

THE ROLE OF ERROR-SENSITIVITY IN MOTOR ADAPTATION

by

David J. Herzfeld, B.S., M.S.

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

Baltimore, Maryland

May, 2015

© David J. Herzfeld 2015

All Rights Reserved

# Abstract

When we experience an error in a motor task, we adapt our next movement to partially compensate. The process of adaptation can be modeled as  $u^{(n+1)} = \alpha u^{(n)} + \eta^{(n)} e^{(n)}$  where  $u^{(n)}$  is the motor command on trial  $n$ ,  $\alpha$  is a decay factor,  $e^{(n)}$  is error, and  $\eta^{(n)}$  represents the subjects' sensitivity to the experienced error. Here, we explore the rules that govern the value of  $\eta^{(n)}$  as well as the brain-regions that are responsible for its evaluation.

In Chapter 2, we begin with a puzzle: in motor learning tasks, humans are able to modulate how much they learn from a given error. In some conditions, they learn a large amount, but in other conditions they learn only a small amount. That is, the brain selects how much it is willing to learn from error. We suggest that 'error-sensitivity' is modulated by the history of previous errors.

What brain region is responsible for determining the amount subjects are willing to learn from an error? Adaptation is critically dependent on the cerebellum, as demonstrated by patient and lesion studies. In Chapter 3 we use transcranial direct current stimulation (tDCS) to alter the function of the cerebellum, and observe its effects on error-sensitivity. We find that increasing the excitability of the cerebellum via anodal tDCS increases the rate of learning, while decreasing

cerebellum excitability via cathodal tDCS decreases error-sensitivity. That is, we suggest the cerebellum is responsible for determining how much subjects are willing to learn from a motor error.

How does the cerebellum accomplish the task of adaptation? It has been proposed that the firing rates of the principal cells of the cerebellum, Purkinje (P-)cells, should encode movement kinematics. Yet, this has remained a long standing puzzle, as no clear encoding of movement kinematics has been found. How the cerebellum learns has been difficult to approach because the problem of encoding remains unresolved. In Chapter 4 we approach this problem from a new direction: we propose that the cerebellum is composed of micro-clusters of P-cells, organized based on their preference for error. When the cells are organized in this manner, a clear encoding of kinematics emerges.

*Thesis Advisor:* Reza Shadmehr, Ph.D.

*Secondary Reader:* John Krakauer, M.D.

*Thesis Committee:* John Krakauer, M.D., Amy Bastian, Ph.D., P.T.

# Acknowledgements

This research was furthered, in part, by the National Institute of Biomedical Imaging and Bioengineering (grant number, T32EB003383), the National Institute of Neurological Disorders and Stroke (grant numbers R01NS078311 and F31NS090860), and by the National Eye Institute (grant numbers R01EY019258 and R01EY023277).

This work would not have been possible without the support of my dissertation advisor, Dr. Reza Shadmehr. His invaluable guidance and support both scientifically and personally, have been crucial to my success as a doctoral student. Over the past five years, our relationship has grown from that of a mentor and mentee to a truly collaborative enterprise. Many of the fundamental advances that we have achieved are largely a result of this collaborative relationship.

I am indebted to Robijanto Soetedjo and Yoshiko Kojima from the University of Washington National Primate Center. They kindly provided all of the Purkinje cell data used in the analysis of cerebellar saccades (Chapter 4). Their data and expertise has provided significant headway in our understanding of the cerebellum and its role in motor learning. Their valiant data collection and analysis efforts will continue to push the limits of knowledge in motor control for years to come.

I would also like to specifically thank my colleague in the Laboratory for Computational Motor Control, Pavan A. Vaswani. His critical thinking abilities and open mind have significantly extended the scope of my thesis work. Specifically, his

advice was invaluable in the implementation of the error-sensitivity hypothesis presented in Chapter 2.

I am extremely grateful to Allison Murawski for her patience, dedication, and support. She is continually a source of inspiration and her companionship has made my time in Baltimore one of the happiest of my life.

