman and Shadmehr, 2000; Donchin et al., 2003; Smith et al., 2006; Joiner and Smith, 2008; Lee and Schweighofer, 2009; Sing et al., 2009; Sing and Smith, 2010) should be near zero.

2.4.2 Increased environmental variability leads to increases in GF but not MF control

We found that the average GF profiles displayed in each block systematically increased with variability. This is evident in both the raw GF profiles (Fig 2.2D) and baseline-subtracted GF profiles which show the variability-induced changes from the zero-variability baseline blocks (Fig 2.2E). The pattern of variability-induced GF change was similar in shape for all non-zero variability levels, but its amplitude grew systematically increased as σ increased. Changes in GF began at movement onset and peaked near the peak velocity point, when the effect of the LF from the viscous curl force-field perturbation would be the strongest (due to the velocity-dependent FF perturbation). The amplitude of these GF changes scaled with the variability level, peaking at 1.5 N for the high-variability blocks (Fig. 2.2E).

In order to quantify the magnitude of variability-induced changes in both GF and MF, we computed adaptation coefficients for each movement for both GF and MF. These adaptation coefficients were based on linear regressions of the corresponding GF or MF force profile onto the LF profile (Scheidt et al., 2000; Smith et al., 2006; Joiner and Smith, 2008; Gonzalez Castro et al., 2011; Joiner et al., 2011; Yousif and Diedrichsen, 2012; Sing et al., 2013; Gonzalez Castro et al., 2014; Wu et al., 2014) (see Materials and Methods). We found that subjects exhibited significantly greater GF adaptation in the low, medium, and high-σ blocks compared to the zero-σ blocks ($p<0.001$ in all three cases, one-tailed paired t-tests, Fig. 2.2F). Moreover, the GF adaptation was greater in both the medium and high-σ blocks than the low-σ blocks ($p<0.01$ in both cases), and greater in the high-σ blocks than the medium-σ blocks ($p<0.05$, one-tailed paired t-tests).

In contrast, we found no consistent changes in MF control with variability. The subject-averaged MF profiles were essentially identical for different amounts of variability (Fig. 2.2G), and the variability-related changes in MF shown in Fig. 2.2H were small. Correspondingly, the MF adaptation coefficients were not significantly different from zero (Fig. 2.2I). In particular, note that the high-σ environments display MF levels that are essentially the same as the zero-σ environment ($p>0.48$, two-tailed paired t-test Fig. 2.2G-H).

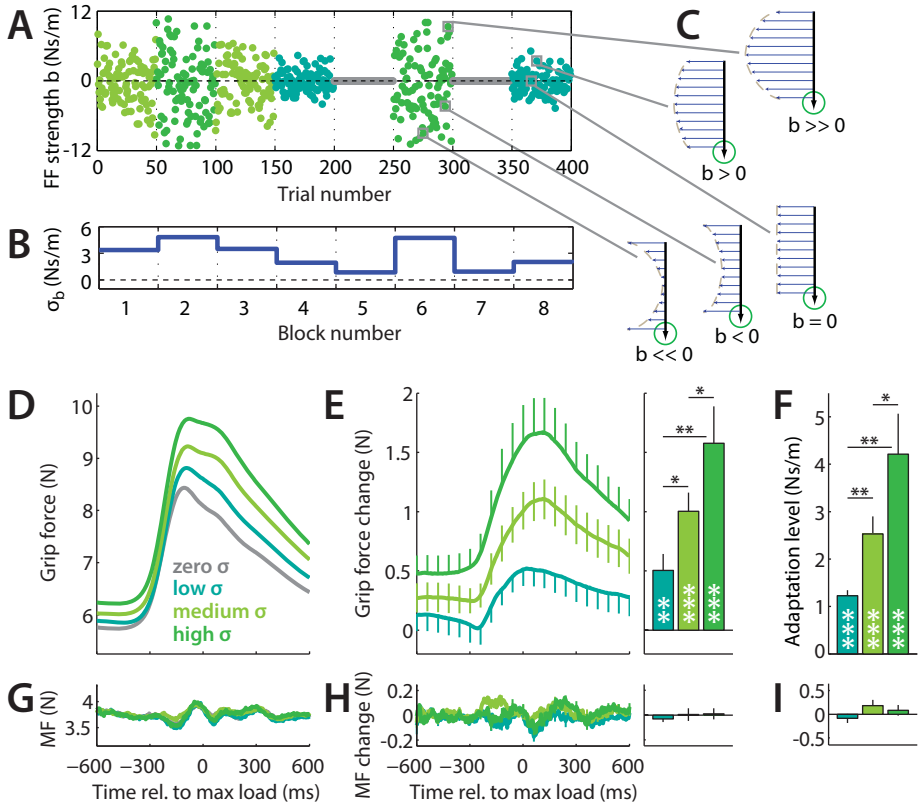


Figure 2.2. Grip forces, but not manipulatory forces, scale with the level of environmental variability.

(A) Design of Experiment 2.1. Following a baseline period (not shown) participants were exposed to environments which randomly changed in variability from one 50-trial block to the next, as illustrated. Specifically, the viscosity (*b*) of a velocity-dependent force-field perturbation was drawn from a normal distribution with mean zero and standard deviation (σ) that assumed one of four values (see Materials and Methods). (B) Standard deviation (σ) of the values of *b* for each block, illustrating the four distinct variability levels. (C) Example force-fields (blue arrows) presented to subjects on individual movements (black arrows) for different values of b. Note that a force offset F0 prevents even a strongly negative b from reversing the direction of the net load force. (D) Mean GF profiles for each of the four variability levels, showing increasingly stronger forces as σ increases. (E) Changes in GF for the three nonzero variability levels compared to the zero σ variability level. Error bars indicate SEM across subjects. The right panel shows the average amplitude of these changes for a short time window (±50ms) around the peak speed point. (F) Baseline-referenced adaptation levels (mean±SEM) for the GF profiles in E, showing significant increases in GF levels with variability. * $p<0.05$, ** $p<0.01$, *** $p<0.001$. Red indicates comparisons with the baseline level. Note how the separation between different variability groups is even more clear than in the right panel of 2E, in spite of the fact that here adaptation levels are based on the transient changes in GF during movement, and thus these comparisons do not benefit from the separation between pre-movement levels suggested in 2E. (G-I) Same as D-F but for MF.

These findings indicate that when the mean dynamics are held constant, GF but not MF control systematically depends on the variability of environmental dynamics.

2.4.3 GF control is even more sensitive to changes in the LF variability than the mean LF

Since previous work maintained that GF is controlled based on the expected LF, we wanted to know how the strength of the dependence of GF on LF variability observed in Experiment 2.1, compares to the dependence of GF on the expected value of the LF. To that end, we estimated the sensitivity (i.e. the slope of the relationship) that GF control displays with respect to both the standard deviation (σ_{LF}) and the mean (μ_{LF}) of the LF. Across the four different variability levels (zero, low, medium, and high) in Experiment 2.1, we found the relationship between subject-averaged GF levels and σ_{LF}, to be approximately linear, with a strong positive correlation between GF levels and σ_{LF} (R^2>0.99, slope = 1.02, see the squares and solid line in Fig. 2.3B). When we estimated the GF/σ_{LF} sensitivity for all 11 individual participants, we found positive slopes for all 11, with an average value of 1.02±0.21 (mean±SEM across subjects, $p<0.001$, one-tailed paired t-test, Fig. 2.3B,D). Note that these slopes are all dimensionless because the levels of GF adaptation and σ_{LF} are both quantified in the units of force-field strength (Ns/m).

To put the amplitude of the GF/σ_{LF} sensitivity into perspective, we estimated the GF/μ_{LF} sensitivity by measuring GF responses to systematic changes in the LF environment. To that end, we presented subjects with a deterministic step change in the environmental dynamics, changing the viscosity coefficient (*b*) of a viscous curl FF to a constant value of either +7.5 Ns/m (Experiment 2.2a, leftward, n=18) or -7.5 Ns/m (Experiment 2.2b, rightward, n=18) for 240 trials following a 160-trial baseline period with *b*=0(Fig. 2.3A). Note that, like in Experiment 2.1, a leftward static offset force was present throughout to ensure that LFs were reduced in magnitude rather than reversed in direction during rightward -FF trials (Fig. 2.2C). We found highly significant and nearly symmetric changes in GF during asymptotic adaptation (last 60 trials) in Experiments 2.2a and 2.2b (2.46±0.51 Ns/m, p=0.0002 and -2.36±0.74 Ns/m, p=0.0055, respectively, two-tailed paired t-tests, Fig. 2.3C). These changes in GF correspond to sensitivity estimates (slopes) of 0.33±0.07 and 0.31±0.10 for individuals in Experiments 2.2a and 2.2b, respectively, and

0.32±0.06 for the combined data (Fig. 2.3C,D). Remarkably, these values are 3-fold lower than the sensitivity of GF to changes in variability observed in Experiment 2.1 (1.01±0.19, compare panels B and C in Fig. 2.3), indicating LF variability to be a stronger determinant of GF control than mean LF.

We conducted three additional experiments to determine the robustness of the finding that GF control is more sensitive to σ_{LF} than to μ_{LF}. We first checked whether the stability of GF adaptation in the 50-trial blocks from Experiment 2.1 was sufficient for accurate GF/σ_{LF} sensitivity estimates. In Experiment 2.1b, (see Fig. 2.3A), variability was modulated from one block to the next like in experiment 2.1, but with longer (100 vs 50 trials) and correspondingly, fewer (8 vs 16) blocks. Experiment 2.1b yielded sensitivity estimates of 1.15±0.17 across individuals ($p<10^{-6}$, one-tailed paired t-test), similar to the ones obtained from Experiment 2.1 (see Fig. 2.3B,D), suggesting GF/σ_{LF} sensitivity estimates to be stable for blocks of 50 trials or greater.

In a second additional experiment (Experiment 2.2c), we checked whether the GF/μ_{LF} sensitivity estimates from Experiments 2.2a and 2.2b would be consistent with the sensitivity observed from a smaller change in μ_{LF}. Because Experiments 2.2a and 2.2b were based on changes in μ_{LF} with a magnitude (7.5 Ns/m) that was greater than the σ_{LF} changes examined in experiments 2.1 or 2.1b (up to 4.1 Ns/m), examining a smaller change in μ_{LF} allowed us to compare GF/σ_{LF} and GF/μ_{LF} sensitivity estimates using similar displacements for σ_{LF} and μ_{LF}. This also allowed us to examine the linearity of the relationship between GF and μ_{LF}. Experiment 2.2c mirrored 2.2a, except that the FF amplitude (*b*) was halved (+3.75Ns/m vs +7.5Ns/m). The results showed an asymptotic adaptation level in the last 60 trials of the 240-trial FF exposure that was almost exactly half of that observed in experiment 2.2a (1.16±0.42 vs 2.46±0.51 Ns/m). This corresponds to nearly identical GF/μ_{LF} sensitivity estimates for Experiments 2.2a, 2.2b, and 2.2c (0.33±0.07, 0.31±0.10, and 0.31±0.11 Ns/m, respectively), suggesting a fixed level of GF/μ_{LF} sensitivity up to at least 7.5Ns/m, in line with a linear relationship between GF and μ_{LF}. Comparison of the GF/σ_{LF} sensitivity estimates from the high-σ blocks in Experiment 2.1 and the GF/μ_{LF} sensitivity estimate from 2.2c revealed the former to be significantly greater than the latter (1.04±0.21 vs

0.31±0.11, p=0.0043, one-tailed unpaired t-test, see right panel of Figure 2.3D). This is in line with the idea that GF is substantially more sensitive to the variability than the expected value of the LF environment, even when displacements for σ_{LF} and μ_{LF} are similar (4.1Ns/m and 3.75Ns/m, respectively).

In a third additional experiment (Experiment 2.1c), we checked whether the GF/σ_{LF} sensitivity estimates from Experiments 1 and 1b would be consistent with the sensitivity observed when GF variability was centered around a non-zero mean. We performed this experiment immediately after Experiment 2c in the same session with the same participants, also allowing us to compare GF/σ_{LF} and GF/μ_{LF} sensitivity estimates in the same participants. Note that we used a value of +3.75Ns/m for μ_{LF} in 1c to maintain continuity with the FF training in 2.2c. Again, we found GF to be more sensitive to σ_{LF} than μ_{LF} (0.78±0.11 vs 0.31±0.11, p=0.0013, one-tailed paired t-test, see right panel of Figure 2.3D), consistent with our previous findings.

Analogous analysis of the MF data from Experiments 2.2a, 2.2b and 2.2c revealed highly significant asymptotic adaptation of MF levels to changes in μ_{LF} (Exp. 2.2a (b = +7.5Ns/m): 6.66±0.27 Ns/m; Exp. 2.2b (b = -7.5Ns/m): -7.14±0.15 Ns/m; Exp. 2.2c (b = +3.75Ns/m): 3.15±0.25 Ns/m, $p<10^{-6}$ in all three cases, two-tailed paired t-tests, Fig. 2.3F). Note the nearly symmetric increases and decreases in the adaptive responses for Experiments 2.2a and 2.2b. These MF levels correspond to MF/μ_{LF} sensitivity estimates of 0.89±0.04, 0.95±0.02, and 0.84±0.07, individually (Fig. 2.3G). The similarity between these estimates corresponds to the nearly linear relationship between MF and μ_{LF} observed in Figure 2.3F. In contrast, the overall MF/σ_{LF} sensitivity we observed in Experiments 2.1, 2.1b and 2.1c was much smaller and not different from zero (0.002±0.020, p>0.9, two-tailed paired t-test, Fig. 2.3E,G) in line with the lack of MF modulation observed in Figure 2.2G-I when variability was increased.

2.4.4 Grip force control maintains a safety margin based on statistical confidence

Remarkably, our findings show not only that GF control is sensitive to the standard deviation of environmental dynamics, but that this sensitivity is higher ($p<10^{-6}$, two-tailed unpaired t-test between

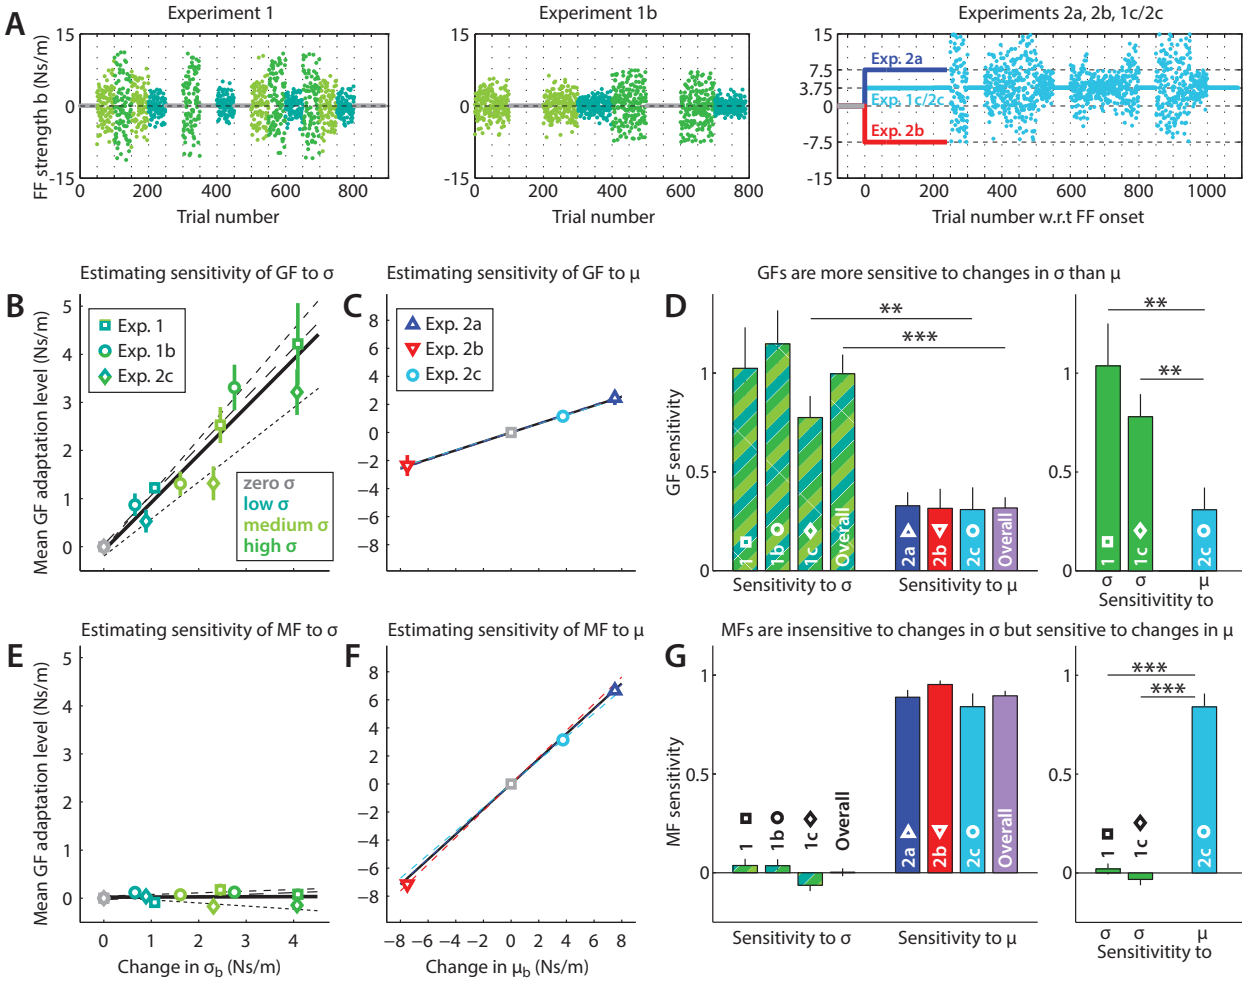


Figure 2.3. Grip forces are three fold more sensitive to changes in the variability than changes in the expected value of environmental dynamics.

(A) Comparison between the design of Experiments 2.1, 2.1b, 2.2a, 2.2b and 2.1c/2.2c. Examples of training schedules are shown. The duration of variability blocks in Experiment 2.1b is double that of Experiment 2.1. Experiments 2.2a and 2.2b present a step increase or decrease of the FF, respectively. Experiment 2.1c/2.2c begins with a step increase in the FF which persists for 240 trials (Experiment 2.2c) and is then followed by the addition a variable FF component (Experiment 2.1c) in an analogous fashion to Experiment 2.1. (B) Baseline-referenced adaptation levels (mean±SEM) for GF for each of the four types of variability blocks and each experiment (2.1,2.1b and 2.1c) plotted against the corresponding increases in the standard deviation of *b* with respect to the baseline. The black lines indicate linear fits to the data (thin dashed lines, in order of increasing density: Experiment 2.1, 2.1a and 2.1c; thick solid line: overall data from all three experiments). The three groups display similar slopes for GF vs. σ and indicate that GF is sensitive to environmental variability. (C) Baseline-referenced adaptation levels (mean±SEM) for GF for experiments 2.2a, 2.2b and 2.2c, plotted against the corresponding changes in the mean of *b*

Figure 2.3 (continued)

with respect to the baseline. The colored lines indicate linear fits to the data for each experiment (the black line indicates the linear fit to the overall data from all three experiments). Note that these fits lie on top of each other indicating that the sensitivity of GF to the μ is linear over the range of changes in μ that we examined. The aspect ratio is the same as in (B), illustrating that the sensitivity of GF to μ (i.e. the slope of the fits) is substantially smaller than the sensitivity of GF to σ (the slope of the fits in (B)). (D) Grip forces show significantly greater sensitivity to the standard deviation, σ, compared to the mean, μ, of environmental dynamics. In the left panel, we calculate sensitivities to σ based on the slopes in 3B. In the right panel, we calculate sensitivities based on the high-variability blocks from Experiments 2.1 and 2.1c, where the magnitudes of the baseline-referenced changes in σ were similar to magnitude of the baseline-referenced changes in μ from Experiment 2.2c (3.75 Ns/m). (E-G). Manipulatory forces are not sensitive to changes in the standard deviation, but only to the mean of environmental dynamics. The format of panels E-G parallels the format of panels B-D.

Experiments 2.1, 2.1b and 2.1c vs. Experiments 2.2, 2.2b and 2.2c) than the sensitivity to the environment's mean dynamics as illustrated in Figure 2.3 (2.82±0.27 fold higher, estimated using bootstrap when combining the data from Experiment 2.1, 2.1b, 2.1c, 2.2a, 2.2b and 2.2c). Since the GF/σ_{LF} sensitivity acts to scale the GF safety margin against slip by the LF variability, the ratio between the GF/μ_{LF} sensitivity and the GF/σ_{LF} sensitivity can be used to determine the statistical confidence associated with the observed safety margin (see Materials and Methods). In particular, the ratio of 2.82±0.27 that we observed corresponds to GF control in a variable environment that is 2.82 standard deviations above that required for the expected value of dynamics. For normally distributed environmental perturbations, like the ones administered in experiments 1 and 1b, this corresponds to a z-score of 2.82±0.27, which amounts to greater than 99% confidence interval against slip. Unfortunately, the precise frequency of slippage incidents is difficult to measure experimentally as these incidents usually result in micro-slips which subjects rarely notice and often correct almost instantaneously (Johansson and Westling, 1984, 1987). In fact, we were not able to detect any instances of loss of grasp from slip during movement in our experiments. The low likelihood of overt slippage we observed is consistent with the high statistical confidence against slip that we report here.

2.4.5 A computational model for understanding how GF control responds to both the mean and variability of load force dynamics.

We next investigated how the adaptation of grip force levels to environmental variability might manifest following a sudden but enduring change in the environment as in experiment 2.2. We hypothesized that the motor system might estimate the current levels of the mean and the variability of the LF environment based on a window of recent experience. If this is the case, a recent sudden change in the environment should increase the motor system's estimate of variability because some trials in the window would be at the previous perturbation level and others would be at the new perturbation level. This increase should wane after prolonged exposure when most or all of the trials in the recent memory window would be associated with the new environment. The Additive Combination of Expected Value and Uncertainty (ACEVU) model, illustrated in Figure 2.4, posits that the overall adaptive response simply reflects the sum of a mean-driven and a variability-driven component, based on windowed estimates of the expected value and the variance of the environment.

While the exact shape of the memory window is not crucial for the model, we used an exponentially decaying memory window to estimate both the mean and the variability of the environment because it describes an estimation process which weighs past experiences less and less as they become more distant. This window shape is also compatible with linear state-space models of trial-by-trial adaptation with finite retention (retention coefficients < 1), it captures the roughly exponential learning curves that have been observed for motor adaptation, and a recent study found evidence that subjects track the variance of a visuomotor perturbation using a exponential memory window (Landy et al., 2012).

Online estimates of the mean and variability of the environment for the ACEVU model in response to a step change in the environment are illustrated in Figure 2.4. Following the step, the mean estimate changes monotonically, asymptotically approaching a level appropriate for the new environment (Fig. 2.4A) in a symmetric manner that is specific for positive and negative steps. In contrast, the variability estimate increases transiently as discussed above, but in an identically positive manner after either step

change (Fig. 2.4B). Thus a critical feature of the ACEVU model is that the mean-driven component of GF adaptation produces opposite responses for opposite changes in the environment (i.e. positive versus negative FFs), whereas the variability-driven component produces identical responses which cannot be specific to the direction of the step change. This is the case because the variance of a signal is equal to the variance of its opposite. Panels C and D show how estimates of the expected value and variability would combine in contributing to the adaptation of GF and MF throughout the time course of this adaptation. In these predictions, MFs and GFs both respond to changes in the mean of the environment as previous studies have shown (Flanagan and Wing, 1993, 1997; Scheidt et al., 2001; Takahashi et al., 2001a; Mawase and Karniel, 2012); however, only GFs respond to changes in the variability, based on the results of Experiment 2.1.

2.4.6 The ACEVU model predicts non-monotonic adaptation

The superposition of a specific μ-driven and a nonspecific σ-driven component of the adaptive response in the ACEVU model, results in a series of counterintuitive predictions for GF adaptation. In combining separate contributions from μ-driven and σ-driven adaptation, our model predicts that positive perturbations (increases in LF) will yield adaptation curves that rise especially rapidly but overshoot their asymptotic levels before settling (Fig. 2.4C). According to the model, the rapid rise will occur because during initial adaptation, the task-specific μ-driven and nonspecific σ-driven components will both contribute to an increase in GF levels. However, the σ-driven contribution will fall off as training proceeds, and an overshoot will result if this fall off outpaces the continued rise of the task-specific μ-driven component.

Perhaps even more counterintuitive is a second prediction also illustrated in Figure 2.4C: that *decreases in LF will initially lead to increases in GF*. This will occur if the increase in GF due to the nonspecific σ-driven component of adaptation to a negative perturbation outweighs the reduction in GF due to the task-specific μ-driven component. Prolonged training with a decreased LF will then eventually lead to a

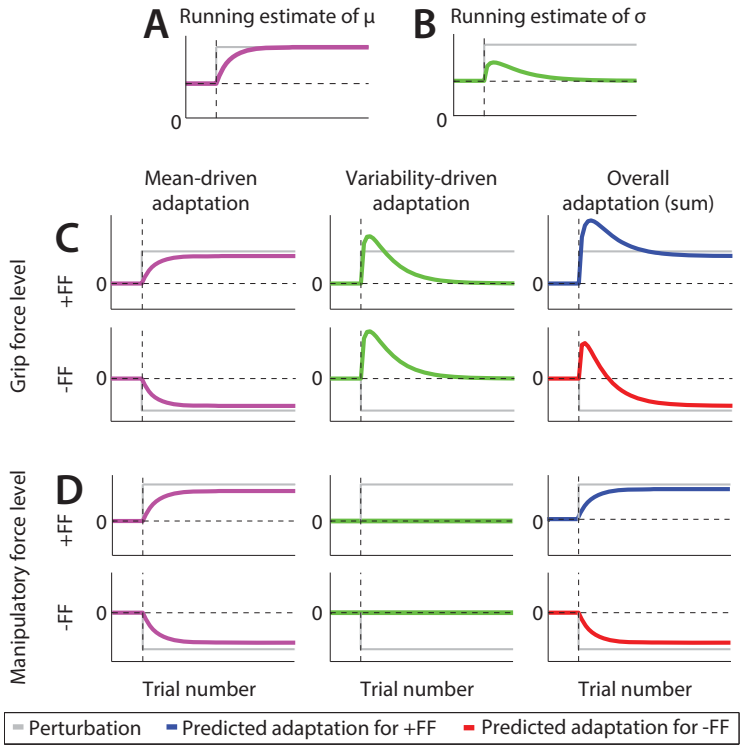


Figure 2.4. Predictions of a model for the adaptation of grip and manipulatory forces based on distinct mean-driven and variability-driven components.

(A,B) Running estimates of mean and variability during a step change in environmental dynamics. These estimates weigh recent memories more strongly than older ones using an exponentially-shaped memory window. The mean estimate (A) will tend monotonically towards the post-step value, whereas the standard deviation estimate (B) will transiently increase, rising rapidly following the step change and then receding as memories that include the pre-step environment fade away. (C) Mean-driven and variability-driven components of the adaptive GF response following a positive (+FF) or negative (-FF) change in environmental dynamics. Critically, whereas the mean-driven component will follow positive or negative changes in environmental dynamics, the variability-driven component will be positive for changes in either direction. Thus our Additive Combination of Expected Value and Uncertainty (ACEVU) model, which incorporates these two components, predicts asymmetrical, non-monotonic GF adaptation curves. A positive FF would elicit overshooting adaptation (blue) due to complementary mean-driven and variability-driven components, whereas a negative FF would elicit bimodal adaptation (red) due to opposing mean-drive and variability-driven components. (D) Same as (C) but for MFs. As MFs are insensitive to the variability of environmental dynamics, the variability-driven component will be zero throughout adaptation leading to monotonic adaptation curves resulting in predictions of symmetrical, monotonic adaptation for the ACEVU model, for both positive and negative changes in environmental dynamics.

reduction of GF levels towards those appropriate for the environment as the nonspecific σ-driven component subsides. Thus the ACEVU model remarkably predicts an initial adaptive response in the "wrong" direction followed by a reversal resulting in appropriately decreased GF levels when decreased LFs are encountered.

Unlike GF control, nonspecific σ-driven adaptation should not contribute to MF control according to the results of Experiment 2.1 which show MF to be insensitive to σ_{LF}. Consequently our model predicts that MF adaptation curves will be fully determined by task-specific μ-driven adaptation and thus be monotonically directed towards an asymptotic level appropriate for each new environment be it positive or negative as shown in Figure 2.4D.

2.4.7 Grip forces adapt to an additive combination of uncertainty and the expected value

To examine how grip forces and manipulatory forces actually evolve following a sudden change in physical dynamics, we analyzed the full time course of the adaptation data from experiment 2.2. For the leftward +FF group, we found that the GF learning curve displayed a clear overshoot. They reached a maximum value at about 10 trials after exposure before settling at a significantly lower asymptotic adaptation level lower as shown in the blue GF profiles in Figure 2.5A-B and the blue learning curve in 2.5C (adaptation levels of 7.90±1.15 Ns/m versus 2.46±0.51 Ns/m for the first 20 trials vs of adaptation versus the last 60, $p<0.001$, two-tailed paired t-test, Fig. 2.5D). Individually, 16 out of the 18 participants in the +FF group displayed higher average GFs in the first 20 trials than in the last 60 of the FF training epoch.

For the rightward –FF group, the initial adaptation resulted in seemingly inappropriate increases in GF above baseline which persisted for about 40 trials before the adaptive responses decreased below baseline levels where they slowly settled to asymptote as shown in Fig. 2.5C-E The adaptation level during the first 20 trials was significantly greater than baseline, whereas the adaptation level during the last 60 was significantly lower than baseline as shown in the red GF profiles in Figure 2.5A-B and the red learning

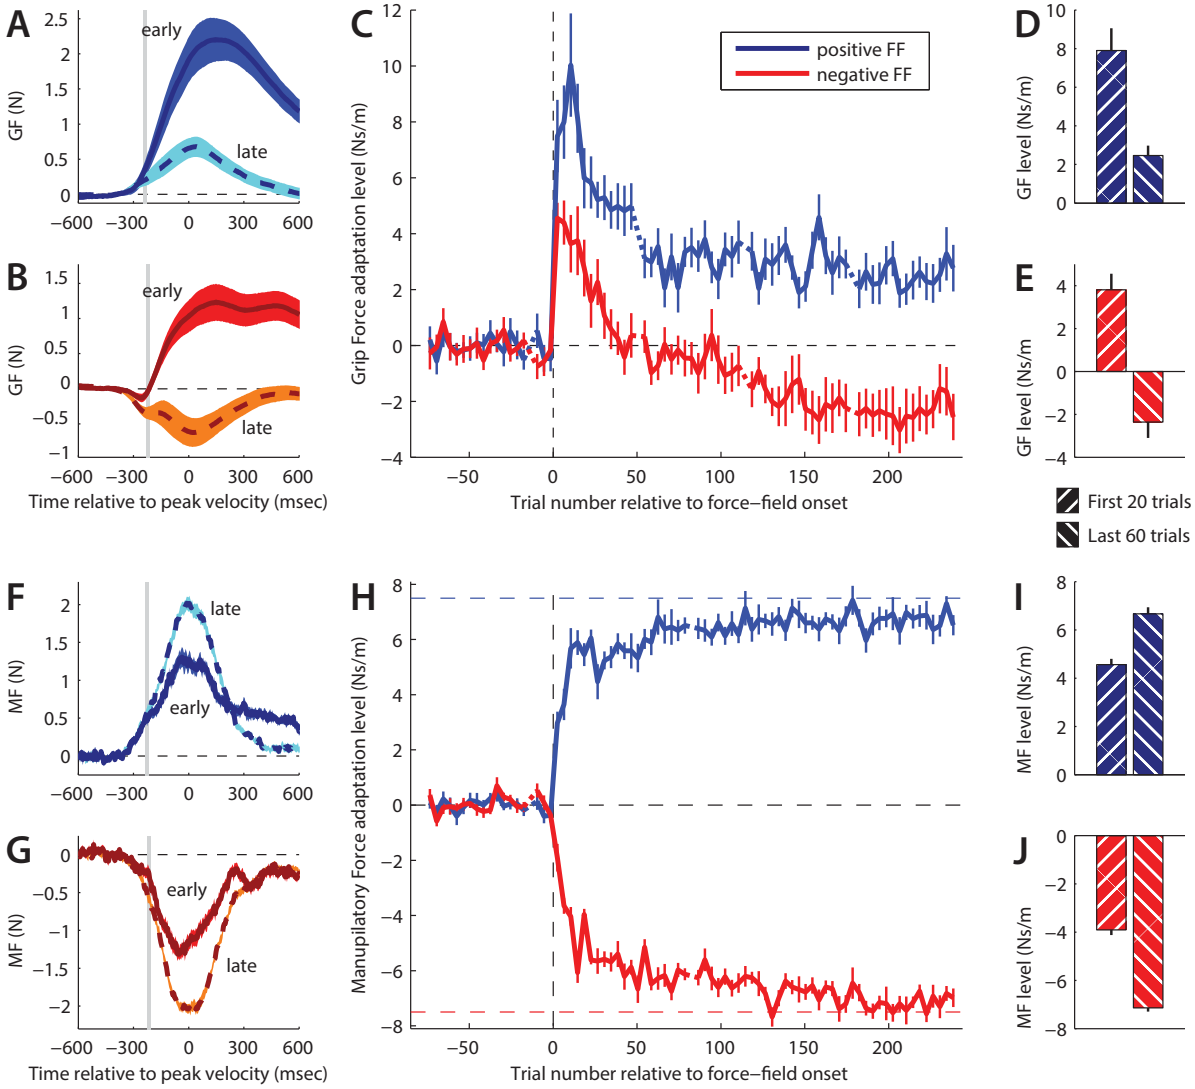


Figure 2.5. Grip forces, but not manipulatory forces, exhibit asymmetric, non-monotonic adaptation curves following opposite perturbations.

(A,B) Baseline-referenced GF profiles early and late (mean of first 20 vs last 60 trials) following the onset of positive (A, leftward) and negative (B, rightward) force-field perturbations. Note that in B, the early response is opposite of what is appropriate for the applied force-field, whereas the late response is appropriately directed. The shaded areas represent SEM across subjects. The gray vertical line represents the 95% confidence interval for movement onset. (C) Corresponding GF adaptation curves based on the adaptation coefficient (see Materials and Methods). GFs initially increase in both force-fields, but eventually asymptote at levels that are roughly equal and opposite. This corresponds to overshooting adaptation for the +FF group and bimodal adaptation for the –FF group as predicted by the model in 2.4C. (D,E) GF levels for the early and late (mean of first 20 vs last 60 trials) stages of adaptation. (F-J) Same as A-E but for MF, illustrating symmetric, monotonic adaptation patterns as predicted by the model in 2.4D.

curve in 2.5C (adaptation levels of +3.80±0.75 Ns/m vs -2.36±0.74 Ns/m, $p<0.0001$ and $p<0.01$, respectively, two-tailed paired t-tests, see Fig. 2.5E). Individually, the GF level was higher than baseline in 16 out of 18 subjects during the first 20 trials, and lower than baseline in 15 out of 18 subjects during the last 60 trials. Both the direction reversal observed in the –FF data (red), and the pattern of overshoot and subsequent fall off observed in the +FF data (blue) were remarkably well predicted by the ACEVU model displayed in Figure 2.4 which combines nonspecific σ-driven adaptation and FF-specific μ-driven adaptation.

2.4.8 Manipulatory forces adapt chiefly to the expected value of environmental dynamics

In contrast to GF adaptation, we found MF adaptation to be characterized by highly symmetric, monotonic responses to increases and decreases in LF, as evidenced by the MF learning curves displayed in Figure 2.5H and the force profiles displayed in Figure 2.5F-G. Accordingly, the amplitude of the average MF adaptation level during the first 20 trials was significantly lower in magnitude than the average adaptation level during trials 181-240 (leftward FF: 4.56±0.24 Ns/m versus 6.66±0.27 Ns/m, $p<10^{-4}$; rightward FF: -3.92±0.21 Ns/m versus -7.14±0.15 Ns/m, $p<10^{-6}$, two-tailed paired t-tests, Fig. 2.5I-J). The highly symmetric learning curves indicate that there is little nonspecific adaptation, with MF adaptation highly specific to the experienced environment. This is consistent with the idea that the nonspecific component of the adaptive response arises from σ-driven adaptation, which we found to be minimal for MF in Experiment 2.1.

2.4.9 Robustness of results using different metrics

To ensure that our observations were not specifically related to the adaptation coefficient metric we used to quantify FF adaptation, we repeated the analysis using an even simpler metric: the GF observed when the effect of *b* on the LF produced by the imposed environment was maximal (Sing and Smith, 2010; Pekny et al., 2011; Joiner et al., 2013). We found that the learning curves derived from this measure were essentially identical in shape to those displayed in Figure 2.5C,H ($R^2>0.97$ for the +FF and –FF GF learning curves and for the +FF and –FF MF learning curves).

2.4.10 Decomposition of grip forces into perturbation specific and nonspecific components

Based on the ACEVU model, task-specific adaptation should differentially affect the adaptive responses to two opposite perturbations, whereas the nonspecific adaptation will be common to both. We used this fact in order to decompose the experimentally observed adaptive responses to the +FF and –FF environments into specific (differential) and nonspecific (common) components (see Methods). Specifically, we applied this differential vs. common component analysis to the GF profiles (GF(t)) in Experiment 2.2a,b to determine the specific and nonspecific components of these profiles (Figure 2.6A,C). This analysis reveals that over the course of adaptation, the amplitude of the nonspecific common-mode GF profile rises from baseline during initial adaption and wanes as adaptation proceeds. Despite the changes in amplitude, the shape of the common-mode GF profile remains remarkably constant as training progresses (colored lines in Fig. 2.6A-B). The extent of this shape consistency is apparent when amplitude-normalized force profiles are compared to each other in Figure 2.6B. Critically, this shape is strikingly similar to the shape of the GF profiles observed in Experiment 2.1 when the effects of environmental variability were measured (R^2 values between 0.92 and 0.99 for the 6 common-mode profiles shown normalized in Fig 2.6B, and $R^2 = 0.98$ for the mean of these). Note that the GF profile from Experiment 2.1 is shown as the dark gray dashed line in 6B for reference. The close match between the shapes of the force profiles for the transiently occurring common-mode component of the step response from Experiment 2.2 and the variability-driven adaptive response directly measured in Experiment 2.1 is consistent with the idea that the transient nonspecific component the GF learning curve reflects the motor system's estimate of environmental variability as predicted by the ACEVU model.

On the other hand, the task-specific differential mode component of the grip force profile assumes a shape very close to the shape of the viscous LF profile (Fig. 2.6C-D). This similarity increases with training, with R^2 values increasing to 0.91 by the end of the training period ($R^2 = 0.93$ for the average profile). Thus the shapes of the force profiles associated with the specific and nonspecific components of the GF adaptation data from Experiment 2.2 were more closely matched to the mean LF profiles observed in

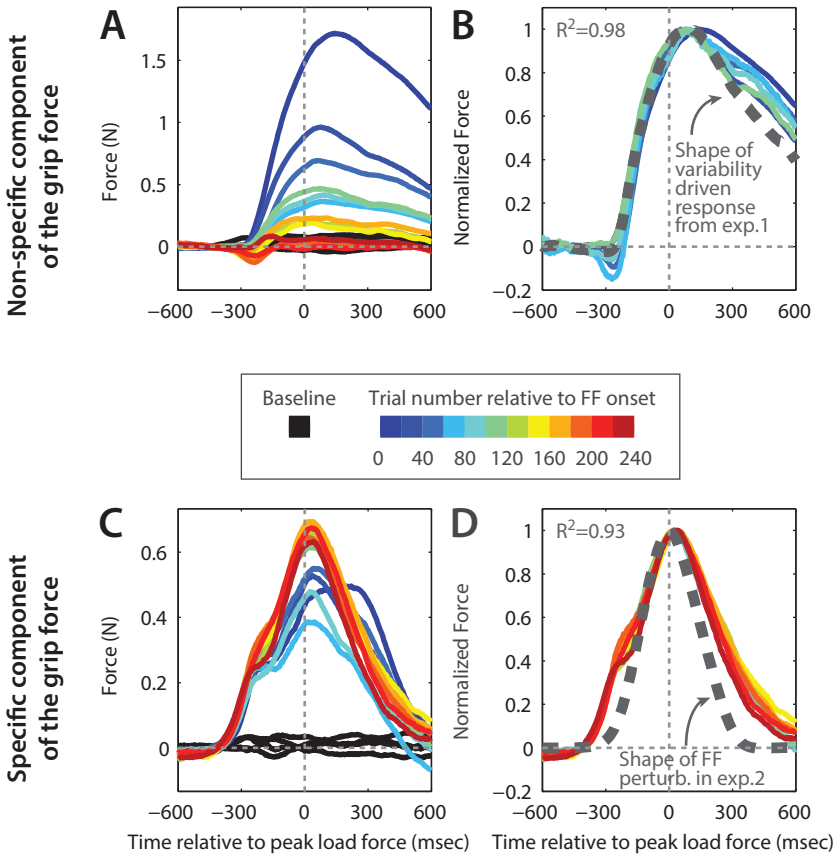


Figure 2.6. Distinct shapes of the specific and non-specific components of grip force profiles during learning reflect the shape of mean-driven and variability-driven responses, respectively.

GF profiles were averaged within bins of 20 trials and across subjects. The FF-specific (non FF-specific) component was obtained by taking, for each such bin, the differential (common) mode of GF profiles from the positive and negative FF experiments. (A) Non-specific components of GFs (warmer colors indicate later parts of training). (B) Normalized profiles for the non-specific components for the first six bins after perturbation onset. The grey dashed curve shows the normalized shape of the mean variability-driven change in GF shown in Fig. 2.2E. Its similarity to the average shape of the non-specific GF components (R^2=98%) is in line with the idea that non-specific GF adaptation is driven by environmental variability. (C) FF-specific components of GFs (color scheme the same in all panels). (D) Normalized profiles for the specific components last eight bins. The grey dashed curve shows the shape of the ideal compensatory force pattern from Experiment 2. Its similarity to the average shape of the specific GF components (R^2=93%) is in line with the idea that specific GF adaptation is driven by the expected value of the environment.

Experiment 2.2 and the variability-driven changes in GF observed in Experiment 2.1, respectively, than to each other (R^2 values of 0.93 and 0.98, respectively, versus an R^2 of 0.50). This finding provides direct experimental support for the theoretical links posited in the ACEVU model between task-specific adaptation and expected value estimation and between nonspecific adaptation and variability estimation.

2.4.11 Learning curves for the specific and nonspecific components of GF and MF adaptation

We then leveraged our knowledge about the shapes of the force profiles associated with the nonspecific (Fig. 2.6A-B) and task-specific (Fig. 2.6C-D) components of adaptation to estimate the learning curves associated with each component. Instead of the standard practice of singly regressing the adaptive responses onto the shape of the viscous load force profile to estimate an adaptation coefficient, as in Fig. 2.5, we simultaneously regressed the adaptive responses we observed onto the shapes of both the common-mode and differential-mode components (see Materials and Methods), to obtain learning curves for non-specific (green) and task-specific (purple) adaptation, respectively, for both GF (Fig. 2.7A) and MF (Fig. 2.7B). This analysis reveals that the task-specific differential-mode component of adaptation (purple) displays a monotonic increase towards an asymptote, following a similar time-course for both GF and MF (see Fig. 2.7C for a direct comparison). In contrast, the nonspecific common-mode component (green) for the GF is transient (it displays a rapid increase towards a sharp peak and then a gradual decay toward baseline) and present only during GF adaptation. This is consistent with the mirror-opposite MF learning curves observed for the leftward and rightward FF groups in Figure 2.5I which suggest fully-specific MF learning and with the finding in Experiments 2.1 and 2.1b that MFs are unrelated to environmental variability. Remarkably, we find that the shapes of these task-specific and nonspecific components of the GF learning curves (Fig 2.7A-B) closely match the predictions of the ACEVU model (Fig 2.4) for distinct mean-driven and variability-driven responses based on exponentially-windowed running estimates of these quantities.

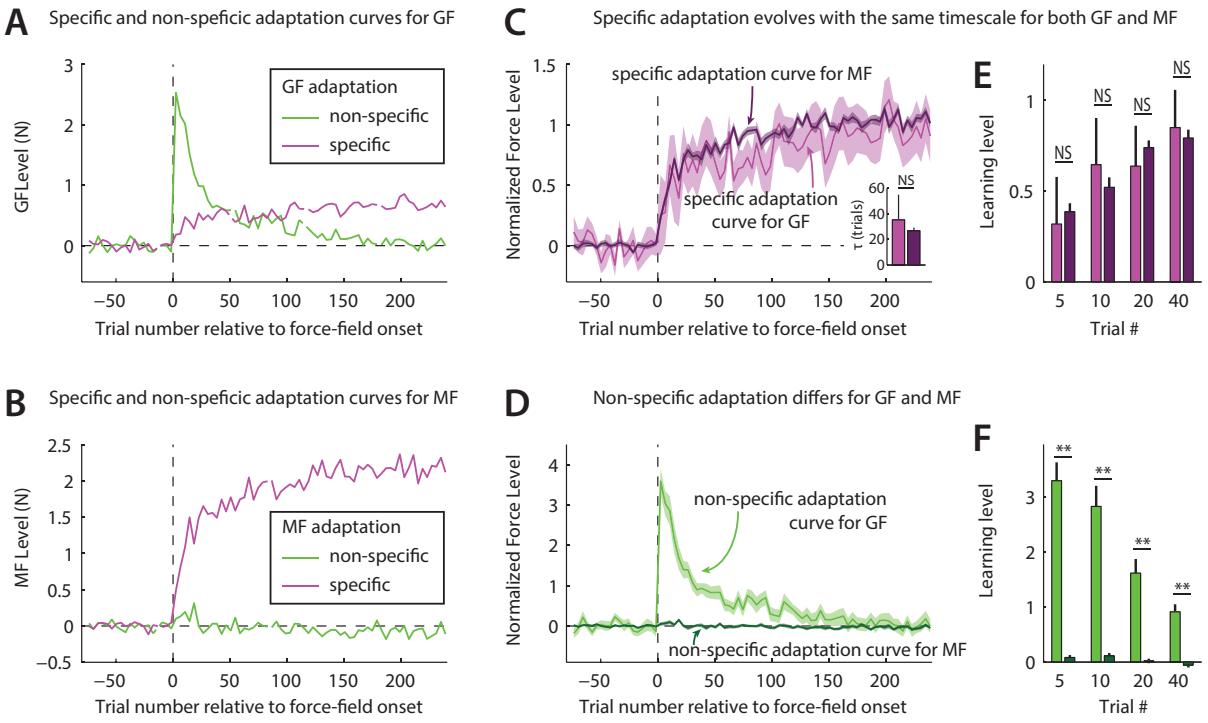


Figure 2.7. Adaptation curves for specific and non-specific components.

(A,B) Estimates of the specific (magenta) and non-specific (green) components of the learning curves for GF (A) and MF (B). The shapes of these components match the predictions in Figure 2.4C-D. (C) Comparison of the learning curves for the specific components of GF and MF, showing that they display similar time-courses of adaptation. Adaptation curves were normalized by the overall asymptotic adaptation level (last 60 trials). Shading indicates SEM across subjects. The inset shows the time constants of adaptation for these two specific components, (mean ±SEM estimated using bootstrap). (D) Same as (C) but for the non-specific component. (E) Comparison of the adaptation levels for the specific components of GF (light) and MF (darker) in 5-trial windows centered around trials 5, 10, 20 and 40. None of the differences are significant ($p>0.5$ in all cases). (F) Same as (E) but for the non-specific component. All differences are highly significant ($p<10^{-6}$ in all cases).

The decomposition of the learning curves into task-specific and nonspecific components provides the opportunity to compare the rate of task-specific GF and MF adaptation. We observed a strong similarity between the learning curves for the task-specific (purple) adaptation for GF and MF ($r = 0.84$), suggesting similar adaptation rates (see Fig. 2.7C which compares asymptote-normalized task-specific learning curves). Correspondingly, fitting single exponentials to these learning curves revealed no difference between the time constants for task-specific GF vs. MF adaptation (35.1±19.6 vs. 26.6±2.6 trials,

p=0.334, estimated using bootstrap, inset in Fig. 2.7C). Moreover, we found no differences between normalized adaptation levels for GF and MF at trials 5, 10, 20 and 40, estimated using 5-trial bins ($p>0.5$ in all cases, two-tailed paired t-test, see Fig. 2.7E). Together these findings indicate that the FF-specific components of GF and MF adapt at essentially the same rate.