Finally, I would like to thank my family and friends, who provided continued support during my years in Baltimore. *No man is an island* ~ John Donne.

## TABLE OF CONTENTS

<b>Abstract</b> . . . . .	<b>ii</b>
<b>Acknowledgements</b> . . . . .	<b>iv</b>
<b>Table of Contents</b> . . . . .	<b>vi</b>
<b>List of Figures</b>	
<b>1 Introduction</b>	
1.1 Background . . . . .	2
1.1.1 Previous models of error-based learning . . . . .	2
1.1.2 Previous experimental evidence that error-sensitivity is not static	3
1.1.3 Previous experimental evidence for modulation of error-sensitivity	5
1.2 Specific aims . . . . .	6
1.2.1 Determine the rules that govern control of error-sensitivity . . . . .	7
1.2.2 Identify the locus of error-sensitivity in the human brain . . . . .	8
1.2.3 Purkinje cells in the cerebellum encode movement kinematics necessary for motor learning . . . . .	9
1.3 Significance . . . . .	11
<b>2 A memory of errors in sensorimotor learning</b>	
2.1 Introduction . . . . .	13
2.2 Materials and Methods . . . . .	13
2.2.1 Experiment 1 . . . . .	14
2.2.2 Experiment 2 . . . . .	19
	vi

2.2.3	Model of error-sensitivity . . . . .	20
2.2.4	Experiment 3 . . . . .	22
2.2.5	Experiment 4 . . . . .	24
2.2.6	Data collection and statistical analysis . . . . .	26
2.3	Results . . . . .	27
2.3.1	History of errors alters error-sensitivity . . . . .	27
2.3.2	Error-sensitivity is local to the experienced error . . . . .	29
2.3.3	Computational model of error-sensitivity . . . . .	31
2.3.4	Experimental tests of error-sensitivity model predictions . . . . .	32
2.3.5	Error-sensitivity model explains savings and meta-learning . . . . .	36
2.4	Discussion . . . . .	39
2.4.1	Why does learning from error saturate with large perturbations? . . . . .	40
2.4.2	Why does error-sensitivity depend on the mean of the perturbation distribution? . . . . .	44
2.4.3	Why does error-sensitivity depend on the sequential order of the perturbation distributions? . . . . .	47
2.4.4	Model explains data attributed to structural learning . . . . .	48
2.4.5	Model explains savings following washout . . . . .	51
2.4.6	Model explains savings in a gait-adaptation experiment . . . . .	52
2.4.7	Model explains the limited range of savings . . . . .	53
2.4.8	Model explains savings that was attributed to reinforced repetition . . . . .	54
2.4.9	Why do gradual perturbations sometimes produce savings, and sometimes not? . . . . .	55
		vii

2.5	Conclusion . . . . .	56
<b>3</b>	<b>Contributions of the cerebellum to acquisition of motor memories</b>	
3.1	Introduction . . . . .	59
3.2	Materials and Methods . . . . .	61
3.2.1	Experiment 1: cerebellar stimulation . . . . .	62
3.2.2	Experiment 2: motor cortex stimulation . . . . .	63
3.2.3	Behavioral procedures . . . . .	63
3.2.4	Data collection and statistical analysis . . . . .	67
3.3	Results . . . . .	69
3.3.1	Effects of tDCS in the null field . . . . .	70
3.3.2	Effects of tDCS on reaction time . . . . .	71
3.3.3	Effect of cerebellar stimulation on learning from error . . . . .	73
3.3.4	Robustness of statistical results . . . . .	76
3.3.5	Effects of motor cortex stimulation on learning from error . . . . .	77
3.3.6	Effects of cerebellar stimulation on stability of the motor memory . . . . .	78
3.3.7	Effects of motor cortex stimulation on stability of the motor memory . . . . .	81
3.3.8	Effect of stimulation on generalization . . . . .	82
3.3.9	Effect of stimulation on over-night retention . . . . .	83
3.4	Discussion . . . . .	86
3.4.1	Feedback control . . . . .	86
3.4.2	Learning from error . . . . .	87
3.4.3	Functional stages of motor memory . . . . .	89



3.4.4	Retention . . . . .	91
3.4.5	Generalization . . . . .	93
3.4.6	Limitations . . . . .	93
3.5	Conclusion . . . . .	95
<b>4</b>	<b>Encoding of action by the Purkinje cells of the cerebellum</b>	
4.1	Introduction . . . . .	97
4.2	Materials and Methods . . . . .	98
4.2.1	The dataset . . . . .	98
4.2.2	Statistical analysis . . . . .	100
4.2.3	Complex-spikes . . . . .	100
4.2.4	Population response . . . . .	102
4.3	Results . . . . .	104
4.3.1	P-cell population response predicts saccade speed . . . . .	104
4.3.2	P-cell responses, organized by error, predict saccade kinematics as a gain-field . . . . .	111
4.3.3	Individual Purkinje cells do not predict saccade kinematics . . . . .	114
4.3.4	Anatomical distribution of Purkinje cells supports population hy- pothesis . . . . .	121
4.4	Discussion . . . . .	123
4.4.1	The importance of bursting and pausing cells in the population . . . . .	124
4.4.2	The importance of organizing the P-cells via their complex-spike properties . . . . .	126
4.5	Conclusion . . . . .	127

<b>5</b>	<b>Conclusions &amp; Future Directions</b>	
5.1	Future directions . . . . .	133
5.1.1	Cerebellar organization provides clues for determining the neuro- physiological basis of adaptation . . . . .	134
	<b>Bibliography</b> . . . . .	<b>139</b>
	<b>Curriculum Vitae</b> . . . . .	<b>156</b>

## LIST OF FIGURES

2.1	History of error alters error-sensitivity . . . . .	29
2.2	Error-sensitivity is a local function of experienced errors . . . . .	30
2.3	Theoretical model and Experiment 3 . . . . .	33
2.4	Comparison of model and experimental results for Exp. 3 . . . . .	35
2.5	Saving occurs only when previously experienced errors are re-visited .	37
2.6	Additions groups in Experiment 4 . . . . .	39
2.7	Model accounts for a large body of experimental results . . . . .	43
2.8	Detailed model explanations of previous experimental results . . . . .	46
3.1	Experimental protocol and effects of tDCS on feedback control . . . . .	64
3.2	Reach kinematics and measures of error-dependent learning during cerebellar and motor cortex stimulation . . . . .	75
3.3	Force in error-clamp blocks and measures of decay in motor output during cerebellar or motor cortex stimulation . . . . .	79
3.4	Effects of tDCS on generalization and retention . . . . .	84
4.1	A population of burst and pause P-cells together predict eye speed in real-time . . . . .	105
4.2	Firing rates of individual P-cells as a function of saccade amplitude and peak speed . . . . .	106
4.3	Determination of complex-spike (CS) properties of P-cells . . . . .	108

4.4	Complex spikes encode direction of error, not direction of the saccade that preceded the error . . . . .	109
4.5	A cluster of P-cells, organized by their complex-spikes, produced a population response that predicted in real-time the motion of the eye	111
4.6	Population response of P-cells predicted saccade speed and direction in real-time as a gain-field . . . . .	113
4.7	The simple-spike population response of P-cells, organized by their complex-spike properties correlates with motion of the eye in real-time	116
4.8	Mean and peak/trough firing rates of the burst and pause cells were poorly modulated by saccade direction . . . . .	118
4.9	Population of P-cells predicted the real-time speed of the eye better than activity of individual cells . . . . .	120
4.10	Change in saccade direction was associated with a change in the time of pause cells . . . . .	120
4.11	CS-dependent organization of the P-cells . . . . .	122
4.12	The population response was sensitive to the fraction of pause and burst cells that composed a cluster of P-cells . . . . .	125
4.13	Gain-field encoding of saccade kinematics in the population response of P-cells disappeared if the P-cells were organized by their simple-spike activity . . . . .	129

# 1 INTRODUCTION

When we experience an error during a motor task, we adapt our next movement to partially compensate. Even when perturbations are random, trial-to-trial changes in the motor output are still present (Donchin et al., 2003; Izawa et al., 2008; Marko et al., 2012), demonstrating that on every movement the brain learns from error. One potential mathematical formulation of this trial-to-trial learning is via “state-space” models (Thoroughman and Shadmehr, 2000; Smith et al., 2006).