2.5 Discussion

The current results show that the variability of LF dynamics plays a crucial role in GF control, in contrast to previous work which primarily focused on how the predictive control of GF depends on the expected value of environmental dynamics(Johansson and Westling, 1984; Westling and Johansson, 1984; Johansson and Westling, 1988; Flanagan and Wing, 1993, 1997; Flanagan et al., 2003). In experiment 2.1, we found that modulating load force variability has no effect on MF control but can drive large, systematic changes in GF control (Fig. 2.2). In fact, we found GF control to be 3-fold more sensitive to the standard deviation of the environment than its mean (Fig. 2.3), suggesting that GF control is effectively programmed to provide 3-sigma statistical confidence against slip and that internal estimates of environmental variability may play an even more important role in GF control than internal estimates of expected value.

We next examined how variability-driven GF control might affect the responses to sudden deterministic (step) changes in the environment. The ACEVU model we created to instantiate our hypothesis that GF adaptation is based on a superposition of distinct adaptive responses to the mean and variability of the LF environment predicts surprisingly unusual asymmetric, and nonmonotonic learning curves following sudden changes in load (Fig. 2.4). Experimental data directly confirm these unusual predictions, demonstrating adaptation that begins in the wrong direction when the load decreases, and overshoots its asymptote when the load increases (Fig. 2.5). Moreover, we used the ACEVU model to specifically decompose GF adaptation into the force profile components associated with σ-driven and μ-driven learning and to determine the corresponding learning curves (Figs. 2.6-2.7). Interestingly, we found distinct shapes for these force profile components that were predicted by the GF responses to environmental variability directly measured in Experiment 2.1 and the shape of the expected dynamics, respectively (Fig. 2.6). This decomposition also revealed on one hand that σ-driven nonspecific learning is present in GF adaptation but absent in MF adaptation, but on the other that μ-driven task-specific learning proceeds at remarkably similar rates for GF and MF adaptation (Fig. 2.7), indicating that

previously reported differences in adaptation rate for GF vs MF arose from the failure to take variability-driven GF control into account.

2.5.1 Previous ideas about safety margins for grip force control

Our results challenge the longstanding idea that, in precision grip tasks, the motor system maintains a constant GF/LF ratio (Johansson and Westling, 1984; Westling and Johansson, 1984), resulting in a safety margin that is a fixed fraction of the expected LF dynamics and is thus insensitive to LF uncertainty, analogous to the idea of a fixed safety factor like that often used in engineering design(ACI Committee 318, 2005; NASA, 2008). In contrast, our data show that GF is 3-fold more sensitive to LF variability than to the expected LF, resulting in GF/LF ratios that are readily modulated by environmental uncertainty rather than remaining fixed.

The sensitivity of GF control to environmental variability may have been difficult to appreciate in previous studies due to a focus on static paradigms that are likely to incur substantially smaller uncertainty than paradigms that focus on transient dynamics during rapid motion. For example, the predominant paradigm for studying GF control has involved the maintenance of a single static posture during slow lift and hold tasks(Johansson and Westling, 1984; Westling and Johansson, 1984; Johansson and Westling, 1987, 1988; Cole and Johansson, 1993). However, even in these tasks, the role of variability in driving GF control may have been underappreciated. In particular, the finding that the safety margin for GF control scales linearly with the LF in a static hold task(Westling and Johansson, 1984; Cole and Johansson, 1993), may be a variability-driven rather than an expected value driven effect. Although variability was not explicitly manipulated, both the force variability associated with motor output(Jones et al., 2002) and the perceptual variability associated with weight estimation(Holway and Pratt, 1936; Weber, 1978) scale linearly with the load. In fact, Weber's Law, which maintains that perceptual uncertainty is generally proportional to intensity, was originally based on studies of weight perception(Weber, 1978). These observations make it clear that in static tasks the safety margin for GF

control is simultaneously proportional to weight estimate variability, motor output variability, and the LF itself. It would be interesting to determine which of these effects dominate.

2.5.2 Previous findings comparing the rates of grip force and manipulatory force adaptation

A key finding from Experiment 2.2 is that GF adaptation is not consistently faster than MF adaptation, as previous work has suggested(Flanagan et al., 2003). We find that the environment-specific mean-driven response that is present in both MF and GF control adapts at essentially identical rates for MF and GF. The apparent difference between the rates of overall GF and MF adaptation observed when only LF increases were examined has been taken as general evidence for a GF rate advantage and used to support the idea that a forward model for predicting the sensory consequences of environmental dynamics adapts faster than an inverse model(Flanagan et al., 2003). However, the current findings suggest this conclusion to be problematic. Instead, we show that faster adaptation for GF compared to MF is confined to instances in which the LF *increases*, with the opposite observed when LF *decreases*. The mechanism here is that variability-driven responses act to increase GF whenever the environment changes, regardless of whether that change would require smaller or larger GFs. Thus, this nonspecific variability-driven adaptive response *synergizes* with environment-specific mean-driven adaptation when LF increases, but *interferes* with it when LF decreases, leading to dramatic differences in adaptation rates. However, this dichotomy does not occur with MF adaptation, because, MF is insensitive to environmental variability. Our findings also provide explanations for previous results showing that GFs are paradoxically increased following initial exposure to microgravity(Augurelle et al., 2003) and that real-time feedback responses can lead to increased GF following sudden, unexpected increases or decreases in LF(Danion, 2007). These results mirror the current findings, and suggest that variability-driven changes in GF control govern real-time feedback responses as well as feedforward motor adaptation. Moreover, the finding that environment-specific MF and GF adaptation proceed at essentially identical rates suggests a unified mechanism for the

adaptation of internal models for the expected environmental dynamics in the control of manipulatory and grip forces.

2.5.3 Uncertainty vs Variability

Experiments 2.1, 2.1b, and 2.1c modulated environmental variability to increase the motor system's uncertainty about the environment. Variability and uncertainty, though related, are distinct qualities. Variability is easy to rigorously define (it can be measured as statistical variance: the second-order moment of the difference between the value of a random variable and its mean). However uncertainty is trickier: it is related to the unpredictable part of the variability. Thus, to estimate uncertainty, the properties of the predictor must be understood. This would allow uncertainty to be defined in an analogous way to variability, as the statistical variance of the difference between actual and predicted values of the environment rather than the variance of actual values alone. Our view is that trial-by-trial predictions of the load force occur and these predictions can be fairly well approximated by running estimates of the mean forces experienced over a window of previous movements. If this is the case, our conception of force variability estimation that takes into account a running estimate of the mean in the ACEVU model is, in spirit, an estimate of uncertainty rather than a true variability estimate, which would be based on a single fixed estimate of the mean. Further work may be required to fully dissociate uncertainty and variability, but the substantial transience of the nonspecific effects predicted by the ACEVU model critically depend on uncertainty estimation, suggesting that safety margin control is based on uncertainty estimates.

2.5.4 Uncertainty-driven control of grip forces and muscular co-contraction

Uncertainty-driven safety margins, like the ones we observed in the control of GF, may explain previous observations about the control of co-contraction of antagonist muscle pairs. Like GF control, antagonist muscle control consists of specific and nonspecific components: reciprocal activation (which drives the exerted MF and is *specific* to the expected environmental dynamics) and co-contraction (which has no effect on the net MF and is *nonspecific* to environmental dynamics)(Feldman, 1980; Franklin and

Wolpert, 2011). Co-contraction increases the mechanical stiffness of the limb in resisting unexpected perturbations, allowing for improved control in uncertain environments(Franklin et al., 2003; Franklin et al., 2004; Milner and Franklin, 2005; Franklin et al., 2008). While not providing a frank safety margin that *prevents* errors from unexpected perturbations, the increased ability to reduce errors stemming from unexpected perturbations that is afforded by co-contraction make it like a "soft" safety margin against unexpected perturbations, raising the possibility that co-contraction may be controlled in a fashion similar to the mechanism we describe for GF. Indeed, co-contraction increases immediately after detection of an unexpected load(Smeets et al., 1990) and more generally while moving in dynamic environments of increased trial-to-trial variability(Takahashi et al., 2001a), but only transiently following step changes in environmental dynamics(Milner and Cloutier, 1993; Thoroughman and Shadmehr, 1999; Franklin et al., 2003; Franklin et al., 2004). All of these effects are consistent with uncertainty-driven control.

2.5.5 Impact of neurologic disease on uncertainty-driven grip force control

Internal estimates of uncertainty for GF control may reflect variability in the environment itself, imprecise sensory information about the environment, or imprecise control of motor output. Thus the current results predict that increased sensory or motor uncertainty will also increase GF safety margins leading to consistently higher GF, because it should be the total uncertainty that determines the statistical confidence associated with the safety margin. This may explain why both healthy subjects under anesthesia(Westling and Johansson, 1984; Cole and Abbs, 1988; Nowak et al., 2001) and patients with a range of sensorimotor dysfunction exhibit elevated GF safety margins. In particular, elevated GFs have been reported in patients with multiple sclerosis(Iyengar et al., 2009), Parkinson's disease(Fellows et al., 1998; Wenzelburger et al., 2002), carpal tunnel syndrome(Zhang et al., 2011; Afifi et al., 2012), peripheral deafferentation (Nowak et al., 2003), and cerebral stroke(Hermsdorfer et al., 2003; Dubrowski et al., 2005; Anens et al., 2010).

2.5.6 Uncertainty estimation and metacognition

Estimates of uncertainty – or, inversely, confidence – about the state of the environment related to one's decisions or actions have traditionally been classified as metacognitive abilities – abilities which encompass "what we know about what we know"(Metcalfe and Shimamura, 1994; Nelson, 1996) (Hacker, 1998; Metcalfe, 2008). While such meta-cognitive abilities have been considered to be a very recent product of evolution exclusive to humans and, to a limited extent, primates(Metcalfe, 2008), recent work on decision-making has found a number of behavioral and neural correlates of decision confidence in rodents (e.g.(Foote and Crystal, 2007; Kepecs et al., 2008; Terrace and Son, 2009; Kepecs and Mainen, 2012)). This should not come as a surprise when considering motor control given the longstanding hypothesis that the motor system relies on unconscious access to internal estimates of variability and uncertainty for motor planning. In fact, it has been widely hypothesized that movements are specifically planned to minimize the task-relevant component of motor variability, possibly in combination with effort(Harris and Wolpert, 1998; Todorov and Jordan, 2002; Izawa et al., 2008; Izawa and Shadmehr, 2008; Braun et al., 2009). Moreover, the strength of online movement corrections has been shown to depend both on the certainty associated with sensory information and the certainty of one's prior beliefs about the environment in a manner consistent with Bayesian integration(Kording et al., 2004; Kording and Wolpert, 2004; Franklin et al., 2012). This integration often occurs without awareness as simple neural computations can provide the required variability estimates(Ma et al., 2006; Kepecs et al., 2008; Kiani and Shadlen, 2009; Kepecs and Mainen, 2012; De Martino et al., 2013). For example, variability estimates like the ones used to implement the ACEVU model can arise from a simple trial-by-trial update rule, whereby the current estimate of variability can be expressed as a weighted average of the previous variability estimate and a rectified version of current sensory information. In Methods, we show this is the case by deriving a simple algorithm for trial-by-trial updating of an internal estimate of environmental variability that requires no memory except for that of the previous variability estimate, yet exactly computes variability over an exponential memory window. Together, these findings suggest that uncertainty estimates need not be the exclusive domain of high-level cognitive endeavors; in fact, our

work illustrates that they are indeed an integral part of the machinery involved in performing basic motor actions that do not require explicit specification, like controlling grip forces. Moreover, the exquisite sensitivity of GFs to environmental variability identifies GF control as a valuable paradigm for probing the nervous system's estimates of uncertainty. Going forward, this establishes a new tool for better understanding the mechanisms underlying the estimation of uncertainty in the nervous system about which little is currently known.

2.6 References

ACI Committee 318 (2005) Building code requirements for structural concrete (ACI 318-05) and commentary (ACI 318R-05). Farmington Hills, Mich.: American Concrete Institute.

Afifi M, Santello M, Johnston JA (2012) Effects of carpal tunnel syndrome on adaptation of multi-digit forces to object texture. Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology 123:2281-2290.

Anens E, Kristensen B, Hager-Ross C (2010) Reactive grip force control in persons with cerebellar stroke: effects on ipsilateral and contralateral hand. Exp Brain Res 203:21-30.

Augurelle AS, Penta M, White O, Thonnard JL (2003) The effects of a change in gravity on the dynamics of prehension. Exp Brain Res 148:533-540.

Braun DA, Aertsen A, Wolpert DM, Mehring C (2009) Learning optimal adaptation strategies in unpredictable motor tasks. J Neurosci 29:6472-6478.

Burge J, Ernst MO, Banks MS (2008) The statistical determinants of adaptation rate in human reaching. Journal of vision 8:20 21-19.

Cole KJ, Abbs JH (1988) Grip force adjustments evoked by load force perturbations of a grasped object. J Neurophysiol 60:1513-1522.

Cole KJ, Johansson RS (1993) Friction at the digit-object interface scales the sensorimotor transformation for grip responses to pulling loads. Experimental brain research 95:523-532.

Danion F (2007) The contribution of non-digital afferent signals to grip force adjustments evoked by brisk unloading of the arm or the held object. Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology 118:146-154.

Danion F, Sarlegna FR (2007) Can the human brain predict the consequences of arm movement corrections when transporting an object? Hints from grip force adjustments. J Neurosci 27:12839-12843.

De Martino B, Fleming SM, Garrett N, Dolan RJ (2013) Confidence in value-based choice. Nat Neurosci 16:105-110.

Donchin O, Francis JT, Shadmehr R (2003) Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. J Neurosci 23:9032-9045.

Dubrowski A, Roy EA, Black SE, Carnahan H (2005) Unilateral basal ganglia damage causes contralesional force control deficits: a case study. Neuropsychologia 43:1379-1384.

Feldman AG (1980) Superposition of motor programs--I. Rhythmic forearm movements in man. Neuroscience 5:81-90.

Fellows SJ, Noth J, Schwarz M (1998) Precision grip and Parkinson's disease. Brain 121 (Pt 9):1771-1784.

Flanagan JR, Wing AM (1993) Modulation of grip force with load force during point-to-point arm movements. Exp Brain Res 95:131-143.

Flanagan JR, Wing AM (1997) The role of internal models in motion planning and control: evidence from grip force adjustments during movements of hand-held loads. J Neurosci 17:1519-1528.

Flanagan JR, Vetter P, Johansson RS, Wolpert DM (2003) Prediction precedes control in motor learning. Curr Biol 13:146-150.

Foote AL, Crystal JD (2007) Metacognition in the rat. Curr Biol 17:551-555.

Franklin DW, Wolpert DM (2011) Computational mechanisms of sensorimotor control. Neuron 72:425-442.

Franklin DW, So U, Kawato M, Milner TE (2004) Impedance control balances stability with metabolically costly muscle activation. J Neurophysiol 92:3097-3105.

Franklin DW, Osu R, Burdet E, Kawato M, Milner TE (2003) Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. J Neurophysiol 90:3270-3282.

Franklin DW, Burdet E, Tee KP, Osu R, Chew CM, Milner TE, Kawato M (2008) CNS learns stable, accurate, and efficient movements using a simple algorithm. J Neurosci 28:11165-11173.

Franklin S, Wolpert DM, Franklin DW (2012) Visuomotor feedback gains upregulate during the learning of novel dynamics. J Neurophysiol 108:467-478.

Fu Q, Zhang W, Santello M (2010) Anticipatory planning and control of grasp positions and forces for dexterous two-digit manipulation. J Neurosci 30:9117-9126.

Goldsmith M (1956) Pablo Casals talks of modern music. In: The News and Courier, p 9C. Charleston, S.C.

Gonzalez Castro LN, Monsen CB, Smith MA (2011) The binding of learning to action in motor adaptation. Plos Comput Biol 7:e1002052.

Gonzalez Castro LN, Hadjiosif AM, Hemphill MA, Smith MA (2014) Environmental consistency determines the rate of motor adaptation. Curr Biol 24:1050-1061.

Hacker DJ (1998) Metacognition: Definitions and empirical foundations. Metacognition in educational theory and practice:1-24.

Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. Nature 394:780-784.

Hermsdorfer J, Hagl E, Nowak DA, Marquardt C (2003) Grip force control during object manipulation in cerebral stroke. Clinical Neurophysiology 114:915-929.

Holway AH, Pratt CC (1936) The Weber-ratio for intensitive discrimination. Psychol Rev 43:322-340.

Iyengar V, Santos MJ, Ko M, Aruin AS (2009) Grip force control in individuals with multiple sclerosis. Neurorehabil Neural Repair 23:855-861.

Izawa J, Shadmehr R (2008) On-line processing of uncertain information in visuomotor control. J Neurosci 28:11360-11368.

Izawa J, Rane T, Donchin O, Shadmehr R (2008) Motor adaptation as a process of reoptimization. J Neurosci 28:2883-2891.

Johansson RS, Westling G (1984) Roles of glabrous skin receptors and sensorimotor memory in automatic control of precision grip when lifting rougher or more slippery objects. Experimental brain research 56:550-564.

Johansson RS, Westling G (1987) Signals in tactile afferents from the fingers eliciting adaptive motor responses during precision grip. Exp Brain Res 66:141-154.

Johansson RS, Westling G (1988) Programmed and triggered actions to rapid load changes during precision grip. Exp Brain Res 71:72-86.

Joiner WM, Smith MA (2008) Long-term retention explained by a model of short-term learning in the adaptive control of reaching. J Neurophysiol 100:2948-2955.

Joiner WM, Brayanov JB, Smith MA (2013) The training schedule affects the stability, not the magnitude, of the interlimb transfer of learned dynamics. Journal of Neurophysiology 110:984-998.

Joiner WM, Ajayi O, Sing GC, Smith MA (2011) Linear hypergeneralization of learned dynamics across movement speeds reveals anisotropic, gain-encoding primitives for motor adaptation. J Neurophysiol 105:45-59.

Jones KE, Hamilton AF, Wolpert DM (2002) Sources of signal-dependent noise during isometric force production. J Neurophysiol 88:1533-1544.

Kawato M (1999) Internal models for motor control and trajectory planning. Curr Opin Neurobiol 9:718-727.

Kepecs A, Mainen ZF (2012) A computational framework for the study of confidence in humans and animals. Philosophical transactions of the Royal Society of London Series B, Biological sciences 367:1322-1337.

Kepecs A, Uchida N, Zariwala HA, Mainen ZF (2008) Neural correlates, computation and behavioural impact of decision confidence. Nature 455:227-231.

Kiani R, Shadlen MN (2009) Representation of confidence associated with a decision by neurons in the parietal cortex. Science 324:759-764.

Kording KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. Nature 427:244-247.

Kording KP, Ku SP, Wolpert DM (2004) Bayesian integration in force estimation. J Neurophysiol 92:3161-3165.

Krakauer JW, Ghilardi MF, Ghez C (1999) Independent learning of internal models for kinematic and dynamic control of reaching. Nat Neurosci 2:1026-1031.

Landy MS, Trommershauser J, Daw ND (2012) Dynamic estimation of task-relevant variance in movement under risk. J Neurosci 32:12702-12711.

Lee JY, Schweighofer N (2009) Dual adaptation supports a parallel architecture of motor memory. J Neurosci 29:10396-10404.

Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. Nat Neurosci 9:1432-1438.

Marsaglia G (1965) Ratios of Normal Variables and Ratios of Sums of Uniform Variables. J Am Stat Assoc 60:193-204.

Mawase F, Karniel A (2012) Adaptation to sequence force perturbation during vertical and horizontal reaching movement-averaging the past or predicting the future? Front Syst Neurosci 6:60.

Metcalfe J (2008) Evolution of metacognition. Handbook of metamemory and memory:29-46.

Metcalfe J, Shimamura AP (1994) Metacognition : knowing about knowing. Cambridge, Mass.: MIT Press.

Milner TE, Cloutier C (1993) Compensation for mechanically unstable loading in voluntary wrist movement. Exp Brain Res 94:522-532.

Milner TE, Franklin DW (2005) Impedance control and internal model use during the initial stage of adaptation to novel dynamics in humans. J Physiol 567:651-664.

NASA (2008) Structural Design and Test Factors of Safety for Spaceflight Hardware. In. Washington, DC.

Nelson TO (1996) Consciousness and metacognition. Am Psychol 51:102-116.

Nowak DA, Glasauer S, Hermsdorfer J (2003) Grip force efficiency in long-term deprivation of somatosensory feedback. Neuroreport 14:1803-1807.

Nowak DA, Hermsdorfer J, Glasauer S, Philipp J, Meyer L, Mai N (2001) The effects of digital anaesthesia on predictive grip force adjustments during vertical movements of a grasped object. European Journal of Neuroscience 14:756-762.

Pekny SE, Criscimagna-Hemminger SE, Shadmehr R (2011) Protection and Expression of Human Motor Memories. Journal of Neuroscience 31:13829-13839.

Quaney BM, Rotella DL, Peterson C, Cole KJ (2003) Sensorimotor memory for fingertip forces: evidence for a task-independent motor memory. J Neurosci 23:1981-1986.

Saijo N, Gomi H (2012) Effect of visuomotor-map uncertainty on visuomotor adaptation. J Neurophysiol 107:1576-1585.

Scheidt RA, Dingwell JB, Mussa-Ivaldi FA (2001) Learning to move amid uncertainty. J Neurophysiol 86:971-985.

Scheidt RA, Reinkensmeyer DJ, Conditt MA, Rymer WZ, Mussa-Ivaldi FA (2000) Persistence of motor adaptation during constrained, multi-joint, arm movements. J Neurophysiol 84:853-862.

Shadmehr R, Mussaivaldi FA (1994) Adaptive Representation of Dynamics during Learning of a Motor Task. Journal of Neuroscience 14:3208-3224.

Sing GC, Smith MA (2010) Reduction in Learning Rates Associated with Anterograde Interference Results from Interactions between Different Timescales in Motor Adaptation. Plos Comput Biol 6.

Sing GC, Orozco SP, Smith MA (2013) Limb motion dictates how motor learning arises from arbitrary environmental dynamics. J Neurophysiol 109:2466-2482.

Sing GC, Joiner WM, Nanayakkara T, Brayanov JB, Smith MA (2009) Primitives for motor adaptation reflect correlated neural tuning to position and velocity. Neuron 64:575-589.

Smeets JBJ, Erkelens CJ, Vandergon JJD (1990) Adjustments of Fast Goal-Directed Movements in Response to an Unexpected Inertial Load. Experimental brain research 81:303-312.

Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. PLoS Biol 4:e179.

Takahashi CD, Scheidt RA, Reinkensmeyer DJ (2001a) Impedance control and internal model formation when reaching in a randomly varying dynamical environment. Journal of Neurophysiology 86:1047-1051.

Takahashi CD, Scheidt RA, Reinkensmeyer DJ (2001b) Impedance control and internal model formation when reaching in a randomly varying dynamical environment. J Neurophysiol 86:1047-1051.

Terrace HS, Son LK (2009) Comparative metacognition. Curr Opin Neurobiol 19:67-74.

Thoroughman KA, Shadmehr R (1999) Electromyographic correlates of learning an internal model of reaching movements. J Neurosci 19:8573-8588.

Thoroughman KA, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. Nature 407:742-747.

Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. Nat Neurosci 5:1226-1235.

Wasserman L (2004) All of statistics : a concise course in statistical inference. New York: Springer.

Weber EH (1978) E. H. Weber : the sense of touch. London ; New York: Academic Press for Experimental Psychology Society.

Wenzelburger R, Zhang BR, Poepping M, Schrader B, Muller D, Kopper F, Fietzek U, Mehdorn HM, Deuschl G, Krack P (2002) Dyskinesias and grip control in Parkinson's disease are normalized by chronic stimulation of the subthalamic nucleus. Ann Neurol 52:240-243.

Westling G, Johansson RS (1984) Factors influencing the force control during precision grip. Exp Brain Res 53:277-284.

Wolpert DM, Kawato M (1998) Multiple paired forward and inverse models for motor control. Neural Networks 11:1317-1329.

Wolpert DM, Miall RC, Kawato M (1998) Internal models in the cerebellum. Trends Cogn Sci 2:338-347.

Wu HG, Miyamoto YR, Gonzalez Castro LN, Olveczky BP, Smith MA (2014) Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. Nat Neurosci 17:312-321.

Yokoi A, Hirashima M, Nozaki D (2011) Gain Field Encoding of the Kinematics of Both Arms in the Internal Model Enables Flexible Bimanual Action. Journal of Neuroscience 31:17058-17068.

Yousif N, Diedrichsen J (2012) Structural learning in feedforward and feedback control. Journal of Neurophysiology 108:2373-2382.

Zhang W, Johnston JA, Ross MA, Smith AA, Coakley BJ, Gleason EA, Dueck AC, Santello M (2011) Effects of carpal tunnel syndrome on adaptation of multi-digit forces to object weight for whole-hand manipulation. PloS one 6:e27715.

Chapter 3: Safety margin control is based on a metric of environmental variability robust to outliers

3.1 Summary

In the previous chapter we found that, when environment dynamics vary randomly from one trial to the next, GFs scale with the standard deviation of these dynamics. Since dynamics were normally distributed, this also amounted to the maintenance of a set confidence against slip. Here, we compare GF responses in cases where dynamics are non-normally, asymmetrically distributed, and find that the motor system does not compute a frank confidence margin against slip, but instead computes a measure of environmental variability by which it scales the GF safety margin. We proceed to investigate how this variability measure might be computed, and find that it is likely based on a low-order norm or rectified deviations over a memory window, in a way that discounts outliers.

3.2 Introduction

In Chapter 2, we found that, under randomly-varying normally-distributed environmental dynamics, the safety margin for grip force control scales with the variability of these dynamics in a way that maintains a set confidence against slip. From our finding that grip forces scaled linearly with both the mean and the standard deviation of environmental dynamics, we can take the ratio between the GF/σ_{LF} sensitivity over the GF/μ_{LF} sensitivity in order to determine the statistical confidence associated with the observed safety margin. The idea here is that this sensitivity ratio expresses the safety margin in units of standard deviation, σ, beyond the mean, μ. Within the context of the normal distribution from which environmental dynamics were drawn in Experiment 2.1, this amounts to a set z-score and thus a fixed amount of confidence against slip – in Chapter 2, we found the ratio between the σ and μ sensitivities to be about 2.8, providing a >99% confidence against slip. What is the mechanism of this uncertainty-driven safety margin control? Here, we will examine two distinct possibilities: if the motor system indeed tries to maintain a set confidence against slip, or if it instead scales grip forces with a measure of environmental variability in a way that is equivalent to a set confidence against slip if the dynamics are normally distributed.

These two mechanisms would be served by qualitatively different forms of memory about uncertainty. The maintenance of a frank confidence margin $\alpha\%$ could be based on a memory of recent "worst-case" load forces, LF_a, which would scale GFs (Equation 3.1):

(Equation 3.1) $$GF = k \cdot LF_{max} \approx k \cdot LF_a$$

Alternatively, the maintenance of a variability estimate, $\hat{V}$, could be based on a memory of rectified deviations of dynamics over a recent time window, in a way similar to our analysis in Section 2.2.3, paragraph "Computational model of combined mean-driven and variability-driven adaptation". This model would amount to the maintenance of a safety margin scaled by $\hat{V}$ on top of the GF appropriate for the mean environment (Equation 3.2):

(Equation 3.2) $$GF = k_\mu \cdot LF_\mu + k_v \cdot \hat{V}$$

As mentioned above, for a Gaussian distribution of dynamics, these two approaches can't be distinguished, as a set confidence interval would map onto a specific scaling of the variability metric. For example, for $\alpha = 95\%$ and the metric $\hat{V}$ being an estimate of the standard deviation, σ, of environmental dynamics, we have (Equation 3.3):

(Equation 3.3) $$LF_{95} = LF_\mu + 1.645\, LF_\sigma$$

Thus, to resolve the question of which mechanism is used to estimate uncertainty we need to consider their predictions when different distributions are involved. For example, in Figure 3.1 we have two distributions of environmental dynamics that are not symmetrically distributed. Both have the same mean and standard deviation, so the "variability measure" hypothesis would predict equal grip force (GF) safety margins for these two cases. However, if the motor system were employing the "set confidence" mechanism, safety margins would be drastically different. A 95% confidence interval would amount to a GF safety margin of only 0.7σ beyond the mean for the top distribution, but it would amount to a safety margin more than twice as strong – 1.6σ – for the bottom distribution. Here we will compare GF responses to similar asymmetrical distributions of dynamics.

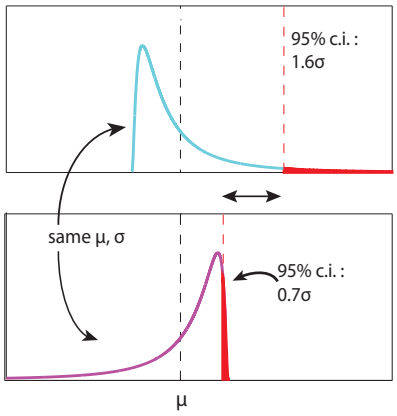


Figure 3.1: Example distributions of environmental dynamics.

The top (cyan) and bottom (magenta) oppositely-skewed distributions share the same mean, μ, and standard deviation, σ (in fact, any measure of their variability based on rectified deviations from the mean would be equal). Correspondingly, the hypothesis that grip force control is based on a measure of variability, predicts that both distributions of environmental dynamics result in the same GF safety margin. In contrast, the hypothesis that the GF safety margin is based on a direct estimate of the confidence against slip, predicts higher GFs for the positively-skewed distribution (cyan) compared to the negatively-skewed one (magenta). For example, maintaining a 95% confidence against slip for these distributions requires a safety margin equal to 1.6σ in the top distribution but only 0.7σ in the bottom distribution.

3.3 Materials and Methods

3.3.1 Participants and Ethics statement

A total of 30 healthy individuals without known neurological disease participated in this experiment. All subjects were naïve to the purpose of the experiments and provided informed consent. All experiment protocols were approved by the Harvard University Committee on the Use of Human Subjects in Research.

3.3.2 Task and analysis

Our experiment setup, the form of the dynamical environments and the statistical analyses we implemented were the same as what used for Experiment 2.1 (see Figure 2.1, section 2.3). After a practice period of 300 trials, we presented subjects 15 50-trial blocks with different variability conditions. Five different types of blocks (5x3) were presented: four types of variability blocks and baseline blocks. During each of the four types of variability blocks, we exposed participants to viscous force-field dynamics of strength (*b*). Critically, as in chapter 2, we applied a leftward static force offset of 4N (F_0 in Equation 2.1, repeated below as Equation 3.4) on top of the time-varying FF perturbation in order to ensure that the net direction of the force would always remain leftward even for negative values of *b* (see Section 2.3.2, Figure 2.2C).

(Equation 3.4)
$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} F_0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 & b \\ -b & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix}$$

In this experiment, *b* was varied from one trial to the next based on one of four different distributions, all with mean zero and standard deviation σ = 3.6 Ns/m. Specifically, depending on the type of block, *b* was either drawn from a Gaussian distribution, a balanced bimodal distribution, a bimodal distribution with a left, negative skew (9:1), and a bimodal distribution with a right, positive skew (1:9, see Figure 3.3AC). During the fifth type of blocks, baseline blocks, *b* was set to zero throughout. These five types of blocks were each presented 3 times in random order, but without back-to-back repetition of the same block type.

3.4 Results

To answer these questions, we designed an experiment (Experiment 3.1) whereby we tested the grip force (GF) responses of N = 30 subjects under different conditions of environmental variability. As in Chapter 2, subjects performed rapid point-to-point reaching movements while grasping a small object with physical dynamics that could change randomly from one trial to the next. They first experienced a practice period of 300 trials to familiarize themselves with the basic task, and were then exposed to 15 50-trial blocks with different variability conditions. Five different types of blocks (5x3) were presented: four types of variability blocks and baseline blocks. During each of the four types of variability blocks, the strength (*b*) of a viscous curl force-field environment (see Equation 3.1 in the methods) was varied from one trial to the next by drawing *b* from one of four different distributions, all with mean zero and standard deviation σ = 3.6 Ns/m. Specifically, depending on the type of block, *b* was either drawn from a Gaussian distribution, a balanced bimodal distribution, a bimodal distribution with a left, negative skew (9:1), and a bimodal distribution with a right, positive skew (1:9, see Figure 3.3AC). During the fifth type of blocks, baseline blocks, *b* was set to zero throughout. These five types of blocks were each presented 3 times in random order (without however allowing for the same type of block to be repeated back-to-back). Similarly to Chapter 2, in order to give subjects time to adapt to each type of block, we focused our analysis on the last 70% of the data from each block.

3.3.1 The control of the safety margin for grip force is based on sensitivity to a measure of environmental variability

We first compared GF responses during the two types of skewed dynamical environments in order to investigate whether GF control was programmed to maintain safety margin strictly based on confidence, regardless of the type of the distribution of environmental dynamics – for example, keeping a memory for the "worst case" (highest load) trials (in the sense of a soft maximum) – or, was instead programmed based on a fixed sensitivity to the variability of environmental dynamics – for example, a fixed $\Delta GF/\Delta\sigma$. These two cases would be exactly equivalent under normally-distributed environmental dynamics (e.g.

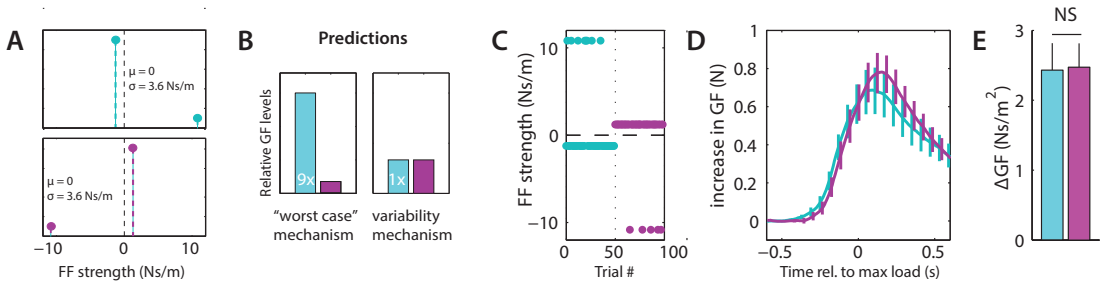


Figure 3.2: The motor system adjusts the safety margin for grip forces based on a measure of variability, instead of directly computing confidence against slip.

A: The two bimodal distributions of environmental dynamics used to elucidate between the "set confidence" and "variability measure" mechanisms of uncertainty estimation.

B: Predictions of the two mechanisms for the environmental dynamics in A. Under the "set confidence" mechanism, GF responses would be 9 times stronger for the positively-skewed distribution (cyan), since the strongest 10% of its values are 9 times stronger (referenced to the mean) compared to the strongest 10% of the values of the negatively-skewed distribution (magenta). Under the "variability measure" mechanism, GF responses would be the same for both distributions, since, like in the example shown in Figure 3.1, σ (and, any measure of their variability based on rectified deviations from the mean) is equal.

C: Example of the training schedule under each one of these two distributions. In Experiment 3.1, each distribution was presented in three 50-trial blocks in random order [repetition?] (see Figure 3.3 for a sample schedule of Experiment 3.1 complete with all types of variability blocks).

D: Average baseline-referenced grip force profiles for each type of block, illustrating very similar responses for both positively and negatively skewed distributions. Errorbars indicate SEM.

E: Grip force levels for each type of block, showing no difference between the two distributions, leading us to reject the hypothesis that the motor system directly computes confidence against slip, and instead conclude that the safety margin adjustments we observe are based on a measure of variability.

Experiment 2.1), as in a normal distribution the ratio of the sensitivity to σ over the sensitivity to μ can determine the confidence margin (see Section 2.4.4). However, this is not the case for any type of distribution. For the two oppositely skewed bimodal distributions in Experiment 3.1, for example, both of which have the same mean and standard deviation, the hypothesis that the GF safety margin is programmed based on a fixed sensitivity to variability would predict that GF responses to each of these two environments are equal (Figure 3.2B). In contrast, the idea that the GF safety margin confers the same statistical confidence against slip would predict that the responses to the right-skewed bimodal distribution would be 9 times greater above baseline compared to the left-skewed bimodal distribution.

The idea here is that, in order to maintain a confidence level above 90%, the motor system would have to increase GF levels based on the strongest 10% of possible environmental dynamics. However, for the right-skewed distribution, the strongest 10% of b values is nine times higher than the strongest 10% of b values for the left-skewed distribution. We found that responses to the two kinds of skewed distributions were remarkably similar (adaptation coefficients of 2.43±0.38 Ns/m vs. 2.47±0.34 Ns/m above baseline for the positively vs. negatively skewed distribution, p>0.9, Figure 3.2DE), indicating that the motor system scales the GF safety margin based on a fixed sensitivity to a measure of the variability of environmental dynamics.

3.3.2 Grip force control is based on a variability measure that is robust to outliers

We proceeded to investigate the details of the variability measure most closely explains the observed sensitivity of GF to variability, and, specifically, how sensitive this variability measure is to outliers. For this variability metric, we consider the p-norms (Horn and Johnson, 1991) of the deviation of the observed viscous load force, b, from its mean, $\bar{b}$ (Equation 3.5).

(Equation 3.5) $$I_p = \sqrt[p]{E|b - E(b)|^p}$$

Our question thus becomes, which is the value of n which best explains the sensitivity of GF to variability? For p = 1, I_p becomes equal to the average of absolute deviations from the mean, whereas, for p = 2, I_p becomes equal to the standard deviation, σ. Low values for p mean that the metric is more robust to outliers in the distribution of b, whereas higher values of p will give more and more weight to the strongest values of the absolute deviation $|b - E(b)|$, making the metric more sensitive to these outliers. At the limit $p \to \infty$, the metric will equal the maximum absolute deviation (Horn and Johnson, 1991):

(Equation 3.6) $$I_\infty = \max\{|b - E(b)|\}$$

Different values of p give correspondingly different predictions for the relative sizes of the GF safety margin for the four different distributions we examine here. For $p = 2$, we expect GF levels to be the same in all four cases: $I_2 = \sigma$, and σ is the same in all four distributions. For $p = 1$, GF levels will be the highest for the balanced bimodal distribution and the lowest for the skewed distributions. The reason is that, the low-order metric I_1 will discount outliers more than I_2, and thus reduce the variability estimate for the skewed distributions since for these distributions variability is mainly driven by the outliers. In contrast, for the balanced bimodal distribution – which in a qualitative sense is the least outlier-heavy relative to

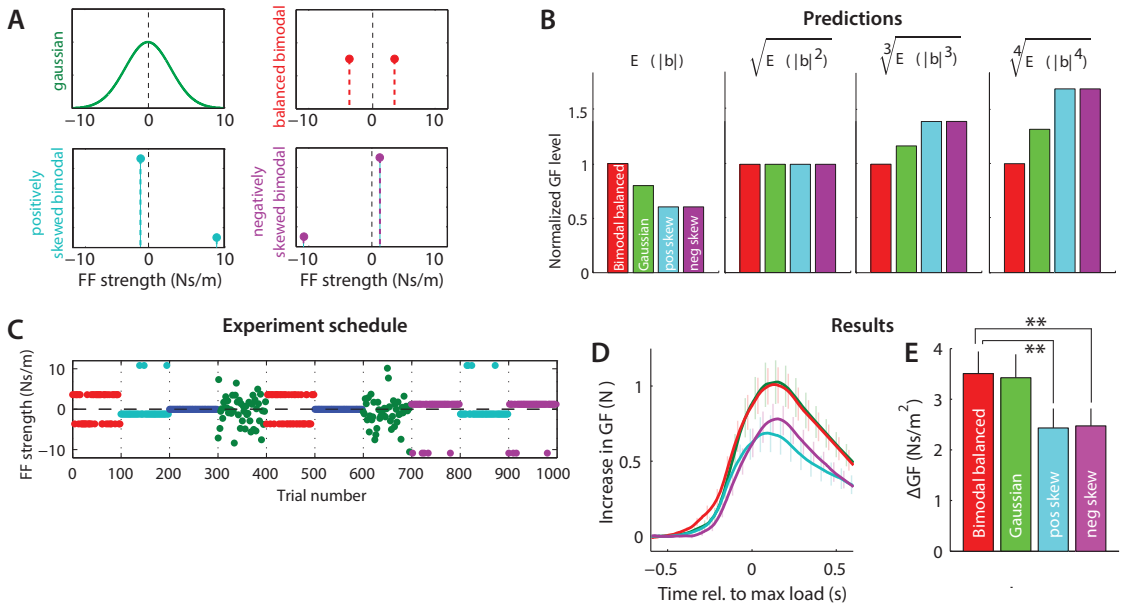


Figure 3.3: Variability estimation by the motor system relies on a low-order measure that is robust to outliers.

A: The four different distributions of dynamics presented in Experiment 3.1, all with the same μ and σ: a Gaussian distribution, like the ones used in Experiment 2.1 (green), a balanced bimodal distribution (fair coin toss, red) and the two skewed, asymmetric distributions described in Figure 3.2 (cyan and magenta).

B: Predictions for the relative size of the GF safety margin across the four different distributions of dynamics, based on different variability metrics. Here, p-norms (Equation 3.6) are used as variability metrics. For p = 2, the p-norm is equal to σ and, since all distributions have the same σ, all responses are equal. Higher values of p, which indicate higher sensitivity to outliers, increase responses to the skewed distributions over the balanced bimodal distribution and vice versa.

C: Sample training schedule for Experiment 3.1, showing the five types of blocks (four variability blocks corresponding to the four different distributions in A and a baseline block).

Figure 3.3 (continued).

D: Average baseline-referenced grip force profiles for each type of block, indicating that responses to the balanced bimodal and Gaussian distributions are higher than responses to the two skewed distributions. Errorbars indicate SEM.

E: Grip force levels for each type of block. Responses to the balanced bimodal distribution are higher than responses to both skewed distributions, in line with a low-order ($p<2$) variability metric that discounts outliers.

its σ – a low-order metric will still result in high variability estimates and thus higher GF levels compared to the other distributions. For values of p greater than 2, the variability metric will weigh outliers more and more, thus resulting in a higher GF levels for the outlier-heavy skewed distributions (see Figure 3.3B).

To investigate the behavior of variability estimates with respect to outliers in environmental dynamics, we compared GF responses across the four different variability conditions in Experiment 3.1. We found that these GF responses were the strongest for the balanced bimodal group than either of the two skewed distribution groups (baseline-referenced GF adaptation levels of 3.51±0.44 Ns/m for the balanced vs 2.43±0.38 Ns/m and 2.47±0.34 Ns/m for the positively and negatively skewed distributions respectively, both $p < 0.01$, two-tailed paired t-tests), suggesting that the GF variability metric that determines the safety margin tends to discount outliers. The responses for the Gaussian distribution were similar to that for the balanced bimodal (3.42±0.46 vs. 3.51±0.44 Ns/m, respectively, $p>0.87$, two-sided paired t-test). Correspondingly, fitting the variability metric models on our data revealed that the value of p that best explains the data is equal to 1.30.

The relationship of variability estimation in the motor system and higher order moments of the distribution of the states of the environment

A way to systematically characterize the presence of outliers in a distribution is by considering its high-order moments, like kurtosis. Kurtosis is the fourth standardized moment of a distribution, and is defined as the ratio between the fourth moment about the mean, and the square of the variance (DeCarlo, 1997):

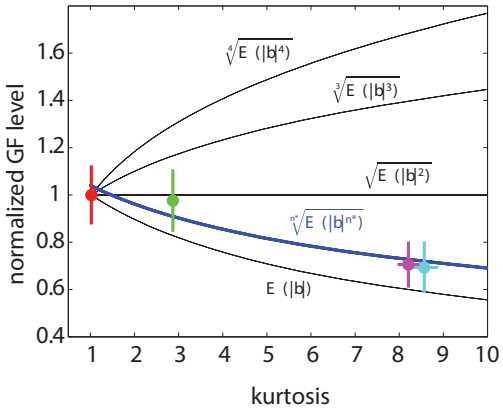


Figure 3.4: The variability measure used by the motor system discounts high-order moments.
Plotted here are grip force responses to each type of variability block, normalized by the average response to the balanced bimodal blocks, against the kurtosis of environmental dynamics in each type of block. The black curves indicate the predicted relationship between kurtosis and GF responses if the variability metric used to drive these responses were calculated based on each one of the first 4 p-norms, whereas the blue line indicates the responses based on the norm that optimally fits the data (p* = 1.30). The two skewed distributions, which have the highest kurtosis, also have the weakest GF responses, showing that the low-order metric used to calculate variability discounts higher-order moments.

(Equation 3.7) $$k_b = \frac{E[(b-\mu_b)^4]}{E[(b-\mu_b)^2]^2} = \frac{\mu_{b,4}}{\sigma_b^4}$$

The minimum value kurtosis can take is 1, which is also the kurtosis of the balanced bimodal distribution. The kurtosis of a Gaussian distribution is 3, whereas the kurtosis of a 9:1 skewed bimodal distribution, like the positively and negatively skewed ones in Experiment 3, is equal to $(9^2 + 9^{-1})/10 = 8.1111$. The relationship between kurtosis and the value of the variability metrics we considered for different values of p is shown in Figure 3.4. For $p < 2$, the variability metric decreases with kurtosis, in line with a metric robust to outliers, whereas for $p > 2$ it increases, in line with metric sensitive to outliers. As shown above, our data are in line with a variability metric of order < 2, and thus a metric that discounts high-order statistical moments like kurtosis, making it robust to outliers.

3.3.3 Size of the memory window for estimating variability

The sparseness and specificity of outliers during the positively and negatively skewed bimodal distribution blocks enables us to investigate the duration of the effect these outliers have on GF and thus allows us to assess the size of the memory window that is used for estimating variability. Specifically, we focused on data right before and right after a positive or a negative outlier for the negatively and positively skewed distributions, respectively. The ACEVU model (section 2.4.5) predicts that positive and negative outliers that have the same magnitude, will give rise to two types of responses: a mean-driven response, specific to the sign of the outlier and thus opposite between positive and negative outliers, and a variability-driven response, non-specific to the sign of the outlier, that will be the same (Figure 2.5A). Given the higher sensitivity to the variability of environmental dynamics compared to the mean, the ACEVU model predicts that the variability-driven response might overpower the mean-driven one resulting in an overall positive response in both cases (rightmost set of panels in Figure 2.5A).

We focused our analysis on the early part of the movement (50-200ms from onset) in order to avoid any effects of feedback-driven grip force corrections that might occur due to the strong outliers, and more closely focus on the predictive component of the response. Correspondingly, we regressed grip force profiles onto the shape of the variability-driven component (as in Figure 2.7), but isolated this regression to the part of the GF profile 50-200ms from movement onset. The results are plotted in Figure 2.5C. We found that both positive and negative outliers prompted an increase in this GF metric for the immediate trials, in spite of the fact that negative outliers considerably decreased the strength of the load force (early force levels of 0.54 0.20 vs. 1.68 0.25, right before vs. right after a positive outlier trial, respectively, 0.60 0.16 vs. 1.35 0.23, right before vs. right after a positive outlier trial, both $p < 10^{-4}$, 2-sided paired t-test). This is in line with the ACEVU model, and the finding in Experiment 2.2 whereby both sudden increases and decreases in the load led to increases in GF (Figure 2.5C). GF levels continued to decline with repeated training, returning to baseline in about 10 trials. Note that, in Figure 2.5C, it appears that GF levels drop below their pre-outlier baseline. This is because we only display trials for which no other

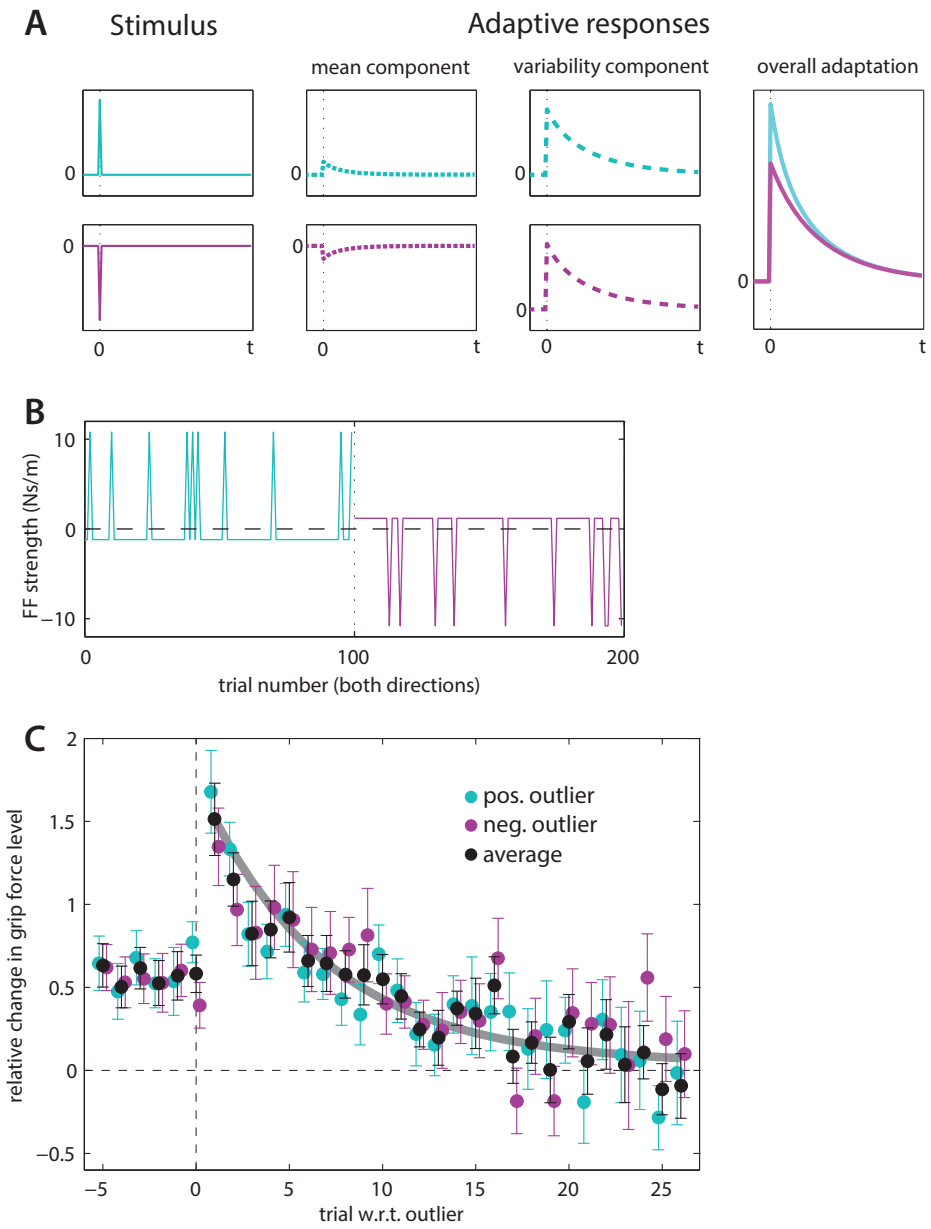


Figure 3.5: Size of the memory window for estimating variability.