To explain this idea, suppose that on trial  $n$  a perturbation  $x^{(n)}$  is imposed on motor command  $u^{(n)}$ , so that the sensory consequences that are observed by the learner are  $y^{(n)} = u^{(n)} + x^{(n)}$ . The learner has a prediction about the sensory consequences  $\hat{y}^{(n)}$ , and updates its belief about the perturbation state,  $\hat{x}^{(n)}$ , from the prediction error  $e^{(n)} = y^{(n)} - \hat{y}^{(n)}$ . This error signal is fundamental to adaptation, producing learning, as evidenced by changes in the motor command,  $u^{(n)}$  in the next trial (Thoroughman and Shadmehr, 2000). However, error is not the only factor that affects how much is learned from trial-to-trial. An equally important factor is error-sensitivity. This sensitivity represents the fraction of the error that is compensated. Mathematically, the belief about the state of the environment on trial  $n$ , written as  $\hat{x}^{(n)}$ , can be related to the belief on trial  $n + 1$  as

$$\hat{x}^{(n+1)} = \alpha \hat{x}^{(n)} + \eta^{(n)} e^{(n)}, \quad (1.1)$$

where  $\eta^{(n)}$  represents error-sensitivity. Eq. 1.1 (representing a single state model), and its extensions to multiple timescales of memory, have been successfully applied to a variety of motor adaptation tasks (Smith et al., 2006; Fine and Thoroughman,

2006; 2007; Kording et al., 2007).

## 1.1 Background

In Eq. 1.1, we can quantify the amount of trial-to-trial learning, or adaptation, from trial  $n$  to  $n + 1$  as the change in the subject’s belief about the perturbation state:

$$\begin{aligned}\Delta\hat{x}^{(n+1:n)} &= \hat{x}^{(n+1)} - u^{(n)} \\ \Delta\hat{x}^{(n+1:n)} &= \hat{x}^{(n+1)} - \alpha\hat{x}^{(n)} - \eta^{(n)}e^{(n)}\end{aligned}\tag{1.2}$$

That is, the learner updates their belief about the environment based on two separate factors,  $\alpha\hat{x}^{(n)}$ , which is a decay of the motor memory back to its baseline level, and  $\eta^{(n)}e^{(n)}$ , which represents the amount of learning due to an error. In this section, we review the relevant literature, describing how the brain determines how much to learn from a given error.

### 1.1.1 Previous models of error-based learning

Earlier models of learning featured a constant ‘learning rate’ parameter, which corresponds to constant error-sensitivity. In these models, we can replace  $\eta^{(n)}$  with  $\eta$ , indicating that the value of the error-sensitivity parameter does not depend on the trial number,  $n$ . For instance, classical conditioning experiments were modeled using sensory prediction errors (Rescorla and Wagner, 1972). In these studies, a single learning-rate variable was used to update model weights based on the error in the current trial. More recently, behavioral results from saccade adaptation studies and reaching studies have been modeled via ‘single-state’ or ‘two-state’ models in which one or two error-sensitivity terms are used, but both are kept constant throughout training (Thoroughman and Shadmehr, 2000; Smith et al., 2006).

Therefore, such models do not vary error-sensitivity. If we were to plot the amount

of adaptation,  $\Delta\hat{x}^{(n+1:n)}$ , as a function of the error size experienced in trial  $n$ , we would see that the function is linear, where the slope of the line is the error sensitivity,  $\eta$ . That is, these models predict that the learning should adapt their movements based on the same fraction of the error, regardless of the error's magnitude. We will see in Chapter 2 that human subjects do not necessarily follow this pattern (see also Section 1.1.2).