A: Predictions of the ACEVU model on grip force responses to positive (cyan) and negative (pink) outliers. Left: example positive and negative stimuli with the same magnitude. These outliers lead to both a mean-driven response, specific to the sign of the outlier and thus opposite between positive and negative outliers (second column), and a variability-driven response, non-specific to the sign of the outlier, that will be the same (third column). Given the higher sensitivity to the variability of environmental dynamics compared to the mean, the ACEVU model predicts that the variability-driven response might overpower the mean-driven one resulting in an overall positive response in both cases (rightmost column).

Figure 3.5 (continued)

B: Example of how trial-to-trial environmental dynamics behave in the skewed bimodal blocks in Experiment 3.1.

C: Analysis of GF responses to determine the size of the memory window for variability estimation. Here, we focused on the two skewed-distribution variability blocks, whereby outliers in the strength of load dynamics were relatively rare (10% of trials) and had the same value. This enabled us to average together the grip force responses before and after these outliers in order to systematically investigate the duration of their effect. The GF adaptation level that is plotted here was calculated based on regression of the GF profile onto the shape of the variability-driven component (Figure 2.6), but focusing only on the early part of each movement (50-200 ms after movement onset) in order to avoid any effects of feedback responses to the strong outliers. We found that GF responses increased by a similar amount right after both a negative and a positive outlier, in line with the non-specificity of variability-driven responses described in Chapter 2. The elevated GF responses then gradually subsided over the course of the next trials with a time constant of 6.8 trials, which provides an estimate for the size of the time window used in estimating variability.

outlier has occurred yet; given that outliers were about 10% of the data, the average time between outliers was about 10 trials. Beyond that point, we are analyzing data for which outliers are more distant than average, allowing the adaptation level – and the variability estimate – to drop further.

We proceeded to estimate the speed of decay of the effects of outliers in GF control, which would give us an estimate for the size of the memory window within which environmental variability is calculated. Since both positive and negative outliers result in similar increases in GF levels, we averaged those two cases together and fit a single exponential to the data in order to determine the time constant of decay, which we found equal to 6.8 trials (95% confidence interval estimated using bootstrap: [3.5 13.4] trials).

3.5 Discussion

In order to elucidate the mechanisms by which uncertainty is estimated in the motor system, we investigated the adjustments of the grip force safety margin under variable dynamics that do not follow Gaussian, or even symmetric distributions. We first examined whether the motor system directly computes a safety margin against slip, or instead relies on a metric of environmental variability to scale this safety margin by comparing responses to distributions with the same variability metrics but 9-fold different confidence intervals against slip. Remarkably, we found similar responses in both cases, indicating that safety margin modulation relies on a variability metric. We proceeded to formulate this variability metric as a norm of deviations of environmental dynamics from the mean, and investigate the order of this norm which would show how sensitive this estimate is to outliers in dynamics. We found that the variability metric is based on a low-order norm which tends to discount outliers and high-order moments in the distribution of dynamics, providing robust estimates of variability.

Future work could further investigate the mechanisms of variability estimation in the motor system by probing grip force responses under different conditions. For example, would a greater degree of importance in avoiding slip – say, holding a more expensive vs. a cheap object – increase the safety margin, and if so, would that happen through an increase in sensitivity to variability, or, could variability estimation itself be altered in a manner more sensitive to outliers? Another interesting research direction would be to investigate the decay of the motor memories associated with variability estimates. We predict that the passage of time will lead to a transient increase in grip forces, in contrast with manipulatory forces where the passage of time leads to their decay back towards baseline (Sing et al., 2009) – also see Chapter 4. The reason behind the increase of grip forces after a time delay would be the heightened uncertainty about environmental dynamics that is driven by that delay. Finally, future work could investigate where this variability estimation for motor control is localized in the brain, and see if the areas involved are similar to areas usually associated with uncertainty representations in cognitive tasks, such as, for example, the orbitofrontal cortex (Critchley et al., 2001; Hsu et al., 2005).

3.6 References

Critchley HD, Mathias CJ, Dolan RJ (2001) Neural activity in the human brain relating to uncertainty and arousal during anticipation. Neuron 29:537-545.

DeCarlo LT (1997) On the meaning and use of kurtosis. Psychological methods 2:292.

Horn RA, Johnson CR (1991) Matrix analysis, [Reprinted with corr.] Edition. Cambridge Cambridgeshire ; New York: Cambridge University Press.

Hsu M, Bhatt M, Adolphs R, Tranel D, Camerer CF (2005) Neural systems responding to degrees of uncertainty in human decision-making. Science 310:1680-1683.

Sing GC, Najafi B, Adewuyi A, Smith MA (2009) A novel mechanism for the spacing effect: Competitive inhibition between adaptive processes can explain the increase in motor skill retention associated with prolonged inter-trial spacing. In: Advances in Computational Motor Control. Chicago, IL.

Chapter 4: Temporally-stable and temporally-labile memories of motor adaptation

4.1 Summary

Successful and efficient motor control relies on the formation and retrieval of motor memories associated with the skill at hand. Here, we assess the temporal stability of motor memories formed through training on a visuomotor rotation task. We identify two distinct components of adaptation based on the persistence of their underlying memories: a temporally-labile component that decays over a few seconds, and a temporally-stable component that remains constant over the course of at least 20 minutes. While the amount of motor adaptation that is temporally-stable is relatively consistent between different measurements within the same training session, it can vary widely both across different individuals and different training sessions for the same individual. We proceed to investigate longer-term motor memories by testing for retention one day after initial training. We find that the temporally-stable adaptation at the end of the training session can predict retention 24 hours later, in line with previous work on force-field adaptation. Using brain stimulation and a reanalysis of patient data from previous adaptation studies, we proceed to investigate where these temporally-labile and temporally-stable components of adaptation might be formed in the brain. Our results suggest that the primary motor and posterior parietal cortices might be involved in temporally-labile adaptation, whereas the cerebellum might be involved in temporally-stable adaptation. Furthermore, we investigate the directional generalization pattern of these motor memories, which may provide a fingerprint of how the underlying memory is encoded in the brain, and we find that the temporally-stable adaptation itself consists of two components: one that is more specific to the trained movement direction (locally-generalizing) and displays weak second-day retention, and one that transfers widely across the workspace (globally-generalizing) and displays a considerably more robust second-day retention. Finally, we find that globally-generalizing adaptation is robust to delays in sensory feedback during training, whereas locally-generalizing adaptation is not.

4.2 Introduction

μετὰ χρόνου πᾶσα μνήμη

any memory implies a time elapsed

Aristotle, On Memory and Reminiscence

Memory manifests itself as the ability to retrieve information after the passage of time. While we do not intend to dwell on Aristotle's statement above (Bloch and Aristotle, 2007), we note that the relationship between memory and time is rather one-sided: memories generally decay with the passage of time. For example, the reader will be less likely to recall Aristotle's words above accurately in half a minute as they are right now. The deterioration of declarative memories has been widely studied in the past 100 years (Luh, 1922; Peterson and Peterson, 1959; Rubin and Wenzel, 1996; Wixted, 2004), whereas the deterioration of motor memories has more recently been the subject of studies involving motor learning and adaptation (Shadmehr and Brashers-Krug, 1997; Krakauer et al., 2005; Richardson et al., 2006; Joiner and Smith, 2008; Sing et al., 2009).

Yet, there appears to be no clear consensus as to the mechanisms behind the decay of memories. Let's consider the question: now, after about half a minute into this introduction, how easy it is to recall Aristotle's aphorism? Is it harder because time has passed, because of all the information that might have interfered with the memory of it, or a combination of both? This question illustrates the difference between two distinct theories of memory decay: the "decay theory" posits that memories can indeed decay simply by the passage of time, whereas the "interference theory" rejects the possibility of pure passive, time-dependent decay, and instead posits that memories deteriorate only with active interference from other (Conrad, 1957; Brown, 1958; Altmann and Gray, 2002; Lewandowsky et al., 2009).

Motor memories, on the other hand, were found to decay both passively with time, and actively, through the removal or the reversal of the training stimulus (Scheidt et al., 2000; Davidson and Wolpert, 2004;

Criscimagna-Hemminger and Shadmehr, 2008; Kitago et al., 2013; Vaswani and Shadmehr, 2013). Recent work systematically examined the decay of force-field learning within a session and identified two distinct adaptation components: a labile component that decays rapidly with time, and a temporally-stable component that only decays with experience (Sing et al., 2009; Sing, 2011). However, whether this distinction is a general characteristic of motor learning or specific to the types of dynamical environments investigated in (Sing et al., 2009) remains unknown.

To assess long-term retention beyond a single session, previous studies usually relied on the presence of savings – the faster relearning of a previously learned stimulus – as an indicator of memory consolidation (Ebbinghaus et al., 1913; Krakauer and Shadmehr, 2006; Richardson et al., 2006; Landi et al., 2011). However, direct measurement of both temporally-stable adaptation, temporally-labile adaptation, and adaptation displaying long-term (1-day) retention revealed that temporally-stable adaptation can predict, and thus act as a gateway to, long-term retention (Joiner and Smith, 2008; Sing et al., 2009). In contrast, long-term retention could not be predicted by either overall or temporally-labile or overall adaptation.

Here, we investigate the temporal stability of motor memories by isolating the temporally-labile vs. temporally-stable components of motor learning in a visuomotor adaptation task. We first study the relative contribution of each component to overall adaptation, measure how quickly temporally-labile adaptation decays, and investigate differences in temporal stability of adaptation both across individuals and also for different sessions for the same individual. We proceed to investigate the brain areas associated with labile and stable adaptation, using transcranial magnetic stimulation and a reanalysis of previous adaptation data from patients. We then trace the shape of the memories underlying temporally-stable adaptation by probing the extent to which they generalize across different movement directions, in turn allowing us to investigate how each part of this shape relates to longer-term retention. Finally, we examine the extent to which the formation of these temporally-stable motor memories relies on the close temporal correspondence between motion and sensory feedback.

4.3 Materials and Methods

A total of 121 subjects participated in the studies described in this chapter (16 in Experiment 4.1, 46 in Experiment 4.2, 25 in Experiment 4.3, and 54 in Experiment 4.4). All subjects were naïve to the purpose of the experiments and provided informed consent. Each subject took part in only one of the above experiments. The protocols for Experiments 4.1, 4.3 and 4.4 were approved by the Harvard University Committee on the Use of Human Subjects in Research. The protocol for Experiment 4.2, which involved transcranial magnetic stimulation (TMS), was approved by the Institutional Review Board in Beth Israel Deaconess Medical Center, where the study was conducted.

4.3.1 Apparatus

The experimental setup we used was the same as in recent previous work (Brayanov et al., 2012; Brayanov, 2013), illustrated in Figure 4.1. Subjects sat against a table on which we placed a digitizing tablet (Wacom Intuos 3). Using full grasp, they held a handle containing a digitizing pen the position of which could be tracked by the tablet, and used this handle to make movements on the tablet. We minimized friction between the handle and the tablet by covering the bottom of the handle with a thin sheet of Teflon. Their position was sampled at 200Hz, with a resolution of 0.005 mm and an accuracy of 0.25 mm and processed using custom routines in Matlab which utilized the Psychophysics Toolbox (Brainard, 1997). About 8 inches above the tablet (around the level of the subjects' shoulder), we horizontally mounted a 120 Hz 23-inch LCD monitor which provided visual cues on subjects' movement and targets, at the same time preventing direct vision of their arm and hand.

4.3.2 General task

Subjects were instructed to make fast, 9-cm point-to-point reaching movements on the tablet as prompted by visual cues on the screen. Targets were displayed against a black background and indicated by colored circles 1cm in diameter, whereas the subject's hand position was represented by a 2.5 mm diameter cursor. Movements were rapid: subjects were usually trained to strive for a movement duration of less

than 250ms or 300ms (see specific procedures), based on audiovisual feedback about movement speed following each trial.

During baseline blocks, the cursor feedback displayed followed hand motion. During training, however, cursor feedback was rotated by 30 degrees with respect to the hand motion, based on a visuomotor rotation (VMR) paradigm (Krakauer et al., 2000; Krakauer et al., 2005; Hadipour-Niktarash et al., 2007; Huang et al., 2011; Brayanov et al., 2012; Shmuelof et al., 2012; Kitago et al., 2013) – see Figure 4.1. In all experiments, the sign of this 30° VMR was balanced across subjects. Moreover, the experiments also contained trials in which the visual feedback of the cursor position was withheld altogether (referred as NVF trials throughout the chapter).

4.3.3 Data analysis

To assess the amount of adaptation to the trained VMR, we measured the direction of hand motion on each trial. In movements where the cursor was visible, this was calculated as the angle between hand and target direction 150ms after movement onset, whereas in movements without visual feedback (e.g. in the generalization blocks in Experiments 4.3 and 4.4, see Specific Procedures), we calculated this angle at the end of the movement.

Statistical Comparisons

We used paired t-tests for statistical comparisons between adaptation levels across conditions within groups, and unpaired t-tests for comparisons between adaptation levels and sensitivity estimates across groups. We also used 1-sided ANOVAS to investigate effects across multiple groups, and F-tests to assess the significance of linear fits in the data. We used bootstrapping for analyses where we expected the distribution of the data to be highly non-normal or where large amounts of noise impeded reliable estimation of the mean values of model parameters, Specifically, we used bootstrapping for ratio estimation, as the ratio between random variables is known to produce a distribution which can be far from normal, even when the random variables themselves are normally distributed (Marsaglia, 1965), as for example in section 4.4.5. We also used bootstrapping for estimating the time constants of decay using

exponential fits (section 4.4.1). We chose this method instead of performing a t-test between parameters derived from fits to individual subject data, because the high noise in these individual data leads to low confidence about the corresponding individual parameters. The statistical tools that were used in each case are outlined in detail in the results section along with the corresponding analysis.

4.3.4 Specific procedures

Experiment 4.1

We recruited 16 subjects (10 males, age = 22.9±6.9 years, mean ± std, 2 left-handed) each of whom participated in 3 or 4 experiment sessions, usually on consecutive days. The first session started with a 100-trial familiarization block to introduce subjects to the task, followed by a 200-trial baseline period. During both familiarization and baseline, veridical visual feedback was provided. After baseline, subjects were trained on a 30 degree VMR (clockwise or counter-clockwise) for 200 trials and then entered the testing phase of the experiment. In the testing phase, subjects were made to wait for a time delay Δt, had their adaptation remaining after this Δt measured in a no-visual-feedback (NVF) trial, and were then repeatedly retrained and retested for a total of 8 different amounts of Δt: 0, 3, 6, 10, 20, 30, 60 and 120 seconds. Within each measurement block, these delays were presented once in random order. There were a total of 8-12 such measurement blocks in each session. After this testing phase, subjects were presented with 200 trials with veridical visual feedback in order to wash out the trained VMR. This washout period concluded the first session (See Figure 4.2 for a detailed diagram of the experimental protocol). Subsequent sessions were identical to the first, with the differences that (i) there was no familiarization period (ii) the baseline period consisted of 100, instead of 200 trials, and (iii) the sign of the VMR was reversed compared to the previous session. We reversed the sign of the VMR between consecutive sessions in order to minimize the chance that there was any transfer of adaptation from the previous day, which could affect our results.

Experiment 4.2

We recruited 46 subjects split into four different groups based on the brain area that underwent rTMS (16 in the occipital pole group; 9 in the pre-frontal cortex, Area 46, group; 13 in the posterior parietal cortex, PPC, group; and 8 in the primary motor cortex, M1, group). These experiments were conducted in collaboration with the Berenson-Allen Center for Non-invasive Brain Stimulation at the Beth Israel Deaconess Medical Center.

Subjects were trained on a 30 degree visuomotor rotation (clockwise or counter-clockwise, balanced across subjects) for 180 trials. About 20 trials into the training period, we started interspersing three kinds of events among the training trials: (i) no-visual-feedback (NVF) trials, (ii) 3.5s delays, and (iii) 30s delays, with about 18% of the trials during the training block being event trials (11 of each kind out of 180 trials). During the time delays the subject remained still on the starting position, waiting for the target.

TMS procedures

After the baseline, and right before the onset of the training period, we performed 1 Hz repetitive transcranial magnetic stimulation (rTMS) on the subjects, on one of the four brain areas: the primary motor cortex (M1), the inferior parietal lobule (IPL) of the posterior parietal cortex (PPC), the area 46 in the prefrontal cortex (PFC) and the occipital pole. To extend the effects of rTMS throughout the training period, we chose to perform stimulation for 20 minutes, which was upper limit of our estimate of how long the training period would take. Our choice was in line with previous work which found that the effects of 15 minutes of rTMS lasted for at least 15 minutes (Chen et al., 1997). We note that the training period duration eventually ranged from about 14 to 20 minutes. However, because of the expected wearing off of the rTMS effects, we focused our analysis on the first 120 trials in the data (first 2/3 of the training period, which included 7 events of each kind).

To determine the stimulation strength for each subject, we used single-pulse TMS on the M1 in order to find his or her motor threshold – the minimum stimulation strength that is capable of eliciting a muscle

response (here, a finger twitch on the contralateral hand, while the hand was resting on its back). The stimulation strength was then set equal to 90% percent of this motor threshold.

We determined the loci of stimulation for the IPL and PFC groups using computer-aided stereotactical guidance based on an MRI brain scan of each subject . For the M1 group, we determined the locus of stimulation as the point on which single-pulse stimulation gave the stronger responses (during the process of finding the motor threshold). Finally, for the occipital group, the stimulating coil was placed directly on the occipital pole.

Experiment 4.3

We recruited N = 25 subjects for this study (8 males, age 23.9±4.7, 6 left-handed). Subjects first experienced a short familiarization period, during which they made reaching movements to one of 19 different targets on the workspace, visiting each target 6 times in random order. Targets were arranged 9cm away from the center starting position, 135 degrees to either side of the trained movement direction (0° , at 12 o' clock), in increments of 15 degrees (Figure 4.14B). In line with the standard sign convention for angles, clock-wise directions are defined as negative whereas counter-clockwise directions are defined as positive. During familiarization, continuous visual feedback of the hand position was presented through a cursor on the screen.

In the next phase, we measured the subjects' baseline generalization pattern using no-visual-feedback (NVF) trials (6 per target) interspersed with trials displaying full visual feedback (12 per target, total length of this baseline period: 342 trials, with two intervening rest breaks, see Figure 4.14AB). Subjects then experienced 120 training trials on the 0° direction with a 30° VMR (balanced across subjects between clock-wise and counter-clockwise). After a 1-minute break, we then measured the generalization of the trained VMR to all 19 movement directions. Each movement direction was presented three times in random order, in a block where all movements were without visual feedback, save for a small area … mm around the center position. After this testing block, subjects experienced a 60-trial retraining block, and then the session was concluded with a second testing block that had the same structure as the first one.

Subjects then returned after about ~24 hours and were immediately retested for generalization in two back-to-back testing blocks (Figure 4.14A).

Experiment 4.4

We recruited 54 subjects (18 males, age 22.1±5.1, 5 left-handed) split in 5 different groups, with each group experiencing a different latency between hand motion and cursor visual feedback: 25ms (the minimum possible under our setup), 50ms, 87ms, 300ms and 600ms. After a baseline period, during which subjects experienced the visual latency but not any visuomotor rotation, subjects were trained on a 30 degree VMR for 120 trials (balanced across subjects within each group, Figure 4.17A). During the training periods subjects also experienced 60 second waits every 30 trials in order to assess the temporally-stable component of adaptation (see Figure 4.17B).

Measurement of baseline latency between hand and cursor motion

Before running the experiment, we used a Hero 3 (GoPro, Inc.) camera to take high-speed (240 fps) videos of both hand and cursor while the hand was moving in an oscillatory fashion on the tablet. We analyzed these videos to extract the two signals of hand motion of cursor motion, and then performed a cross-correlation analysis to estimate the delay between them, which corresponded to the latency of the system. We measured this latency for different adjustments of the experiment setup, both with respect to hardware (for example changes in the input device, computer monitor, or the video card) and the software (such as video drivers or adjustments of the Matlab Psychophysics toolbox through which the experiment was ran). Optimizing for these parameters, we achieved an overall hand-to-cursor latency of 24.6±0.4 ms (mean±SD across 8 different measurements).

Modification of reward scheme

In this experiment, like in other experiments in this chapter, subjects were rewarded for reaching the target quickly, with a high-pitch sound after the completion of the movement. Specifically, they were rewarded if the movement took less than 300ms in Experiment 4.1, and less than 250ms in Experiments 4.2 and 4.3. For this experiment, however, we expected that the high visual latencies would make

reaching the center of the target considerably harder, and a fixed time limit for reward had the danger of making reward rare enough to not be relevant. We thus modified our reward scheme to maintain a reasonable likelihood of reward, using a dynamic reward scheme which adjusted the maximum movement time that would result in reward, t_{min}, in each trial, based on the subjects' success rate within recent movements:

(Equation 4.1) $$t_{min} = 0.4 - 0.2 \times (fraction\ of\ the\ last\ 40\ trials\ rewarded) \quad [s]$$

4.4 Results

4.4.1 Visuomotor adaptation consists of distinct temporally-stable and temporally-labile components

We designed an experiment to measure the temporal stability of motor learning, whereby 16 participants were trained to asymptote on a 30 degree visuomotor rotation (VMR), were then tested for retention after a variable delay Δt that could range from 0 to 120 seconds, and then repeatedly retrained and retested (Figure 4.1, 4.2). Specifically, after a baseline period of 200 trials, participants were trained on a 30° VMR for another 200 trials by the end of which they had reached asymptote adaptation (27.9±0.3° for the last 40 trials, mean±SEM, Figure 4.1). They then entered the testing period, which consisted of 8-12 blocks, each one testing for retention after one of 8 different delays (Δt): 0, 3, 6, 10, 20, 30, 60 and 120s. After the end of the testing period, subjects experienced 200 washout trials (where the VMR perturbation was removed and veridical visual feedback was provided instead) which concluded the session (Figure 4.2). Subjects returned for a total of 3-4 similar sessions, each on a different day (usually in consecutive days), and with the sign of the trained VMR switching between consecutive sessions in order to minimize the chance that transfer of adaptation or savings could occur from the previous session.

We found a clear decay of adaptation with the passage of time: from an initial value of 28.6±0.2° for Δt = 0, adaptation drops down to 24.5±0.8° in Δt = 30 seconds. However, there is little further decay at Δt = 60 or 120s (23.9±0.8° and 23.6±1.0°, respectively, with $p = 0.28$ between these two datapoints) indicating that the time-dependent decay of VMR is incomplete, with the adaptation remaining after 60 or 120 seconds showing temporal stability. Correspondingly, the overall decay curve illustrates two distinct components of adaptation. The first component displays rapid decay with time, which is well fit (R^2 = 99.4%) with an exponential with a time constant τ of only 16.5±2.2s (mean±SEM estimated using bootstrap, see Materials and Methods). This component accounts for only 17% of the initial adaptation; the remaining 83 % was instead stable with time. This result exposes a key dichotomy in visuomotor adaptation between temporally-labile and temporally-stable adaptation components.

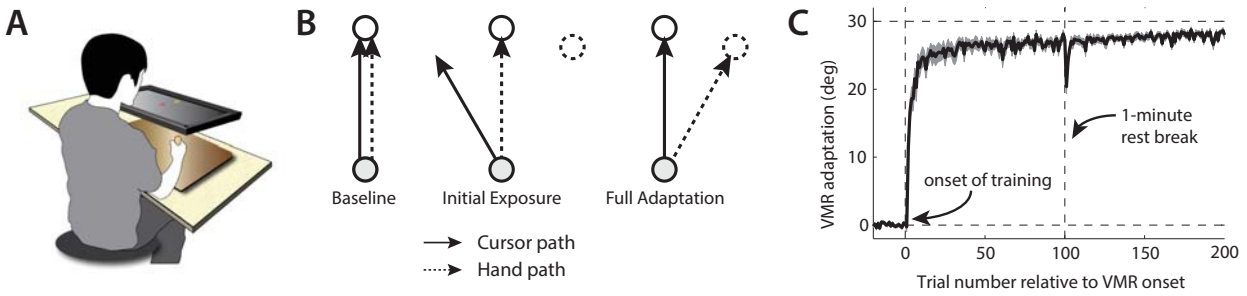


Figure 4.1: Learning a visuomotor rotation task.

A: Experiment setup. Subjects made point-to-point reaching movements on a tablet and observed visual feedback on a screen mounted above.

B: Visuomotor rotation paradigm. During baseline (left), the cursor follows the hand motion. During training, the cursor trajectory is displaced by 30° about the hand movement direction (in this example, counter-clockwise), resulting in a 30° error for the initial exposure (middle). Full adaptation will have the hand move in a direction 30° clockwise about the target, so that the cursor moves directly towards the target (right).

C: Learning curve from Experiment 4.1. To estimate the amount of adaptation at each trial, we measured movement direction 150ms after onset. A 30° movement direction would indicate full adaptation (data from both clockwise and counter-clockwise are included, flipped appropriately). Adaptation is fairly rapid, approaching asymptotic levels 50 trials into training. A short rest break (~1 min) in the middle of the training session resulted in a partial decay of adaptation, followed by rapid relearning. Shading indicates SEM.

The dichotomy between temporally-labile and temporally-labile components of adaptation is not unique to visuomotor adaptation. Previous work found similar timecourses of decay for a different motor task, that of force-field (FF) adaptation (Sing et al., 2009; Sing, 2011). When tested for retention after a similar gamut of different Δt's, subjects exhibited rapid (τ = 19.7s) but also incomplete (25% of initial adaptation) decay, showing distinct temporally-labile and temporally-stable components that amount for similar fractions of zero-delay adaptation as the corresponding ones our study of VMR decay (Figure 4.2). Importantly, we note that the (Sing et al., 2009) study included a decay measurement after a 20 minute Δt, which showed nearly identical decay as the data after the 2 minute Δt (18% vs. 20% decay for the 20 vs. 2 minute Δts, p=0.65), showing that the stability of this component extends at least in the order of tens of minutes. The time course of decay for the temporally-labile component of FF adaptation was found to be similar to the one observed for VMR adaptation (τ = 19.7 vs. 16.5s, respectively). Moreover,

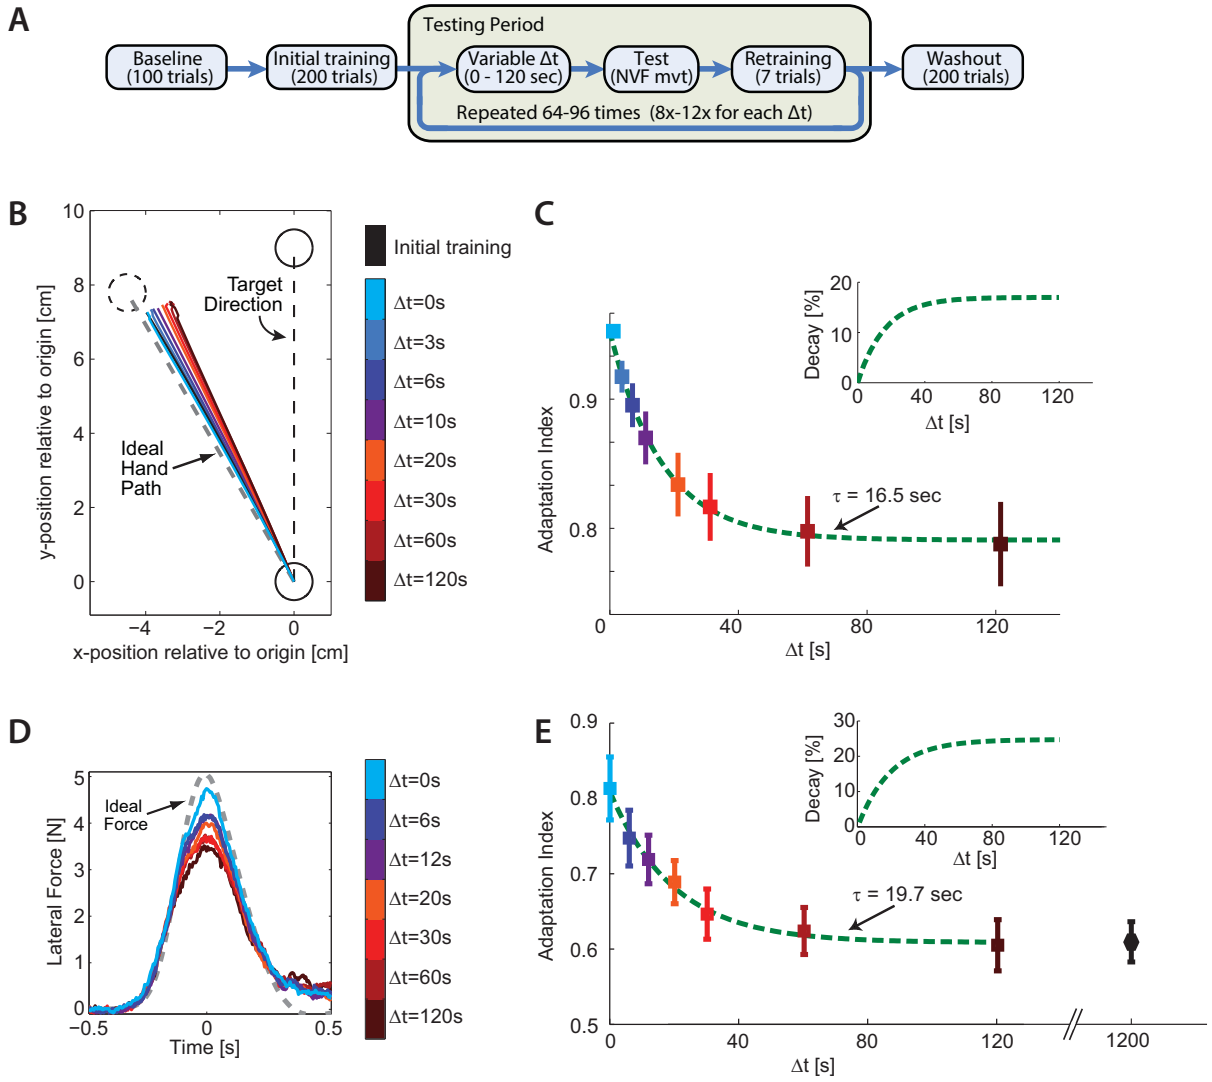


Figure 4.2: Temporally-labile and temporally-stable components of motor adaptation.

A: Training schedule for Experiment 4.1. In each session, a baseline period was followed by 200 training trials, after which subjects entered the testing period. During testing, subjects waited for a variable Δt ranging for 0-120s, were then tested for retention of adaptation, and then repeatedly retrained and retested. Each testing block consisted of probing retention for 8 different values of Δt in random order, and each session consisted of 8-12 such blocks. After the testing period, subjects experienced a 200-trial washout period which concluded the session. Subjects returned for a total of 3-4 sessions.

B: Decay of visuomotor adaptation. After initial training, subjects' average handpath (black curve) is very close to ideal (gray dashed line). As Δt increases, handpaths deviate more and more from ideal, indicating decay. However, there is little change in deviation after 30-60s, suggesting that time-dependent decay levels off and is thus partial.

Figure 4.2 (continued)

C: Decay curve for visuomotor adaptation. The adaptation index is based on the movement direction (with 1 amounting full adaptation, 30°). Decay is rapid (fit with a time constant τ of 16.5s) and incomplete: only about 20% of overall adaptation is temporally-labile (inset), suggesting that the remainder is temporally-stable.

D: Decay of force-field adaptation. Figure modified from (Sing et al., 2009). In this force-field adaptation task, subjects are tasked to counter novel environmental dynamics by producing an ideal force pattern (gray dashed curve). Increased Δt results in weaker and weaker force compensation, showing decay. However, there is little change in force production after 60s, suggesting partial decay in a manner analogous to the VMR decay pattern in B.

E: Decay curve for force-field adaptation. The adaptation index is based on regression of experienced forces onto the ideal ones (with 1 amounting to perfect match and full adaptation). Decay, as in C, is rapid (time constant τ of 19.7s) and incomplete: only about 25% of adaptation is temporally-labile. In fact, force adaptation measured 20 minutes after training (black datapoint) is indistinguishable from the one measured after 2 minutes, suggesting that the remaining temporally-stable adaptation is robust to further passage of time. Together, C and E suggest that the decomposition of motor adaptation in distinct temporally-stable and temporally-labile components might be a general characteristic of motor adaptation.

this rather small difference might be explained by slight differences in experiment design: whereas our study took 8-12 measurements of the full decay curve, the (Sing et al., 2009) study only used 6. Analysis of our measurement-to-measurement data (see section 4.2.1 below) suggests that the time constant of decay, τ, slightly decreases as the number of repeated measurements increases (Figure 4.3). To more directly compare the decay time constant we obtain here vs. the time constant for FF adaptation we isolated our analysis to the first 6 measurements, which revealed a $\tau = 19.5 \pm 3.1$s (mean±SEM estimated using bootstrap), even closer to the 19.7s observed for FF (Figure 4.3I). Together, our analyses identify similar temporally-labile and temporally-stable components for different types motor adaptation, suggesting a common architecture for motor adaptation.

Evolution of temporally-labile and temporally-stable components with training

We then investigated the effects of repeated training on both the amount of temporally-stable adaptation and the speed of decay of temporally-labile adaptation. To see how these components evolve during a session, we averaged together decay curves, from all sessions, that had the same measurement block number (in total, each session had 8-12 measurement blocks). Comparing these decay curves for earlier

vs. later measurement blocks shows that adaptation decays faster for later measurements compared to earlier ones (Figure 4.3A, block numbers increasing from blue to red). This increase in the speed of decay can be more clearly seen when we account for the different levels of asymptote adaptation (measured at the $\Delta t=0$ trials) slowly increased during a session ($p<0.0001$, see Figure 4.3D) by expressing decay as the difference between retention and the asymptote adaptation for that particular measurement block (Figure 4.3B).

To systematically examine changes in both the speed and overall amount of decay, we characterized decay curves using exponential fits that contained two free parameters: the depth of the decay, D, and the time constant of decay, τ (Equation 4.2)

(Equation 4.2) $$retained\ adaptation\ after\ a\ delay\ \Delta t: \qquad r(t) = r_0 - De^{-\frac{\Delta t}{\tau}}$$

In Equation 4.2, r_0 refers to the adaptation level when the delay Δt applied was zero (indicating asymptote adaptation). Thus, when $\Delta t = 0$, the level of the temporally-labile adaptation is D, and the level of temporally-stable adaptation is $r_0 - D$.

Because estimates of exponential fits are generally susceptible to noise in the data, we used a bootstrapping analysis to estimate τ and D for each block number (see Materials and Methods). Our analysis revealed that, in line with the increase in decay speed shown in Figure 4.3BC, the time constant τ decreased with time, from about 25 seconds to 10-20 seconds (Figure 4.3E). In contrast, changes in the overall amount of decay, D, did not show systematic trends, with the possible exception of a stronger decay during the first measurement compared to the rest (Figure 4.3F). In line with the higher decay during this first measurement, the decay of adaptation during the rest break between the two initial training blocks which preceded the testing (Figure 4.1C) was also high at 7.0±0.9°, suggesting that decay (and thus the temporally-labile component of adaptation) might indeed be stronger for the first 100-200 trials of adaptation but then tends to level off at a somewhat lower level.

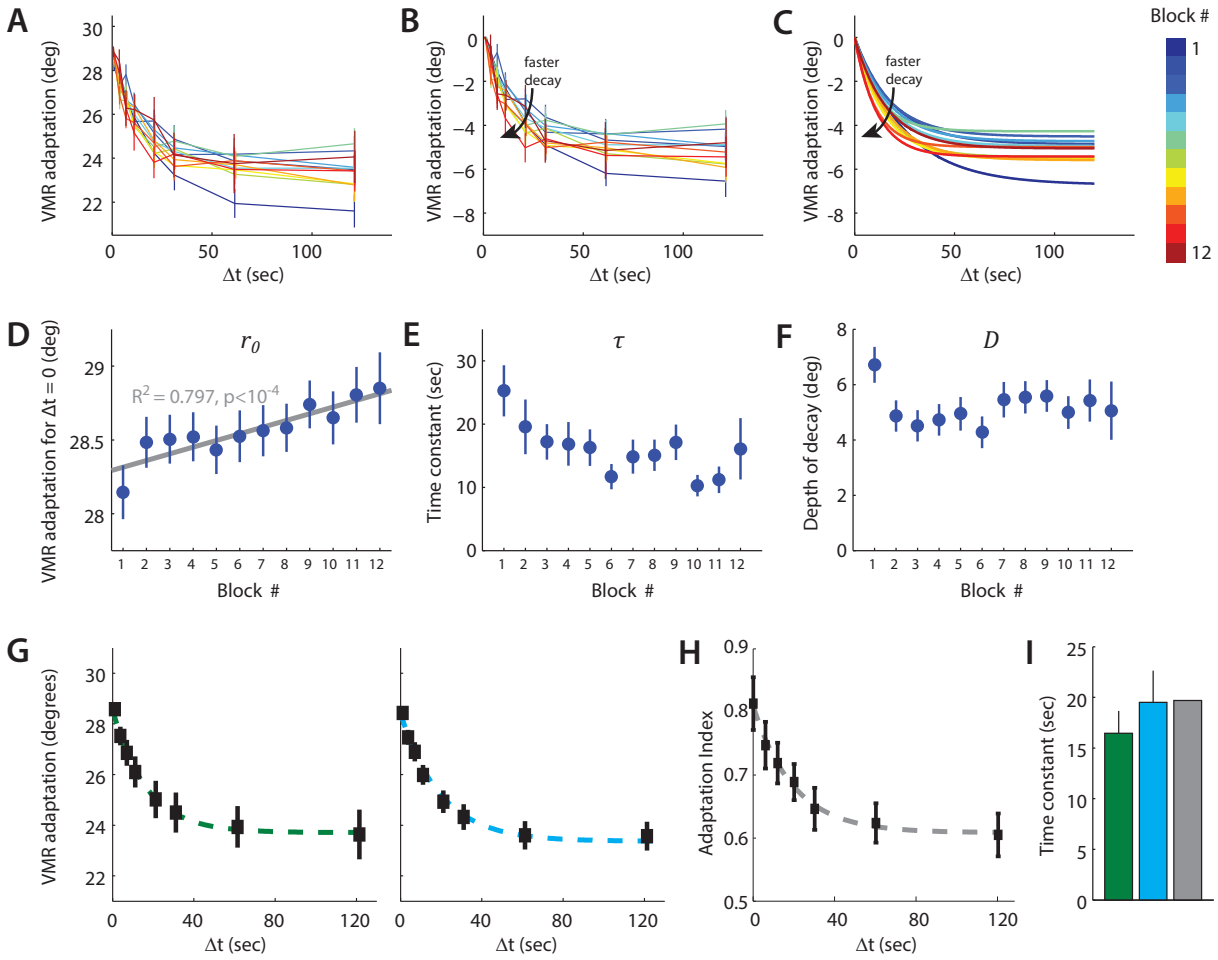


Figure 4.3: Effects of the length of testing on decay measurements.

A: Decay curves grouped block number (1-12, going for dark blue to dark red, see legend on the right of panel C). Errorbars indicate SEM.

B: Same as (A), but curves are referenced to the adaptation level at Δt = 0, correcting for the fact that the Δt = 0 level increased with block number as shown in panel D. Whereas decay curves are similar in depth (with possibly the block 1 curve being slightly, but not clearly, deeper), they appear to decay faster as testing progresses to higher block numbers.

C: Exponential fits of the data in B, indicating the increase in the speed of decay.

D: Evolution of the asymptote adaptation level at Δt = 0 (*r0* in Equation 4.2) with block number, showing a slight but significant increase. Errorbars indicate SEM.

E: Evolution of the time constant, τ, of decay curves with block number, showing a decrease in τ which amounts to an increase in the speed of decay as testing progresses. The time constants were estimated using exponential fits through a bootstrapping procedure, which also gave estimates of the SEM (errorbars).

Figure 4.3 (continued)

F: Evolution of the depth of decay, *D*, of decay curves with block number, showing little modulation apart from a slightly deeper decay on the first block. Errorbars indicate mean±SEM estimated using the same bootstrapping procedure as in E.

G: Left: subject-averaged decay curve and exponential fit (same as Figure 4.2C), which uses all 8-12 blocks that were available for each subject. Right: same as left, but using only the first 6 blocks in order to better compare with the force-field data in panel H / Figure 4.2E, which only had 6 testing blocks. In line with panel E, decay is slightly slower. Dashed lines indicate exponential fits, and errorbars indicate SEM.

H: Subject-averaged decay curve for force-field adaptation (same as Figure 4.2E) and exponential fit.

I: Time constants of decay for each case, showing the similarity of the speed of decay across the two paradigms (VMR and FF adaptation). The two left bars indicate mean±SEM estimated using bootstrap, whereas the gray bar indicates the time constant for FF adaptation decay as reported in (Sing et al., 2009).

4.4.2 Intra- vs. inter- individual differences in time-driven decay: moods and personalities in the decay of motor adaptation

In the previous section we focused our analysis on data that were averaged both across subjects and across different sessions, which allowed us to reliably trace the timecourse of decay. However, behind the subject-averaged data shown in Figure 4.2, also lies a considerable amount of variability across, first, the 16 subjects, second, the 3-4 sessions each subject participated in (for a total of 60 sessions), and third, the 8-12 different measurements in each session. For example, Figure 4.4A shows decay curves from individual sessions, against the subject-averaged data. Which are the sources of this large amount of variability in decay?

Here, we systematically compare the magnitude of differences in retention both between individuals and between different sessions for a given individual against the magnitude measurement-to-measurement differences within each session. We first examine differences on how deep this decay is (i.e. differences in the amount of adaptation that was temporally-labile, D in equation 4.2), and then proceed to examine differences in the speed of decay (i.e. differences in the time constant τ of this temporally-labile adaptation).

Differences in the amount of decay

The session averaged data shown in Figure 4.2 displayed considerable variability, with some sessions showing little to no decay but others displaying more than 50% decay. To systematically examine these differences, we characterized the decay curve on each measurement (each block which cycled through all 8 possible Δts once) by a single scalar D indicating the depth of the decay through fitting the data on Equation 4.2 with the time constant τ fixed at the subject average (16.5s) in order to avoid the effects of overfitting and focus on differences on the amount of decay, D.

Some subjects displayed consistently different amounts of decay than others. For example, for the two subjects shown in Figure 4.4B, all sessions from the high-decay subject (thin blue lines) show considerably higher decay than any of the sessions of the high-decay subject (thin red lines). We systematically compare the amount of decay across all subjects in Figure 4.4C, where we find a stunningly wide gamut of decay behavior with 10-fold differences in some cases: whereas some subjects showing little or no decay, others subjects show average decay close to 10 degrees. In Figure 4.4C, each colored datapoint indicates the average decay for each subject, with the colored errorbars indicating the standard error across different sessions. Note that the size of most of these within-subject, across-session errorbars is much smaller than the range of inter-individual differences, illustrating that a large fraction of session-to-session differences can be explained by this subject-specific behavior. Correspondingly, a majority of the session-to-session variability (56%) was due to differences between individuals (1-way ANOVA, $F(15,44)=3.71$, $p<0.001$). In contrast, only 3% of the session-to-session variability was due to session order ($p>0.4$) as shown in Figure 4.4D. We refer to the systematic differences in motor performance that we observe across individuals as “motor personalities”.

Moreover, our data revealed a large variability in decay across different sessions by the same subject. Figure 4.4C illustrates data from two sessions (black and gray) of an example subject (the subject in blue from Figure 4.4B) which are consistently different across different measurements which make up each session. Data from each one of these measurement blocks (i.e. a single measurement for each different

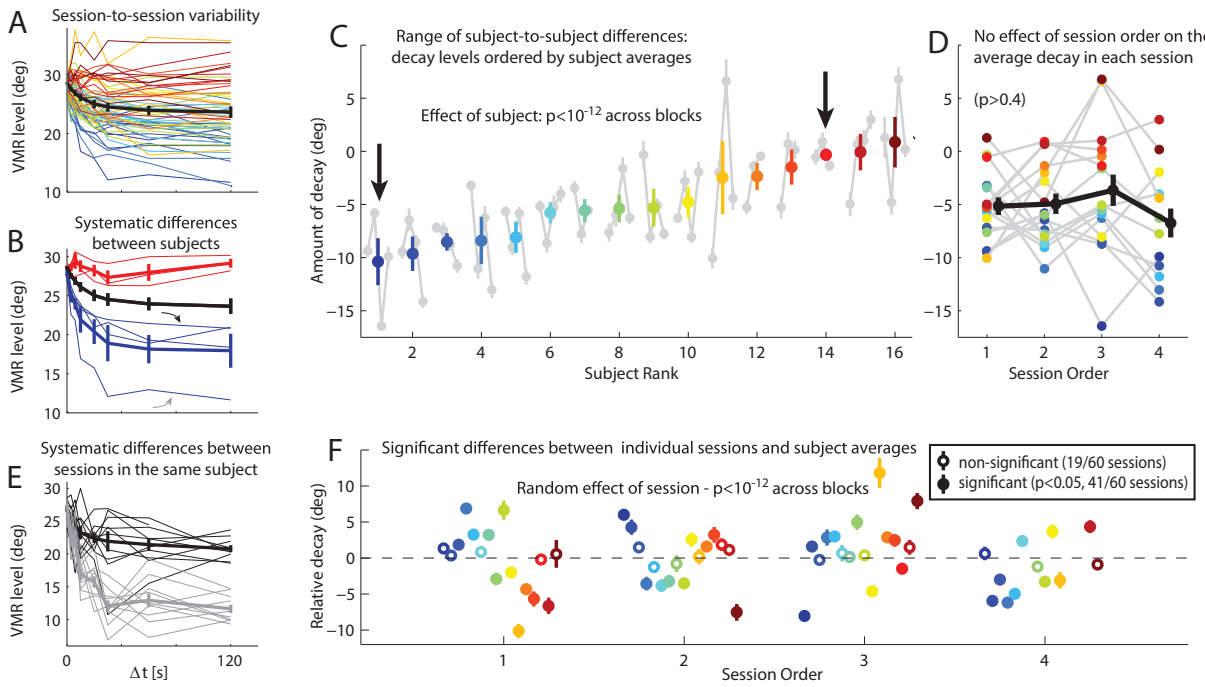


Figure 4.4: The structure of variability in decay.

A: The thick black curve shows the average decay curve for all subjects (as in Figure 4.2C), whereas the thin curves (colored by subject as in C) are the average decay curves for each session (60 in total), illustrating the large spread between sessions.

B: Black line shows subject-averaged decay data as in Figure 4.2C. The thick colored lines indicate mean behavior for 2 individual subjects (identified by the arrows in panel c). The decay curves for each session, the thin colored lines, show that the differences between these subjects are consistent.

C: Session averages grouped by subject (gray) with their corresponding subject averages (color). Subjects are ranked by their average decay level, to highlight the wide range of decay differences across individuals. The two arrows point out the individuals displayed in B.

D: Session averages ordered by session number. The thick line indicates the average decay for each session #, showing no effect of session order. The colors correspond to the data in C.

E: Systematic differences between the decay curves from different sessions in the same subject. The thick lines show session averages for two sessions from the blue subject in b and the thin lines show decay curves from all blocks in these two sessions.

F: Differences between each session's decay level and the mean decay from the remaining sessions from the same subject, illustrating the random effect of session in session-to-session variability. Errorbars indicate SEM across blocks, filled circles indicate the sessions where a significant ($p<0.05$) deviation from the mean of the other sessions was observed.

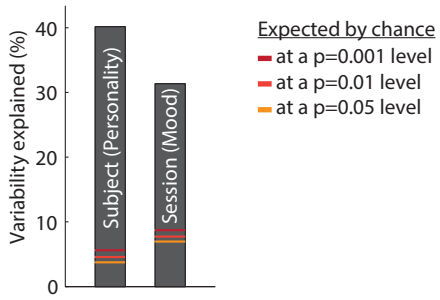


Figure 4.5: Block-to-block decay variability explained by motor personalities and motor moods.
The gray bars show the percentage of the block-to-block decay variability explained by differences between subjects (effect of "motor personality") and differences between different sessions of each subject (effect of "motor mood"). The colored marks indicate the percentages that can be expected by chance (given a confidence level), showing that the effects observed are far greater and highly significant.

Δt), represented by the thin lines, show clear separation between the two sessions. These consistent within-subject, across-sessions differences were not isolated to a few subjects, but were, instead, a characteristic of the dataset: 41 out of the 60 sessions we studied displayed significantly different ($p<0.05$) decay levels than the other sessions from the same subject (Figure 4.4F). We refer to the systematic differences in motor performance that we observe across different sessions for the same individual as "motor moods".

To compare the relative importance of these "moods" vs. "personalities" in the decay of motor adaptation, we performed a nested ANOVA which revealed that, whereas person-to-person differences account for 40% of the variability in single-block decay levels ($F(15, 647) = 29.0$, $p<10^{-12}$), an additional 31% is explained by session-to-session differences within individuals ($F(59, 603) = 11.24$, $p<10^{-12}$). The fraction of block-to-block variability explained by either of these effects is far greater than what could be expected by chance, as summarized in Figure 3.y. Remarkably, our data reveal that "motor moods" are almost as important as "motor personalities" (31% vs. 40%) in explaining differences in decay, and consequently the differences in the composition of the acquired motor memories with respect to temporally-labile and temporally-stable components.

Factors that might account for the different motor "moods" and "personalities" in the decay of motor adaptation

The breadth of inter-subject and intra-subject differences invited us to consider the factors that might determine them. We first investigated whether the different "motor personalities" we observed could be explained by demographic information such as age, gender, height or weight. However, neither of these factors was able to explain the inter-subject variance (no factor was significant, all $p>0.4$; % of inter-subject variance explained: age: 0.3%, gender: 7.0%, height: 3.3%, weight: 3.3%, Figure 4.6). We proceeded to investigate the factors that could explain the different "motor moods", testing for properties of the task such as the sign of the rotation, the number of days passed since the last session, the session order, or circumstances of the task like the time of the day. We only found a small and marginal effect of the VMR sign (explaining 7.0% of the variance associated with the inter-session-to-session "mood" effect, 0.8±0.4 degrees more decay for adaptation to the clock-wise force-field, $p = 0.042$), with other effects, such as the time of day, days since last session, and session order were not significant (all $p>0.27$, Figure 4.6). Thus, further work will be needed to understand the cause of the motor mood and motor personality differences that we demonstrate here.

Block-to-block trends

Finally, we investigated whether there were any trends between one measurement block to the next, and whether these within-session trends could to some extend predict decay from one session to the next. Linear fits within revealed that block-to-block decays within a session showed significant ($p<0.05$) trends in 19/60 sessions (14/44 sessions for which next-session data were available, Figure 4.12). For 9 out of these 14 sessions, the slope sign predicted whether the average decay for the next session would be higher or lower (Figure 4.12). This observation –though rather weak – suggests the possibility that slow trends within a session might, in some cases, extend to the next. This prompted us to further investigate by directly examining within-subject, across-session trends. Specifically, we examined linear fits across all measurement blocks and across all sessions from each subject, and found significant ($p<0.01$) trends in 9

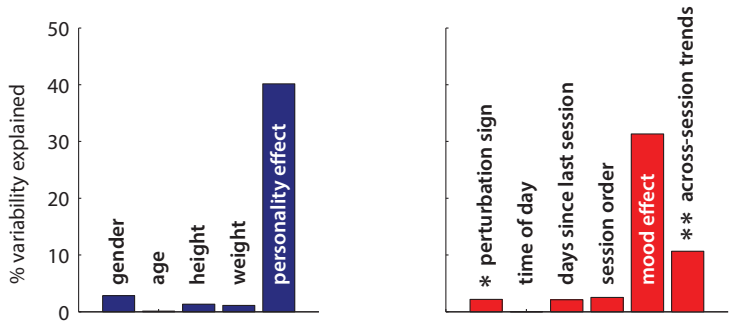


Figure 4.6: Factors investigated to explain the different motor personalities and motor moods.

Left: factors that we tested to see if they account for the different motor personalities – inter-individual differences in decay, showing no significant effects. Right: factors that we tested for the different motor moods – inter-session differences in decay within the same individual. Beyond a marginal effect of rotation sign, we found no significant effects. However, a considerable fraction of variability could be explained by trends in measurement to measurement decay across consecutive sessions (rightmost bar).

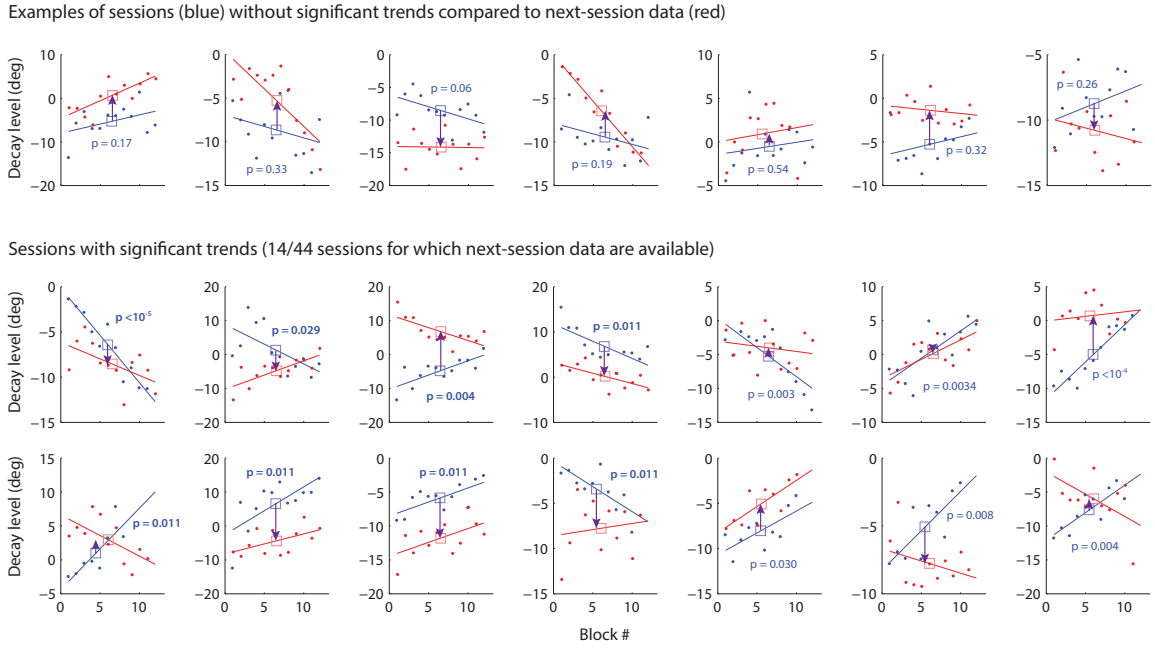


Figure 4.7: Systematic trends in session-to-session decay measurements.

Plotted are examples of within-session trends for consecutive sessions (earlier session in blue, next session in red). Lines indicate best fits, whereas the squares indicate mean decay for each session. Top row: examples of cases where there was no significant trend in the earlier session; Bottom two rows: the cases where there was a significant trend in the earlier session, where in 9 cases the slope of the earlier session predicted whether the mean decay in the next session would be higher or lower.

out of 16 subjects. These across-session fits explained 10.66% of the overall variance in the data ($F(632,31) = 24.76$, $p<10^{-12}$), far more than any of the session effects that we examined (Figure 4.11).

4.4.3 The neural circuits behind the temporally-labile process of adaptation

To investigate the neural substrates underlying both temporally-stable and temporally-labile processes, we first need to consider which brain areas might be involved with motor learning in general. The primary motor cortex (M1) and the posterior parietal cortex (PPC) were found to be involved with motor memories. Recordings from the M1 in macaque monkeys found cells that changed their tuning towards the direction of the compensatory force during a perturbation task (Arce et al., 2010). In human studies, it was found that anodal TDCS - which increases excitability (Filmer et al., 2014) - applied on the M1 can enhance the acquisition of motor adaptation and its retention between successive training sessions (Galea and Celnik, 2009; Reis et al., 2009) but not long-term retention (Reis et al., 2009). Disruption of the M1 using transcranial magnetic stimulation (TMS) sped up the decay of motor adaptation (Hadipour-Niktarash et al., 2007) whereas rTMS-induced inhibition of the M1 disrupted retention of motor skill improvements across consecutive training sessions (Muellbacher et al., 2002) and the second-day retention of motor adaptation (Richardson et al., 2006) – see (Reis et al., 2008) for an extensive review of the effect of M1 TMS on motor learning. Disruption of the PPC (specifically, the contralateral inferior parietal lobule, IPL) reduced the asymptotic level of motor adaptation to a viscous force-field (i.e. similar to the one described in Chapter 2), whereas disruption of the occipital pole did not (Della-Maggiore et al., 2004). Moreover, neurons in the macaque PPC displayed activity compatible with the encoding of the current motion state with little or no delay, in line with the output of a forward model for motion state (Mulliken et al., 2008).

On the other hand, areas like the prefrontal cortex, are rather involved with high-order functions such as monitoring reward, directing attention, goal selection or goal-directed motor planning rather than motor adaptation (Ridderinkhof et al., 2004). Finally, the occipital pole might present a likely control site

because the visual cortex, which lies around this stimulation site, is unlikely to be directly involved in motor adaptation.

To investigate the brain areas that might be involved in temporally-labile vs. temporally-stable adaptation, we designed an experiment where we compared the decay of temporally-labile adaptation when the activity of different brain areas is disrupted. We recruited 46 subjects which, before being trained on a 30° VMR, underwent a 20-minute session of repetitive 1Hz TMS (rTMS), which known to transiently inhibit brain excitability (Chen et al., 1997; Maeda et al., 2000; Mottaghy et al., 2000; Valero-Cabre et al., 2007). Subjects were split into four groups, each one of which experienced TMS in one of four different brain areas: the primary motor cortex (M1, N=8), the inferior parietal lobule (IPL) in the posterior parietal cortex (PPC, N=13), the pre-frontal cortex (Area 46, N=9), and the occipital pole (N=16) as a control site (hence the larger number of subjects).

Measurement of the decay curve

About 20 trials into the training session, we begun to repeatedly measure decay of VMR adaptation either after a short ($\Delta t = 3.5$ s) or a long ($\Delta t = 30$ s) delay. These time delays were selected in order to allow us to probe the speed and depth of the decay, based on our results from Experiment 4.1 which found a time constant of $\tau = 16.5$s. A guiding concept was that a potential rTMS-induced impairment specific to the temporally-labile process could decrease the time constant τ of decay, resulting in much greater decay after a short Δt when compared to control data. In contrast, the remaining adaptation measured after the long Δt (and thus the amount that decayed over the Δt, assuming similar adaptation levels before the delay) would be mainly based on the temporally-stable component, and thus would not be affected by an impairment in the temporally-labile process (see Figure 4.8). In short, impairment in the temporally-labile component would manifest as faster, but not deeper decay, whereas impairment in the temporally-stable component would manifest as deeper, but not faster decay.

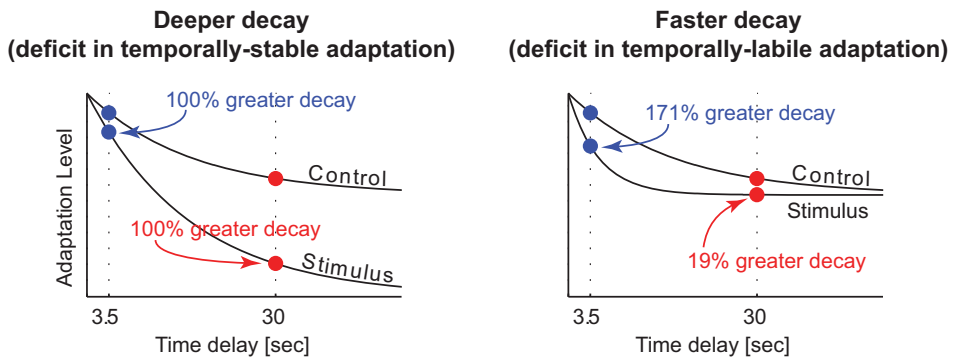


Figure 4.8: Measurement of deeper vs. faster decay of adaptation.

An increase in the depth of decay for the stimulation group compared to controls (left), which would suggest a weakening of the temporally-stable adaptation component, would increase both early (3.5s) and late (30s) decay by the same fraction. An increase in the speed of decay (right), which would suggest a weakening of the temporally-labile adaptation component, would strongly increase early decay but result in weaker relative increases in late decay.

Passive vs. active decay during a regular inter-trial interval

Previous data from a VMR task that were collected in our lab suggested that, when visual feedback is withheld during a NVF trial, the subsequent trial displays systematically lower adaptation. This is in line with the time-dependent decay: the absence of visual feedback during NVF movements prevents additional training, effectively increasing the time between training trials by one inter-trial interval (about 3.5s, equal to the short Δt). However, the fact that the motor system is actively engaged in motion during that time, might increase decay compared to the case of passively waiting for the same amount of time to pass. In order to elucidate the potential effect of additional active decay over single trial intervals, we also interspersed trials without visual feedback (NVF trials) throughout the training block, and compared decay after both NVF and $\Delta t = 3.5$s trials. We found similar decay after NVF trials vs. 3.5s trials for the three testing groups (M1: 3.2±0.7° vs. 2.6±0.7°, IPL: 2.0±0.5° vs. 2.3±0.7°, Area 46: 1.1±0.8° vs. 0.9±0.8°, all $p>0.3$, 2-tailed paired t-tests), suggesting that any effect of active decay due to the additional movement was minimal. The occipital pole group showed some difference though that was marginal (2.5±0.5° vs. 1.2±0.4°, $p = 0.06$).

The M1 and PPC might be involved in temporally-labile adaptation

We found that asymptote adaptation was similar across the four groups (M1: 26.8±0.8°, IPL: 26.0±0.3°, Area 46: 26.8±0.4°, Occipital pole: 25.9±0.5°, all $p>0.11$ between groups, Figure 4.9A-D). Correspondingly, an ANOVA across different groups did not reveal any significant effect ($p = 0.45$). In addition, temporally-stable adaptation, as measured after the longer time delay $\Delta t = 30s$ was also similar across these groups (M1: 22.3±1.3°, IPL: 21.4±1.1°, Area 46: 22.7±0.7°, Occipital pole: 21.9±0.6°, all $p>0.38$ between groups, ANOVA: $p = 0.80$ for effect of group, Figure 4.9E). As temporally-stable adaptation in either non-control group (M1, IPL, Area 46) did not differ compared to controls (Occipital), it is unlikely that either one of the other three brain areas is involved in temporally-stable adaptation.

However, there were differences in the decay after a short time delay ($\Delta t = 3.5s$), whereby both the M1 and IPL groups decayed about twice as much as controls (2.6±0.7° and 2.3±0.7°, vs. 1.2±0.4°, respectively, $p = 0.034$ for M1 vs. controls and $p = 0.068$ for IPL vs. controls, 1-tailed unpaired t-test, Figure 4.9E). In contrast, the 3.5s decay for the Area 46 group was similar to controls (0.9±0.8° vs. 1.2±0.4°, $p = 0.72$, 2-sided unpaired t-test). The faster, but not deeper, decay in the M1 and IPL groups is in line with a disruption of the temporally-labile process (Figure 4.8), suggesting that a network involving these two brain areas is responsible for temporally-labile motor adaptation.

4.4.4 The neural circuits behind the temporally-stable component of adaptation

Experiment 4.2 did not yield any evidence that either the M1, the IPL, or the prefrontal cortex are involved in the temporally-stable component of motor adaptation, suggesting that temporally-stable adaptation might primarily involve different circuits. Here we hypothesize that temporally-stable adaptation could involve the cerebellum, a brain area widely suggested to be a location where the internal representations of the state of the environment are formed and stored during motor adaptation (Wolpert et al., 1998; Imamizu et al., 2000; Imamizu et al., 2003; Kawato et al., 2003; Ito, 2008; Cullen and Brooks, 2014). The theory behind these internal models, posits that the motor system maintains a model of the state of the environment, and compares the sensory input resulting from an action with what would be

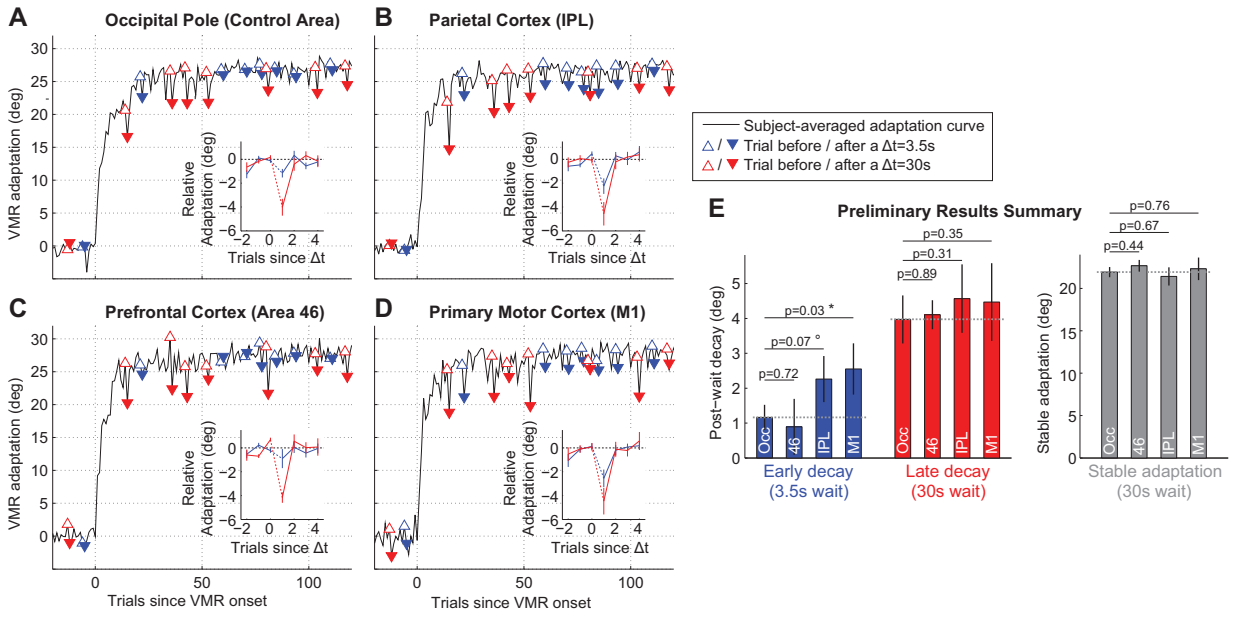


Figure 4.9: The primary motor and posterior parietal cortices might be involved in temporally-labile adaptation.

A: Average learning curves for the Occipital Pole (controls) group, indicating adaptation before (empty, upwards-pointing triangles) and after (filled, downwards-pointing triangles) a Δt (blue: 3.5s, red: 30s), and illustrating the extent of time-dependent decay. Inset: adaptation around the two types of Δt trials (three trials before to four trials after), showing the breadth of decay and almost immediate recovery. Adaptation is referenced to the average of the two pre-Δt trials (trials -1 and 0), whereas errorbars indicate SEM.

B-D: Same as (A), but for the posterior parietal cortex (IPL), prefrontal cortex (Area 46) and primary motor cortex (M1) groups. Note, in the insets, how decay after the long Δt (30s) appears similar across the four groups, however, the decay after the short Δt (3.5s) is greater for IPL and M1.

No different temporally-stable adaptation (remaining after 30s Δt) or temporally-labile adaptation (decay after 30s Δt).

E: Summary of the results from Experiment 4.2. Late decay (red, Δt = 30s) was similar across all groups, and, correspondingly, the temporally-stable components (gray) for each group were also similar, suggesting that neither area is involved in temporally-stable adaptation. However, early decay (blue, Δt = 3.5s) was twice as strong in the primary motor and posterior parietal cortex stimulation groups, compared to the occipital pole and the prefrontal cortex groups. This faster, but not deeper, decay suggests that the M1 and PPC might be involved in the formation and maintenance of temporally-labile motor memories.

predicted by this model, providing a mechanism for updating that model and thus enabling motor learning. In line with the idea that the cerebellum is the locus of these state estimates, previous research found that disruption of the cerebellum, using single-pulse TMS right after the onset of movement, in turn disrupted the estimation of the state of the environment (Miall et al., 2007).

Moreover, cerebellar damage has been shown to produce deficits in motor learning (Sanes et al., 1990; Smith and Shadmehr, 2005; Tseng et al., 2007; Criscimagna-Hemminger et al., 2010; Gibo et al., 2013), in line with the idea that the cerebellum is involved in motor learning. Interestingly, however, the degree of impairment varies considerably from one study to another, with widely ranging results reported even within the same task. For example, three recent force field adaptation studies in reaching arm movements have reported very different levels of impairment (patient performance when compared to controls) ranging from 30-70%, as shown in Figure 4.10. (Smith and Shadmehr, 2005) reported a substantial 75% deficit in force field adaptation, but two more recent studies (Criscimagna-Hemminger et al., 2010; Gibo et al., 2013) found only 30-40% learning deficits in this task.

In this section, we suggest that these differences can be explained by, and provide evidence for, the idea that the cerebellum is essential for temporally-stable adaptation. If this were the case, paradigms with longer inter-trial spacing might be more sensitive to cerebellar damage, as this damage would lead to greater decay of adaptation in the time interval between movements. Interestingly, the (Smith and Shadmehr, 2005) study was based on an 8-movement-direction paradigm whereas (Gibo et al., 2013) and (Criscimagna-Hemminger et al., 2010) employed only one or two movement directions, respectively. Correspondingly, the inter-trial intervals between consecutive movements in the same direction were 4-fold greater in the (Smith and Shadmehr, 2005) study, as shown in Figure 4.10D.

To examine our idea that temporally-stable adaptation involves the cerebellum – and thus that cerebellar damage reduces the temporal stability of motor memories – we reanalyzed data from these three previous studies. We took advantage of the fact that inter-trial intervals could range widely within the same study, as shown in Figure 4.11 for the Criscimagna-Hemminger and Gibo studies. Comparing adaptation after a

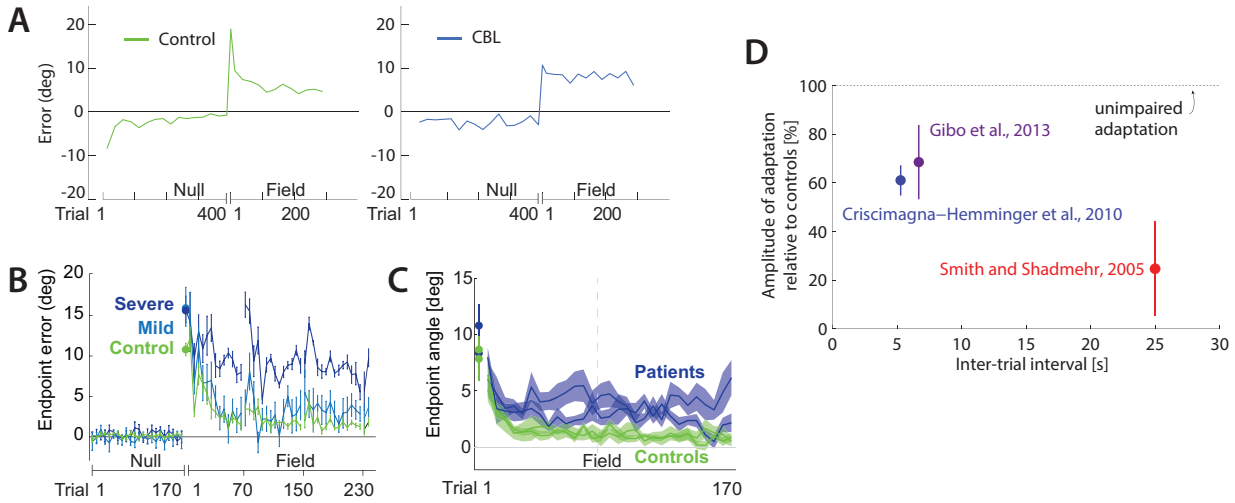


Figure 4.10: Previous studies show a wide range of motor learning deficits in cerebellar patients.

A: Force-field adaptation in controls (green, left) vs. cerebellar patients (blue, right) in the (Smith and Shadmehr, 2005) study from which this panel is adapted. Whereas controls are able to reduce their initial error after force-field introduction, illustrating motor learning, patients do not reduce their initial error, showing a near-complete deficit in learning.

B: Force-field adaptation in the (Criscimagna-Hemminger et al. 2010) study from which this panel is adapted, showing considerably less pronounced deficits for cerebellar patients when compared to A. Specifically, even severe cerebellar patients (dark blue, with an ICARS score of ≥40) are able to reduce their error by about 50%, illustrating motor learning; at the same time, the performance mild cerebellar patients (light blue) is close to that of controls (green).

C: Force-field adaptation in the (Gibo et al. 2013) study from which this panel is adapted, showing a considerable amount of motor learning, in line with the results of (Criscimagna-Hemminger et al. 2010) and in contrast to (Smith and Shadmehr, 2005). Here, two patient groups (blue) reduce their initial error, amounting to motor learning are compared to controls (green).

D: Stronger deficits might be influenced by longer inter-trial time intervals. Here, the deficit in motor adaptation for patients is expressed as a % of control performance, and is plotted against the median inter-trial interval for each study. Errorbars indicate SEM.

short vs. a long inter-trial interval within the same study enables us to systematically investigate decay, and estimate decay curves analogous to the ones shown in Figure 4.2.

Patient performance is more sensitive to increased inter-trial intervals when compared to controls

We first focused our analysis on the Criscimagna-Hemminger study as it contained more data over a wider range of inter-trial intervals (ITIs, Figure 4.11). In this study, 15 patients with cerebellar disease, along with 15 age-matched controls, were trained to hold a robotic manipulandum and use it to make out-and-back movements through a target. During training, the robot produced a viscous force-field, which perturbed the motion laterally to the movement direction, with forces that scaled with movement velocity. That dynamical perturbation was very similar to the one used in Experiment 2.1, 2.2 (Chapter 2) and 3 (Chapter 3), and had a force-field strength *b* equal to 11 Ns/m. Training consisted of 240 trials, separated by three rest breaks (trials 70, 150, and 230) that ranged from 30 seconds to three minutes long. The experiment consisted of two arms: one where the force-field was introduced gradually, and one where the force-field was introduced abruptly, which we analyze first.

The data from the abrupt force-field training block showed that controls displayed a considerable error when they experienced this force-field perturbation the first time (gray curve in Figure 4.12A); however, with training (during the last 75% of training trials), they learned to minimize this error and move almost directly to the target, amounting to near-complete asymptote adaptation. Figure 4.12A shows, with the darkest green, the average hand path when the ITI since the last training trial – and thus the time-dependent decay – was minimal: when the preceding training movement was on the same direction. A similar analysis of patient behavior shows comparable results: while patients do not learn as well as controls (as illustrated by the higher deviation on asymptote, dark blue in Figure 4.12B), they have also reduced most of their error by asymptote adaptation – for trials that were immediately preceded by trials on the same direction, meaning any decay of temporally-labile adaptation was minimal.

In contrast, patient and control behavior drastically diverge when we compare training trials that occurred after increasingly higher ITIs. We divided the data into a total of 4 bins based on the number of intervening movements between each trial and the immediately preceding trial that was on the same

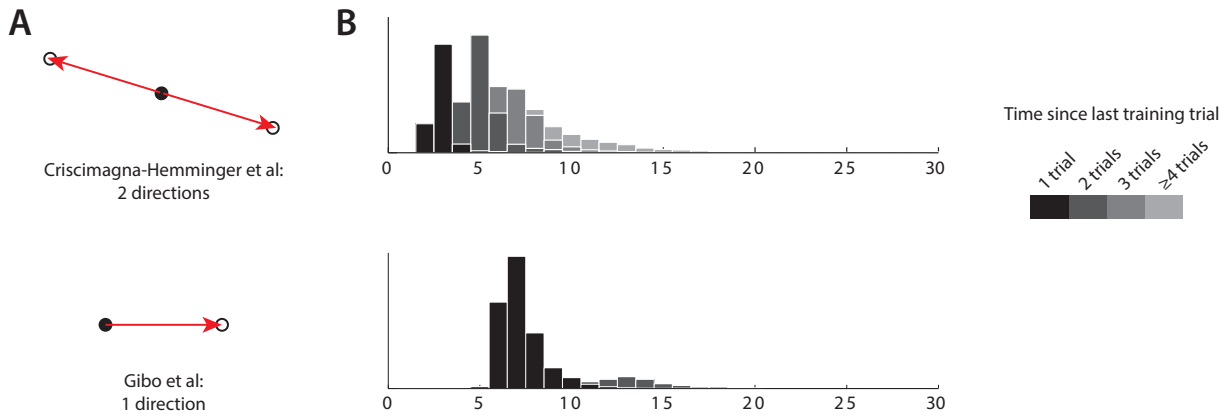


Figure 4.11: Distribution of inter-trial intervals in the (Criscimagna-Hemminger et al., 2010) and (Gibo et al., 2013) studies.

A: Illustration of the different movement directions for the (Criscimagna-Hemminger et al., 2010) and (Gibo et al., 2013) studies. The two different targets in (Criscimagna-Hemminger et al., 2010), which were presented in random order, meant a wide distribution of inter-trial intervals depending on how many trials passed since the last time the same target was visited. For (Gibo et al., 2013), the presence of only one movement direction meant a narrower distribution of inter-trial intervals; however, ~10% of trials were error-clamp (EC) trials, which effectively doubled the inter-trial interval, as they probed forces without providing additional training.

B: Distributions of inter-trial intervals for the two studies, colored based on the number of intervening trials between training trials on the same direction.

movement direction: 0 (when the previous movement was on the same direction, 1, 2, or 3 or more. We added a 5th bin for trials that immediately followed rest breaks, and had a considerably longer ITI than the other four cases (Figure 4.12C). We found that, for both controls and patients, increased ITIs result in the handpaths deviating more and more from ideal, showing decay of adaptation (from the darker to the lighter colors in Figure 4.12AB). Deviations were highest after a break, in line with the considerably longer ITIs (the lightest green and blue colors in Figure 4.12AB). For controls, the handpath deviations after these longest-ITI trials were still considerably smaller than the initial deviations, suggesting that most of adaptation did not decay with time. In contrast, for patients, deviations, got increasingly larger with longer ITIs, with deviations after a break being almost as large as the initial ones (Figure 4.11B), suggesting that almost all of their adaptation decayed with time.

Estimation of decay curves

To more systematically investigate the decay curves for both controls and patients, we quantified decay by computing a simple measure of adaptation for each trial: the fraction of the initial angular error, e_{init}, that each subject compensated for during each trial *k* (Equation 4.3):

(Equation 4.3) $$y(k) = \frac{e_{init} - e(k)}{e_{init}} * 100 \quad [\%]$$

The adaptation metric $y(k)$ would equal 0% when the angular error during the trial was as bad as the initial one ($e_{init} = e(k)$), and would equal 100% when the angular error during the trial was zero. The analysis of how this adaptation metric evolves with increased ITI mirrored the qualitative comparisons made in Figure 4.12C: while the adaptation of cerebellar patients after the shortest ITIs was less than a third less than the adaptation of controls (66.0±5.9% vs. 94.3±1.9% $p<10^{-4}$), the passage of time led to a much deeper decay whereby the adaptation of patients was three quarters less than controls (18.0±10.6% vs. 68.1±4.9%, $p<10^{-4}$ for patients vs. controls after breaks).

The adaptation deficits of cerebellar patients are specific to temporally-stable adaptation, suggesting that temporally-stable adaptation critically involves the cerebellum

We proceeded to break adaptation down into temporally-labile and temporally stable components, by taking temporally-stable adaptation to be equal to the amount remaining after a break, and temporally-labile adaptation as the difference between overall adaptation after the minimum ITI and temporally-stable adaptation (Figure 4.12C). We found that temporally-stable adaptation was severely impaired for patients (18.0±10.6% vs. 68.1±4.9%, $p<10^{-4}$) showing that the cerebellum is crucial for the temporal stability of motor memories. In contrast, temporally-labile adaptation was not impaired – it was, in fact, slightly higher for patients than controls (48.0±9.9% vs. 26.2±5.1%, p = 0.045, 2-sided unpaired t-test), suggesting that the motor system partially compensated for the reduced temporally-stable component with a stronger temporally-labile one.

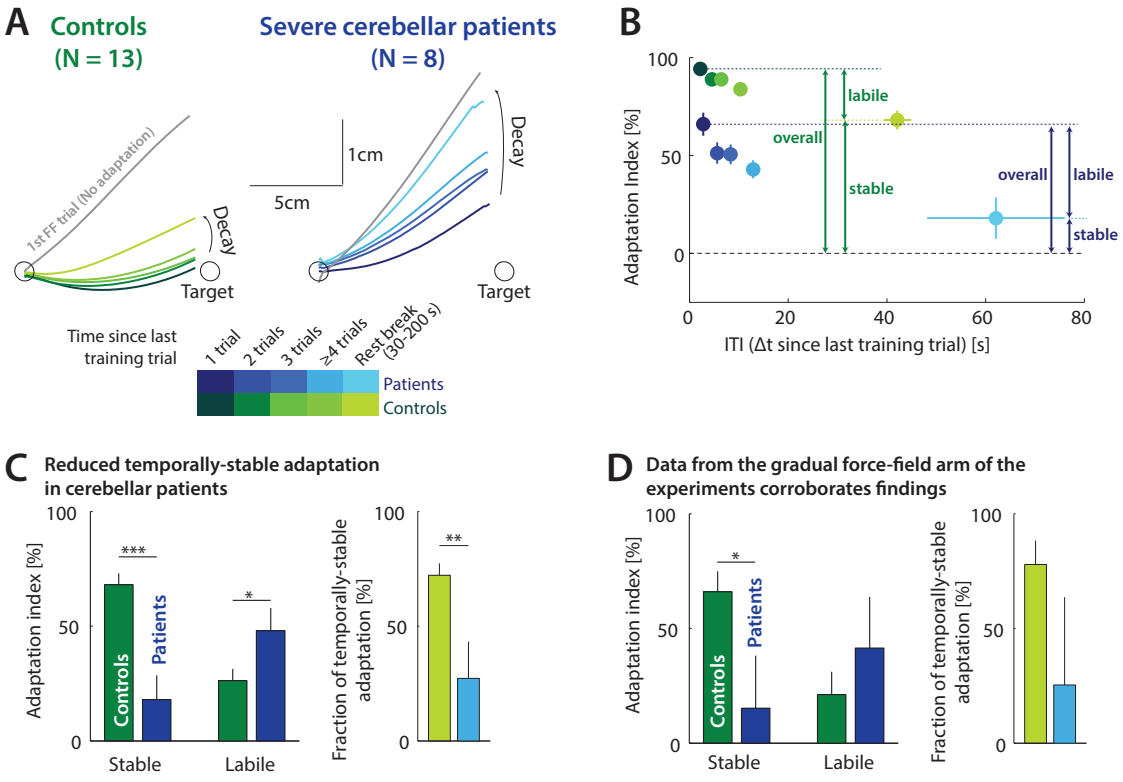


Figure 4.12: The motor learning deficits of cerebellar patients are specific to temporally-stable adaptation.

A: Average paths from the Criscimagna-Hemminger et al. (2010) study for each type of inter-trial interval, as was defined by the distance to the most recent trial on the same direction (1, 2, 3 or ≥4 trials, or after a rest break). Controls are shown on the left; severe cerebellar patients are shown on the right (ICARS index > 40). Note that patients are able to considerably reduce errors compared to the initial exposure to the force-field (gray) when the ITIs are low. However, while both patients and controls show increasing deviations (decaying performance) with increased ITI, these deviations increase much faster for the cerebellar patients. While the performance of controls after a rest break (which indicates temporally-stable adaptation) is still much better than the performance during the initial force-field exposure, in contrast, the performance of patients after a rest break is nearly as bad as during initial force-field exposure, indicating higher temporal volatility of motor memories.

B: Decay of adaptation with increased ITI. Note how controls' adaptation decay to a nonzero level, in line with the data shown in Figure 4.2 whereas patients' adaptation decays close to zero. The adaptation index was calculated as the % improvement from the angular error that was experienced during initial force-field exposure (Equation 4.3).

C: *Left panel:* Temporally-stable and temporally-labile components for patients and controls. The temporally-labile component was calculated based on the difference between overall adaptation (measured during consecutive same-direction learning trials, darkest shades in A and B) and adaptation after a rest break, whereas the temporally-stable component was calculated based on the adaptation after a break. Temporally-stable adaptation is drastically reduced in cerebellar patients, with controls displaying three-fold more stable adaptation; however, temporally-labile

Figure 4.12 (continued)

adaptation is not reduced, but in fact slightly higher in patients. This result shows that the motor learning deficits of cerebellar patients are specific to temporally-stable adaptation. *Right panel:* Fraction of overall adaptation that is temporally-stable. In line with B, the motor adaptation for controls is about three times more stable compared to patients (72±5% for controls vs. 27±16% for patients, $p<0.01$).

D: Same as C, but for data form the gradual arm of the Criscimagna-Hemminger et al., 2010 experiment, showing similar results.

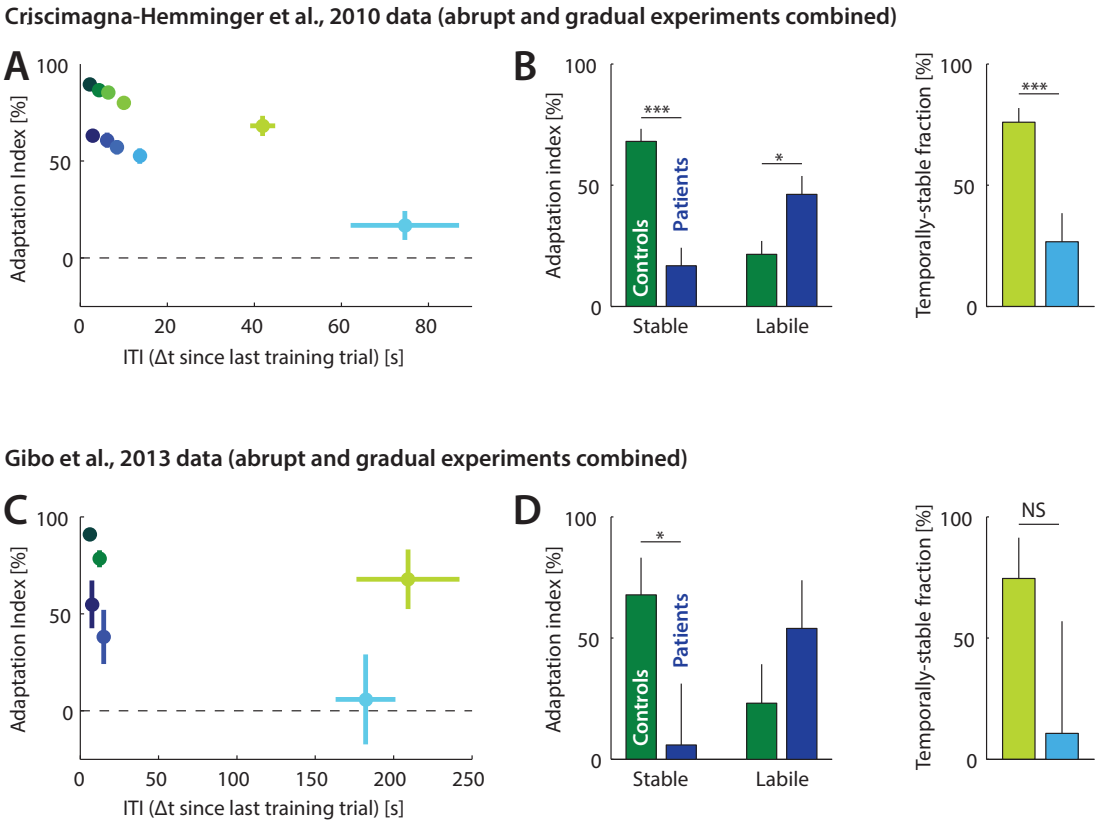


Figure 4.13: The motor learning deficits of cerebellar patients are specific to temporally-stable adaptation – summary of the analysis.

A: Decay of adaptation with increased ITI in the Criscimagna-Hemminger et al., 2010 experiment, in the same format as 4.12B, but averaging together results from both the abrupt and gradual arm of the experiment.

B: Temporally-stable and temporally-labile components for patients and controls (left panel) and fraction of overall adaptation that is temporally-stable (right panel) in the same format as 4.12C, but encompassing results from both the abrupt and gradual arms of that experiment together. These results further illustrate that the motor learning deficits of cerebellar patients are specific to temporally-stable adaptation.

C-D: Analogous to A-B but for the Gibo et al. (2013) study showing similar results.

Together, these results show that the impairments in motor learning cerebellar patients are specific to temporally-stable adaptation. Correspondingly, when we calculated the fraction of adaptation that was temporally-stable, we found that for controls, temporally-stable adaptation is the bulk of overall adaptation – 72%, similar to previous work on force-field adaptation with healthy participants (Sing et al., 2009) and our own data for VMR adaptation (section 4.1). In contrast, we found that, for patients, temporally-stable adaptation was only 27% of total.

An analogous analysis of the gradual force-field adaptation arm of the (Criscimagna-Hemminger et al., 2010) experiment provided further evidence for this idea. As shown in Figure 4.12D, temporally-stable adaptation, was also severely impaired for patients compared to controls (15.2±23.0% vs. 66.0±8.9 %, p=0.029), whereas temporally-labile adaptation was not (41.4±22.3% vs. 21.1±10.0 %, p=0.36, 2-sided paired t-tests). We thus summarized the results of this experiment by combining abrupt and gradual data in Figure 4.13AB.

The analysis of the (Gibo et al., 2013) data mirrors these results. Because in that study there was only one movement direction, we binned the data based on only three types of trials. The first bin consisted of trials preceded immediately by a training trial. These trials were the most common (90% of the training trials) and generally had the smallest ITI compared to the other trial types. The second bin consisted of trials that were preceded by an error-clamp trial, a type of trial used to measure force responses while providing no additional training (see Chapter 2), thus effectively doubling the inter-trial interval. These post-error-clamp trials were about 10% of the training trials. Finally, the third bin consisted of trials after a break which by far had the longest ITI (Figure 4.13C). Since there was only one break during training, this bin consisted of only one trial per subject.

While we had to base our estimates on a smaller number of trials compared to the first study (with only one post-break trial available instead of four), a fact which consequently resulted in noisier estimates, our results on average agreed with our analysis of the Criscimagna-Hemminger study. Specifically, we found that time-dependent decay for patients was much deeper compared to controls, from an initial overall

adaptation level of 54.8±12.3% to a temporally-stable adaptation level of 5.9±25.3% for patients, compared with 90.9±2.1% overall decaying to 67.8±15.3% stable for controls. Correspondingly, this temporally-stable adaptation was significantly weaker for patients compared to controls ($p = 0.018$, 1-tailed unpaired t-test) whereas the temporally-labile adaptation was not compromised (53.9±19.9% for patients vs. 23.2±16.0% for controls, $p = 0.24$). In short, we observed near-complete decay for patients, and much smaller decay for controls, in line with our analysis of the previous study.

Reanalysis of the Smith and Shadmehr, 2005 data

Finally, we conducted a reanalysis of the Smith and Shadmehr data, which had on average larger inter-trial intervals. Long inter-trial intervals generally would prevent temporally-labile adaptation to build up – the median ITI was about 30s, which is more than 1.5 times as long as the time constant of decay for temporally-labile adaptation – making harder to assess its decay as we did for the other two studies. However, we took advantage of a special kind of trials – catch trials (Thoroughman and Shadmehr, 2000) – which were interspersed with force-field training trials. During catch trials, the force-field was switched off, effectively providing a strong training signal towards the opposite direction. The subsequent opposite learning can thus result in increased error for the next force-field trials. Indeed, we see a clear increase in error after a catch trial for controls (6.51±0.35° vs. 8.83±0.41°, $p < 10^{-4}$) – but not a clear increase for patients (9.78±0.94° vs. 1.30±1.41°, $p = 0.49$, Figure 4.14A). This lack of difference corresponds to a near absence of overall opposite learning, in line with the general learning deficit for patients in the study.

However, we were able to break this single-trial opposite learning down into temporally-labile and temporally-stable adaptation in order to examine whether results from the Smith and Shadmehr study agree with the other two studies, in particular, with our finding that temporally-stable, but not temporally-labile adaptation is affected by cerebellar damage. From the force-field trials that followed catch trials, we isolated the occasions (13 out of the 40 post-catch FF trials) on which the FF trial immediately followed

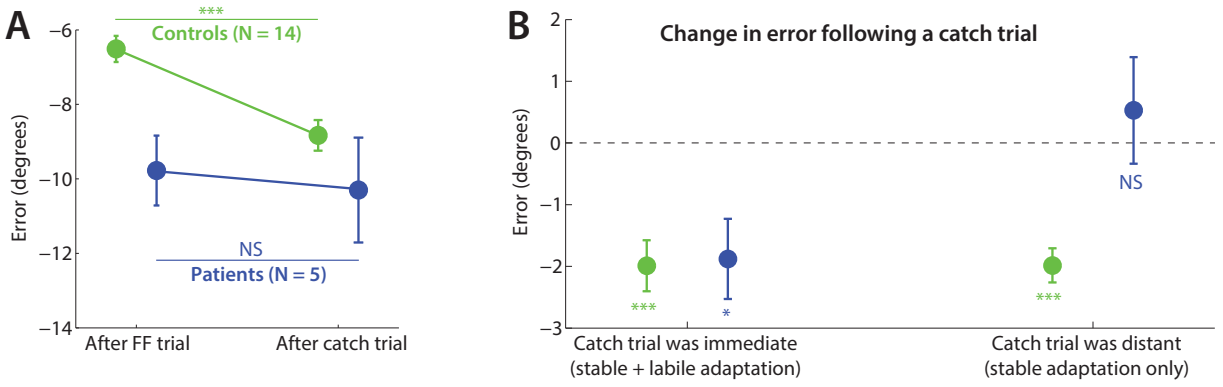


Figure 4.14: Temporally-labile adaptation for cerebellar patients in the Smith and Shadmehr, 2005 experiment but not temporally-stable adaptation.

A: Catch trials lead to an increased error in overall adaptation for controls (green) but not patients (blue). Here, we analyze adaptation before and after catch trials, during which the force-field was switched off, effectively providing a strong training signal towards the opposite direction. The subsequent opposite learning can thus result in increased error for the next force-field trials, providing a way to assess whether the mechanisms of learning are intact. The stronger error for controls indicates the presence of single-trial overall adaptation, whereas its absence in patients is in line with the deficits in overall adaptation that were observed in that study.

B: Temporally-labile but not temporally-stable adaptation for patients illustrates that their deficits were specific to the temporally-stable component. Plotted here is the change in error between catch and post-catch trials (negative numbers denote increased error). To probe temporally-labile adaptation, we isolated the post-catch trials which immediately followed the catch trial – i.e. without any intervening movements to other targets that would increase ITI. We found that, for patients, these trials displayed stronger error in line with the presence of motor adaptation. However, for post-catch trials for which the catch trial was more distant in time – 2 or more intervening movements, amounting to an increased ITI – there was no adaptation, showing no temporally-stable motor adaptation. Together, these results show that the adaptation deficits displayed by patients in the Smith and Sharmehr, 2005 study were specific to temporally-stable adaptation.

the catch trial – when there were no intervening movements to other targets – resulting in a relatively small ITI (~3.5 s). We found increased errors compared to adaptation following non-catch trials (1.88±0.53° increase, $p = 0.022$, 1-tailed paired t-test, Figure 4.14B) indicating that the inverted error on the previous trial resulted in learning. In contrast, when the catch trial was more distant in time – 2 or more intervening movements, 21 out of the 40 post-catch trials – there was no similar increase in error (-0.65±0.86° increase, $p = 0.71$, 1-tailed paired t-test, Figure 4.14B), indicating that there was no stable

learning. Thus, while these patients exhibited learning immediately after these catch trials, this learning was only labile. This corroborates the analysis of the previous datasets, as it shows no stable adaptation but a degree of labile adaptation.