Most recently, the Kalman-filter framework has been used to model adaptation (van Beers, 2009). In these models, error-sensitivity changes during learning based on the uncertainty that the learner has with regard to its prediction  $\hat{y}^{(n)}$  versus observation  $y^{(n)}$ . For example, if the uncertainty associated with  $\hat{y}$  is large, error-sensitivity is large, making the learner learn more from the prediction error. From a theoretical perspective, this approach is attractive, but one problem is that in the Kalman framework uncertainty rapidly converges to the values described by the noises in the state update and measurement equations. It is also unclear how the learner should estimate these noises. In addition, the error-sensitivity term used in these models, termed the Kalman gain, is the same for all error magnitudes, which we will show is not the case for human motor adaptation.

### **1.1.2 Previous experimental evidence that error-sensitivity is not static**

As described in the previous section, most mathematical models of motor learning assume that error-sensitivity,  $\eta$ , is fixed both in time as well as a function of error. In this section, we outline previous experimental evidence suggesting that these assumptions do not fully describe the behavior of human subjects.

Marko et al. (2012) provide compelling evidence that error-sensitivity is not fixed as a function of error magnitude. That is, subjects do not learn the same fraction of the error for every error magnitude that they experience. Briefly,

subjects held the handle of a robotic manipulandum, and were asked to “shoot” to a target presented directly in front of them. In some trials, the robot applied forces on the subject’s hand. These forces were small, medium, or large, resulting in proprioceptive errors of varying magnitudes. The authors then assessed the amount of trial-to-trial learning via Eq. 1.2. They found that subjects learned a relatively high fraction of the error when the error was small compared to when the proprioceptive error was large. That is  $\eta$  for a small error was significantly larger than  $\eta$  for a large error. The author’s repeated this experiment, providing subjects with visual rather than proprioceptive perturbations, and found similar results. Together, these results suggest that error sensitivity is not constant as a function of error magnitude, but rather decays for large errors.

Reanalysis of previous psychophysical results also support the assertions made by Marko et al. (2012). For instance, Wei and Körding (2009) performed a reaching experiment in which the subject’s reach was visually perturbed in a direction perpendicular to the reach direction. The author’s used several visual perturbation magnitudes and could therefore assess how much the subjects deviated from a perfectly straight movement. Similar to the result of Marko et al. (2012), subjects did not learn the same from all perturbation magnitudes, rather subjects learned fractionally more from the small visual perturbations.

Similar results were obtained by Fine and Thoroughman (2006), where the authors perturbed the reaching movements of healthy human subjects via single force pulses of various magnitudes. They also found that as a percentage, the error-sensitivity was not constant, but large for small errors and small for large errors. That is, error-sensitivity fell as a function of the size of the perturbation (and, similarly, error).



### 1.1.3 Previous experimental evidence for modulation of error-sensitivity

Previous computational models of motor learning suggest that error-sensitivity is fixed as a function of error magnitude and also as a function of trial. In the previous section, we found that multiple studies in humans suggest that error-sensitivity is not constant, but changes as a function of error magnitude. In this section, we review experimental evidence suggesting that error-sensitivity can also change as a function of trial. That is, error-sensitivity may not be fixed over the course of an experimental session.

Previous work by Smith and Shadmehr (2004) found that people change their error-sensitivity as a function of history of previous perturbations. In that task, volunteers ( $n = 5$ ) reached to a target while holding a robotic arm that produced force perturbations. The force perturbation on trial  $n + 1$ , written as  $f^{(n+1)}$ , was described by:  $f^{(n+1)} = af^{(n)} + \epsilon$  where  $\epsilon \sim \mathcal{N}(0, \sigma^2)$ . When  $a$  was close to 1.0, the perturbations in trial  $n$  and  $n + 1$  were highly correlated. Similarly, when  $a$  was near -1.0, the two perturbations were anti-correlated. When there is a high correlation between the perturbations in two successive trials, learning the perturbation in the first trial would benefit performance in the next trial. However, when there is low or negative correlation between two trials, learning from error in one trial will not help next-trial performance. The authors found that in the case where the perturbations were correlated ( $a = 0.9$ ), the subjects learned from a given error significantly more than when the same error was experienced in an anti-correlated ( $a = -0.9$ ) environment.

Together with the previous section, these results suggest that error-sensitivity is likely not fixed as a function of trial - our model must contain  $\eta^{(n)}$  rather than  $\eta$ . In addition, it is likely that error-sensitivity is dependent on the

magnitude of the experienced error:  $\eta^{(n)}(e^{(n)})$ .

## 1.2 Specific aims

When we experience an error during a motor task, we alter the motor commands on the next trial in an attempt to partially compensate for the error. This partial compensation depends on two important variables: error, and error-sensitivity. Error is the difference between predicted and observed sensory consequences of motor commands. Error-sensitivity is the “learning rate” that determines how much we learn from error. Previous motor control studies have assumed that error-sensitivity is independent of error. However, recent results suggest that people change their error-sensitivity as a function of previous perturbations as well as a function of perturbation magnitude.

Suppose that, in principle, the brain could modulate the amount it learns from an error. How might this be done? Using behavioral studies, we suggest that the brain modulates error-sensitivity by keeping a history of recent errors. That is, in contrast to current models of learning in which the objective is to predict a sensory outcome, here we suggest that the brain also stores a memory of errors that were experienced in previous movements.

Using behavioral psychophysics, we first identify the rules that the brain uses to systematically modulate how much it is willing to learn from error. Next, we identify the primary brain region responsible for control of this “learning rate” parameter, by modulating the brain’s activity using non-invasive stimulation. Finally, we analyze neurophysiology from the cerebellum, the first step in a process to determine the neural mechanisms underlying cerebellar-dependent learning.

### 1.2.1 Determine the rules that govern control of error-sensitivity

When we make a movement error, we modify our motor output in the next trial to compensate for the error. Previous models of motor adaptation have assumed that error-sensitivity is independent of error. However, this assumption is inconsistent with experimental data as well as our results. These results suggest that the brain may have a method for adjusting error-sensitivity based on both the history of past errors and the size of the errors. To account for the behavioral data, we propose a theoretical framework in which the brain monitors the history of prior errors, and then modulates error-sensitivity using this stored memory (see Chapter 2). This theory links the history of experienced errors and error-sensitivity. In this computational model, error is encoded in the nervous system via a set of basis elements, where each basis element has a preferred error. The model predicts that error-sensitivity for a particular error magnitude would generalize to near-by error magnitudes, similar to how motor output generalizes to nearby target directions. We then provide experiments which identify the rules that govern control of error-sensitivity during motor learning in humans.

Understanding error-sensitivity is important because it provides insight into the phenomena of savings. Savings refers to the observation that when a subject practices a task with perturbation (A), and then the perturbation is removed (i.e., extended washout), they exhibit faster re-learning of (A). Remarkably, savings of (A) is present even when washout is followed by training in (B), a perturbation in the opposite direction. In addition, training in (A) and then washout can produce savings in (B), exhibiting a form of ‘meta-learning’. Current models of learning cannot account for these fundamental observations (Zarahn et al., 2008; Mawase et al., 2014). Here, we suggest that modulation of error-sensitivity may underlie both the phenomena of savings and meta-learning, and show that in theory, the

equations can account for these results. The new idea that emerges is that when we are better at compensating for a perturbation than before, it is not because we recall the actions we had produced or the perturbation we had estimated, but because we recall the errors we experienced during that learning process. By manipulating the history of experienced errors, thereby manipulating error-sensitivity, we show that the process of savings can be disrupted or enhanced. We provide experiments that directly test the hypothesis that a memory of errors can largely account for the phenomena of savings and meta-learning.