4.4.5 Temporally-stable adaptation predicts second-day retention

So far, we contrasted the temporally-labile component of motor adaptation – decaying over the course of tens of seconds – with the temporally-stable component – which persists at least over the course of minutes. How persistent is the latter over longer timescales, and how does it relate to long-term retention? Previous work on force-field adaptation found that retention 24 hours later scaled with the amount of temporally-stable adaptation present at the end of training, and was equal to about half of temporally-stable adaptation (Sing et al., 2009). Moreover, the same study suggested that the temporally-stable component of FF adaptation corresponds to the slow-learning, high-retention process ("slow process") of a two-state model of motor adaptation (also see Discussion). Importantly, this slow process was also found to predict retention after 24 hours (Joiner and Smith, 2008; Shadmehr et al., 2010a).

Here, we first investigate whether the ability of the temporally-stable component to predict long-term retention is also true for visuomotor adaptation tasks. We examine this relationship both for subject-averaged behavior, analogous to what was done in previous FF studies, but also extend our analysis to investigate how temporally-stable adaptation predicts long-term retention across individuals. We proceed to compare the shape of these two memories, based on their pattern of generalization across different movement directions, which provides a fingerprint of the underlying memory in line with the directional tuning of neuronal elements associated with motor learning (Shadmehr, 2004) – also see section 1.3.

To investigate the long-term retention of VMR adaptation, we trained N=25 subjects to asymptote performance (120 trials) on a 30 degree VMR. We then gave a break that was at least 1 minute long in order to allow the temporally-labile component to decay almost completely (based on the time constant of

decay from paragraph 4.4.1, τ = 16.5s, 1 minute ≈ 3.6 τ, corresponding to >97% decay of temporally-labile adaptation). After the break, we measured the directional generalization of the temporally-stable adaptation component in a block that consisted of repeated probe trials to 19 different movement directions, arranged from -135 to 135 degrees about the trained direction in increments of 15 degrees (see Figure 4.15B) . During these probe trials, visual feedback, which could drive further error-dependent learning, was withheld. This process was repeated after a 60-trial retraining block, and we then brought subjects back the following day and immediately tested them for generalization (see Figure 4.15A).

Temporally-stable adaptation predicts second-day retention

The temporally-stable component, measured for movements on the trained direction, contributed 60.7±3.7% (mean±SEM across subjects) of the asymptotic adaptation during training (16.5±1.0 out of 27.3±0.3 degrees). Importantly, about half of this temporally-stable component were retained when tested one day later (8.8±1.2 out of 16.5±1.0°, or 51.5±5.2%, see Figure 4.15E,F). We note that the fraction of the temporally-stable component of VMR that displays long-term retention is similar to the one observed for FF adaptation (40-50% (Joiner and Smith, 2008) or 60% (Sing et al., 2009) also depending on the length of the training block), showing that the apparent function of the temporally-stable process as a gateway to long-term retention (Joiner and Smith, 2008; Sing et al., 2009) is not a characteristic of only one specific learning paradigm, but could instead be a more general characteristic of motor learning.

Moreover, analysis of individual subject data shows that the temporally-stable component can also predict long-term retention on the individual level, whereby subjects with weak temporally-stable adaptation on day 1 generally display weaker retention on day 2. Specifically, we found a significant ($p = 0.0001$, R^2=0.50) correlation between day 1 temporally-stable adaptation and day 2 retained adaptation (Figure 4.15G), further suggesting that long-term motor memories derive from temporally-stable adaptation.

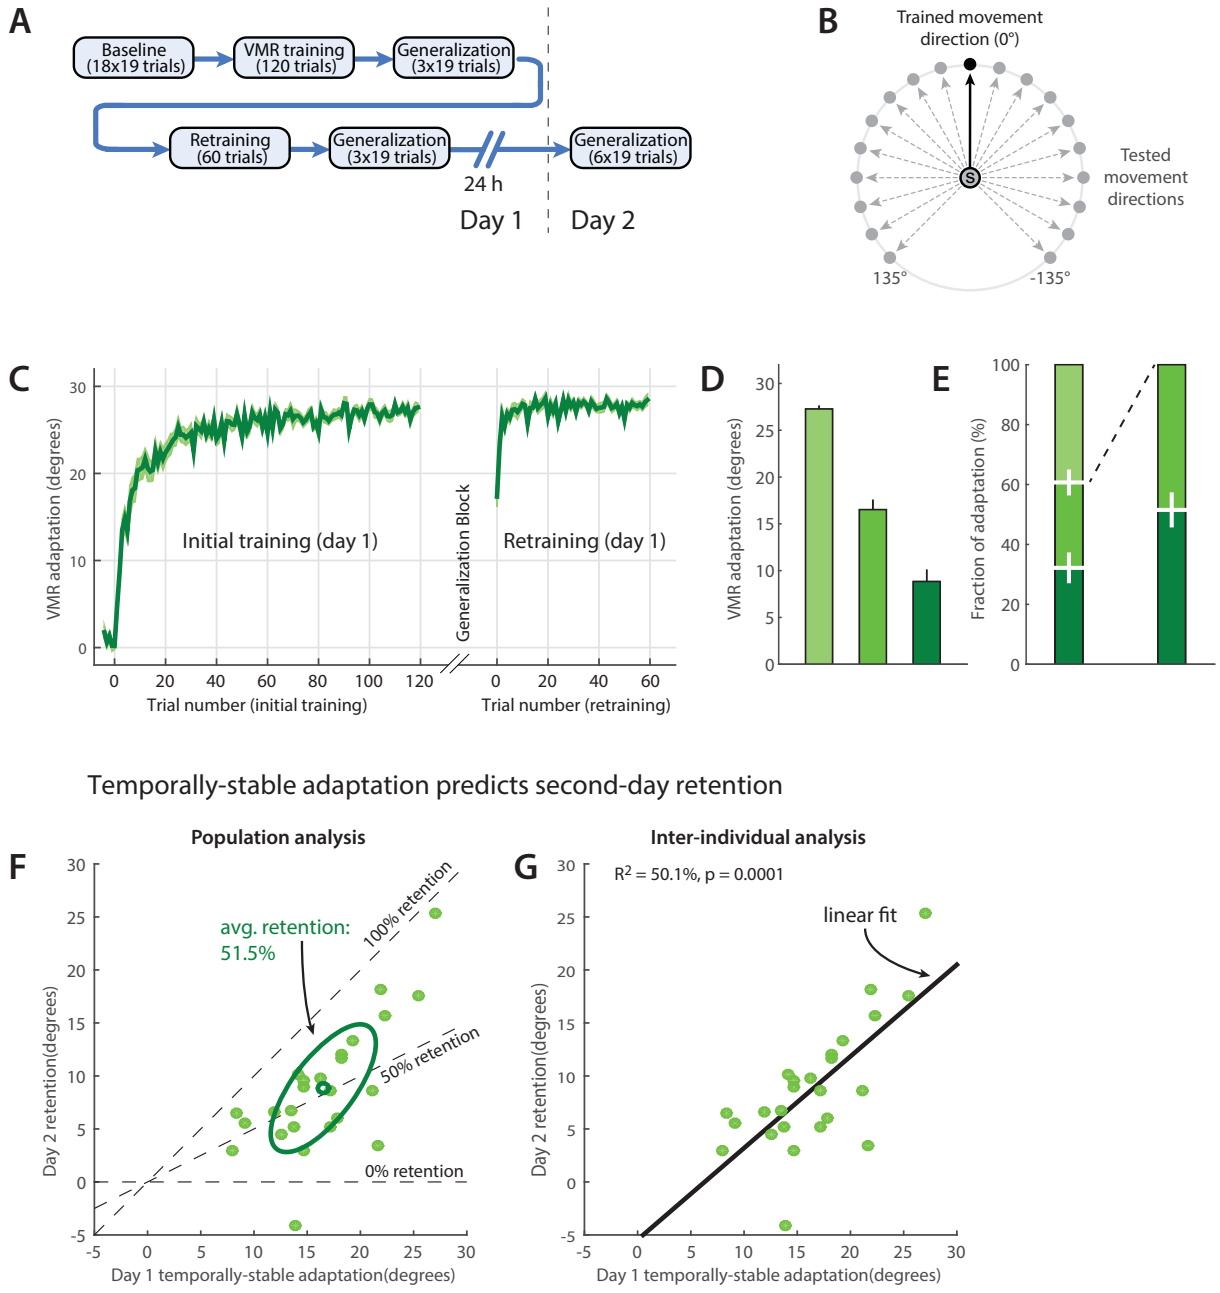


Figure 4.15: Long-term retention of visuomotor adaptation.

A: Illustration of the trained movement direction (black, 0°) and the testing / generalization directions (gray, from -135° to 135° in 15° increments) for Experiment 4.3. During baseline and generalization blocks, these 19 different movement directions were presented in random order.

B: Experiment design for Experiment 4.3. After a baseline period, during which subjects visited each one of the 19 different targets 18 times (6 without and 12 with visual feedback), subjects entered the training phase whereby they were trained on a 30° VMR (CW or CCW, balanced across subjects) for 120 trials. They were then tested for generalization in a block which contained 3 movements to each direction, all without visual feedback. They were

Figure 4.15 (continued)

then retrained for 60 trials and then retested for generalization, which concluded the day 1 experiment schedule. Subjects returned a day later and were retested for generalization.

C: Learning and relearning curves from Experiment 4.3, showing that overall adaptation asymptotes quickly, over the course of a few tens of trials. We thus quantified overall adaptation by taking the last 40 trials of both the training and retraining sessions.

D: Comparison of overall adaptation on day 1 (light green), temporally-stable adaptation measured on day 1 (green) and adaptation retained on day 2 (dark green).

E: The composition of overall (left) and temporally-stable (right) adaptation. Roughly 40% of overall adaptation is temporally-labile, whereas out of the remaining 60% of temporally-stable adaptation, about 50% is retained on day 2.

F. Population analysis of Day 1 temporally-stable adaptation vs. Day 2 retention, showing an average retention of 51.5%.

G: Same data as in (F), focusing on the significant ($p = 0.0001$) linear relationship (black line) between Day 1 temporally-stable and Day 2 retained adaptation across subjects, showing how temporally-stable adaptation, measured only 1-2 minutes after training, can predict long-term retention a day later.

Globally-generalizing stable adaptation is more strongly retained than the locally-generalizing one

We then wanted to examine how different components of the memory traces associated temporally-stable adaptation might contribute to long-term retention. Previous research (Brayanov et al., 2012; Brayanov, 2013) found that these memories are characterized by two components based on how they generalize across different movement directions: one component that is specific to the trained direction, and thus generalizes narrowly about the trained direction (local component) and a component that generalizes uniformly throughout the workspace (global component). In order to systematically characterize each subject's directional generalization pattern (what we refer to as the generalization function), we fit their data using a Gaussian curve (corresponding to the local component) with an offset (corresponding to the global component) (Brayanov et al., 2012; Brayanov, 2013) as in Equation 4.4:

(Equation 4.4) $$x(\theta) = L * e^{-\frac{\theta^2}{2\sigma^2}} + G$$

where $x(\theta)$ gives the generalization function x at an angle θ away from the trained direction, L and σ correspond to the strength and the width of the local component, and G corresponds to the strength of the global component.

In line with previous work, we found that our generalization data for the temporally-stable process displayed distinct local and global components, with the bulk of temporally-stable adaptation contained within the local component (11.1±0.7° vs. 4.5±1.0° for the local and global components, respectively, $p<10^{-6}$, or 67.2±2.4% of stable adaptation for the local component, ratio estimated using bootstrap, Figure 4.14A, right panel and Figure 4.15C). However, when we examined our next-day data, we found that the contributions of the global and local components to the retained generalization were instead roughly equal (5.4±0.5° vs. 3.7±0.5° for the local and global components, respectively, $p<0.001$, or 54.9±5.4% of stable adaptation for the local component, ratio estimated using bootstrap, Figure 4.16B, right panel and Figure 4.14C). This change was the result of the differential retention of the global and local components: while 68.7±8.9% of the first-day global component was retained on the next day (4.5±1.0° on day 1 vs. 3.7±0.5° on day 2, $p=0.39$), only 36.0±7.8% of the corresponding local component was retained (11.1±0.7° on day 1 vs. 5.4±0.5° on day 2, $p<10^{-5}$), leading to a change in the shape of the generalization function (Figure 4.16A,B). In short, the globally-generalizing component of temporally-stable adaptation is much more robust to decay, than the locally-generalizing one and thus contributes more to long-term retention.

Both local and global generalization components of long-term retained adaptation are predicted from temporally-stable adaptation

The power of the total temporally-stable component to predict long-term retention of individual subjects extends to both the local and global components. When we compared the strength of these components on day 1 vs. day 2, we found significant relationships in both cases ($p = 0.0004$, $R^2=0.43$ for the local component, Figure 4.17C, and $p = 0.0044$, $R^2=0.30$ for the global component, Figure 4.17D). This ability of the temporally-stable component to predict, not only the magnitude but also the shape of generalization for long-term retention – with the shape itself being a fingerprint of the underlying motor memory

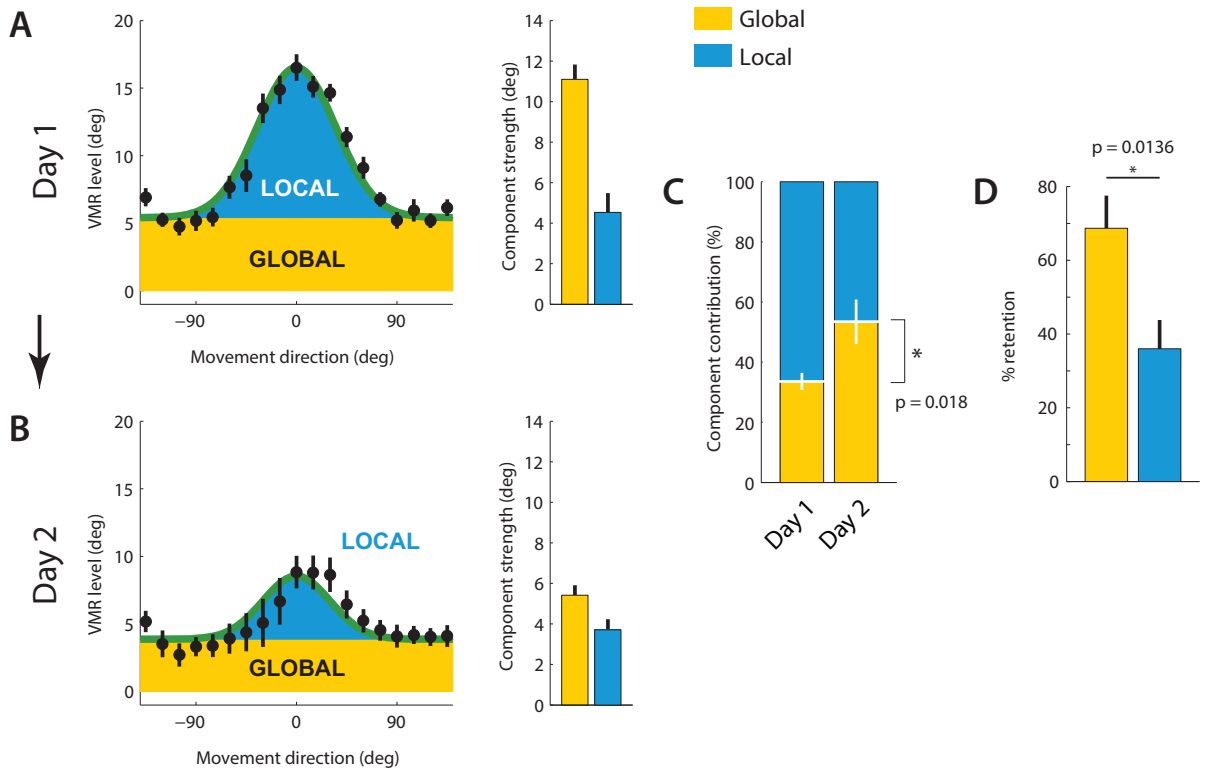


Figure 4.16: The globally-generalizing component of motor adaptation is better retained than the locally-generalizing one.

A: Generalization pattern for the temporally-stable component as measured ~1 minute after training on Day 1 (left panel). The generalization function is characterized by a locally-generalizing component, fit with a gaussian function (green curve), and a globally-generalizing component (the offset of the green curve). The local component is considerably larger than the global one (right panel).

B: Generalization pattern for the memory retained ~24 hours after training. As in (A), the generalization function is characterized by a local and global component. Whereas both components are reduced compared to Day 1 (A), the local component is reduced more strongly so that the Day 2 local and global components are similar (right panel).

C: % contribution of the local and global component to the retained memory on Day1 vs. Day 2. The stronger reduction in local adaptation compared to global adaptation means that the contribution of the latter is increased in Day 2. Errorbars represent SEM.

D: Day 2 retention of the local and global components as a % of their strength on Day 1. The retention of the global component is about two times stronger (68.7% vs. 36.0%).

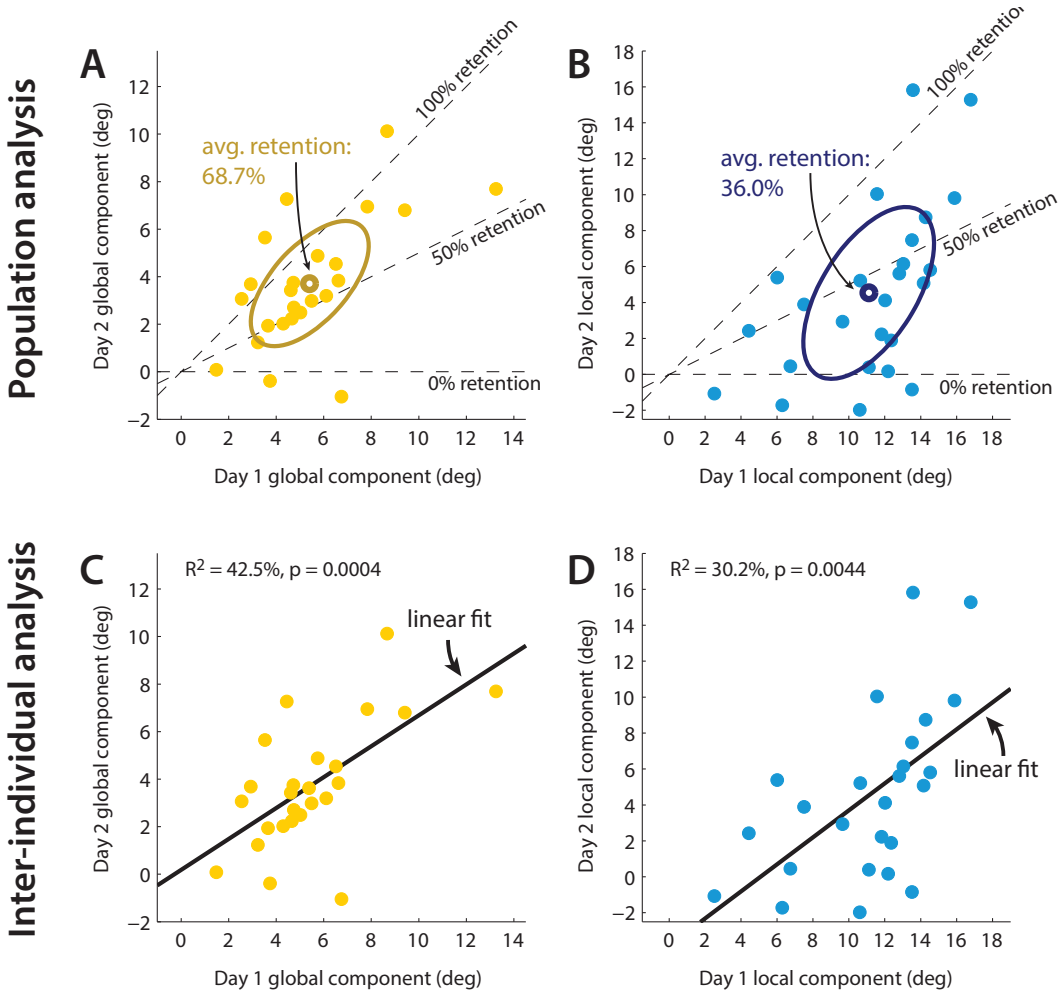


Figure 4.17: The individual components of temporally-stable memory predict the corresponding components of long-term retention.

A: Scatter plot of the local components of generalization for each individual subject, for Day 1 (temporally-stable adaptation) against Day 2 (long-term retention). This population analysis shows an average retention of 68.7% for the local component.

B: Same population analysis as in A, but for the of Day 1 temporally-stable adaptation vs. Day 2 retention, showing an average retention of 36.0% for the global component.

C: Same data as in (A), focusing on the significant ($p = 0.0004$) linear relationship (black line) between the Day 1 and Day 2 locally-generalizing components across subjects. This result shows how the locally-generalizing component of temporally-stable adaptation can predict the locally-generalizing adaptation component of long-term retention a day later.

D: Same data as in (C), focusing on the significant ($p = 0.0044$) linear relationship (black line) between the Day 1 and Day 2 globally-generalizing components across subjects. Thus, the globally-generalizing component of temporally-stable adaptation can predict the globally-generalizing adaptation component of long-term retention a day later.

(Shadmehr, 2004) – suggests that long-term retention is the manifestation of the same memory trace as the temporally-stable component as measured only a minute after training. This, in turn, provides further evidence for the idea that long-term motor memories derive from the slow-learning, temporally-stable process (Joiner and Smith, 2008; Sing et al., 2009; Shadmehr et al., 2010a; Sing, 2011).

Moreover, since local and global components display different capacities for long-term retention, our result implies that measuring both the local and global components of temporally-stable adaptation should provide a better overall prediction of an individual's long-term retention of adaptation, compared to simply measuring the total level of temporally-stable adaptation. Correspondingly, a training paradigm designed to enhance the global portion of temporally-stable adaptation could result in more robust long-term retention. Further work could identify how this could be achieved and examine the extent of the potential benefits.

4.4.6 The temporal stability of motor adaptation is sensitive to the temporal correspondence between action and feedback[2]

The finding that temporally-stable adaptation can serve as a gateway to long-term retention, invites us to investigate the conditions that might affect the fraction of motor learning that is due to temporally-labile vs. temporally-stable adaptation. Here, we investigate the effects of delays in visual feedback in the temporal stability of motor memories. Delays in visual feedback may have important consequences for motor learning, as this learning is often driven by the comparison between observed and expected sensory feedback. While the motor system routinely deals with the relatively large delays inherent in sensory feedback loops – which can range from 100 – 250 ms for visual feedback (Katayama and Kawato, 1993)

[2] This work was conducted in collaboration with Katherine E. Mallett, and a large part of the results presented in this section is also contained in her senior thesis, Mallett KE (2014) Visual Feedback Latency Reduces the Retention of Visuomotor Learning and Alters Its Internal Representation. In. Cambridge, MA: Harvard University.

– previous studies (Kitazawa et al., 1995; Kitazawa and Yin, 2002) have suggested that (further) external delays in the presentation of endpoint-only visual feedback reduce both the rate and amount of visuomotor adaptation, even for delays as short as 50ms. However, when visual feedback is continuously presented rather than being relegated to the endpoint, the consequences of visual latencies on the rate, amount and temporal stability of adaptation are unknown.

Visual latencies during training reduce temporally-stable, but not temporally-labile or overall adaptation

We thus examined the effect of delays in continuous visual feedback on both the learning curve and the retention of motor adaptation by training and testing subjects on a VMR under different amounts of visual latency between 25-600ms. Following a 240-trial baseline period, we trained 5 groups of subjects for 120 trials in a single movement direction on a ±30° VMR. Each group of 10-12 participants experienced a different visual latency (25, 50, 87, 300 or 600ms) present in both the baseline and training periods. The smallest latency we tested corresponded to the base delay inherent in our experimental setup, which we optimized down to 25ms using the combination of a low-input lag, fast-response 120Hz LCD display and latency-optimized experimental software with asynchronous graphics output. We also inserted 60 second waits every 30 training trials in order to probe the temporally-stable component of adaptation (on trials 10, 40, 70, and 100, see Figure 4.18A). Training was followed by another 60s wait period to let the temporally-labile component of adaptation wash out. After this period, subjects were tested for retention along the trained movement direction as well as generalization in 18 more different target directions (for a total of 19 movement directions from -135° to +135° about the trained direction in 15° increments, similar to Experiment 4.3, see Figure 4.14B). After testing, subjects experienced a 60-trial retraining block (with the visual latency present) and were then retested (Figure 4.18A).

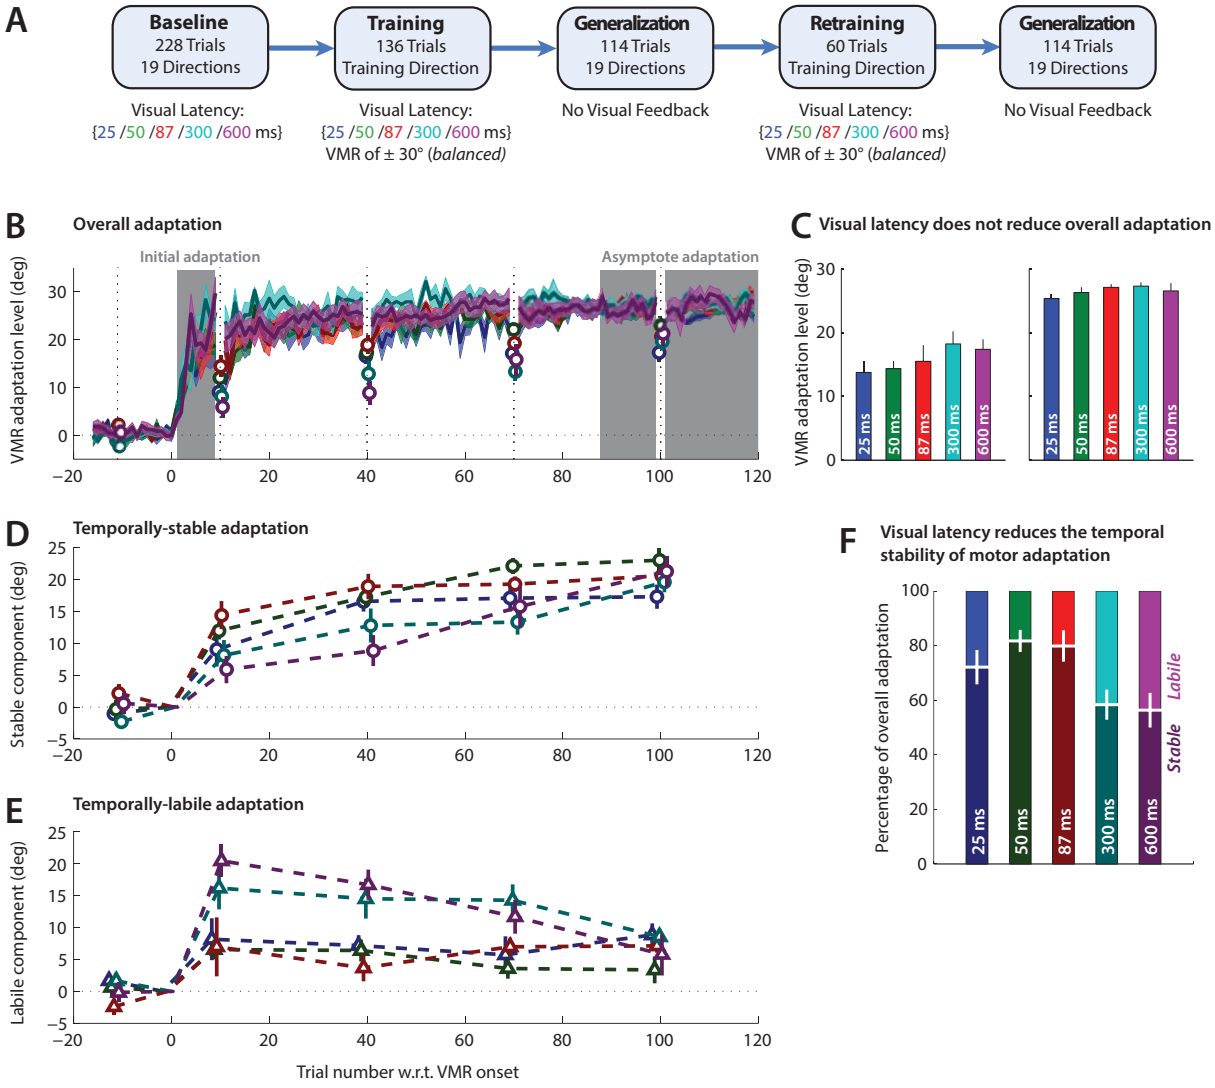


Figure 4.18: Overall learning is not reduced by increased visual latencies, but temporally-stable learning is reduced.

A: Schedule for experiment 4.4. Before training, subjects experienced a baseline period, whereby we probed movements to 19 different directions and at the same time introduced a visual latency between hand and cursor. Subsequently, subjects were trained on a 30° visuomotor rotation on a single movement direction, tested for generalization, retrained and retested.

B: Learning curves for the 5 latency groups, illustrating that both the speed and the asymptotic level of visuomotor adaptation were not compromised by visual latency. Specifically, neither early adaptation (trials 2-10) or asymptote adaptation (last 30 trials) displayed any clear trends with increased visual latency (C).

D: Temporally-stable adaptation during training, measured after 1-minute waits. In contrast to the robust overall adaptation even with increased visual latency shown in B, temporally-stable adaptation appears somewhat lower for

Figure 4.18 (continued)

the high-latency groups (300ms and 600ms), especially early during training. This hints that the temporal stability of motor memory might be reduced by latencies in visual sensory feedback during its learning.

E: Temporally-labile adaptation during training, measured as the difference before vs. after 1-minute waits, appears considerably increased for the high-latency groups (300ms and 600ms) especially early during training. This suggests that temporally-labile adaptation is robust to visual delays during training, and in fact compensates for the reduction in temporally-stable adaptation.

F: The percentage of stable as a fraction of overall adaptation is reduced with visual latency, showing how increased latencies shift the balance from temporally-stable to temporally-labile memories. Darker shades indicate temporally-stable adaptation, whereas lighter shades indicate temporally-labile adaptation. Errorbars indicate SEM.

Surprisingly, we found that visual feedback latency had little effect on the VMR learning curves, as neither early (trials 1-10) nor asymptotic learning (trials 100-120) showed any significant effect of visual feedback latency (13.8±1.7, 14.3±1.2, 15.5±2.5, 18.2±2.0, and 17.3±1.6° for early adaptation for the 25, 50, 87, 300 and 600ms groups, respectively, and 25.4±0.6, 26.3±0.8, 27.1±0.5, 27.3±0.6, and 26.5±1.2° for asymptotic learning, $p>0.39$ for the effect of latency in both cases using ANOVA, see Figure 4.18BC). However, temporally-stable adaptation, measured after the 60s waits, appeared to be reduced during early learning (trials 10 and 40) for the high-latency groups (300 and 600ms, cyan and magenta in Figure 4.18D). In contrast, the temporally-labile adaptation for these two groups was considerably higher around the same trials (Figure 4.18E). Correspondingly, the balance between temporally-stable and labile adaptation shifted, with 72.1±6.3, 81.7±4.0, and 79.9±5.8% of adaptation being temporally stable for the 25, 50, and 87ms groups, but only 58.4±5.5 and 56.3±6.3% for the 300 and 600ms groups ($p = 0.0002$ for the first three vs. the two high-latency groups, unpaired two-tailed t-test, Figure 4.18F). Our results thus suggest that high latencies in visual feedback during training specifically target the temporally-stable component of adaptation.

Visual latencies during training specifically target the locally-generalizing component of temporally-stable adaptation

Visual feedback latency also profoundly affected the shape of generalization of motor learning. All groups displayed distinct global and local components of the generalization across movement directions which we estimated by fitting the data from each subject on Equation 4.4 (Figure 4.19A). However, the amplitude of the local component, and thus the fraction of the generalization it accounted for, was specifically reduced by feedback latency (local components of 13.3±1.6°, 12.0±1.1°, 8.1±1.0°, 6.1±0.7°, and 4.6±1.0° for the 25, 50, 87, 300 and 600ms groups, respectively, $p < 10^{-5}$ for a linear fit to the data, , $p < 10^{-5}$ for the effect of latency group using ANOVA, Figure 4.19CD), whereas the global component was similarly affected (global components of 5.3±0.8, 5.4±1.0, 7.3±0.7, 5.1±0.7, and 5.0±0.8°, $p = 0.27$ for a linear fit to the data, $p = 0.22$ for the effect of latency group using ANOVA, Figure 4.19B). Our results amounted to reductions in the local component (with respect to the 25ms minimum latency group) of 39%, 54% and 66% for the 87, 300ms, and 600ms groups respectively (all $p<0.01$, unpaired two-tailed t-tests). Overall, our findings indicate that visual feedback latency determines both the retention of learning (Figure 4.19F) and its internal representation, reflected in the pattern of generalization (Figure 4.18A).

We note that, whereas the local component for the 600ms group was reduced by 66%, it was still significantly greater than zero ($p<10^{-3}$) and not significantly different than the local component for the 300ms group (4.6±1.0° for the 600ms group vs. 6.1±0.7° for the 300ms group, $p = 0.248$). There is thus a possibility that the local component might asymptote at a non-zero level when latency increases further, which would in turn imply that part of this local component is unaffected by increased visual latency (see Figure 4.19C for a plot of the local component against the corresponding visual feedback latency).

As a concluding note, we point out that the sensitivity of temporally-stable adaptation to very small delays that we observe, can be especially important when comparing results across different laboratories, each of which may utilize its own experiment setup with its own native feedback delay. For example, the LCD projectors used in some experimental setups can have latencies as large as 100ms.

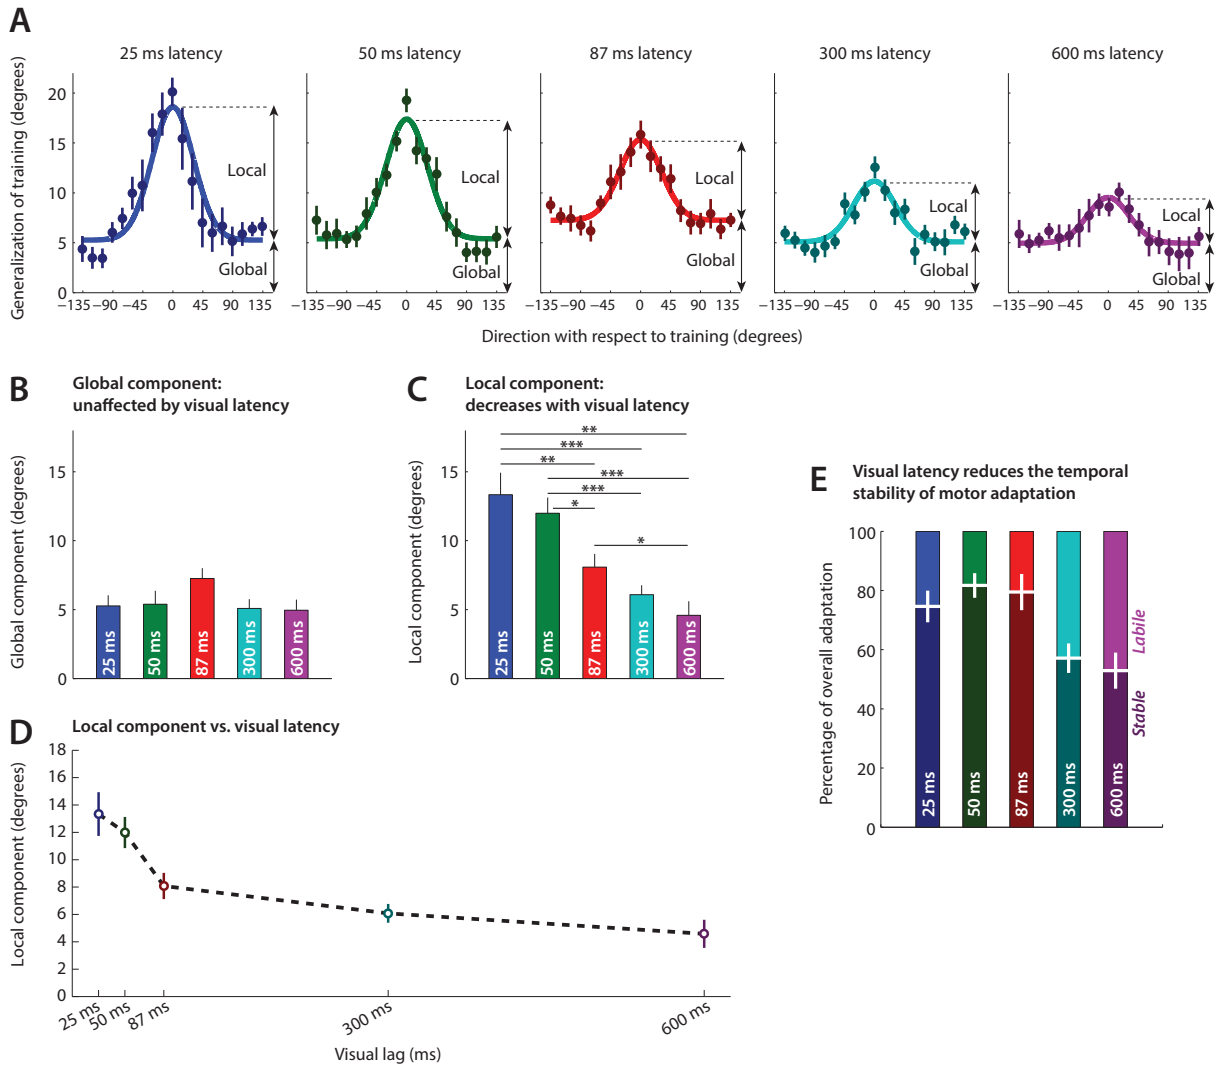


Figure 4.19: Visual latency during training suppresses the locally-generalizing component of temporally-stable adaptation.

A: Generalization functions for the 5 visual feedback latency groups in Experiment 4.4, illustrating how increased latency reduces the temporally-stable component of motor adaptation (as measured on the training direction, 0°) but appears to have little effect on generalization across the directions distal to training. This corresponds to a reduction in the local, but not the global component of generalization. Solid curves indicate Gaussian fits to the generalization curve (Equation 4.4), whereas errorbars indicate SEM.

B: Global component of generalization for each latency group, illustrating no effect of visual latency.

C: Local component of generalization (mean±SEM) for each latency group, illustrating its reduction with increased visual latency. * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

D: The local component of generalization (same data as in B) plotted against the corresponding visual latency. The reduction in the local component with latency appears to stabilize after 300ms.

Figure 4.19 (continued)

E: The percentage of stable as a fraction of overall adaptation is reduced with visual latency, showing how increased latencies shift the balance from temporally-stable to temporally-labile memories. Same as 4.18F, but including the temporally-stable data from the generalization blocks (along the previously trained, 0° direction). Darker shades indicate temporally-stable adaptation, whereas lighter shades indicate temporally-labile adaptation. Errorbars indicate SEM.

4.5 Discussion

Here, we employed a visuomotor adaptation task and found that this motor adaptation consists of two distinct components: a temporally-labile component, that decays fast with the passage of time (τ = 16.5s) and accounts for about 20-25% of asymptote adaptation, and a temporally-stable component which persists well beyond that timeframe and accounts for about 75-80% of adaptation (paragraph 4.4.1, Experiment 4.1). The amount of adaptation that decays, and thus the balance between temporally-stable and temporally-labile adaptation, could vary widely not only between different subjects, but also across different training sessions from the same subject – in fact, inter-individual vs. intra-individual differences account for comparable fractions of the variance in our measurements (40% vs. 30%, paragraph 4.4.1, Experiment 4.1). The speed of decay (time constant) did not show substantial differences across subjects or sessions. In contrast, in another experiment we found that the speed of decay for temporally-labile adaptation increased through rTMS-induced inhibition in the primary motor and posterior parietal cortices. This result suggests that the temporally-labile component is likely based on a network encompassing these two areas (paragraph 4.4.2, Experiment 4.2). Reanalysis of data from previous studies on the motor adaptation of cerebellar patients revealed that the temporally-stable component is likely based on a network involving the cerebellum (paragraph 4.4.3).

We proceed to find that the temporally-stable component of visuomotor adaptation can predict longer-term retention over the span of 24 hours (paragraph 4.4.4, Experiment 4.3), in line with previous work on force-field adaptation (Joiner and Smith, 2008; Sing et al., 2009; Sing, 2011). We investigate the pattern of how temporally-stable adaptation generalizes across different movement directions and found that it

consists of two components: one that generalizes broadly across the workspace, and one that generalizes narrowly around the trained direction. We find that both the broad and narrow generalization components of second-day retention can be predicted from the temporally-stable component, suggesting that temporally-stable adaptation, which can be measured one minute after training, represents the same memory as long-term retention. Finally, we find that temporally-stable visuomotor adaptation can be greatly reduced by visual feedback delays (paragraph 4.4.5, Experiment 4.4), suggesting that the formation temporally-stable motor memories is sensitive to the close temporal correspondence between action and feedback.

4.5.1 The stable-labile dichotomy and state-space models of motor adaptation

The dichotomy between temporally-stable vs. temporally-labile adaptation that we observe mirrors dichotomies between different processes in state-space models of motor adaptation. These models usually consist of a "fast" process – which learns fast from error but also retains poorly – and one (Smith et al., 2006) or multiple (Lee and Schweighofer, 2009) "slow" processes, which learn more slowly but show better retention. State-space models can explain a variety of observations in motor learning, like, among others, savings (Smith et al., 2006), anterograde interference (Sing and Smith, 2010), or even more general learning phenomena like the spacing effect (Ebbinghaus et al., 1913; Sing et al., 2009). Our findings are in line with Sing et al., 2009 which suggest that the fast adaptation process maps onto the temporally-labile component, and, correspondingly, that the slow adaptation process maps onto the temporally-stable component. Specifically, we find slower acquisition for the strongly-retaining temporally-stable adaptation, but faster early acquisition for the weakly-retaining temporally-labile adaptation (e.g. Figure 4.18C-E). For example, the timecourse of temporally-labile adaptation during learning in Experiment 4.4 (Figure 4.18E for the high-latency groups where labile adaptation is magnified), is non-monotonic displaying an early increase before asymptoting to lower levels, in contrast with the monotonic temporally-stable adaptation (Figure 4.18D). These shapes timecourses are predicted by the interactions of a fast and slow process (Smith et al., 2006; Sing et al., 2009). Following (Sing et al.,

2009), our results also support the idea that the fast-learning process decays passively with time, whereas the slow-learning, stable process decays only with trial-to-trial experience.

4.5.2 Time-dependent vs. experience-dependent decay

Measuring decay within a single session suggests that the temporally-stable component does not experience time-dependent decay, as retention remains essentially unchanged between 1 and 2 minutes (section 4.4.1) and even between 2 and 20 minutes (Sing et al., 2009). Yet, when measured after 24 hours (see paragraph 4.4.4), the remaining memory is only about 52% of its value on the day before, showing that this temporally-stable component underwent further decay during that time. A similar partial retention of the temporally-stable component was found in force-field adaptation (Joiner and Smith, 2008; Sing et al., 2009). Why would that be? One explanation could be that what we deem as a temporally-stable component is, instead, another temporally-labile component that instead decays over a much longer time scale – with a time constant of days, instead of seconds (a decay of 50% over 24 hours would imply a τ of approximately 34-35 hours). However, another explanation would be that this further decay is driven by experience. Data from previous work on the decay of force-field adaptation shows that the temporally-stable component can decay actively when the same movements are produced in the absence of error (Criscimagna-Hemminger and Shadmehr, 2008; Sing et al., 2009). This decay, when error is absent, is also something predicted by linear state-space models of motor adaptation that assume a non-trivial (i.e. <1) trial-to-trial retention factor for the stable/slow component (Smith et al., 2006; Joiner and Smith, 2008; Sing et al., 2009; Shadmehr et al., 2010a). It is very likely that, the multitude of reaching movements produced by the subjects during the 24 hours between experiment sessions, which are free of the visuomotor or force-field perturbations experienced in the lab, lead to a partial unlearning and/or experience-driven decay of adaptation. It would be interesting to investigate whether this mechanism of partial deadaptation can explain the differences between the decay of the global and local components: for example, a stronger decay in the local component might be due to a greater amount of interfering natural motion happening along the trained direction, which in our study was the forward one.

4.5.3 Implications for rehabilitation

The idea that the temporally-stable process, but not the temporally-labile one, can predict the formation of long-term motor memories, as supported by our results and previous work (Joiner and Smith, 2008; Sing et al., 2009), might help improve rehabilitation procedures where the goal is focused on long-term retention rather than short-term, within-session gains. Since both stable and labile processes are fed from the same motor error they both act to reduce, they effectively act to inhibit each other (Smith et al., 2006; Sing et al., 2009). Thus, a disruption of the temporally-labile process will lead to an increase in the temporally-stable process, and thus in stronger long-term retention. This disruption could result from the manipulation of inter-trial intervals, the increase of which allows temporally-labile adaptation to decay (Sing et al., 2009). The rTMS-induced disruption of areas involved in the temporally-labile process that we observe here can also be used to that effect, without the need to increase inter-trial intervals and thus extend the training session. Further research could investigate how the temporary disruption of brain areas involved in temporally-labile adaptation, like the M1 and PPC, can be thus employed to maximize long-term retention.

4.5.4 Possible reasons on why visual latencies reduce the temporal stability of motor adaptation

Our finding that the temporally-stable component adaptation is reduced by visual latencies during training might relate to the idea that delays in feedback about an action might cause a reduction in the intrinsic reward associated with completing that action. Here, an increased delay in sensory feedback – which will cause a corresponding delay between the time an action would appear completed and the time the corresponding motor command had been issued – might temporally discount the reward associated with completing the motion (Shadmehr et al., 2010b; Haith et al., 2012). It has been suggested that reward can cause the reinforcement of a motor plan, enhancing the stability of the memory associated with that plan (Shmuelof et al., 2012). Thus, we might expect that visual delays would reduce the stability of motor

memory, specifically targeting the temporally-stable, but not the temporally-labile, component of adaptation.

4.5.5 Reverse hierarchy theory and the robustness of the globally-generalizing component to increased visual latencies

Our finding that the globally-generalizing adaptation component is more robust against increased visual latencies, compared to the locally-generalizing component, might be explained by the reverse hierarchy theory. This theory posits that high-level components of learning are easier to learn and generalize widely; whereas low-level components, which are harder to learn, generalize narrowly (Ahissar and Hochstein, 1997, 2004; Ahissar et al., 2009). The reverse hierarchy theory was originally based on the encoding of visual features in the brain, whereby the encoding of basic, low-level features of a visual scene, is very specific (i.e. generalizes narrowly), whereas the encoding of high-level features generalizes widely. Extending this idea to our experiment, we expect that the more basic, low-level components of motor learning – which are more likely to rely on a close temporal correspondence between hand motion and visual feedback – will generalize narrowly; in contrast, high-level components of motor learning – which should be robust to visual feedback delays – will generalize widely. This would in turn make local generalization more susceptible to visual feedback delays, but should affect global generalization much less, if at all – explaining our current results. This idea is further supported by the previous finding that globally-generalizing VMR adaptation is learned much faster than the locally-generalizing one (Brayanov, 2013).

4.6 References

Ahissar M, Hochstein S (1997) Task difficulty and the specificity of perceptual learning. Nature 387:401-406.

Ahissar M, Hochstein S (2004) The reverse hierarchy theory of visual perceptual learning. Trends in cognitive sciences 8:457-464.

Ahissar M, Nahum M, Nelken I, Hochstein S (2009) Reverse hierarchies and sensory learning. Philosophical transactions of the Royal Society of London Series B, Biological sciences 364:285-299.

Altmann EM, Gray WD (2002) Forgetting to remember: the functional relationship of decay and interference. Psychol Sci 13:27-33.

Arce F, Novick I, Mandelblat-Cerf Y, Vaadia E (2010) Neuronal correlates of memory formation in motor cortex after adaptation to force field. The Journal of neuroscience : the official journal of the Society for Neuroscience 30:9189-9198.

Brainard DH (1997) The psychophysics toolbox. Spatial vision 10:433-436.

Brayanov JB (2013) Internal Representations for the Generalization of Motor Memories.

Brayanov JB, Press DZ, Smith MA (2012) Motor memory is encoded as a gain-field combination of intrinsic and extrinsic action representations. The Journal of neuroscience : the official journal of the Society for Neuroscience 32:14951-14965.

Brown J (1958) Some tests of the decay theory of immediate memory. Quarterly Journal of Experimental Psychology 10:12-21.

Chen R, Classen J, Gerloff C, Celnik P, Wassermann EM, Hallett M, Cohen LG (1997) Depression of motor cortex excitability by low-frequency transcranial magnetic stimulation. Neurology 48:1398-1403.

Conrad R (1957) Decay theory of immediate memory.

Criscimagna-Hemminger SE, Shadmehr R (2008) Consolidation patterns of human motor memory. The Journal of neuroscience : the official journal of the Society for Neuroscience 28:9610-9618.

Criscimagna-Hemminger SE, Bastian AJ, Shadmehr R (2010) Size of error affects cerebellar contributions to motor learning. Journal of neurophysiology 103:2275-2284.

Cullen KE, Brooks JX (2014) Neural Correlates of Sensory Prediction Errors in Monkeys: Evidence for Internal Models of Voluntary Self-Motion in the Cerebellum. Cerebellum.

Davidson PR, Wolpert DM (2004) Scaling down motor memories: de-adaptation after motor learning. Neuroscience letters 370:102-107.

Della-Maggiore V, Malfait N, Ostry DJ, Paus T (2004) Stimulation of the posterior parietal cortex interferes with arm trajectory adjustments during the learning of new dynamics. The Journal of neuroscience : the official journal of the Society for Neuroscience 24:9971-9976.

Ebbinghaus H, Ruger HA, Bussenius CE (1913) Memory; a contribution to experimental psychology. New York,: Teachers College, Columbia University.

Filmer HL, Dux PE, Mattingley JB (2014) Applications of transcranial direct current stimulation for understanding brain function. Trends in neurosciences 37:742-753.

Galea JM, Celnik P (2009) Brain polarization enhances the formation and retention of motor memories. Journal of neurophysiology 102:294-301.

Gibo TL, Criscimagna-Hemminger SE, Okamura AM, Bastian AJ (2013) Cerebellar motor learning: are environment dynamics more important than error size? Journal of neurophysiology 110:322-333.

Hadipour-Niktarash A, Lee CK, Desmond JE, Shadmehr R (2007) Impairment of retention but not acquisition of a visuomotor skill through time-dependent disruption of primary motor cortex. The Journal of neuroscience : the official journal of the Society for Neuroscience 27:13413-13419.

Haith AM, Reppert TR, Shadmehr R (2012) Evidence for hyperbolic temporal discounting of reward in control of movements. The Journal of neuroscience : the official journal of the Society for Neuroscience 32:11727-11736.

Huang VS, Haith A, Mazzoni P, Krakauer JW (2011) Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. Neuron 70:787-801.

Imamizu H, Kuroda T, Miyauchi S, Yoshioka T, Kawato M (2003) Modular organization of internal models of tools in the human cerebellum. Proceedings of the National Academy of Sciences of the United States of America 100:5461-5466.

Imamizu H, Miyauchi S, Tamada T, Sasaki Y, Takino R, Putz B, Yoshioka T, Kawato M (2000) Human cerebellar activity reflecting an acquired internal model of a new tool. Nature 403:192-195.

Ito M (2008) Control of mental activities by internal models in the cerebellum. Nature reviews Neuroscience 9:304-313.

Joiner WM, Smith MA (2008) Long-term retention explained by a model of short-term learning in the adaptive control of reaching. Journal of neurophysiology 100:2948-2955.

Katayama M, Kawato M (1993) Virtual trajectory and stiffness ellipse during multijoint arm movement predicted by neural inverse models. Biological cybernetics 69:353-362.

Kawato M, Kuroda T, Imamizu H, Nakano E, Miyauchi S, Yoshioka T (2003) Internal forward models in the cerebellum: fMRI study on grip force and load force coupling. Progress in brain research 142:171-_188.

Kitago T, Ryan SL, Mazzoni P, Krakauer JW, Haith AM (2013) Unlearning versus savings in visuomotor adaptation: comparing effects of washout, passage of time, and removal of errors on motor memory. Frontiers in human neuroscience 7:307.

Kitazawa S, Yin PB (2002) Prism adaptation with delayed visual error signals in the monkey. Experimental brain research 144:258-261.

Kitazawa S, Kohno T, Uka T (1995) Effects of delayed visual information on the rate and amount of prism adaptation in the human. The Journal of neuroscience : the official journal of the Society for Neuroscience 15:7644-7652.

Krakauer JW, Shadmehr R (2006) Consolidation of motor memory. Trends in neurosciences 29:58-64.

Krakauer JW, Ghez C, Ghilardi MF (2005) Adaptation to visuomotor transformations: Consolidation, interference, and forgetting. Journal of Neuroscience 25:473-478.

Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. Journal of Neuroscience 20:8916-8924.

Landi SM, Baguear F, Della-Maggiore V (2011) One week of motor adaptation induces structural changes in primary motor cortex that predict long-term memory one year later. The Journal of neuroscience : the official journal of the Society for Neuroscience 31:11808-11813.

Lee JY, Schweighofer N (2009) Dual adaptation supports a parallel architecture of motor memory. J Neurosci 29:10396-10404.

Lewandowsky S, Oberauer K, Brown GD (2009) No temporal decay in verbal short-term memory. Trends in cognitive sciences 13:120-126.

Maeda F, Keenan JP, Tormos JM, Topka H, Pascual-Leone A (2000) Modulation of corticospinal excitability by repetitive transcranial magnetic stimulation. Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology 111:800-805.

Mallett KE (2014) Visual Feedback Latency Reduces the Retention of Visuomotor Learning and Alters Its Internal Representation. In. Cambridge, MA: Harvard University.

Marsaglia G (1965) Ratios of Normal Variables and Ratios of Sums of Uniform Variables. J Am Stat Assoc 60:193-204.

Miall RC, Christensen LO, Cain O, Stanley J (2007) Disruption of state estimation in the human lateral cerebellum. PLoS biology 5:e316.

Mottaghy FM, Krause BJ, Kemna LJ, Topper R, Tellmann L, Beu M, Pascual-Leone A, Muller-Gartner HW (2000) Modulation of the neuronal circuitry subserving working memory in healthy human subjects by repetitive transcranial magnetic stimulation. Neuroscience letters 280:167-170.

Muellbacher W, Ziemann U, Wissel J, Dang N, Kofler M, Facchini S, Boroojerdi B, Poewe W, Hallett M (2002) Early consolidation in human primary motor cortex. Nature 415:640-644.

Mulliken GH, Musallam S, Andersen RA (2008) Forward estimation of movement state in posterior parietal cortex. Proceedings of the National Academy of Sciences of the United States of America 105:8170-8177.

Reis J, Schambra HM, Cohen LG, Buch ER, Fritsch B, Zarahn E, Celnik PA, Krakauer JW (2009) Noninvasive cortical stimulation enhances motor skill acquisition over multiple days through an effect on consolidation. Proceedings of the National Academy of Sciences of the United States of America 106:1590-1595.

Reis J, Robertson EM, Krakauer JW, Rothwell J, Marshall L, Gerloff C, Wassermann EM, Pascual-Leone A, Hummel F, Celnik PA, Classen J, Floel A, Ziemann U, Paulus W, Siebner HR, Born J, Cohen LG (2008) Consensus: Can transcranial direct current stimulation and transcranial magnetic stimulation enhance motor learning and memory formation? Brain stimulation 1:363-369.

Richardson AG, Overduin SA, Valero-Cabre A, Padoa-Schioppa C, Pascual-Leone A, Bizzi E, Press DZ (2006) Disruption of primary motor cortex before learning impairs memory of movement dynamics. The Journal of neuroscience : the official journal of the Society for Neuroscience 26:12466-12470.

Ridderinkhof KR, van den Wildenberg WP, Segalowitz SJ, Carter CS (2004) Neurocognitive mechanisms of cognitive control: the role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. Brain and cognition 56:129-140.

Sanes JN, Dimitrov B, Hallett M (1990) Motor learning in patients with cerebellar dysfunction. Brain : a journal of neurology 113 (Pt 1):103-120.

Scheidt RA, Reinkensmeyer DJ, Conditt MA, Rymer WZ, Mussa-Ivaldi FA (2000) Persistence of motor adaptation during constrained, multi-joint, arm movements. Journal of neurophysiology 84:853-862.

Shadmehr R (2004) Generalization as a behavioral window to the neural mechanisms of learning internal models. Human movement science 23:543-568.

Shadmehr R, Smith MA, Krakauer JW (2010a) Error correction, sensory prediction, and adaptation in motor control. Annual review of neuroscience 33:89-108.

Shadmehr R, Orban de Xivry JJ, Xu-Wilson M, Shih TY (2010b) Temporal discounting of reward and the cost of time in motor control. The Journal of neuroscience : the official journal of the Society for Neuroscience 30:10507-10516.

Shmuelof L, Huang VS, Haith AM, Delnicki RJ, Mazzoni P, Krakauer JW (2012) Overcoming motor "forgetting" through reinforcement of learned actions. The Journal of neuroscience : the official journal of the Society for Neuroscience 32:14617-14621.

Sing GC-W (2011) Multiple Timescales and Motion-Dependence in Motor Learning and Memory: Harvard University.

Sing GC, Smith MA (2010) Reduction in learning rates associated with anterograde interference results from interactions between different timescales in motor adaptation. PLoS computational biology 6.

Sing GC, Najafi B, Adewuyi A, Smith MA (2009) A novel mechanism for the spacing effect: Competitive inhibition between adaptive processes can explain the increase in motor skill retention associated with prolonged inter-trial spacing. In: Advances in Computational Motor Control. Chicago, IL.

Smith MA, Shadmehr R (2005) Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. Journal of neurophysiology 93:2809-2821.

Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. PLoS Biol 4:e179.

Thoroughman KA, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. Nature 407:742-747.

Tseng YW, Diedrichsen J, Krakauer JW, Shadmehr R, Bastian AJ (2007) Sensory prediction errors drive cerebellum-dependent adaptation of reaching. Journal of neurophysiology 98:54-62.

Valero-Cabre A, Payne BR, Pascual-Leone A (2007) Opposite impact on 14C-2-deoxyglucose brain metabolism following patterns of high and low frequency repetitive transcranial magnetic stimulation in the posterior parietal cortex. Experimental brain research 176:603-615.

Vaswani PA, Shadmehr R (2013) Decay of Motor Memories in the Absence of Error. Journal of Neuroscience 33:7700-7709.

Wolpert DM, Miall RC, Kawato M (1998) Internal models in the cerebellum. Trends in cognitive sciences 2:338-347.

Chapter 5: Savings arises from faster acquisition of short-term memories rather than the reemergence of stable long-term memories

5.1 Summary

It has been widely observed that declarative memories and motor skills are easier to reacquire than learn from scratch. This advantage, known as savings, has been assumed to result from the establishment of stable long-term memories formed during initial training. In fact, the presence of savings has even been used as an indicator of whether a motor memory had consolidated. However, recent findings have demonstrated that learning rates can be systematically controlled, providing a possible alternative explanation for savings. Here we investigate whether savings in the human motor system results from the reemergence of latent stable memories, or the more rapid reacquisition of new short-term motor memories. We dissect adaptation into temporally-stable and temporally-labile components, and find no savings in the temporally-stable component which is more resistant to active and passive decay and forms the basis for long-term retention. In fact, surprisingly, we find that relearning for temporally-stable adaptation is even slower than initial learning. In contrast, we find that the temporally-labile component of learning displays savings at a level that not only makes up for slower temporally-stable readaptation, but overrides it to yield robust overall savings. We conclude that savings is restricted to the temporally-labile component of motor adaptation, and thus faster relearning results not from the reemergence of stable long-term memories, but instead from the more rapid acquisition of short-term memories.

5.2 Introduction

Declarative and procedural memories alike are easier to reacquire than learn from scratch. This advantage, known as savings, was first formally defined in Hermann Ebbinghaus's seminal work (Ebbinghaus et al., 1913), in which he observed that relearning a forgotten list of words is faster than learning a completely novel list. Savings has since been demonstrated in a plethora of different paradigms, ranging from cognitive tasks in humans (Nelson, 1985; MacLeod, 1988), operant conditioning in animals (Frey and Ross, 1968; Napier et al., 1992; Medina et al., 2001), and motor tasks in humans including saccade adaptation (Kojima et al., 2004), force-field adaptation (Brashers-Krug et al., 1996; Shadmehr and Brashers-Krug, 1997; Smith et al., 2006), adaptation to a visuomotor rotation (VMR) (Krakauer et al., 2000a; Krakauer et al., 2005a; Krakauer and Shadmehr, 2006; Zarahn et al., 2008; Huang et al., 2011; Kitago et al., 2013) and split-treadmill adaptation (Malone et al., 2011; Mawase et al., 2014)

Previous research generally maintained that savings results from the recall of strong previous memories. In fact, the presence or lack thereof of savings itself is often taken as a litmus test for whether a previously trained memory has been consolidated (Krakauer and Shadmehr, 2006; Richardson et al., 2006; Landi et al., 2011). In line with this general approach, savings was explained as the result of the unmasking of a slower-learning, strong-retention process in a multirate learning model (Smith et al., 2006); context- or relevance- based switching between such multiple slow processes, each specific to a different memory (Lee and Schweighofer, 2009; Berniker and Kording, 2011; Pekny et al., 2011); reverting to the memory of a previous motor plan, which was reinforced by success (Shmuelof et al., 2012) or even mere repetition (Huang et al., 2011). All these different explanations focus on savings as the manifestation of a latent, consolidated motor memory, which is robust to both interference and the passage of time.

Recent work, however, found that the learning rates can be modulated by factors such as the amount of task-relevant variability present before learning (Wu et al., 2014), the balance between sensory

uncertainty and uncertainty about state estimation (Kording et al., 2004; Kording and Wolpert, 2004), prior exposure to perturbations characterized by the same covariance structure between learning parameters (Braun et al., 2009; Turnham et al., 2012), or the trial-to-trial consistency of the learning environment (Yousif and Diedrichsen, 2012; Gonzalez Castro et al., 2014). Highly consistent environments, whereby perturbations tend to persist from one trial to the next and thus confer more predictability to the imposed perturbation, can lead to strong (up to 3x) increases in learning rates. This is a critical finding as far as the study of savings is concerned, since the adaptation paradigms used to study savings usually consist of the same perturbation being active for a large (usually 60-120) number of trials which results in high environmental consistency, which would in turn lead to strongly increased learning rates. This can lead to savings by enabling the faster re-acquisition of a short-term memory of learning, instead of relying on the reemergence of a long-term, stable memory of this learning. In line with this idea, it was suggested that increases in learning rates account for savings within the framework of a linear, time-varying model (Zarahn et al., 2008; Mawase et al., 2014).

Here we investigate the extent to which the increased capacity for faster relearning that characterizes savings is driven by a pull towards latent, previously consolidated memories, or a propensity to more rapidly form a memory similar to what was previously learned. Specifically, we train subjects in a visuomotor rotation paradigm and compare the speed of acquisition and temporal stability of both novel and retrained motor memories. We use time delays to separate adaptation into distinct temporally-stable and temporally-labile components, and first investigate how resistant each component is to active washout when the environment returns to the…. Finally, we estimate and compare savings separately for temporally-stable and temporally-labile adaptation component.

5.3 Materials and Methods

A total of 80 subjects (38 men, age 21.6±4.2 years, 6 left-handed) participated in the studies described in this chapter (20 subjects in each one of Experiments 5.1, 5.2, 5.3 and 5.4). Participants were naïve with respect to the purpose of the experiments and provided informed consent in accordance with the policies of the Harvard University Institutional Review Board.

5.3.1 Apparatus

We used the same experimental setup as in Experiments 4.1-4.4 (Brayanov et al., 2012). Subjects sat in front of an apparatus consisting of a 200Hz digitizing tablet (Intuos 3) positioned below a 120Hz LCD monitor. During the experiment, subjects moved a custom-made handle, which contained a stylus, on top of the tablet allowing us to record hand position. Vision of the hand was occluded by the monitor and subjects instead observed their movement on the through a white cursor representing hand position.

5.3.2 Experiment protocol

Subjects made point-to-point arm reaching movements between the center of the screen and a target 9cm further away from their body. At the end of each movement, they would be rewarded with a bell sound if they had managed to reach the target fast enough. To get the bell sound, subjects had to move and stop at the target within 250ms. Training was isolated to the outward movement direction by removing visual feedback from return movements. Subjects took 1 minute breaks every 200 trials or less (see Figure 5.1), whereas additional 1 minute pauses were inserted every 30 or 40 trials. During these pauses, in contrast to breaks, subjects were required to keep their hand on the starting position and wait for the target to appear before moving.

After familiarizing themselves with the task during a 220-trial baseline block, subjects in Experiment 5.1 entered the main part of the session which consisted of three 80-trial training periods. The first and second training periods were separated by a 40-trial washout period, whereas the second and third training periods were separated by an 800-trial washout period. Experiment 5.2 had a similar schedule with the

difference that the 800-trial washout period came first (between the first and second training periods, see Figure 5.1). Experiment 5.3 consisted of data from the first two sessions of Experiment 4.1 (16 subjects), as well as data from 4 additional subjects who took part in a two-session version of Experiment 4.1. Experiment 5.4 was similar to 5.2 with the difference that the second training episode (after the 800-trial washout period presented after initial training) had a VMR opposite to the one initially trained.

5.3.3 Data analysis

Statistical comparisons

We performed single-sided paired t-tests across subjects to assess the presence of savings for overall, stable and labile adaptation. For all other statistical comparisons two-sided paired t-tests across subjects were implemented, with the exception of the comparisons involving the estimation of washout time constants in Figure 5.2A. For comparison of washout time constants we utilized a bootstrapping procedure (see below) instead of comparing fits to individual subject data, because the high noise in these individual data leads to low confidence about the corresponding individual parameters.

Data inclusion criteria

We performed outlier rejection on the learning curves of each experiment. Specifically, for each trial, we rejected as gross outliers adaptation levels more than 4 inter-quartile ranges (IQRs) away from the subject median. This resulted in the rejection of only 0.2% of trials. Additionally, we removed the first 4 trials after a 1-minute wait from the analysis of overall washout curves, in order to exclude the aftereffects of those 1-minute waits. As a result, washout analysis did not use 10.5% of the trials during the long, 800-trial washout period.

Estimation of washout time constants

To compare the time constants for washout for both temporally-stable and overall adaptation (Figure 5.2A), we focused our analysis on the part of the overall washout data that started from a similar level to the adaptation level of the first temporally-stable measurement trial (which was 10.3±1.1° for Experiment

5.1, and 6.7 ±0.9° for Experiment 5.2). This meant excluding the first 2 washout trials from Experiments 5.1 and 5.2.

To estimate the confidence intervals associated with the above time constants of washout, we utilized a bootstrapping procedure (Efron and Tibshirani, 1994). Specifically, for each one of 10000 bootstrap iterations, we randomly sampled N=20 subjects from each group, and fit their average data with a single-exponential fit (Equation 5.1):

(Equation 5.1) $$x = a + be^{-t/\tau}$$

As the distribution of the fit parameters across bootstrap iterations approaches the distribution of the mean parameters (Efron and Tibshirani, 1994), we were able to use the standard deviation across the parameter values obtained for each iteration as an estimate of the standard error (SEM) of these parameters.

Normalization of adaptation data

To systematically quantify savings and take into account the systematically different baselines between post-long washout vs. post-short washout relearning, and the different baselines between temporally-stable, temporally-labile and overall adaptation, we subtracted baseline adaptation, $x_{baseline}$, and normalized each learning curve x by the distance between baseline and the ideal adaptation level of 30 degrees (equation 5.2). For overall adaptation, the baseline adaptation level was defined as the average of the last 5 trials before training onset, whereas for stable adaptation, the baseline adaptation level was taken from the last after-wait trial, 10 trials before the onset of training (Figure 5.3B).

(Equation 5.2) $$x_{norm} = \frac{x - x_{baseline}}{30 - x_{baseline}} \times 100\%$$

Finally, savings for each type of adaptation were taken to be the % difference in adaptation between initial training and retraining for the same training trial (equation 5.3).

(Equation 5.3) $$s = x_{norm}(retrain) - x_{norm}(initial\ training)\ (\%)$$

5.4 Results

We designed a set of experiments that would allow us to (a) study the speed of washout for temporally stable vs. overall adaptation and (b) investigate whether savings is specifically driven by one or both of these adaptation components. We recruited N = 40 subjects and trained them in a 30° visuomotor rotation (VMR) (Krakauer et al., 2000b; Krakauer et al., 2005b; Zarahn et al., 2008; Huang et al., 2011; Brayanov et al., 2012)(Figure 5.1AB). After 80 trials of initial training, subjects were tested for savings after a short, 40-trial or a long, 800-trial washout period. Subjects were split in two experiments, Experiment 5.1 (N=20) and Experiment 5.2 (N=20), which were identical with the only difference that, in Experiment 5.1, the short, 40-trial washout period was presented first and the 800-trial washout period second, whereas in Experiment 5.2 this order was flipped (Figure 5.1C, for a detailed description see Materials and Methods).

We measured temporally-stable adaptation every 30 or 40 trials by inserting one-minute breaks which would allow for most of the temporally-labile adaptation to decay. Previous data, presented in Chapter 2, suggest that temporally-labile VMR adaptation decays rapidly with time with a time constant τ =15-20s. Thus, the one-minute waits we impose here amount to 3τ-4τ, which sufficient to eliminate 95-98% of labile adaptation and thus effectively isolate the temporally stable adaptation.

5.4.1 Stable adaptation washes out much slower than overall adaptation

We first examined the speed of washout for both overall and stable adaptation during the long, 800-trial washout periods in both Experiments 5.1 and 5.2. During these periods, one-minute wait times were experienced every 40 trials that allowed us to probe the washout curve for the temporally stable process immediately after each wait, as well as the washout curve of the overall adaptation. We found that overall adaptation decayed fast: in both Experiments 5.1 and 2, the overall adaptation level decayed very close to zero only 30-50 trials into the washout period (Figure 5.2A), in agreement to previous work (Zarahn et al., 2008). In fact, very early into the washout period (trials 16-25), overall adaptation already dropped

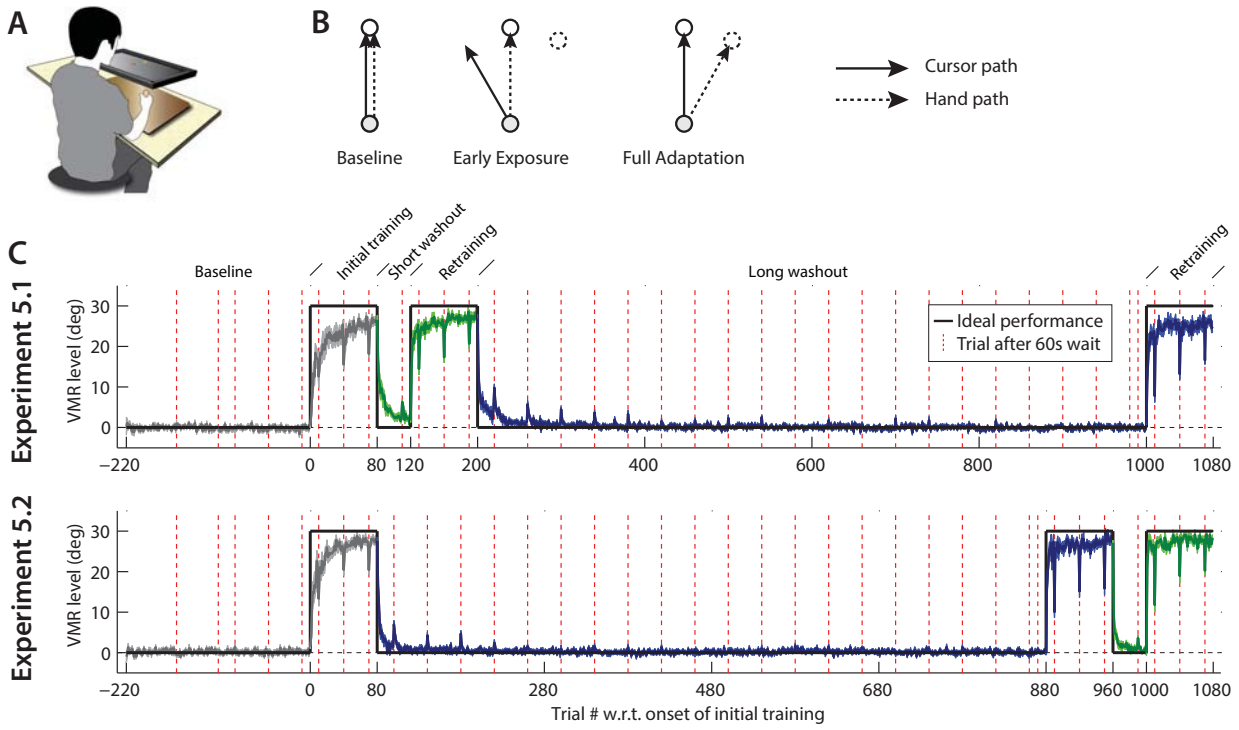


Figure 5.1. Experiment setup and training paradigm.

A: Experiment setup. Subjects made point-to-point reaching movements on a tablet and observed visual feedback on a screen mounted above.

B: Visuomotor rotation paradigm. During baseline (left), the cursor follows the hand motion. During training, the cursor trajectory is displaced by 30° about the hand movement direction (in this example, counter-clockwise), resulting in a 30° error for the initial exposure (middle). Full adaptation will have the hand move in a direction 30° clockwise about the target, so that the cursor moves directly towards the target (right).

C: Top: experiment schedule and overview of data from Experiment 5.1. The experiment consisted of three main parts: a baseline period followed by the initial 80-trial VMR training (average adaptation level shown in gray); a short, 40-trial washout period followed by retraining (green); and a long, 800-trial washout period followed by another 80-trial retraining session (blue). Red dashed lines indicate trials before which 1-minute waits were inserted, in order to isolate the temporally-stable component of adaptation. Note how during the washout periods adaptation peaks during these trials, illustrating a slower washout for temporally-stable vs. temporally-labile adaptation. Bottom: same but for Experiment 5.2, where the long, 800-trial washout period happened first.

below 10% of asymptote learning, and, for trials 51-150 of washout, it dropped below 3% (Figure 5.3B). To systematically quantify the speed of washout of the overall adaptation, we fit single exponentials that revealed time constants of 26.3±6.7 trials for Experiment 5.1 (mean ± SEM estimated through bootstrapping, see Materials and Methods) and 5.7±2.9 trials for Experiment 5.2 (Figure 5.2A).

In contrast, stable adaptation following one-minute waits washed out much more slowly, remaining significantly ($p<0.05$) above zero for at least 140 trials (circles in Figure 5.2A). Early into the washout period (trials 16-25), temporally-stable adaptation was still around 35-50% of asymptote learning, and, for trials 51-150 of washout, it was still at 15-25% of asymptote learning (Figure 5.2A). In line with this much stronger of stable adaptation compared to overall adaptation, exponential fits to the stable washout curve revealed time constants of 105.2±29.7 trials (mean ± SEM) for Experiment 5.1 and 103.2±25.5 trials for Experiment 5.2, more than 3-4 times greater than the time constants for overall washout (Figure 5.2A).

5.4.2 Temporally-stable adaptation leads to a long-term memory

Does the temporally-stable component of adaptation represent the formation of a long-term memory, or does it also decay with time, albeit with a much longer timescale compared to the temporally-labile component described above? To investigate this question, we analyzed data from a different experiment (Experiment 5.3, which comprised of the data from the 16 subjects of Experiment 4.1, as well as data from 4 more subjects, see Materials and Methods) in which participants received prolonged training (704-872 trials) on a 30° VMR and then experienced a 200-trial washout period before coming the next day to receive more washout trials. Our hypothesis was that, if the temporally-stable component were a long-term memory, what is measured on day 1 as temporally-stable adaptation right before the end of the session would be to a significant extent preserved on day 2.

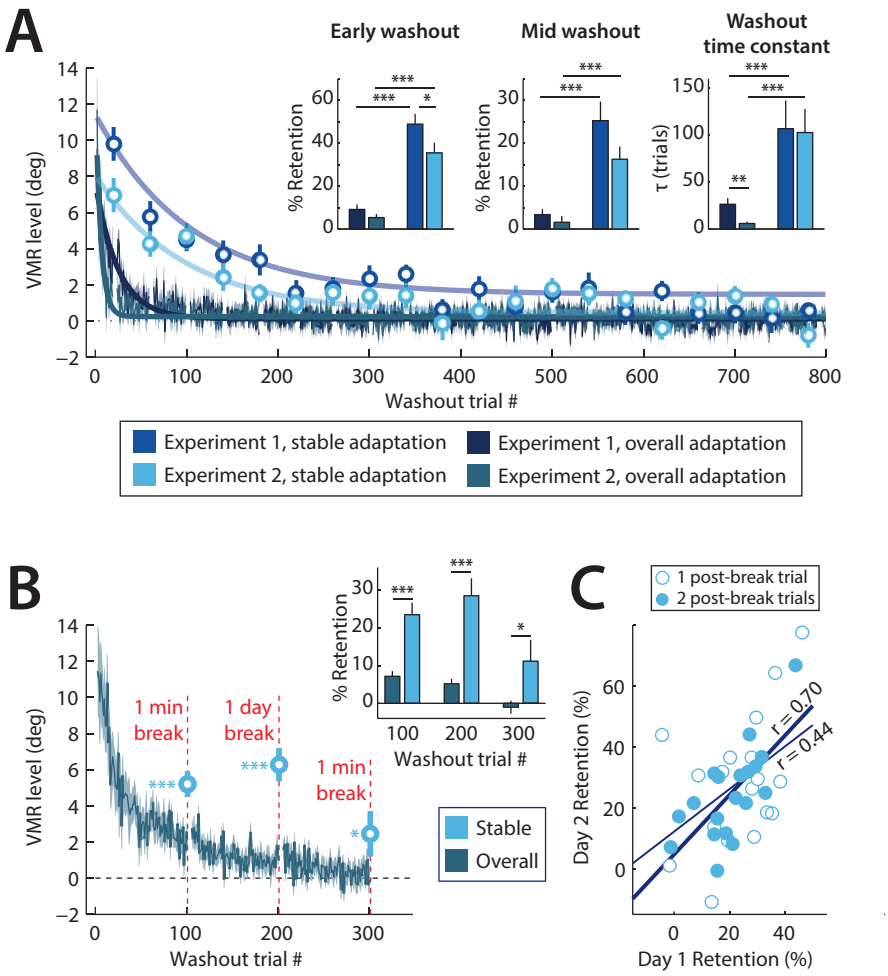


Figure 5.2. Temporally-stable adaptation washes out slowly and exhibits long-term retention.

A: Washout curves for the overall adaptation (darker shades) and temporally-stable adaptation (lighter shades - circles) for both Experiment 5.1 and Experiment 5.2, illustrating the contrast between the fast washout of the overall adaptation and the slower washout of stable adaptation. The thick solid lines indicate exponential fits, whereas errorbars and shading indicate SEM. Inset: Three different metrics quantifying the speed of washout for overall vs. stable adaptation. Left panel: percent retention of asymptote adaptation for trials 16-25 of washout. Center panel: percent retention of asymptote adaptation for trials 51-150 of washout. Both left and center panels show significantly higher retention for the stable adaptation compared to overall adaptation. Right panel: time constants of washout estimated through bootstrapping. Errorbars indicate SEM. *p<0.05, **p<0.01, ***p<0.001.

B: Washout data from Experiment 5.3. Note that, trials 1-200 are part of the first session, whereas trials 201-300 are part of the second session which occurred at least one day later. Note how temporally-stable adaptation is still robust on day 2 (trial 201) as it was on day 1 (trial 101), showing that temporally-stable adaptation leads to the formation of a long-term memory. Inset: % of adaptation that is retained around each washout trial for both stable and overall adaptation, showing robust retention of the temporally-stable component for day2 (trial 201) compared to day 1

Figure 5.2 (continued)

(trial101). To calculate % retention, we compared adaptation to the corresponding stable and overall adaptation levels at asymptote during training.

C: Population analysis of the temporally-stable day 1 component (trial 101) and the retained adaptation on day 2 (trial 201). We found that the strength of the temporally-stable component on day 1 (trial 101) correlated with the retained temporally-stable component on the beginning of day 2 across subjects (trial 201, $r = 0.44$, $p = 0.07$, open circles, fit shown in the thin blue line). When we improved the signal-to-noise ratio of these temporally-stable adaptation estimates by averaging the first 2 trials after a break, we found a more clear and significant relationship ($r = 0.70$, $p = 0.0009$, filled circles, fit shown in the thicker blue line).

We found significant adaptation remaining at the beginning of the next day (6.27±0.91°, trial 201 of washout in Figure 5.2B, $p<10^{-5}$) which was virtually unchanged ($p=0.75$) from the first-day temporally-stable adaptation as measured in the middle of the day 1 washout period (5.84±0.93°, trial 101 of washout in Figure 5.2B). Moreover, the temporally-stable component was still high a further 100 washout trials later during day 2 (3.52±1.59°, trial 301,, $p=0.02$). In other words, increasing the length of one of the breaks after which temporally-stable adaptation is measured, from 1 minute to 1 day, did not appear to result in further decay for the temporally-stable component of adaptation, nor increase the speed of its washout.

Further analysis of Experiment 5.3 corroborated the same general pattern of washout found in Experiments 5.1 and 5.2: a fast washout for overall adaptation and a considerably slower washout for the temporally-stable process. Specifically, temporally-stable adaptation was equal to 25.1±3.5%, 28.5±4.6%,, and 15.9±7.2% of asymptote learning on the three post-break trials (101, 201 and 301, correspondingly) in contrast to overall adaptation which was only equal to 7.1±1.5%, 5.2±1.3%, and -1.0±1.7% of asymptote [as measured during the corresponding 2 preceding trials] (inset in Figure 5.2B).

Together, these results suggest that the temporally-stable component of VMR adaptation reflects the formation of a long-term memory, extending our analysis of Experiment 4.3 in Chapter 4, whereby we found temporally-stable adaptation to predict second-day retention. Here, we also find that the part of temporally-stable adaptation that is resistant to active washout is preserved, almost without further decay,

on the next day, further establishing the role of temporally-stable adaptation as the gateway to long-term memories. Analysis of the temporally-stable component across individuals collaborates these results, with the strength of the temporally-stable component on day 1 (trial 101) correlating with the retained temporally-stable component on the beginning of day 2 across subjects (trial 201, $r = 0.44$, $p = 0.07$, Figure 5.3C). When we improved the signal-to-noise ratio of these temporally-stable adaptation estimates by averaging the first 2 trials after a break, we found a more clear and significant relationship ($r = 0.70$, $p = 0.0009$).

5.4.3 Overall adaptation displays savings even after complete washout of temporally-stable adaptation

A consequence of the washout time constants we found for overall vs. stable adaptation in the previous paragraph is that, whereas the 40 washout trials utilized in (Zarahn et al., 2008) will be enough for near-complete washout of overall adaptation, they will allow for a considerable amount of stable adaptation to remain. This raises the possibility that the savings observed after 40 washout trials might be driven by this residual stable adaptation.

To examine whether previously observed savings were driven by incomplete washout of the temporally stable process of adaptation, and to test for savings separately for the overall vs. stable processes, we analyzed our relearning curves in Experiments 1 and 2 after both short (40-trial) and long (800-trial) washout periods and compared them with the corresponding initial learning curves (green, blue and gray, respectively, Figure 5.3). For each training period, we probed the level of the temporally-stable adaptation using one-minute wait times 10 trials before and 10, 40 and 70 trials after the onset of training.

We found that, in Experiment 5.1, the relearning of overall adaptation after both the short and long washout periods was similarly fast (average adaptation of 19.8±1.1° and 18.5±1.4° for the first 10 learning trials, respectively, $p>0.11$), and faster than initial learning (11.4±1.4°, both $p<10^{-3}$, Figure 5.3A). These savings were similar regardless of the fact that, after the 40-trial washout period there was a strong residual stable component whereas after the 800-trial washout period there was none (Figure 5.3C).

We obtained similar savings in Experiment 5.2 where the order of the short and long washout periods was reversed (average adaptation for the first 10 learning trials: 23.2±0.9° and 21.1±1.2° after the short and long washout vs. 14.0±1.5° for initial learning, both $p<10^{-5}$, Figure 5.3F) indicating that the pattern of savings observed was not dependent on the order the short and long washout periods were presented.

To systematically quantify savings and take into account the systematically different baselines between post-long washout vs. post-short washout relearning, we subtracted baseline adaptation, $x_{baseline}$, and normalized each learning curve x by the distance between baseline and the ideal adaptation level of 30 degrees (equation 5.1). For overall adaptation, the baseline adaptation level was defined as the average of the last 5 trials before training onset, whereas for stable adaptation, the baseline adaptation level was taken from the last after-wait trial, 10 trials before the onset of training (Figure 3B).

(Equation 5.5) $$x_{norm} = 100\frac{x-x_{baseline}}{30-x_{baseline}}[\%]$$

The result of this normalization is shown in Figure 3D. From that, we calculated the raw savings for each retraining trial as the difference between the normalized adaptation at that trial and the normalized adaptation at the corresponding trial during the initial training period (Equation 5.2, Figure 5.3E,J).

(Equation 5.6) $$s_{raw} = x_{norm,retrain} - x_{norm,initial}$$

Moreover, in order to better compare overall with stable savings, we focused our analysis on the trials right before and right after the 60s wait times. We found significant savings for the overall adaptation during early retraining (trial 10, right before the first 60s delay after the onset of training) for both Experiment 5.1 (26.7±7.3% after the short and 23.6±8.3% after the long washout, both $p<0.01$) and Experiment 5.2 (18.2±5.9% after the short and 23.5±7.8% after the long washout, both $p<0.01$).

Figure 5.3. Savings is driven by the temporally-labile adaptation process after both a long and a short washout.

A: Learning curves for Experiment 5.1, showing both overall adaptation and stable adaptation (circles). Errorbars and shading indicate SEM. Overall adaptation is faster for retraining after both a 40-trial washout period (green) and an 800-trial washout period (blue) when compared to initial adaptation (grey). However, there is no clear similar increase in how fast the stable process readapts.

B: Baselines for overall, stable and labile adaptation for initial training (gray) and training after 40 or 800 washout trials (green and blue, respectively). The data show that retraining periods after only 40 washout trials start with significant nonzero baselines. *$p<0.05$, **$p<0.01$, ***$p<0.001$.

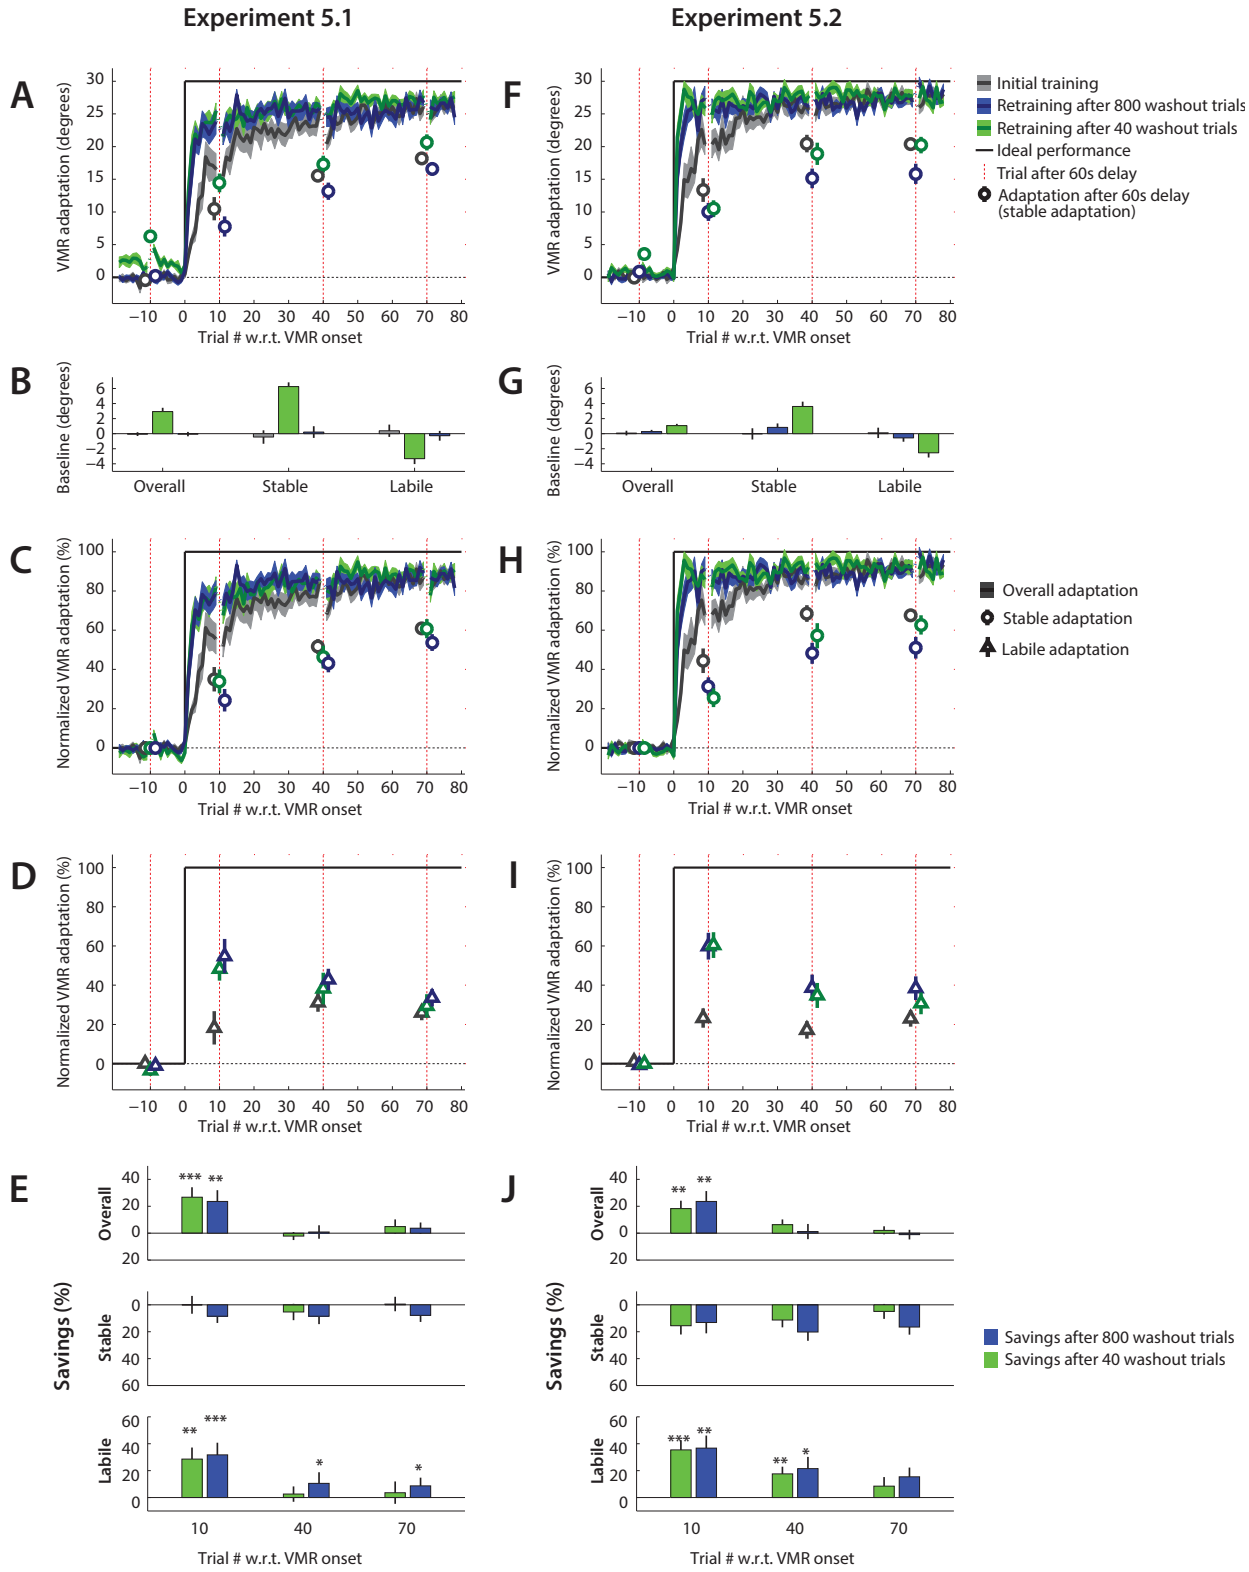
C: Normalized learning curves for overall and temporally-stable adaptation. The raw learning curves for both the overall and stable process shown in A were normalized to take into account the corresponding amounts of baseline learning shown in B. The normalized data show that whereas the overall adaptation curve is faster for retraining after both washout periods, stable adaptation is not higher than during initial adaptation.

D: Normalized learning curve for temporally-labile adaptation. Temporally-labile adaptation after either washout period is higher than initial adaptation, especially for early (trial 10) retraining where it is between 2-3 times higher.

E: Savings for the overall, temporally-stable, and temporally-labile adaptation, illustrating substantial and significant savings early during training (trial 10) for both overall and temporally-labile adaptation, but not for temporally-stable adaptation. This shows that savings are driven by the temporally-labile process of adaptation.

F-H: Same for A-E but for Experiment 5.2 (where the 800-trial washout period occurred first), showing the same result as Experiment 5.1: Savings is driven by the temporally-labile process, with the temporally-stable process displaying no savings.

Figure 5.3 (continued)



5.4.4 Stable adaptation does not display savings

In contrast, stable adaptation was not faster during retraining when compared to initial training (compare the green and blue circles vs. the gray ones in Figure 5.3A,F). Specifically, stable adaptation 10 trials after the onset of training for Experiment 5.1 was 14.5±1.4° after a 40-trial washout and 7.8±1.5° after an 800-trial washout compared to 10.5±1.8° for initial training. For Experiment 5.2, stable adaptation at trial 10 after training onset was 10.5±1.3° after a 40-trial washout and 10.0±1.4° after an 800-trial washout compared to 13.3±1.8° for initial training. Because stable adaptation after the short, 40-trial washout periods began with a significant non-zero baseline (Figure 5.3B,G), we normalized stable adaptation as we did with overall adaptation in the previous paragraph (Equation 5.1). Our normalized data show that readaptation of the stable process was nominally *slower* in both experiments (Figure 5.3C,H) which amounts to no savings (Figure 5.3E,J, trial 10 savings: Experiment 5.1, -0.0±6.6% after the 40-trial washout and -8.5±5.0% after the 800-trial washout; Experiment 5.2, -13.2±8.0% after the 40-trial washout and -15.6±6.5% after the 800-trial washout, $p>0.5$ for savings in all cases). We thus conclude that, whereas relearning after a washout period displays significant overall savings, this does not amount to any savings for the temporally-stable adaptation component. This illustrates that the temporally stable adaptation we observed is not simply the consolidation of a more temporally-labile form of adaptation (as, if that was the case, stable savings would scale with overall savings) but a distinct adaptation component. This observation in turn enabled us to express overall adaptation as the sum of a temporally-stable and a temporally-labile adaptation component.

5.4.5 Savings is driven exclusively by temporally-labile adaptation

To calculate savings for the temporally-labile component of adaptation, we first estimated the level of temporally labile adaptation as the difference between overall and stable adaptation. Specifically, from the overall adaptation measured in trials immediately preceding one-minute waits we subtracted the stable adaptation as measured in the trial right after the one-minute wait. The levels of labile adaptation estimated through these differences – and normalized using Equation 5.1 – are shown in Figure 5.3D,I .

Note that, for relearning after the short, 40-trial washout periods, the baseline used in this normalization for the labile process was negative, resulting from the significantly positive baseline for the stable process (Figure 5.3B,G). We found that labile learning was much faster after both the short and long washout periods in both cases (Trial 10 normalized adaptation – Experiment 5.1: 48.4±6.0% after the 40-trial and 54.8±8.8% after the 800-trial washout period, vs. 18.3±8.5% for initial adaptation; Experiment 5.2: 60.5±6.6% and 59.9±6.8% vs. 23.2±4.9% for initial adaptation, Figure 5.3D,I). When we quantified labile savings as the difference in normalized labile adaptation at each trial between the initial learning and relearning periods, like the method we used to quantify stable savings above, we found clear early (trial 10 of learning) savings of 28.5±8.6% after the 40-trial washout and 31.6±9.1% after the 800-trial washout, both $p<0.01$ for Experiment 5.1 and 35.4±7.2 % after the 40-trial washout and 36.7±9.4% after the 800-trial washout, both $p<0.001$ for Experiment 5.2. In summary, savings was exclusive to temporally-labile adaptation instead of temporally-stable adaptation, indicating that savings arises from a capacity to faster re-acquire short-term memories rather than the manifestation of long-term memories.

We summarized our savings results by focusing on early (trial 10) savings and combining the data from both Experiments 5.1 and 5.2. As illustrated in Figure 5.4A, we found clear raw savings for the overall (22.4±4.2% and 23.4±4.8% after the short and long washout periods, respectively) and temporally-labile (30.9±6.2 and 33.5±6.2%) processes but not for the temporally-stable process (-6.9±5.5 and -9.9±5.4%). Moreover, temporally-labile savings were significantly higher than the corresponding temporally-stable savings ($p<0.005$ both after a long and a short washout period).

A drawback of our simple raw savings metric (Equation 5.6) is that it doesn't take into account the fact that, for subjects with faster initial adaptation, there is little room for further improvement, which would restrict the raw savings metric compared to subjects where initial adaptation is slower – even if the % increase in learning rates was similar in both cases. We thus further normalized raw savings by the maximum possible savings possible given initial adaptation (Equation 5.7):

(Equation 5.7) $$s_{norm} = \frac{x_{norm,retrain} - x_{norm,initial}}{100 - x_{norm,initial}}[\%] = \frac{s_{raw}}{100 - x_{norm,initial}}[\%]$$

Savings calculated under this improved metric are shown in Figure 5.4B, with the results paralleling the non-normalized version: savings for the overall (58.2±10.9% and 62.9±12.9% after the 40- and 800-trial washout periods) and temporally-labile (39.1±7.8% and 42.4±7.9%) processes but not for the temporally-stable process (-12.2±9.3% and -16.8±9.4%).We conclude that savings in VMR adaptation are driven by an increase in the learning of the temporally-labile process, whereas the learning of the temporally-stable process doesn't display any savings and might even slow down to an extent.

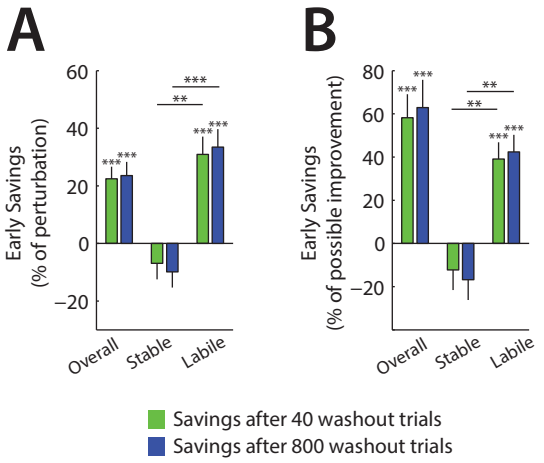


Figure 5.4. Summary of results showing that savings is driven by temporally-labile, but not temporally-stable adaptation.

A: Savings early during training (trial 10) from Experiments 5.1 and 5.2 combined, illustrating strong savings for stable and overall adaptation but not temporally-stable adaptation. Here, savings are calculated based on the simple metric in Equation 5.6. Errorbars indicate SEM. *p<0.05, **p<0.01, ***p<0.001.

B: Same as A, but with savings calculated based on the normalized metric in Equation 5.7, illustrating the same result – robust savings for stable and overall adaptation but not for labile adaptation.

5.4.6 Savings does not transfer to the learning of the opposite perturbation

Here we briefly investigate the mechanism behind the faster re-acquisition of temporally-labile adaptation observed in the previous sections. We will examine two main possibilities based on recent findings. First, environments that are consistent from one trial to the next can lead to increased learning rates and thus higher savings. Critically, this effect is greatly enhanced when the environment repeats the exact same perturbation (Gonzalez Castro et al., 2014) – also see chapter 6. Both conditions are present in the paradigms previously used to study savings: initial training consists of a large number of trials with the same specific perturbation, which in turn presents an environment characterized by both high consistency and repetition. We refer to this repetition mechanism as "perturbation-specific".

An alternative explanation is that this learning rate upregulation is not specific to the exact same perturbation, but to the errors experienced (Herzfeld et al., 2014). This "error-specific" mechanism similarly predicts savings in a train-washout-retrain paradigm like ours: the errors experienced during retraining are similar to the errors experienced during initial training. However, these two approaches make different predictions when, instead of retraining the same perturbation after washout, the opposite perturbation is trained instead. The perturbation-specific model predicts no savings, as the opposite perturbation is novel. In contrast, the error-specific model predicts savings as the errors experienced during the training of the opposite perturbation are very similar to the ones experienced at the beginning of washout after the initial training, when the initially trained perturbation is suddenly removed.

To test between these two ideas, we ran an additional experiment (Experiment 5.4) in which we tested for adaptation to the opposing perturbation after an 800-trial washout period. As with Experiments 5.1 and 5.2, we found no savings for temporally-stable adaptation (-4.6±7.5%, Figure 5.5A-C). However, we didn't find any savings for overall or temporally-labile adaptation either (raw savings of -0.8±7.3% and 2.3±10.0% respectively), in contrast to Experiments 5.1 and 5.2 where the perturbation retrained after washout was the same. Moreover, same-perturbation savings from Experiment 5.1 and 5.2 were both

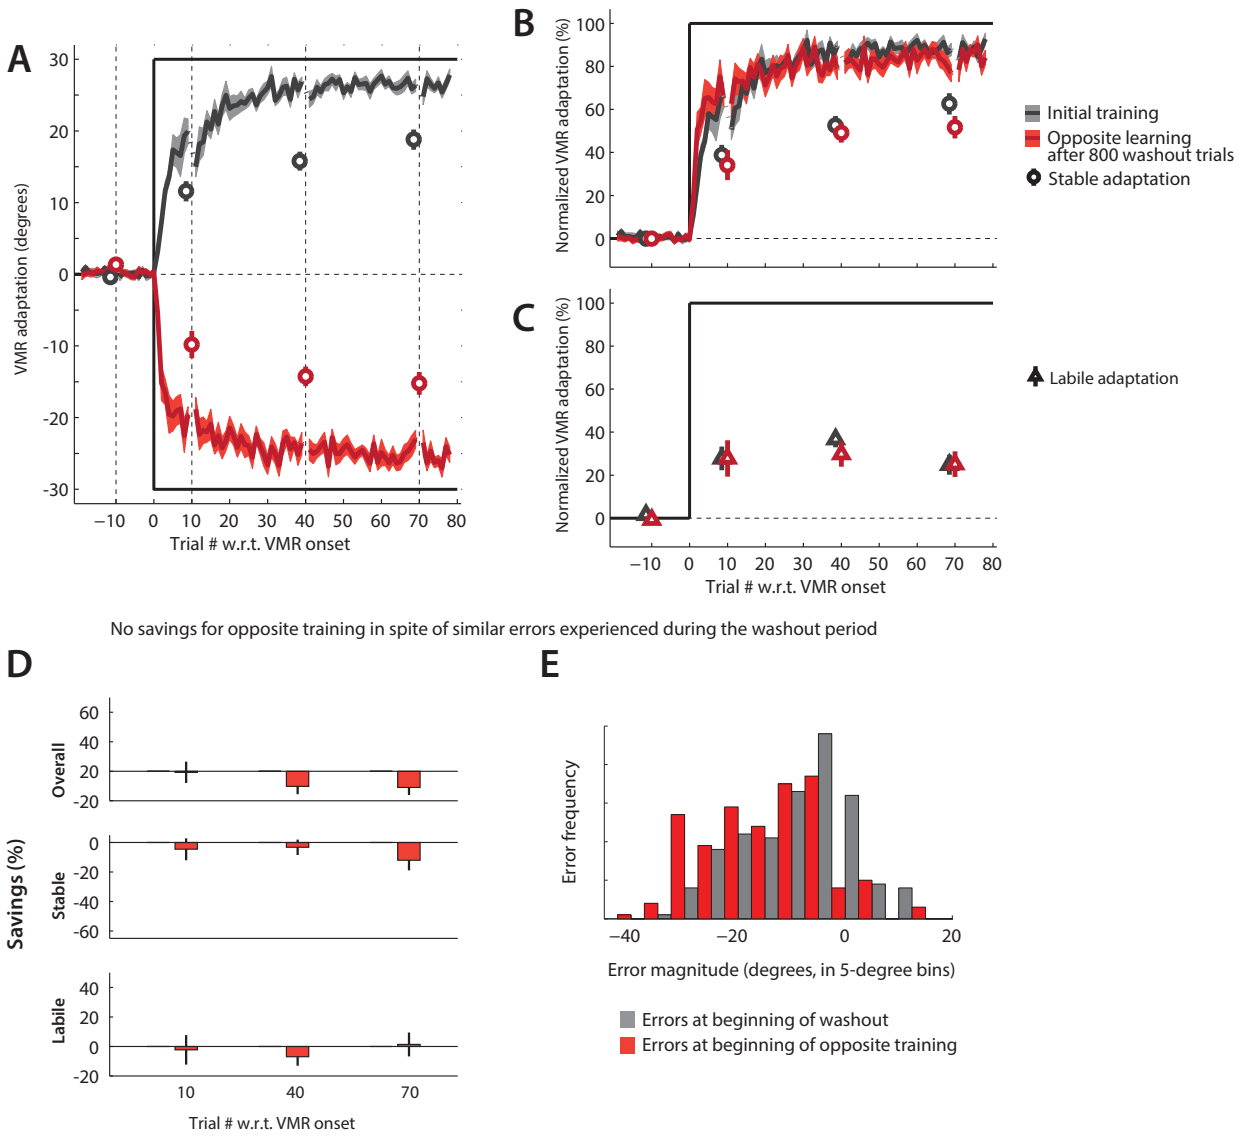


Figure 5.5. Savings does not transfer to the learning of the opposite perturbation.

A: Learning curves from Experiment 5.4, showing the overall and stable adaptation for initial (gray) and opposite learning (red) adaptation. Note how the two learning curves appear to mirror each other. Format analogous to Fig. 5.3A. Shading indicates SEM.

B: Normalized initial (gray) and opposite learning (red) adaptation for the overall and stable components, showing no clear increase in the speed of overall adaptation. Contrast this to the increase in the speed of overall adaptation shown on Figure 5.3CH when the retraining episode presents the same, rather than the opposite, perturbation.

C: Same as B but for the temporally-labile component of adaptation, clearly showing no increase in adaptation speed, in contrast to Figure 5.3DI.

D: Raw savings (Equation 5.6) for overall, stable and labile adaptation, illustrating no opposite savings in contrast to Experiments 5.1 and 5.2.

Figure 5.5 (continued)

E: Distribution of errors from all subjects during the first 10 trials of washout (gray) vs. the first 10 trials of the subsequent opposite learning (red), showing that the range of these errors is very similar. The fact that savings is not present during opposite training, in spite of similar errors being experienced before, suggests that the temporally-labile savings we observe are specific to the perturbation rather than the errors experienced.

higher than opposite-perturbation savings from Experiment 5.4 ($p<0.05$ for trial 10 savings for both overall and labile adaptation, two-tailed unpaired t-tests).

This lack of savings was in spite of the fact that the same errors during early training (first 10 trials) had been experienced before, during the beginning of the first washout period. Figure 5.5E shows the distribution of errors both during the beginning of opposite training and during the beginning of the preceding washout period, showing that their respective ranges are very similar. Thus, our results are in contrast to the error-specific model, and provide support for the perturbation-specific model. Our observations are in line with previous findings in which single-trial adaptation to the opposing FF after prolonged consistent training is in fact reversed (Gonzalez Castro et al., 2014) – also in chapter 6, more specifically Figure 6.5 and 6.6. Overall, these results suggest that the learning rate upregulation that seems to drive temporally-labile savings is specific to the perturbation and not to the errors experienced.

5.5 Discussion

Here we used a visuomotor rotation (VMR) learning task to investigate the extent to which savings motor learning represent the re-emergence of a temporally-stable, previously consolidated memory or instead show an increased capacity for reacquiring adaptation that could be temporally-labile. During both training and washout sessions, we divided adaptive responses into temporally-stable and temporally-labile components using 60-second time delays (Figure 5.1). We first observed that the washout of the temporally-stable component is about 4 times slower than the washout of the temporally-labile component (τ of about 100 trials vs. a τ of about 25 trials or less, Figure 5.2A). Moreover, in a different experiment we found that this washout-resistant, temporally-stable adaptation displayed long-term retention one day later (Figure 5.2B,C). However, having residual temporally-stable adaptation after washout was not necessary for savings: we found similar savings when retraining after both a long, 800-trial washout period which brought both temporally-stable and temporally-labile processes back to baseline, and a short, 40-trial washout period which left a significant temporally-stable residual adaptation component (Figure 5.3). Finally, we investigated savings separately for the temporally-stable and temporally-labile processes, and we found that, remarkably, savings was exclusive to the temporally-labile process (Figures 5.3, 5.4). This indicates that savings is driven by an increased propensity to more rapidly form a memory similar to what was previously learned, but not the re-emergence of a latent temporally-stable memory that was previously consolidated, as previously thought.

5.5.1 Linear time-invariant models do not account for all savings in motor adaptation

Our work shows that linear time-invariant (LTI) models of motor adaptation are not sufficient to explain all observed savings. LTI models assume one or multiple different states of motor adaptation, each with its own learning and retention parameters that are constant. LTI models explain a variety of phenomena in motor control like the two-timescale behavior of learning curves, spontaneous recovery, anterograde interference, and including, under certain conditions, savings (Smith et al., 2006; Sing and Smith, 2010).. Specifically, LTI models predict savings when the apparent unlearning of a stimulus is the result of a

short period of training with the opposite stimulus, a condition which creates two opposing motor memories: a residual memory of the initial stimulus contained in slow-learning strongly-retaining processes and a rapidly acquired memory of the recently trained opposite stimulus contained in the fast-learning weakly-retaining processes (Smith et al., 2006). When the initial stimulus is presented again, the slow processes are still partially adapted to it, whereas the fast processes start to re-adapt rapidly, resulting in savings. However, for the case when the washout of both fast and slow processes is complete – for example, when unlearning is the result of a prolonged washout period where the stimulus is reversed to baseline, as during the 800-trial washout periods in Experiments 5.1 and 5.2 – LTI models predict that, since the system returned to its naïve state, subsequent adaptation should not exhibit savings. This prediction is in contrast with our results which show savings after the 800-trial washout period, during which both temporally-labile and temporally-stable adaptation returned to zero, illustrating that a considerable amount of savings in motor adaptation cannot be explained by the LTI framework.

Previous work similarly showed savings after overall adaptation returned very close to zero within a 40-trial washout period (Zarahn et al., 2008), concluding that LTI models are insufficient to explain savings. While our work validates this hypothesis, we point out that the 40-trial washout period that was utilized, though able to bring overall adaptation very close to zero, would not be sufficient to completely wash out temporally-stable adaptation. Our analysis of the washout curves in Figure 5.2 shows a residual temporally-stable adaptation of about 5°-6° after 40 trials of washout, meaning that savings in overall adaptation after that point will partially be due to the temporally-stable component. In summary, whereas any residual slow-state adaptation to the initial stimulus will result in savings when this stimulus is relearned, here we clearly illustrate that a substantial amount of savings occurs even when this slow-state adaptation has washed out.

5.5.2 What drives savings in the temporally-labile process?

The strong savings we observe for the temporally-labile process signal an apparent increase in its learning rate. What is the mechanism behind this learning rate upregulation? A recent theory, termed structural

learning, posits that learning can be facilitated when the motor system learns the covariance structure between possible learning parameters, allowing for a reduction in the dimensionality of the learning environment (Braun et al., 2009; Landi et al., 2011). For example, experiencing a specific visuomotor rotation might enable the motor system to adapt with an increased learning rate upon the onset of a different rotation, by restricting exploration of the space of all possible visuomotor transformations to a narrower manifold containing only visuomotor rotations. Critically, this learning rate upregulation depends on the consistency of the learning environment (Yousif and Diedrichsen, 2012; Gonzalez Castro et al., 2014) – also chapter 6: consistent environments, which correlate strongly from one trial to the next, can lead up to a 3-fold increase in learning rates, whereas inconsistent ones can lead up to 5-fold reductions compared to baseline.

5.5.3 Using savings as an indicator of memory consolidation can be problematic

Previous work utilized savings as an assay of whether a previously learned was consolidated (Krakauer and Shadmehr, 2006; Richardson et al., 2006; Landi et al., 2011). However, here we find that savings do not necessarily result from the latent reemergence of a previously consolidated motor memory, but instead could reflect an increased propensity to relearn what was previously learned. Our finding thus suggests that using savings as an indicator of memory consolidation might be problematic. For example, under this idea, the savings we observed after the 800-trial washout period would have been taken as an indicator that the longer-term memory associated with the temporally-stable component had consolidated, whereas, in fact, even the more stable component of motor adaptation slowly but clearly washed out to zero during that 800-trial period (Figure 5.2A), and that stable adaptation was in fact weaker during retraining compared to initial training. We thus suggest a degree of separation between the analysis of savings and consolidation be done separately: for example, studies can first examine whether previous-session adaptation had consolidated by measuring retained adaptation using non-training, probe trials (for example, error-clamps) and use these measurements as a starting point from which to analyze the subsequent relearning period and assess savings.

5.6 References

Berniker M, Kording KP (2011) Estimating the relevance of world disturbances to explain savings, interference and long-term motor adaptation effects. PLoS computational biology 7:e1002210.

Brashers-Krug T, Shadmehr R, Bizzi E (1996) Consolidation in human motor memory. Nature 382:252-255.

Braun DA, Aertsen A, Wolpert DM, Mehring C (2009) Motor task variation induces structural learning. Current biology : CB 19:352-357.

Brayanov JB, Press DZ, Smith MA (2012) Motor memory is encoded as a gain-field combination of intrinsic and extrinsic action representations. The Journal of neuroscience : the official journal of the Society for Neuroscience 32:14951-14965.

Ebbinghaus H, Ruger HA, Bussenius CE (1913) Memory; a contribution to experimental psychology. New York,: Teachers College, Columbia University.

Efron B, Tibshirani RJ (1994) An introduction to the bootstrap: CRC press.

Frey PW, Ross LE (1968) Classical conditioning of the rabbit eyelid response as a function of interstimulus interval. Journal of comparative and physiological psychology 65:246-250.

Gonzalez Castro LN, Hadjiosif AM, Hemphill MA, Smith MA (2014) Environmental consistency determines the rate of motor adaptation. Current biology : CB 24:1050-1061.

Herzfeld DJ, Vaswani PA, Marko MK, Shadmehr R (2014) A memory of errors in sensorimotor learning. Science 345:1349-1353.

Huang VS, Haith A, Mazzoni P, Krakauer JW (2011) Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. Neuron 70:787-801.

Kitago T, Ryan SL, Mazzoni P, Krakauer JW, Haith AM (2013) Unlearning versus savings in visuomotor adaptation: comparing effects of washout, passage of time, and removal of errors on motor memory. Frontiers in human neuroscience 7:307.

Kojima Y, Iwamoto Y, Yoshida K (2004) Memory of learning facilitates saccadic adaptation in the monkey. The Journal of neuroscience : the official journal of the Society for Neuroscience 24:7531-7539.

Kording KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. Nature 427:244-247.

Kording KP, Ku SP, Wolpert DM (2004) Bayesian integration in force estimation. Journal of neurophysiology 92:3161-3165.

Krakauer JW, Shadmehr R (2006) Consolidation of motor memory. Trends in neurosciences 29:58-64.

Krakauer JW, Ghez C, Ghilardi MF (2005a) Adaptation to visuomotor transformations: consolidation, interference, and forgetting. The Journal of neuroscience : the official journal of the Society for Neuroscience 25:473-478.

Krakauer JW, Ghez C, Ghilardi MF (2005b) Adaptation to visuomotor transformations: Consolidation, interference, and forgetting. Journal of Neuroscience 25:473-478.

Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000a) Learning of visuomotor transformations for vectorial planning of reaching trajectories. The Journal of neuroscience : the official journal of the Society for Neuroscience 20:8916-8924.

Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000b) Learning of visuomotor transformations for vectorial planning of reaching trajectories. Journal of Neuroscience 20:8916-8924.

Landi SM, Baguear F, Della-Maggiore V (2011) One week of motor adaptation induces structural changes in primary motor cortex that predict long-term memory one year later. The Journal of neuroscience : the official journal of the Society for Neuroscience 31:11808-11813.

Lee JY, Schweighofer N (2009) Dual adaptation supports a parallel architecture of motor memory. The Journal of neuroscience : the official journal of the Society for Neuroscience 29:10396-10404.

MacLeod CM (1988) Forgotten but not gone: savings for pictures and words in long-term memory. Journal of experimental psychology Learning, memory, and cognition 14:195-212.

Malone LA, Vasudevan EV, Bastian AJ (2011) Motor adaptation training for faster relearning. The Journal of neuroscience : the official journal of the Society for Neuroscience 31:15136-15143.

Mawase F, Shmuelof L, Bar-Haim S, Karniel A (2014) Savings in locomotor adaptation explained by changes in learning parameters following initial adaptation. Journal of neurophysiology 111:1444-1454.

Medina JF, Garcia KS, Mauk MD (2001) A mechanism for savings in the cerebellum. The Journal of neuroscience : the official journal of the Society for Neuroscience 21:4081-4089.

Napier RM, Macrae M, Kehoe EJ (1992) Rapid reaquisition in conditioning of the rabbit's nictitating membrane response. Journal of experimental psychology Animal behavior processes 18:182-192.

Nelson TO (1985) Ebbinghaus Contribution to the Measurement of Retention - Savings during Relearning. J Exp Psychol Learn 11:472-479.

Pekny SE, Criscimagna-Hemminger SE, Shadmehr R (2011) Protection and expression of human motor memories. The Journal of neuroscience : the official journal of the Society for Neuroscience 31:13829-13839.

Richardson AG, Overduin SA, Valero-Cabre A, Padoa-Schioppa C, Pascual-Leone A, Bizzi E, Press DZ (2006) Disruption of primary motor cortex before learning impairs memory of movement dynamics. The Journal of neuroscience : the official journal of the Society for Neuroscience 26:12466-12470.

Shadmehr R, Brashers-Krug T (1997) Functional stages in the formation of human long-term motor memory. The Journal of neuroscience : the official journal of the Society for Neuroscience 17:409-419.

Shmuelof L, Huang VS, Haith AM, Delnicki RJ, Mazzoni P, Krakauer JW (2012) Overcoming motor "forgetting" through reinforcement of learned actions. The Journal of neuroscience : the official journal of the Society for Neuroscience 32:14617-14621.

Sing GC, Smith MA (2010) Reduction in learning rates associated with anterograde interference results from interactions between different timescales in motor adaptation. PLoS computational biology 6.

Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. PLoS biology 4:e179.

Turnham EJ, Braun DA, Wolpert DM (2012) Facilitation of learning induced by both random and gradual visuomotor task variation. J Neurophysiol 107:1111-1122.

Wu HG, Miyamoto YR, Gonzalez Castro LN, Olveczky BP, Smith MA (2014) Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. Nature neuroscience 17:312-321.

Yousif N, Diedrichsen J (2012) Structural learning in feedforward and feedback control. J Neurophysiol 108:2373-2382.

Zarahn E, Weston GD, Liang J, Mazzoni P, Krakauer JW (2008) Explaining savings for visuomotor adaptation: linear time-invariant state-space models are not sufficient. Journal of neurophysiology 100:2537-2548.

Chapter 6: Environmental Consistency Determines the Rate of Motor Adaptation

6.1 Summary[3]

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

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

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

[3] Most of this chapter appears in *Gonzalez Castro LN, Hadjiosif AM, Hemphill MA, Smith MA (2014)* Environmental consistency determines the rate of motor adaptation. *Current biology : CB 24:1050-1061.*

In addition, an earlier form of the text and part of the analysis also appears in the dissertation of L. Nicolas Gonzalez-Castro, with whom I collaborated for this research project:

Gonzalez Castro LN. Understanding Generalization, Credit Assignment and the Regulation of Learning Rate in Human Motor Learning. *Harvard University, 2011.*

6.2 Introduction

The human motor system has the remarkable ability to not only adapt its output to reduce motor errors, but also adapt the rate at which this adaptation occurs (Baddeley et al., 2003; Körding et al., 2007; Burge et al., 2008; Wei and Körding, 2010; Malone et al., 2011). However, the mechanisms by which adaptation rates change are still unclear. Previous studies that have examined this phenomenon have posited that adaptation rates are determined by optimal estimation based on the sensory information available to guide learning (Baddeley et al., 2003; Burge et al., 2008; Wei and Körding, 2010; van Beers, 2012). The idea is that based on noisy sensory information about the environment, the adjustments in motor output that occur during motor adaptation represent an ongoing assessment of the motor system's belief about the state of the environment that was most recently experienced. Accordingly, Bayesian inference, in which the relative amplitudes of different types of noise determine how optimal estimates are made, has been suggested as a framework for understanding learning rate modulations. However, several predictions of this theory have not been borne out experimentally (Baddeley et al., 2003; Burge et al., 2008; van Beers, 2012). Here we suggest that motor adaptation rates are primarily determined not by estimation of the state of the environment most recently experienced, but by prediction of the state most likely to be experienced during the next movement, based on the tendency of changes in the environment to persist from one movement to the next.

The contrast between prediction and estimation is made clear in the Kalman filter, a statistically optimal model widely used in the analysis of linear systems(Kalman, 1960). This model uses a two-step process for incorporating new information: estimation and prediction. First, in the *estimation step*, a gain factor *K* (the Kalman gain) is computed using Bayesian inference to determine the statistically optimal weighting for updating the estimate of the current state based on new sensory information based on the relative levels of state noise σ_x and sensory input noise σ_u. In the equations below, $x(n^-)$ and $x(n^+)$ signify the state estimate before and after this update, respectively, and the error, $e(n)$, is the difference between new sensory information and the value predicted for it $x(n^-)$. Second, in the *prediction step*, a prediction

factor *A* (the state transition gain) that models how the system evolves or decays from one state to the next is multiplied by the result of the estimation step to make a prediction about the next state based on the current estimate.

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

$$Prediction\text{: } x([n+1]^-) = Ax(n^+)$$

Putting these two steps together yields:

$$Combined\ estimation\ and\ prediction\text{: } x([n+1]^-) = Ax(n^-) + AKe(n)$$

Closely resembling the form of a learning rule commonly hypothesized to model motor adaptation (Thoroughman and Shadmehr, 2000; Donchin et al., 2003; Smith and Shadmehr, 2005; Smith et al., 2006b; Sing and Smith, 2010):

$$Trial\ to\ trial\ learning\ rule\ for\ motor\ adaptation\text{: } x(n+1) = Ax(n) + Be(n)$$

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

Adaptation rates are also known to be increased when the exposure to a particular perturbation is repeated, a phenomenon termed savings, and recent work has suggested that action repetition is what leads to the improved learning when savings is observed (Huang et al., 2011) Thus here we examine how motor adaptation rates are modulated by Bayesian estimation gain (K), next-trial prediction (A), and repetition. While it is unclear whether a statistically optimal linear model like the Kalman filter can provide insight into the mechanisms that determine how quickly the human motor system learns, as this learning is likely to be neither strictly linear nor statistically optimal, here we suggest that it can. In

particular, we hypothesize that the motor system makes internal estimates of environmental persistence that are highly adaptable and readily modulate the rate at which motor learning proceeds. Accordingly, we demonstrate that experimental manipulation of the extent to which environmental change persists from one movement to the next can upregulate the rate of motor adaptation by 3-fold and downregulate it by 5-fold. In experiments where the persistence and the variability of environmental perturbations were systematically manipulated, we demonstrate a striking double dissociation supporting the hypothesis that prediction is fundamentally different from estimation and is modulated to a greater degree during motor learning: We find that the next-trial motor adaptation rate is predicted by environmental persistence but not variability, whereas the same-trial feedback response strength gain is predicted by environmental variability but not persistence. Moreover, in experiments where persistence and repetition were systematically manipulated, we find that persistence, measured by environmental consistency defined as the lag-1 autocorrelation of the environment, alone can modulate adaptation rates, whereas repetition alone cannot. However the combination of persistence and repetition greatly potentiates the effect of persistence alone leading to even greater adaptation.

6.3 Experimental Procedures

6.3.1 Ethics statement

All experimental participants were naïve to the experimental purpose, provided informed consent and were compensated for their participation. All the experimental protocols were reviewed and approved by the Harvard University Committee on the Use of Human Subjects (CUHS) in research.

6.3.2 General task description

Subjects performed 10cm reaching movements in the horizontal plane with their dominant hands while grasping the handle of a 2-link robotic manipulandum, Figure 6.1A. They were presented with 1cm-diameter circular targets displayed on a vertically oriented LCD monitor. The position of the subject's hand was represented on the LCD monitor by a 3mm cursor. Position, velocity and force at the handle were measured with sensors installed in the manipulandum at a sampling rate of 200Hz. The subjects were instructed to produce fast, continuous movements, and were provided visual feedback throughout the movement. Feedback about the movement time achieved was presented at the end of each movement. Ideal movement times (500±50ms) were signaled by an animation of the target while a chirp sound was played. Here movement time on each trial was defined that the time between when the movement speed first exceeded 0.05m/s and the final time it went below 0.05m/s. For movement completion times that were below or above the ideal range the targets were colored blue and red, respectively and the chirp sound was withheld. The mean peak speed for the movements in all experiments was 0.33±0.04m/s. All movements in the present study occurred in the 90º and 270º directions (Figure 6.1A). In certain movements, the subjects' trajectories were disrupted by velocity-dependent dynamics. This was implemented by a viscous curl force-field (FF) at the handle produced by the motors of the manipulandum, Equation 6.1 and Figure 6.1B.

(Equation 6.1) $$\vec{F}(\vec{v}) = B\vec{v} = \begin{bmatrix} 0 & b \\ -b & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix}$$

In this equation the constant matrix B represents the viscosity associated with this FF. In particular, we constrained B so that it had a curl matrix form with amplitude b as shown above. This curl form resulted in the direction of the force always orthogonal to the direction of the velocity vector. In most of the experiments (P1N1, P1, P1L, P7, P7L and P20) the magnitude of the scalar b was set to 15 N/(m/s) and its sign was set so that one half of subjects experienced a clockwise FF (b=+15 N/(m/s)) while the other half experienced a counterclockwise FF (b =-15 N/(m/s)) during training. In these experiments, the FF blocks or single trials began with movements in the 270º direction. In the RN and RW environments the mean value of b was approximately zero but its specific value was changed from one trial to the next as described below.

We assessed the level of adaptation using error-clamp trials, i.e., we measured the force pattern that subjects produced when their lateral errors were held to near zero values in an error-clamp (Scheidt et al., 2000; Smith et al., 2006b; Joiner and Smith, 2008b; Gonzalez Castro et al., 2011; Joiner et al., 2013), Figure 6.1B. We then regressed these measured force patterns onto the ideal force required to fully compensate the force-field. The slope of this regression was used as the adaptation coefficient (x) that characterized adaptation in that trial(Smith et al., 2006a; Joiner and Smith, 2008a; Yousif and Diedrichsen, 2012).

In all the experiments adaptation rates were calculated by obtaining the difference between the adaptation coefficient (x) for the error-clamp trial following the first FF trial in a measurement cycle (EC_{Post}) and the adaptation coefficient for the error-clamp trial preceding this force-field trial (EC_{Pre}), Equation 6.2 and Figure 6.1B.

(Equation 6.2) $$Adaptation\ Rate = x(EC_{Post}) - x(EC_{Pre})$$

In all the experiments the percent reduction in kinematic error in a given FF trial (e_{FF}) compared to the kinematic error in the first presentation of the FF (e_{Init}) was calculated by subtracting e_{FF}, corrected to account for the amount of adaptation in the preceding trial, from e_{Init}, and normalizing by (e_{Init}), as shown

in Equation 6.3. As adaptation directly leads to corrections in error that are independent of stiffness, we corrected the kinematic error in a given FF trial (e_{FF}) for the amount of adaptation in the preceding trial by dividing e_{FF} by its corresponding gap in adaptation, $1 - x(EC_{Pre})$.

(Equation 6.3) $$\%\ Reduction\ in\ kinematic\ error = \frac{e_{Init} - \frac{e_{FF}}{1 - x(EC_{Pre})}}{e_{Init}}$$

6.3.3 Learning environment experiments

Different learning environments were created to study the environmental modulation of adaptation rate in different experiments, Figure 6.1B. Consistency is defined as the lag-1 autocorrelation of the learning stimulus, i.e., the expected value of the covariance of the magnitude of the force between the current trial and the subsequent trial in the same movement direction, normalized by the overall variance of the force-field environment, Equation 6.4.

(Equation 6.4) $$Consistency: R(1) = \frac{E[(FF_n - \mu_{FF})(FF_{n+1} - \mu_{FF})]}{\sigma_{FF}^2}$$

We note that, for error clamp trials, we used the adaptation coefficient x as a measure of the effective learning stimulus for all environmental consistency and variability calculations.

All experiments began with a 200-trial baseline/familiarization period before the environment being studied was applied. Note also that all blocked experiments were balanced in terms of FF direction with half of the subjects experiencing clockwise curl FFs and the other half experiencing counterclockwise curl FFs. In all experiments, the adaptation rate was calculated as described above (Equation 6.2) for both the 90º and 270º directions. Each separate experiment is described in detail below.

In the anti-consistent learning environment (P1N1; $R(1)$=-0.3) experiment, 21 subjects (mean age=20.7 ± 2.5 years; 7 male) completed two baseline blocks of null movements (100 movements/block) in the 90º/270º directions to obtain familiarity with the task. The subjects were then exposed to 50 environment cycles where they experienced the following sequence of trials in each movement direction: a single positive FF trial (b=15 N/(m/s)), followed by a single negative FF trial (B=-15 N/(m/s)), followed by 11-

13 washout (null) trials. To assess the adaptation rate of the subjects we randomly interspersed error-clamp trials before and after the positive FF trial in a subset of the FF cycles (40%) with the aforementioned sequence.

In the inconsistent learning environment (P1; $R(1)$=-0.05) experiment, 12 subjects (mean age=23.8 ± 6.6 years; 1 male) were exposed to two baseline blocks of null movements (100 movements/block) in the 90º/270º directions to obtain familiarity with the task. The subjects were then exposed to 45 environment cycles where they experienced the following sequence of trials in each movement direction: a single positive FF trial (b=15 N/(m/s)), followed by 10-12 washout (null) trials. Randomly interspersed in a subset of the environment FF cycles (44%) we introduced error-clamp trials before and after the positive FF trial of the aforementioned sequence to assess the subjects' adaptation rate.

In the medium consistency learning environment (P7; $R(1)$=0.74) experiment, 12 subjects (mean age=23.8 ± 6.6 years; 1 male) completed two baseline blocks of null movements (100 movements/block) in the 90º/270º directions to obtain familiarity with the task. The subjects were then exposed to 27 environment cycles where they experienced the following sequence of trials in each movement direction: seven positive FF trials (b=15 N/(m/s)), followed by 15-18 washout (null) trials. Randomly interspersed in a subset of the environment cycles (44%) we introduced error-clamp trials before and after the first FF trial of the aforementioned sequence to assess the subjects' adaptation rate.

In the high consistency learning environment (P20; $R(1)$=0.90) experiment, 28 subjects (mean age=22.8 ± 3.9 years; 15 male) completed two baseline blocks of null movements (100 movements/block) in the 90º/270º directions to obtain familiarity with the task. The subjects were then exposed to 27 environment cycles where they experienced the following sequence of trials in each movement direction: twenty positive FF trials (b=15 N/(m/s)), followed by 28-32 washout (null) trials. Randomly interspersed in a subset of the environment cycles (44%) we introduced error-clamp trials before and after the first FF trial of the aforementioned sequence to assess the subjects' adaptation rate. In addition, for 16 out of these 28

subjects, at the end of the last FF environment cycle we added a single negative FF trial, surrounded by EC trials, to assess the adaptation rate associated with this novel perturbation.

A group of 18 subjects (mean age=20.3 ± 2.7 years; 6 male) completed an extended version of the medium consistency learning environment experiment (P7L; $R(1)$=0.73), modified to include the approximately the same number of FF trials used in the high consistency learning environment experiment (P20) – 539 FF trials in P7L vs 540 FF trials in P20. These subjects completed 2 blocks of 100 null familiarization trials in the 90º/270º directions, and then 77 environment cycles where they experienced the following sequence of trials in each movement direction: seven positive FF trials (B=15 N/(m/s)), followed by 10-14 washout (null) trials. Randomly interspersed in a subset of the environment cycles (45%) we introduced error-clamp trials before and after the first positive FF trial of the aforementioned sequence to assess the subjects' adaptation rate. In addition, at the end of the last FF environment cycle we added a single negative FF trial, surrounded by EC trials, to assess the adaptation rate associated with this novel perturbation.

A group of 12 subjects (mean age=21.8 ± 2.9 years; 4 male) completed an extended version of the inconsistent learning environment experiment (P1L; $R(1)$=-0.45), modified to include the same number of FF trials (540) used in the high consistency learning environment experiment (P20). These subjects completed 2 blocks of 100 null familiarization trials in the 90º/270º directions, and then 540 environment cycles where they experienced the following sequence of trials in each movement direction: one positive FF trials (b=15 N/(m/s)), followed by 1-3 washout (null) trials. Randomly interspersed in a subset of the environment cycles (22%) we introduced error-clamp trials before and after the first positive FF trial of the aforementioned sequence to assess the subjects' adaptation rate.

In the random noise learning environment (RN; $R(1)$=0.02) experiment, 13 subjects (mean age=22.4 ± 3.5 years; 8 male) were exposed to two baseline blocks of null movements (100 movements/block) in the 90º/270º directions to obtain familiarity with the task. The subjects were then exposed to an environment where the FF trials varied randomly, with magnitudes determined according to a normal distribution with

standard deviation of 7.5 N/(m/s). Randomly interspersed in this environment we introduced error-clamp trials (1% of trials) to assess the adaptation rate of the subjects according to the following sequence of trials, (1) EC trial, (2) FF trial of magnitude ± 15 N/(m/s), (3) EC trial, Figure 6.1B.

Finally, in the random walk learning environment (RW; $R(1)$=0.76) experiment, 23 subjects (mean age=22.4 ± 3.5 years; 10 male) were exposed to two baseline blocks of null movements (100 movements/block) in the 90º/270º directions to obtain familiarity with the task. The subjects were then exposed to an environment where the strength of the viscous FF followed a damped random walk:

$$b(k) = Ab(k-1) + w(k)$$

Where $b(k)$ is the strength of the force-field at trial k, $A = 0.88$ was the carryover (retention) coefficient and $w(k)$ was zero-mean random noise with a standard deviation of 2.7 N/(m/s), chosen so that the variance of the overall environment roughly matched the variance of the random noise environment. Note that this is similar to a regular (undamped) random walk for which A would be 1. We used a value for the carryover coefficient (A) smaller than 1 to ensure that the amplitude of b does not grow unbounded. Note that this value is numerically equal to the expected value of the lag-1 autocorrelation ($R(1)$) for an infinitely long sequence $b(k)$. EC-FF-EC measurement triplets, like those described above, were randomly interspersed in this environment with a mean interval of 30 trials to measure the adaptation rate in both the 90º and 270º directions. The perturbation trial in the triplet had a b of +15Ns/m in half the triplets and -15Ns/m for the other half. Note that without interspersing these triplets the $R(1)$ for the RW environment would be close to A = 0.88; however, the fact that the FF perturbations contained within these triplets are uncorrelated to the FF trials that precede or follow them reduces the lag-1autocorrelation from 0.88 to $R(1)$ = 0.76. The amplitude of the FF coefficient b was capped at 24Ns/m.

6.3.4 Data inclusion criteria

In our data analysis grossly irregular trials were excluded. This included movements that were extremely fast (peak velocity > 0.5 m/s) or extremely slow (peak velocity < 0.22 m/s), as well as trials with extremely fast (< 75 ms) or extremely slow (> 2.5 sec) reaction times. This insured that subjects did not initiate movements too quickly, without correctly identifying the location of the target, or too late, indicating that they might have not been attending to the task. Applying these criteria resulted in the inclusion of at least 90% of the trials in each group: 94% in the P1N1 group, 90% in the RN group, 97% in the P1 group, 94% in the P7 group, 93% in the P7L group, and 91% in the P20 group. Furthermore, we excluded from our analyses subjects for which less than one third of their EC triplet data fit the inclusion criteria within the number of trials analyzed. Specifically, from the analyses of the latter half of the data (Figures 2BC, 3 and 5 and Figure 6), 7 subjects were excluded: one from the P7 environment, one from the P20 environment, one from the RN environment and four from the RW environment. From our analysis of the last third of the data (Figure 4B) the same two subjects from the P7 and P20 environments were excluded.

6.3.5 Statistical analyses

Differences in kinematic error during training between the learning environments were assessed with a one-way analysis of variance (ANOVA). We also used ANOVAs to assess differences in initial (first measurement sequence) and late (last half of measurement sequences) adaptation rate and percent reduction in kinematic error between the learning environments. Notice that the number of degrees of freedom (DOF) of the error in the comparison of initial adaptation rate is 66 and not 69 because trials from 3 subjects – one in the P1, one in the P7 and one in the P20 groups – failed at least one of the data inclusion criteria noted above. When significant differences arose in the ANOVAs, hypothesis-based post-hoc comparisons were performed using one-tailed *t*-tests. Comparisons between pairs of groups were performed using unpaired *t*-tests without the assumption of equal variances between groups, whereas within group comparisons were performed using paired *t*-tests.

6.3.6 Calculation of confidence intervals for the responses predicted by the combined-CR model

As the combined-CR prediction for the P7L-Opposite response was based on data from different subject populations (RW and P7L, since the prediction was equal to -RW+(P7L-RW)) we could not directly measure variability across subjects to determine SEM as is the other analyses presented in the paper. Thus we utilized a bootstrapping procedure in order to estimate the SEMs associated with this combined-CR prediction (displayed in Figure 6C). Specifically, we repeatedly randomly sampled with replacement from the RW and P7L subject populations for a total of 10,000 iterations, and used the standard deviation of the corresponding sample averages as a measure of SEM. The same process was used for the combined-CR prediction for the P20-Opposite response (Figure 6D).

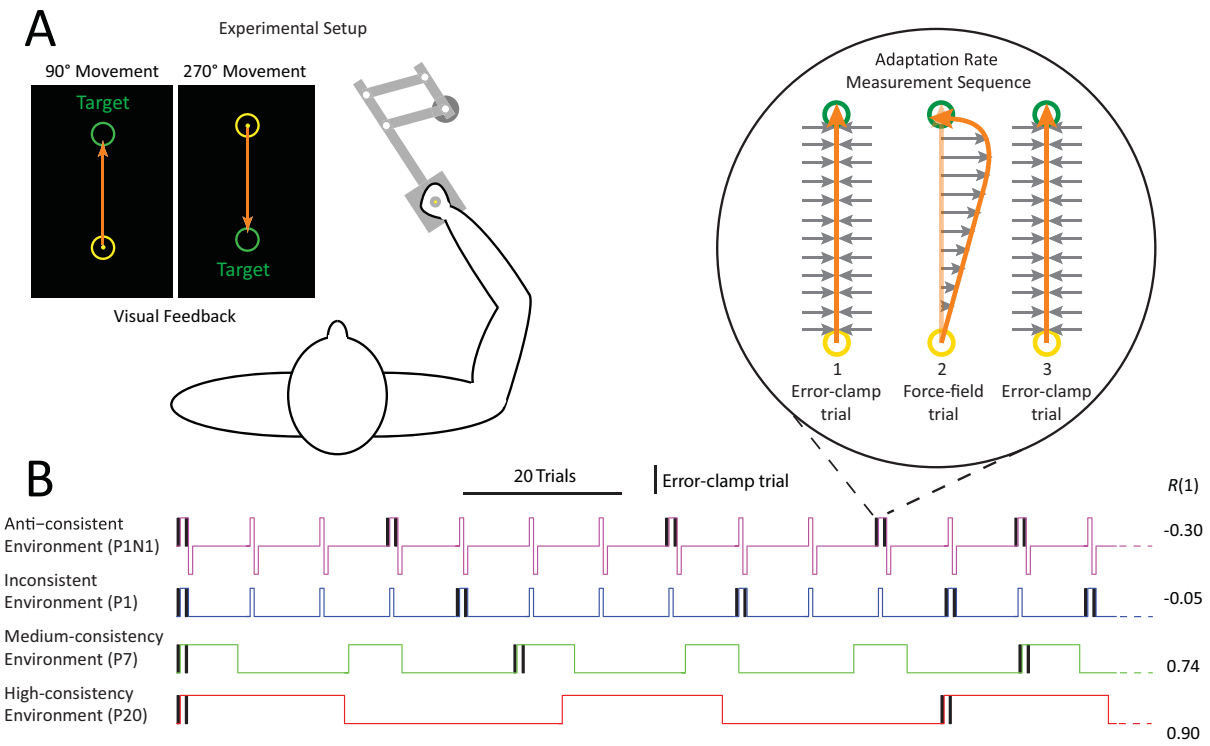


Figure 6.1. Experimental paradigm.

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

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

6.4 Results

We began by systematically manipulating the persistence of the physical environment for movement by cycling different patterns of velocity-dependent force-field (FF) perturbations as illustrated in Figure 6.1B. After performing 200 baseline reaching movements without perturbation (null trials), subjects were exposed to one of four different environments (anti-consistent, inconsistent, medium-consistency, and high-consistency) in which FF perturbation patterns with varying persistence were experienced. We operationally defined consistency as a statistical measure of persistence based on the correlation coefficient between the FF strength in the current trial and the previous trial (the lag-1 autocorrelation, $R(1)$). In all experiments, subjects performed active point-to-point reaching movements in the horizontal plane in two alternating movement directions in the midline, toward and away from the chest, while holding the handle of a robotic manipulandum. The environments were experienced and adaptation rates measured in both movement directions and the results were combined. However, the trial numbers we refer to are within each direction.

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

We designed the experiments so that we could directly compare the adaptation rates induced by each environment. To accomplish this, we estimated the single-trial adaptation rate induced by the first FF trial of each cycle by placing error-clamp (EC) measurement trials (black bars in Figure 6.1B) before and after the FF trial in order to determine the change in motor output associated with the FF exposure (blowout in Figure 6.1B) (Sing et al., 2009; Sing et al., 2013). We refer to this three-trial sequence as a measurement triplet and used the difference between the lateral force output recorded on the two EC trials in each measurement triplet to measure of the change in feedforward motor output induced by exposure to the intervening FF trial. This allowed us to estimate the single-trial adaptation rate identically for all environments by comparing the observed force output change to the change that would be necessary to fully compensate for the FF perturbation.

6.4.1 Environmental consistency can upregulate and downregulate motor adaptation rates

We found that all groups displayed similar adaptation rates on the initial FF trial (one-way ANOVA, $F(3,66)=0.87$, $p=0.46$; 0.089±0.017, mean±SEM for all groups). However, subsequent exposure to different levels of environmental consistency resulted in single-trial adaptation rates that diverged over a range of 20-fold as shown in Figure 6.2A. High environmental consistency resulted in increased adaptation rates, whereas low consistency resulted in decreases, as evidenced by the changes in lateral force output patterns associated with single trial adaptation in the last half of the training period for each group (Figure 6.2B) and adaptation rates estimated from these force output changes (Figure 6.2C). Correspondingly, we found that the adaptation rates were significantly different between groups in the last half of the training period (one-way ANOVA, $F(3,69)=36.46$, $p=3.1\text{x}10^{-14}$). The high-consistency (P20) and medium-consistency (P7) groups both displayed significantly increased adaptation rates (0.303±0.028, $p=1.5\text{x}10^{-8}$ and 0.191±0.033, $p=0.0067$, respectively) relative to the initial rate (0.089±0.017), with the P20 group displaying significantly faster adaptation than the P7 group ($p=0.0075$). In contrast, the low-consistency group (P1) displayed a small but non-significant decrease in adaptation rate (0.058±0.010, $p=0.06$), and the anti-consistent group (P1N1) displayed a significantly

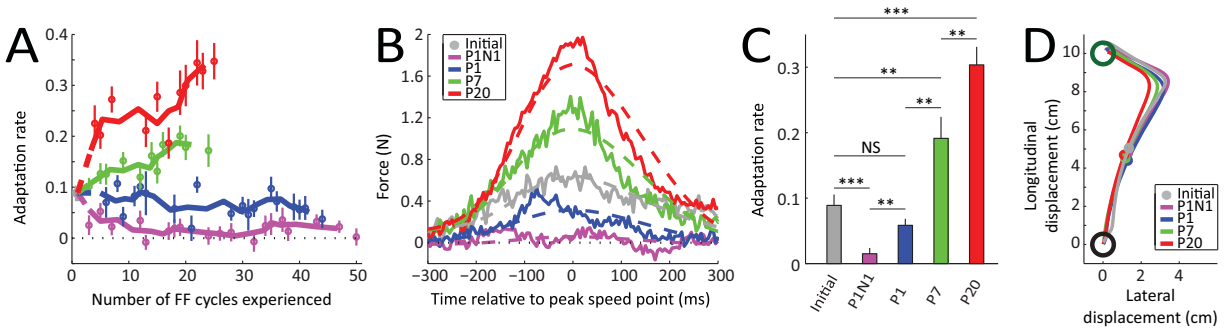


Figure 6.2. Upregulation and downregulation of motor adaptation rates by environmental consistency.

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

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

(C) Average initial adaptation rate for all the different learning environments (gray) compared to the last-half adaptation rates for each environment (colors). Errorbars indicate SEM. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

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

decreased adaptation rate (0.015±0.009) compared to both the initial rate (p=0.0001) and the P1 group (p=0.0014). Because we randomized the duration of the washout period, the first perturbation trial in each cycle was somewhat unpredictable, and subjects showed no evidence of the ability to take advantage of the limited predictability that was present. However, the second (negative) perturbation in P1N1 is *perfectly* predictable since it always follows the first, making the P1N1 environment inherently more predictable than P1, albeit in a negative manner, in line with its negative environmental consistency (R(1)=-0.30). Nevertheless, we observed a near-zero adaptation rate as opposed to a negative one in the

P1N1 data, suggesting that the motor system is unable to make good use of this negative predictability, at least over the time course of our training paradigm. Overall, the data demonstrate that highly-consistent environments upregulate adaptation rates by three-fold whereas anti-consistent environments downregulate adaptation rates by five-fold. Together these results indicate that, adaptation rates can be modulated by more than an order of magnitude over the course of just a few minutes by changes in environmental consistency.

6.4.2 Comparing next-trial adaptive responses with same-trial feedback responses

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

An initial analysis based on data from the P1N1, P1, P7 and P20 environments yielded intriguing but inconclusive results. In short, we used the motor adaptation rate measured during the trial after each perturbation as a measure of next-trial response and the reduction in lateral error that was observed during the perturbation trial itself (with respect to the error observed on the very first perturbation trial) as a measure of the same-trial response – i.e. the feedback response strength. However, the results were unclear. We found that the higher-consistency environments (P7 and P20) which displayed increases in next-trial motor adaptation rates also showed increases in the strength of same-trial feedback responses. But the P7 and P20 environments also had the highest variability, obscuring whether consistency, variability, or both were responsible for the changes in same-trial and next-trial responses. Thus variability and consistency were highly correlated across these four environments ($r=0.82$), preventing them from being dissociated. Interestingly, next-trial responses were somewhat more tightly coupled to changes in environmental consistency ($R^2=0.94$) than variability ($R^2=0.69$), and same-trial responses were

more tightly coupled to changes in environmental variability (R^2=0.97) than consistency (R^2=0.79), hinting but only hinting at a double dissociation.

6.4.3 Next-trial adaptive responses are driven by environmental consistency, whereas same-trial feedback responses are driven by environmental variability

To clearly dissociate the effects of environmental consistency and variability on adaptation rate and feedback response strength, we designed two new environments to have high variability but low consistency. In the first, termed "random noise" (RN), the force-field was randomly varied from one trial to the next (see Figure 6.3A). The second environment, P1-Long (P1L), was similar to P1 but with 2-4 rather than 10-12 null washout trials between the isolated FF trials in order to increase their frequency and thus increase environmental variability. RN's variability (55.4 N^2s^2/m^2) was greater than all the other environments, whereas its consistency was essentially zero ($R(1)$=0.02) due to its random nature. Similarly, P1L's variability was high (46.6 N^2s^2/m^2) whereas its consistency was low ($R(1)$=-0.45). Correspondingly, the inclusion of RN and P1L allows consistency and variability to be dissociated across the environments we studied, breaking the positive correlation between them, dropping their correlation from 0.82 to 0.38 (R^2 of 0.67 vs. 0.14).

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

We found that despite substantially increased variability (55.4 vs. 16.7 N^2s^2/m^2 for RN vs. P1 and 46.6 vs. 33.2 N^2s^2/m^2 for P1L vs. P1N1, see Figure 6.3B,D-E), the RN and P1L environments elicited low adaptation rates, like P1 and P1N1 (0.044±0.008 vs 0.058±0.010, p=0.86 for RN vs. P1, and 0.031±0.008 vs 0.015±0.009, p=0.09 for P1L vs. P1N1). However, this is in line with similar consistency (0.02 vs. -

0.05 for RN vs. P1 and -0.45 vs. -0.30 for P1L vs. P1N1). In contrast, the RN and P1L environments displayed significantly lower adaptation rates (p<0.001 for all pairwise comparisons) than both the P7 and P20 environments (0.191±0.033 and 0.303±0.028) in accordance with substantially lower consistency (0.02 and -0.45 for RN and P1L vs. 0.74 and 0.90 for P7 and P20) despite similar variability (55.4 and 46.6 vs. 45.9 and 53.0 N^2s^2/m^2). These results provide compelling evidence for the idea that environmental consistency rather than variability determines adaptation rates. Correspondingly, the R^2 value for the relationship between adaptation rate and environmental consistency remains essentially unchanged at 90% when the RN and P1L data are included (F(1,4)=35.32, p=0.004, see Figure 6.3D), with similar results observed when considering the consistency of the kinematic error rather than the consistency of the FF environment. However, the R^2 value between adaptation rate and variability drops precipitously from 69% to 16% (F(1,4)=0.76, p=0.43) when we include the RN and P1L experiment data as illustrated in Figure 6.3E. Moreover, a bivariate regression of consistency and variability onto adaptation rate reveals a significant effect of consistency but not variability (F(1,3)=22.3, p=0.018 for consistency, and F(1,3)=0.068, p=0.81 for variability). This bivariate regression accounts for 90% of the variance in the full data set, 0% improvement over consistency alone, but a 74% improvement over variability alone as shown in Figure 6.3F.

When we instead examined changes in feedback response strength, we found that RN and P1L elicited levels (41.8%±6.7% and 35.7%±3.4%) that were far greater than P1 and P1N1 (-6.3%±11.0% and 9.8%±3.2%, p<0.0015 for all pairwise comparisons see Figure 6.3C,G-H), and more like those observed in the P7 and P20 environments (22.2%±9.1% and 21.5%±4.1%). This is in line with levels of environmental variability in RN and P1L (55.4 and 46.6 N^2s^2/m^2) substantially higher than in P1 and P1N1 (16.7 and 33.3 N^2s^2/m^2) but similar to P7 and P20 (45.9 and 53.0 N^2s^2/m^2). However, this result is at odds with similar consistency for RN and P1 (0.02 vs. -0.05) and for P1L and P1N1 (-0.45 vs. -0.30) but reduced consistency for RN and P1L compared to P7 and P20 (0.02 and -0.45 vs. +0.74 and +0.90),

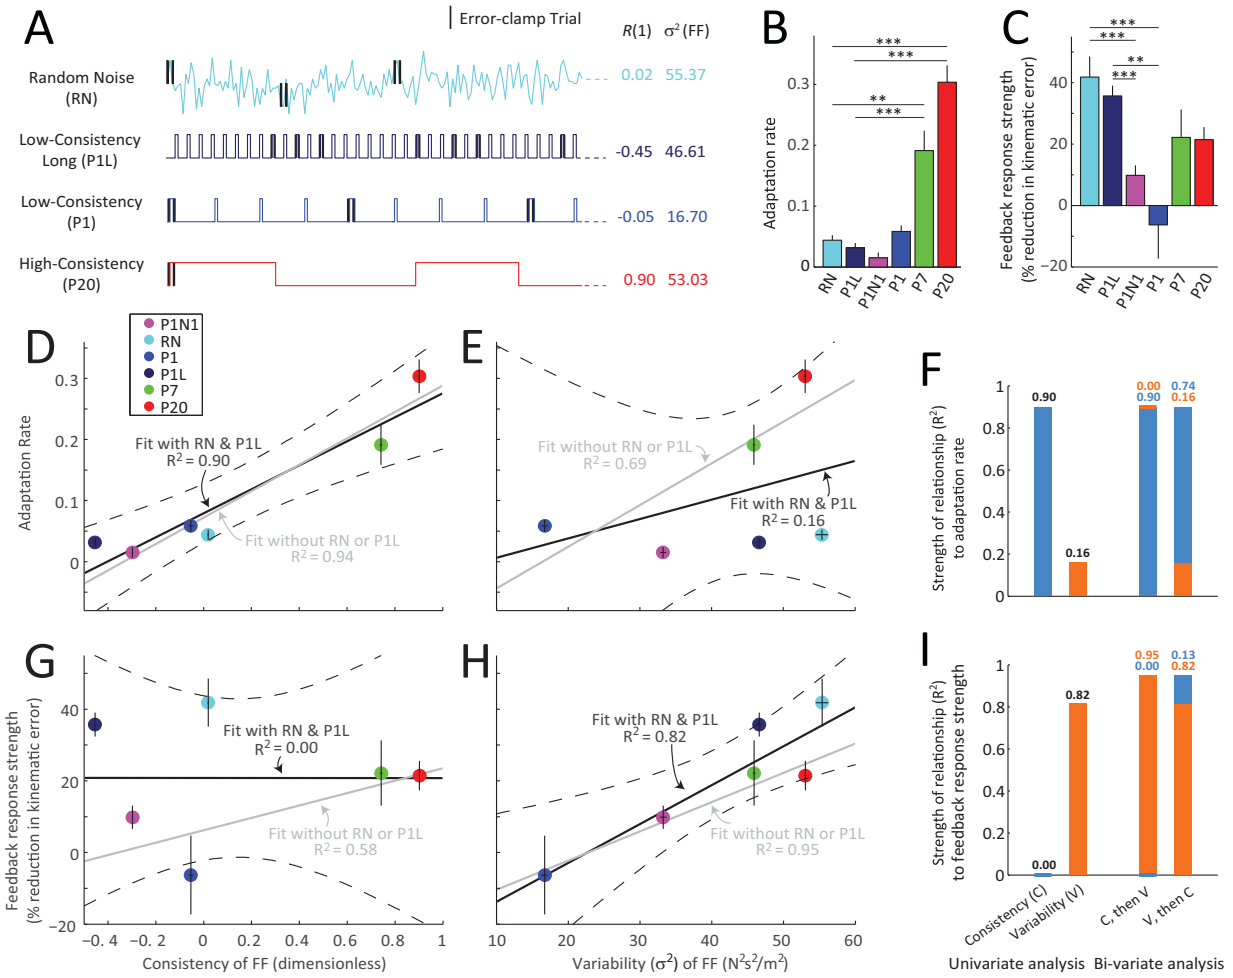


Figure 6.3. Environmental consistency modulates next-trial motor adaptation rates but not same-trial feedback control.

(A) Schematic of learning environments including the random noise (RN) environment. Note that the RN environment was designed to have a high variability like the high-consistency environment (P20) but at the same time a near zero consistency like the inconsistent learning environment (P1).

(B and C) Adaptation rate and feedback response strength (% reduction in kinematic error) in the last half of the P1N1, P1, RN, P7, and P20 environments. Statistical comparisons between RN and the other environments are shown (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$). Errorbars indicate SEM.

(D and E) The relationships between adaptation rate and environmental consistency (D) or variability (E) across experiments. Both before (gray solid line) and after including the RN and P1L data (black solid line), environmental consistency explains a large fraction of the variance in motor adaptation rate, R2>89% in both cases. However inclusion of the RN and P1L data reduces the ability of environmental variability to explain changes in adaptation rate, reducing the R^2 value from 69% to only 16%. The gray dotted horizontal line indicates the initial adaptation rate before exposure to the different environments.

Figure 6.3. (continued)

(G and H) The relationship between feedback response strength and environmental consistency (G) or variability (H) across experiments in a format similar to panels D and E. Both before and after including the RN data, environmental variability explains a large fraction of the variance in feedback response strength, R^2=82% when including the RN /P1L data and R^2=95% without. However inclusion of the RN/P1L data reduces the ability of environmental consistency to explain changes in feedback response strength, reducing the R^2 value from 58% to 0%. Errorbars indicate SEM.

(F and I) Summary of the strength of the relationships (R2) between adaptation rate (F) or feedback response strength (I) and consistency or variability. The first pair of bars summarizes the univariate regression analyses shown in panel D, E, G, and H. The 3rd and 4th bars summarize the corresponding bi-variate regressions. Here the 3rd bar shows the improvement in R^2 when a univariate analysis based on consistency (C, blue) is augmented by variability (V, orange). The 4th bar shows the improvement when a univariate analysis based on V is augmented by C. Errorbars indicate SEM.

suggesting that, unlike next-trial adaptation, same-trial feedback response strength depends on environmental variability rather than consistency. Correspondingly, the R^2 value for the relationship between feedback response strength and environmental variability remains high at 82% when the RN and P1L data are included ($F(1,4)=17.88$, $p=0.013$, see Figure 6.3H). However, the R^2 value between feedback response strength and environmental consistency drops sharply from 58% to 0% ($F(1,4)=0.00$, $p=0.99$) when we include the RN and P1L data, as illustrated in Figure 6.3G. Moreover, a bivariate regression of consistency and variability onto feedback response strength reveals a significant effect of variability but not of consistency ($F(1,3)=59.03$, $p=0.0046$ for variability, and $F(1,3)=8.34$, $p=0.063$ for consistency) .This bivariate regression accounts for 95% of the variance in the full data set, 95% improvement over consistency alone, but only a 13% improvement over variability alone as shown in Figure 6.3I.

Taken together, these findings indicate a striking double dissociation, with motor adaptation rate determined by environmental consistency rather than variability, but feedback response strength determined by environmental variability rather than consistency (Figure 6.3F,I). This finding is in line with a recent study by Yousif and Diedrichsen (Yousif and Diedrichsen, 2012) showing that adaptive

changes in feedback responses could be observed in both consistent and inconsistent environments, whereas changes in feedforward adaptation were only present in consistent environments. The positive relationship between environmental variability and feedback response strength is in line with several previous studies (Takahashi et al., 2001; Osu et al., 2004; Franklin et al., 2008; Selen et al., 2009; Mitrovic et al., 2010) and consistent with the idea that the variance of an environment informs the variance of the prior for Bayesian integration (Körding et al., 2004; Berniker et al., 2010; Chalk et al., 2010; Seydell et al., 2010). Moreover, the extremely tight coupling we find between feedback response strength and environmental variability indicates that the weak relationships we observe (1) between feedback response strength and environmental consistency and (2) between feedforward adaptation rate and environmental variability cannot simply result from the imprecise or noisy characterization of feedback response strength or environmental variability. The specific effect of environmental consistency on adaptation rate suggests that consistency drives changes in the motor system's predictions about the persistence of environmental perturbations from one trial to the next which would affect next-trial adaptation, without driving changes in the motor system's estimation of perturbation size which would affect same-trial feedback responses.

6.4.4 Increases in motor adaptation rate are not explained by the amount of environmental exposure

We next examined whether the improved adaptation levels we observed could be attributed to savings, which refers to the faster relearning of a previously learned task (Krakauer et al., 2000; Medina et al., 2001; Kojima et al., 2004; Smith et al., 2006b; Zarahn et al., 2008). Although our P7 and P20 experiments repeatedly trained the same FF, the durations of individual FF blocks were rather short (20 trials at most) compared to the paradigms in which savings has been demonstrated for motor adaptation (Kojima et al., 2004; Zarahn et al., 2008; Huang et al., 2011), and the partial adaptation achieved in our experiments due to the relatively short FF exposures, precludes the reinforcement of successful performance that has be suggested to be a key ingredient for the instantiation of savings (Huang et al., 2011). In fact, subjects only

learn 50-70% of the force field by the last trial of the training blocks in the P7 and P20 environments where adaptation rate increases were observed (see Figure 6.4C). Correspondingly, we thought it unlikely that savings could be responsible for the faster adaptation observed in the P7 and P20 experiments. Nevertheless, we investigated two experimental manipulations to carefully examine the possibility.

Because the amount of savings should depend on the amount of training, we examined whether the total duration of the exposure to the FF environment could explain the observed increases in adaptation rates as would be predicted by a savings mechanism. This possibility is compatible with the finding that the P20 environment elicited the largest adaptation rates because P20 exposed subjects to a greater number of FF trials than subjects in the other groups. This was the case because the P1 and P7 experiments were designed to match or exceed the P20 experiment in the number of cycles (45, 27, and 27, respectively) rather than FF trials (45, 189, and 540, respectively). We thus examined versions of the medium-consistency and low-consistency environment experiments, P7-Long (P7L) and P1-Long (P1L), that were lengthened in order to match the number of FF trials experienced in the P20 environment (540). We increased the number of cycles from 27 to 77 for the P7L environment and from 45 to 540 for the P1L environment, reducing the number of washout trials somewhat to avoid experiment durations greater than 4 hours (see Figure 6.4A and Experimental Procedures).

The P1L and P7L environments both displayed significantly lower adaptation rates than P20 (0.030±0.009 and 0.252±0.038 vs. 0.333±0.027; mean±SEM for the last third of trials for P1L, P7L and P20, respectively, $p<10^{-8}$ and $p=0.046$ for comparisons against P20). P1L also showed a reduced adaptation rate compared to P1 (0.030±0.009 vs. 0.062±0.010, $p=0.011$), in line with even lower environmental consistency (R(1) of -0.45 for the P1L vs. R(1) of -0.05 for the P1) but in contrast to increased training (540 cycles for P1L vs. 45 cycles for P1). The kinematic error learning curves displayed in Figure 6.4C provide additional evidence for increased adaptation rate in the P20 group compared to P7 and P7L. Comparison of trials 4 through 7 in these normalized learning curves (the

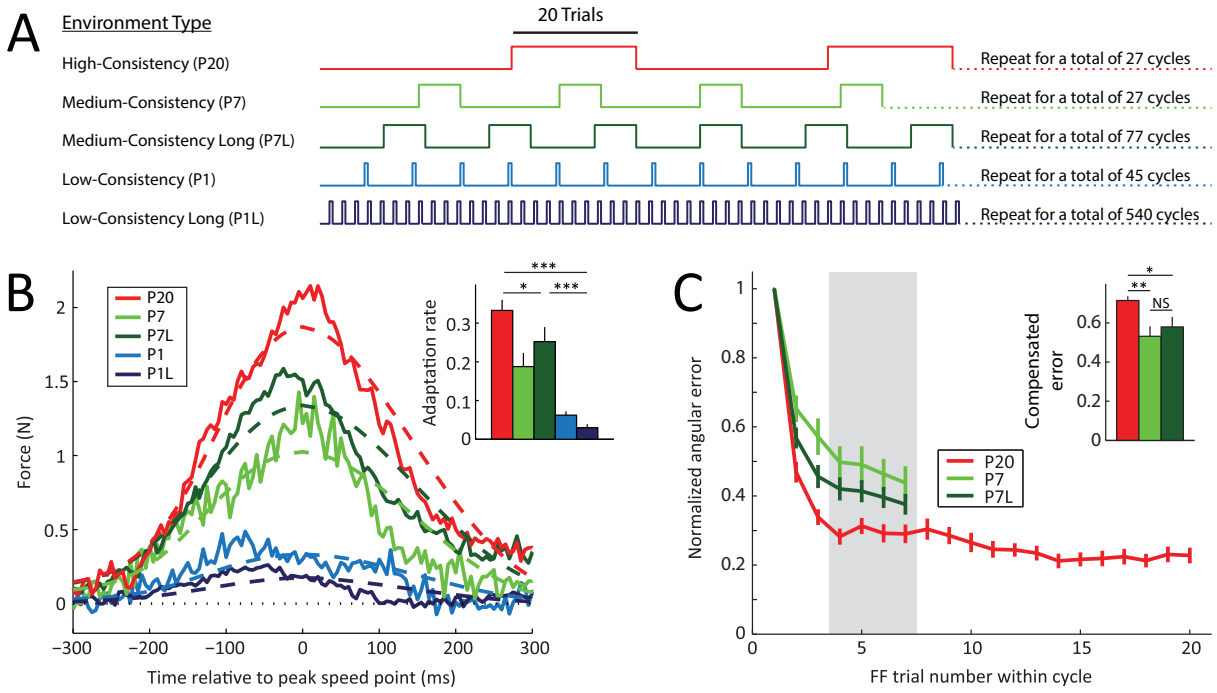


Figure 6.4. Increased repetition does not result in increased adaptation rates in blocked environments.

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

(B) Lateral force profiles comprising the single-trial adaptive response in the P20, P7, P7L, P1 and P1L environments. Note that the subjects in the P7L environment, despite experiencing the same number of FF trials as those in the P20 environment, compensated less than those exposed to the P20 environment. Subjects exposed to the P1L environment compensated even less. The inset shows the average adaptation rates in the P20, P7, P7L, P1 and P1L experiments in the last third of each environment. Errorbars indicate SEM. Notice that subjects in the P20 environment exhibited adaptation rates that were significantly greater than those of subjects in all other environments.

(C) Mean angular error, a second measure of motor adaptation, in late FF cycles in the P20, P7 and P7L environments. The error in trials 4, 5, 6 and 7 (gray shaded region) of the FF cycles was significantly lower in the P20 environment than P7 or P7L. The inset quantifies these differences. Errorbars indicate SEM. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

shaded region) indicates significantly greater kinematic error reduction in the P20 group compared to P7L (p=0.011). These findings indicate that the P20 environment elicits greater adaptation rates than low- or medium-consistency environments when equated based on either the number of cycles (P20 vs. P1 or P7) or the number of trials (P20 vs. P1L or P7L). Thus the amount of exposure cannot account for the adaptation rate differences observed between the low-, medium- and high-consistency environments, as would be predicted by a savings mechanism.

6.4.5 Specificity of adaptation rate increases for the experienced environment dynamics

Although we found that differences in environmental consistency rather than differences in the amount of repetition explain the pattern of learning rate changes observed across different environments (see Figure 6.4 above), we wondered whether repeated exposure to the same dynamics might interact with environmental consistency. In particular, does the relationship between adaptation rate and environmental consistency that we observe across environments require some degree of repetition for the trained dynamics? To investigate this possibility, we began by measuring the adaptation resulting from a single exposure to a completely novel FF – a FF opposite to the one that subjects were repeatedly trained on – after training that would increase adaptation rate. Specifically, we inserted an opposite FF trial at the end of the P7L experiment (P7L-Opposite), as illustrated in Figure 6.5A. The idea was that only adaptation rate upregulation that was entirely independent of repetition could result in full transfer of the adaptation rate increases observed in P7L to the P7L-Opposite trial. However, we found that the adaptive response to the P7L-Opposite trial was in the same direction as the response for the previous FF exposures despite the fact that the P7L-Opposite FF was opposite in direction to these exposures (Figure 6.5B). Thus, the P7L-Opposite response was inappropriately directed, and if an adaptation rate was calculated from it, that rate would be significantly negative in value ($p<0.001$); however, one might question whether an inappropriately-directed response should be characterized by an adaptation rate. Instead, the adaptive

responses elicited by P7L-Opposite trials are more in line with *repetition of the adaptive responses* elicited during the P7L training than *transfer of the response gains* associated with this training, which

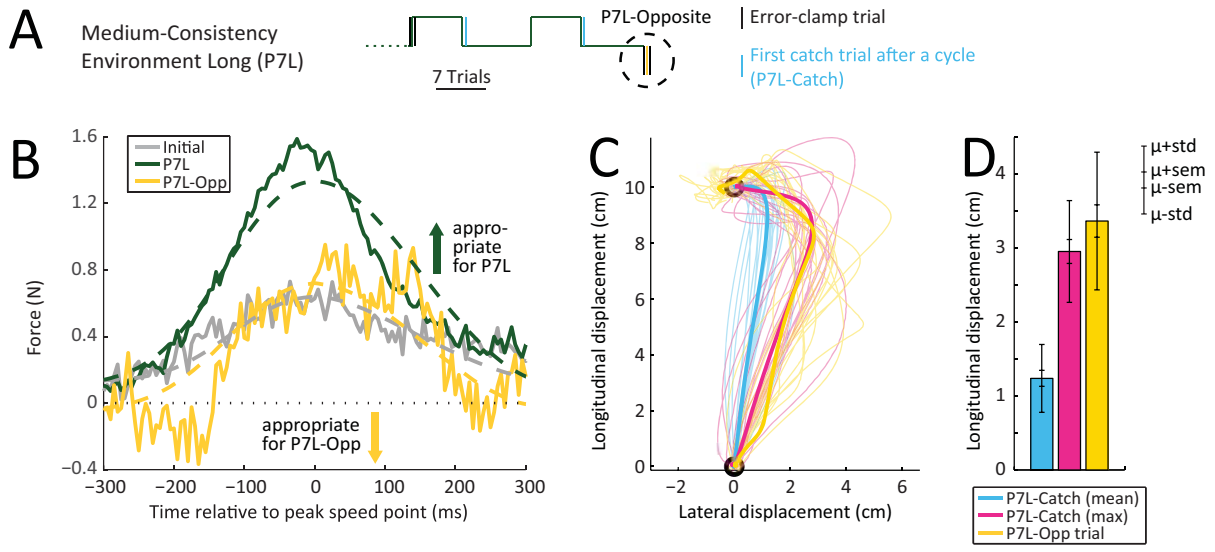


Figure 6.5. Specificity of adaptation-rate increases for the experienced dynamics.

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

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

(C) Hand trajectories during the null trials immediately following FF blocks (P7L-Catch, blue: average trajectories for each subject, pink: trajectories of trials where each subject experienced their maximum deviation) compared to hand trajectories during the single P7L-Opposite trial (yellow). Thin lines indicate individual subject data; thick lines indicate data averaged across subjects. Note that the most deviated P7L-Catch trajectories are similar to the P7L-Opp ones, indicating that, while the FF experienced during the P7L-Opposite trial was completely novel, the errors it elicited were not.

(D) Peak longitudinal displacements / errors experienced during the P7L-Catch and P7L-Opposite trials. While on average each subject experienced weaker errors during P7L-Catch (blue), the largest P7L-Catch errors for each subject (pink) were similar to the errors experienced during P7L-Opposite. Errorbars indicate both SEM (inner notches) and standard deviation (outer notches) across subjects.

were oppositely directed. However, it should be noted that the amplitude of the P7L-Opposite responses falls short of those observed during P7L training (Figure 6.5B, 6.6C-D), suggesting that the P7L-Opposite responses might not be explained by repetition alone.

We next examined whether the inappropriately-directed P7L-Opposite responses were associated with the novelty of the experienced dynamics or of the associated motor errors. By design, the experienced dynamics were completely novel, opposite of the only FF previously experienced. In contrast, motor errors with the same direction as the ones experienced during the P7L-Opposite trial would be expected during at the beginning of the washout period following each FF block. However, because subjects do not fully adapt by the end of each block (Figure 6.4C) the amplitude of the error on the first trial of washout (W1) might be substantially smaller than that experienced during the P7L-Opposite trials. Indeed, we found that the mean W1 trajectory for each participant displayed motor errors that were aligned with, but 63% smaller than the errors experienced during the P7L-Opposite trials (blue vs. yellow in Figure 6.5C-D, with no cases in which mean W1 errors were within 40% of the P7L-Opposite errors). However, W1 trials displayed variable trajectories from one block to the next even within the same individual, and it turns out that the upper limit of this variability (measured as the maximally displaced W1 trial for each individual during the last third of P7L exposure) is largely indistinguishable (8.9% smaller, on average) from the P7L-Opposite motor errors (pink vs. yellow in Figure 6.5D-E, with 83% of cases in which maximum W1 errors were within 25% of the P7L-Opposite errors or higher). Thus, whereas the experienced dynamics during P7L-Opposite were completely novel, the motor errors that participants experienced were similar to errors sometimes experienced at the beginning of each washout block. However, the adaptive response following the maximally displaced W1 trials reduced, rather than increased, the subsequent displacement for all individuals in the P7L experiment, reflecting positive rather than negative adaptation rates for the W1 trials. Thus, the inappropriately-directed P7L-Opposite response cannot be attributed to the motor errors elicited on these trials, and must instead be related to the novel dynamics experienced or the combination of the novel dynamics and the motor errors.

6.4.6 Repetition contributes to, but is not required for consistency-driven adaptation rate upregulation

The adaptation rate reductions observed in the P1L environment compared to P1 or to baseline indicate that repetition without consistency does not increase learning rates. However, the finding that P7L-Opposite responses reflect repetition following exposure to an environment in which repetition and consistency are combined, raises the possibility that repetition might be required for the learning rate increases observed in high-consistency environments. Thus we tested whether learning rate upregulation can occur in a high-consistency environment without repetitious exposure to specific dynamics. In particular, we studied an environment based on a damped random walk (RW) where dynamics randomly varied without repetition from one trial to the next not around zero (as in RN) but around a fixed fraction of the previous trial's dynamics. This resulted in a slowly drifting pattern (see Figure 6.6A) which we parameterized so that consistency ($R(1)$=0.76) was similar to P7 ($R(1)$=0.74) but much higher than RN ($R(1)$=0.02); however, variability (57.7 N^2s^2/m^2) was approximately matched to both RN (55.4 N^2s^2/m^2) and P7 (45.9 N^2s^2/m^2).

We found that the RW environment elicited adaptation rates (0.097±0.012) that were significantly higher than RN (p=0.0004) and the other low-consistency environments, P1 (p=0.0079) and P1L (p=0.00003) as shown in Figure 6.6B, in line with its higher consistency. However its adaptation rate was significantly lower than that observed in P7 (p=0.0094), where a similar consistency was combined with repetition. These findings indicate that environmental consistency alone can elicit learning rate increases; however, when combined with repetition, even greater increases occur.

To examine the mechanism by which consistency and repetition combine in the P7 environment, we can re-examine the P7L-Opposite data in light of the RW results. Here, we consider three possibilities: first, that repetition enables consistency-based learning (C-only hypothesis); second, that consistency enables repetition-based learning (R-only hypothesis); and third, that both consistency and repetition contribute distinct components to the adaptation (combined CR hypothesis). The P7L-Opposite data provide a

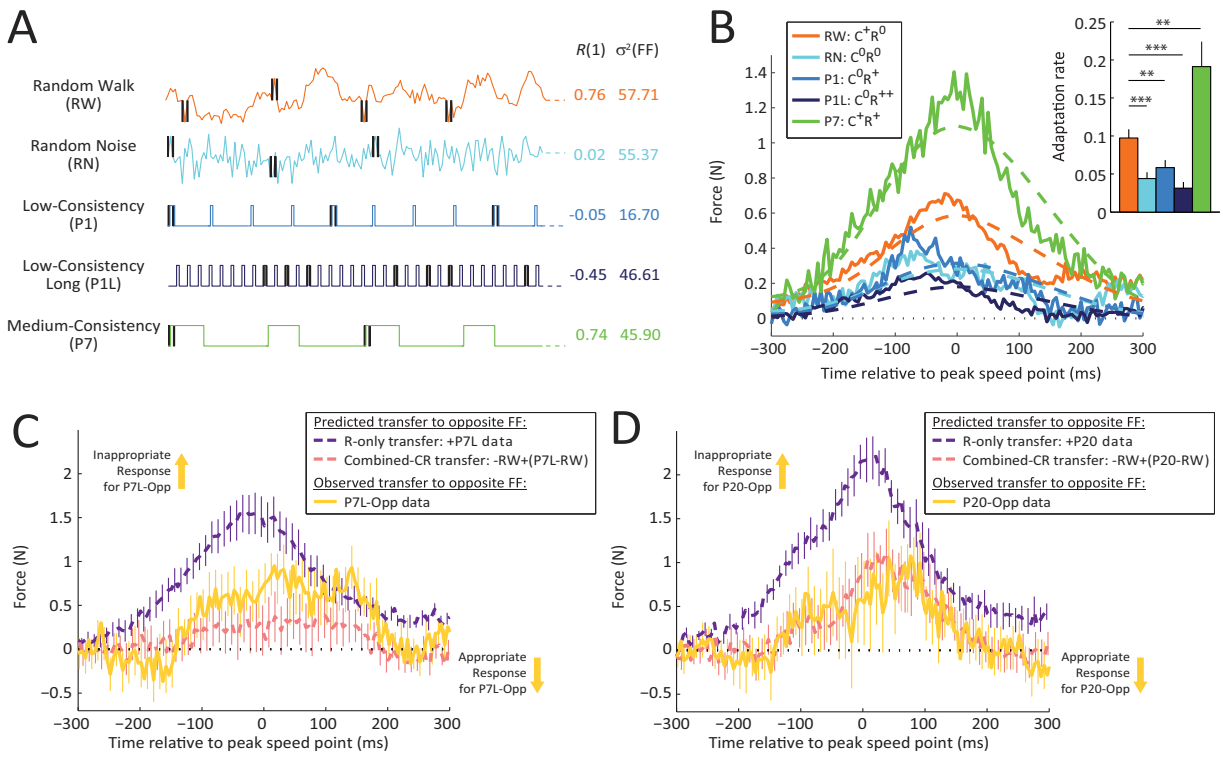


Figure 6.6. Synergistic interaction between repetition and consistency for learning rate upregulation.

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

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

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

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

unique opportunity for examining these hypotheses because consistency-based gain modulations and repetition of the previously learned adaptive responses make opposite predictions, and these predictions can be made without reference to the P7L-Opposite data themselves – instead, deriving the R-only prediction from P7L, and the C-only prediction from RW, and the combined CR prediction from a combination of the two. We can reject the C-only hypothesis, that repetition further enables pure consistency-based learning, because consistency-based adaptation rate transfer to the P7L-Opposite trial would predict an appropriately directed response, which is opposite to what we observe. The R-only hypothesis, that high consistency enables pure repetition-based learning, correctly predicts an inappropriately directed response for P7L-Opposite; however, the amplitude of the predicted response (Figure 6.6C, purple) is considerably larger than the P7L-Opposite data (yellow). In contrast, the combined CR hypothesis, where both consistency and repetition contribute distinct components to the adaptation, appears to explain both the amplitude and direction of the P7L-Opposite response. Here, a flipped version of the RW data (-RW) data provides an estimate of how the consistency-based component of learning should transfer to P7L-Opposite, whereas the difference between the P7L and the RW data (P7L-RW) provides an estimate of the repetition-based component of learning. Figure 6.6C shows that the prediction generated from combining these components (-RW+(P7L-RW); pink) closely explains the P7L-Opposite data, matching these data more closely than the R-only and C-only predictions would. However, it should be noted that, although certainly closer to the combined-CR prediction, the P7L-Opposite data fall in between the R-only and combined-CR predictions, suggesting that this analysis might not be definitive. Thus, to further examine the issue, we examined opposite-FF responses in a group of 16 subjects who experienced a single opposite-FF trial following exposure to the P20 environment. These P20-Opposite data (Figure 6.6D) also match the prediction of the combined CR hypothesis considerably better than the R-only or C-only predictions. Taken together, these data suggest that environments which combine repetition and consistency, elicit a combination of consistency-based and repetition-based adaptation. In the P7 and P20 environments these contributions act in synergy; however in the P7L-Opposite and P20-Opposite data these contributions act in opposition.

6.5 Discussion

We find that experimental manipulation of environmental consistency allows for adaptation rate modulation that spans over an order of magnitude (20x), with highly-consistent environments, in which external disturbances tend to persist from one trial to the next, inducing greater than 3-fold increases in learning rate, and anti-consistent environments, in which these disturbances fluctuate in a negatively correlated manner, inducing 5-fold decreases in learning rate. Moreover, we find that whereas environmental variability determines feedback response strength but not motor adaptation rate, environmental consistency determines motor adaptation rate but not feedback response strength. This striking double dissociation gives insight into the specificity of mechanisms for same-trial control versus next-trial adaptation. However, consistency can be achieved through repetition, and the somewhat surprising results of the P7L-Opposite experiment suggested a key role for repetition in the adaptive responses we observed. Thus we closely examined the interactions between consistency and repetition. We find that even in the absence of repetition, increased consistency elicits highly significant increases in learning (RN vs. RW), however in the absence of consistency, increased repetition fails to elicit increased learning (RN vs. P1 vs. P1L), suggesting that consistency rather than repetition is primarily responsible for the learning rate increases we observe. However, we find that consistency-modulated learning rate increases are greatly magnified when repetition is combined with increased environmental consistency, indicating a strong synergistic effect between repetition and consistency.

The idea that the consistency of the environment modulates the rate of motor adaptation is in line with the idea that forming a *prediction* of the degree to which environmental change should persist from trial to trial is a critical step in determining the rate of adaptation. In a low-consistency or anti-consistent environment the motor system would benefit from a low adaptation rate in order to avoid overlearning from current disturbances that do not positively predict future disturbances. However, in a highly consistent environment, the motor system would benefit from rapid adaptation because disturbances

experienced on one trial would be highly predictive of future disturbances. Our results show that the motor system adjusts its rate of adaptation in a way that is consistent with these ideas.

6.5.1 Interactions between environmental consistency and repetition

Our results reveal a surprising synergy between consistency and repetition in the learning environment. The initial set of experiments, which employed a blocked design in which a single amplitude FF alternated with a null environment (P1N1, P1, P7, and P20), demonstrated large increases in adaptation rate as consistency increased. However, the increased number of FF trials in each block resulted in increased repetition of the FF in the higher consistency environments where the block length was increased (45 vs. 189 vs. 540 FF trial repetitions for the P1, P7 and P20 experiments, respectively). We thus performed the P1L and P7L experiments (540 FF trials each) to determine whether the additional repetition present in the P20 experiment was responsible for the increased adaptation rate it elicited. But in both cases the adaptation rates were smaller than the P20 environment ($p<0.0001$ for P1L, and $p=0.046$ for P7L) and not larger than the corresponding environments with less repetition (P1L actually shows a *lower* learning rate than P1, p=0.011, P7L is not different than P7, p = 0.22). This, suggests that the extremely high learning rate observed in P20 was due to its increased environmental consistency, rather than increased repetition relative to the original P1 and P7 environments, and that increased repetition alone does not improve learning rates, especially when consistency is low (P1 vs P1L). To further study the effect of repetition, we created two environments (RW and RN) in which the amplitude of the FF was randomly varied over a continuous spectrum so that, strictly speaking, no repetition was present except in the probe trials themselves. This allowed us to examine the effect of environmental consistency in the absence of repetition. Here we found that the high-consistency ($R(1)=0.76$, similar to P7), zero-repetition RW environment elicited an adaptation rate that was significantly greater ($p=0.0004$) than observed in the low consistency ($R(1)=0.02$) zero-repetition RN environment, as shown in Figure 6.6. However, the adaptation rate observed for RW was significantly lower ($p=0.009$) than for the P7 environment, which it was matched in consistency, suggesting that the combination of consistency and repetition leads to faster learning than consistency alone.

Together, our results show that increased environmental consistency improves adaptation rates, both when repetition is low (RW vs. RN) and when repetition is high (P20 vs. P7L, and P7L vs. P1L). In contrast, increased repetition fails to increase adaptation rates when consistency is low (P1 vs. RN, and P1L vs. P1), but increased repetition does increase adaptation rates when consistency is high (P7 or P7L vs. RW). The substantial effect of repetition when consistency is high is in line with the inappropriately-directed learning we observe on the P7L-Opposite trial at the end of the P7L experiment, suggesting that repetition may lead to a stereotypic adaptive response when a large error is encountered independent of its direction (Huang et al., 2011). In summary, we find that both consistency and repetition play a role in the modulation of adaptation rates. We show that repetition alone does not lead to an increase in adaptation rate, but consistency alone does lead to an increase in adaptation rate; however, together they lead to an increase in adaptation rate that is about double that elicited by consistency alone, indicating a synergistic effect between these two key features of the learning environment.

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

6.5.2 Attempts to manipulate measurement noise and state noise to alter adaptation rates

Previous studies that have investigated the modulation of motor adaptation rates have focused on the estimation of the identity of the experienced perturbations (Baddeley et al., 2003; Burge et al., 2008; Wei and Körding, 2010). Theoretically this estimation should be based on the appropriate weighting of the uncertainty about the state of the environment (*state noise*, i.e. variance on the *prior* expectation) relative to uncertainty about the fidelity of sensory information (*measurement noise*, i.e. variance on the *likelihood*).

According to optimal Bayesian estimation, the gain on the sensory information should be decreased when the measurement noise rises, and it should increase when the state noise rises (Baddeley et al., 2003; Burge et al., 2008; Berniker et al., 2010; Wei and Körding, 2010). Correspondingly, studies have attempted to manipulate state noise and measurement noise in order to modulate the rate of adaptation (Baddeley et al., 2003; Burge et al., 2008; Wei and Körding, 2010); however, these manipulations have not yielded consistent results. One set of approaches for manipulating state and measurement noise has involved perturbing the visual feedback about hand position, with a sequence that varied from one trial to the next either as white noise or a random walk (Baddeley et al., 2003; Burge et al., 2008). Here the idea was that measurement noise should be independent from one trial to the next like white noise, whereas state noise should accumulate like a random walk, and that the sensorimotor system would leverage this dichotomy by parsing the trial-to-trial variability in sensory information into independent and accumulating components to estimate the levels of measurement noise and state noise. However, there is no compelling reason to believe that this occurs. Instead, we suggest that the sensorimotor system's estimate of measurement noise should be affected by the quality of the sensory feedback rather than its pattern across trials. For example, when a child tracks the motion of a firefly at night, she can clearly identify the location of each flash, despite having little idea where each would appear. Thus, the random pattern of flashes does not indicate high measurement noise. Only a blurred view of a flash would introduce uncertainty about the firefly's current location. Correspondingly, experiments using white noise perturbations to randomize perturbations in an attempt to manipulate the sensorimotor system's estimate

of measurement noise have consistently failed to show the decreases in adaptation rate they predicted (Baddeley et al., 2003; Burge et al., 2008). In contrast, when studies have attempted to increase the motor system's estimate of measurement noise by degrading the intrinsic quality of visual feedback, decreases in adaptation rate have been consistently observed (Burge et al., 2008; Wei and Körding, 2010).

Moreover, studies of sensorimotor integration for same-trial feedback responses have convincingly showed that a white noise perturbation sequence leads to Bayesian integration in which state noise (the prior) rather than measurement noise (the likelihood) is increased by the perturbation variance (Körding et al., 2004; Kording and Wolpert, 2004; Berniker et al., 2010; Chalk et al., 2010; Seydell et al., 2010). In fact, state noise has been taken to fully reflect the variance of the perturbation sequence in these studies. This leaves us with the conclusion that the motor system does not parse trial-to-trial variability in sensory feedback into independent and accumulating components to differentially estimate the levels of measurement noise and state noise, instead taking the entire variability as an estimate of state noise.

But then, how can we explain that studies have found accumulating random walk perturbations to result in faster adaptation rates than white noise perturbations (Baddeley et al., 2003; Burge et al., 2008)? A key observation is the fact that random walk perturbations display a lag-1 autocorrelation that approaches 1 for long duration exposures and are thus, in our parlance, highly-consistent, whereas white noise perturbations display a low-consistency near zero lag-1 autocorrelation. Therefore, changes in environmental consistency can provide a unifying explanation for the current set of results and for increases in adaptation rates previously observed with random walk perturbations. This interpretation is also in line with the finding that in dynamic environments where the spatial complexity is likely to exceed the width of the basis elements for adaptation, learning rates are consistently reduced (Thoroughman and Taylor, 2005), as high spatial complexity could make an environment appear inconsistent.

6.5.3 Implications for retention of motor memories

The idea that forming a prediction of the degree to which environmental change should persist from trial to trial is a critical step in determining the rate of adaptation, also has implications for the retention of

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

6.5.4 Implications for rehabilitation

The findings described in the present study may enable the design of more effective rehabilitation protocols for stroke patients. The first four weeks post-stroke is the time window during which rehabilitation has the greatest impact on the functional recovery of patients (Kwakkel et al., 2003; Kwakkel et al., 2004; Krakauer, 2006). Therefore, an intervention that improves the rate of motor learning during this initial period may help patients achieve better functional recovery, possibly via the design of motor rehabilitation paradigms that upregulate motor learning with high-consistency environments.

6.6 References

Baddeley RJ, Ingram HA, Miall RC (2003) System identification applied to a visuomotor task: Near-optimal human performance in a noisy changing task. J Neurosci 23:3066-3075.

Berniker M, Voss M, Kording K (2010) Learning Priors for Bayesian Computations in the Nervous System. PLoS ONE 5.

Burge J, Ernst MO, Banks MS (2008) The statistical determinants of adaptation rate in human reaching. J Vis 8:20.21-19.

Chalk M, Seitz AR, Series P (2010) Rapidly learned stimulus expectations alter perception of motion. J Vision 10.

Donchin O, Francis JT, Shadmehr R (2003) Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: Theory and experiments in human motor control. J Neurosci 23:9032-9045.

Franklin DW, Burdet E, Peng Tee K, Osu R, Chew C-M, Milner TE, Kawato M (2008) CNS Learns Stable, Accurate, and Efficient Movements Using a Simple Algorithm. J Neurosci 28:11165-11173.

Gonzalez Castro LN, Monsen CB, Smith MA (2011) The binding of learning to action in motor adaptation. PLoS Comput Biol 7:e1002052.

Huang Vincent S, Haith A, Mazzoni P, Krakauer John W (2011) Rethinking Motor Learning and Savings in Adaptation Paradigms: Model-Free Memory for Successful Actions Combines with Internal Models. Neuron 70:787-801.

Joiner WM, Smith MA (2008a) Long-term retention explained by a model of short-term learning in the adaptive control of reaching. Journal of neurophysiology 100:2948-2955.

Joiner WM, Smith MA (2008b) Long-Term Retention Explained by a Model of Short-Term Learning in the Adaptive Control of Reaching. J Neurophysiol 100:2948-2955.

Joiner WM, Brayanov JB, Smith MA (2013) The training schedule affects the stability, not the magnitude, of the interlimb transfer of learned dynamics. Journal of neurophysiology 110:984-998.

Kalman RE (1960) A new approach to linear filtering and prediction problems. Journal of basic Engineering 82:35-45.

Kojima Y, Iwamoto Y, Yoshida K (2004) Memory of learning facilitates saccadic adaptation in the monkey. J Neurosci 24:7531-7539.

Kording KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. Nature 427:244-247.

Körding KP, Ku SP, Wolpert DM (2004) Bayesian integration in force estimation. J Neurophysiol 92:3161-3165.

Körding KP, Tenenbaum JB, Shadmehr R (2007) The dynamics of memory as a consequence of optimal adaptation to a changing body. Nat Neurosci 10:779-786.

Krakauer JW (2006) Motor learning: its relevance to stroke recovery and neurorehabilitation. Curr Opin Neurol 19:84-90.

Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. Journal of Neuroscience 20:8916-8924.

Kwakkel G, Kollen B, Lindeman E (2004) Understanding the pattern of functional recovery after stroke: Facts and theories. Restor Neurol Neurosci 22:281-299.

Kwakkel G, Kollen BJ, van der Grond J, Prevo AJH (2003) Probability of regaining dexterity in the flaccid upper limb - Impact of severity of paresis and time since onset in acute stroke. Stroke 34:2181-2186.

Malone LA, Vasudevan EVL, Bastian AJ (2011) Motor Adaptation Training for Faster Relearning. J Neurosci 31:15136-15143.

Medina JF, Garcia KS, Mauk MD (2001) A mechanism for savings in the cerebellum. J Neurosci 21:4081-4089.

Mitrovic D, Klanke S, Osu R, Kawato M, Vijayakumar S (2010) A Computational Model of Limb Impedance Control Based on Principles of Internal Model Uncertainty. PLoS ONE 5.

Osu R, Kamimura N, Iwasaki H, Nakano E, Harris CM, Wada Y, Kawato M (2004) Optimal impedance control for task achievement in the presence of signal-dependent noise. J Neurophysiol 92:1199-1215.

Scheidt RA, Reinkensmeyer DJ, Conditt MA, Rymer WZ, Mussa-Ivaldi FA (2000) Persistence of motor adaptation during constrained, multi-joint, arm movements. Journal of neurophysiology 84:853-862.

Selen LPJ, Franklin DW, Wolpert DM (2009) Impedance Control Reduces Instability That Arises from Motor Noise. J Neurosci 29:12606-12616.

Seydell A, Knill DC, Trommershauser J (2010) Adapting internal statistical models for interpreting visual cues to depth. J Vis 10:1.1-27.

Sing GC, Smith MA (2010) Reduction in Learning Rates Associated with Anterograde Interference Results from Interactions between Different Timescales in Motor Adaptation. Plos Computational Biology 6:Article No.: e1000893.

Sing GC, Orozco SP, Smith MA (2013) Limb motion dictates how motor learning arises from arbitrary environmental dynamics. Journal of neurophysiology 109:2466-2482.

Sing GC, Joiner WM, Nanayakkara T, Brayanov JB, Smith MA (2009) Primitives for motor adaptation reflect correlated neural tuning to position and velocity. Neuron 64:575-589.

Smith MA, Shadmehr R (2005) Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. J Neurophysiol 93:2809-2821.

Smith MA, Ghazizadeh A, Shadmehr R (2006a) Interacting adaptive processes with different timescales underlie short-term motor learning. PLoS Biol 4:e179.

Smith MA, Ghazizadeh A, Shadmehr R (2006b) Interacting adaptive processes with different timescales underlie short-term motor learning. Plos Biology 4:1035-1043.

Takahashi CD, Scheidt RA, Reinkensmeyer DJ (2001) Impedance Control and Internal Model Formation When Reaching in a Randomly Varying Dynamical Environment. J Neurophysiol 86:1047-1051.

Thoroughman KA, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. Nature 407:742-747.

Thoroughman KA, Taylor JA (2005) Rapid reshaping of human motor generalization. J Neurosci 25:8948-8953.

van Beers RJ (2012) How Does Our Motor System Determine Its Learning Rate? PLoS ONE 7:e49373.

Wei K, Körding K (2010) Uncertainty of feedback and state estimation determines the speed of motor adaptation. Frontiers in Computational Neuroscience 4.

Yousif N, Diedrichsen J (2012) Structural learning in feedforward and feedback control. J Neurophysiol 108:2373-2382.

Zarahn E, Weston GD, Liang J, Mazzoni P, Krakauer JW (2008) Explaining Savings for Visuomotor Adaptation: Linear Time-Invariant State-Space Models Are Not Sufficient. J Neurophysiol 100:2537-2548.