### **1.2.2 Identify the locus of error-sensitivity in the human brain**

When we interact with a novel object, we form a motor memory, which can be recalled the next time the object is encountered. In Chapter 2, we suggest that this motor memory includes a previously undiscovered memory: a memory of errors. This memory of errors modulates the error-sensitivity on each trial, affecting how much a subject learns from an experienced error. What brain region is responsible for modulating error-sensitivity?

Previous research suggests that two brain regions are crucially responsible for the integrity of such a motor memory: the motor cortex (M1) and the cerebellum. Which of these regions, if any, store a memory of errors and modulate the amount a subject is willing to learn?

In Chapter 3, we use transcranial direct current stimulation (tDCS) to alter the function of the cerebellum and motor cortex when healthy subjects are exposed to a novel motor task. Using this experimental manipulation, we can dissociate the roles of the cerebellum and motor cortex during formation of a new motor memory. We show that increasing the excitability of the cerebellum via anodal tDCS increases the rate of learning, while decreasing cerebellar excitability via cathodal

tDCS impairs the ability to respond to sensory feedback and decreases the rate of learning. These results provide evidence that the cerebellum is a crucial neural substrate responsible for modulation of error-sensitivity.

In addition, we suggest a critical role for the human cerebellum in the ability to correct for errors during movement and the ability to learn from that error. Our results suggest that during a reaching movement the errors experienced may be a result of feedback corrections. These corrections may be used in the subsequent learning process. That is, by altering the feedback responses of the subject, the error incorporated into the subsequent motor command via the error-sensitivity term can be modified.

Finally, we find that during the initial part of training, subjects' learning decays quickly, but with further training the decay slows, suggesting that with training the motor memory gained stability. Stimulation of the cerebellum or the motor cortex did not alter these decay patterns. Therefore, long-term storage of motor memory may be the responsibility of a separate region of the brain.

### **1.2.3 Purkinje cells in the cerebellum encode movement kinematics necessary for motor learning**

In Chapter 3 we suggest that the cerebellum is the crucial substrate necessary for storage of a memory of errors. That is, one of the primary roles of the cerebellum during motor learning may be to store the history of prior errors, and thereby modulate error-sensitivity, incorporating a portion of the error into the subsequent motor command.

However, the neural mechanisms in the cerebellum that underlie adaptation remain poorly understood. While the principal cells of the cerebellum, Purkinje cells (P-cells), show some modifications following motor learning (Kojima et al., 2010), changes in the responses of these cells has not been effectively linked to

behavioral adaptation. Given that the integrity of the cerebellum is required for motor adaptation, why have the neural basis of adaptation remained elusive?

Execution of accurate movements depends critically on the cerebellum, indicating that the primary cells of the cerebellum, Purkinje cells, should relate to the kinematic parameters of movement. However, this encoding has remained unclear: Purkinje cells show little consistent modulation with respect to kinematic parameters. Without a clear understanding of this encoding, the problem of learning from error, which entails a change in this encoding, has been even more difficult to solve.

In Chapter 4, we show that the activity of the cerebellum during cerebellar-dependent movements cannot be understood by recording the responses of individual cells. Rather, we suggest that populations of approximately 50 Purkinje cells combine their responses via common projections to a neuron in the deep cerebellar nuclei. We hypothesize that the presynaptic neurons that project to an individual cell in the deep cerebellar nuclei are not selected randomly, but rather share a common preference for error.

That is, imagine that errors divide the cerebellum into anatomical micro-clusters, where each micro-cluster is composed of a relatively small number of Purkinje cells that each prefer a specific error. Further imagine that this population of Purkinje cells project to a common neuron in the deep cerebellar nucleus.

In Chapter 4, we provide evidence that when the responses of Purkinje cells are organized via this hypothesized anatomy, a beautiful encoding of kinematics emerges. The combined population response shows gain-field encoding of speed and direction, similar to the encoding of movements previously reported in posterior parietal cortex (PPC).

We suggest that the neural basis of adaptation can only be understood by using this crucial anatomical organization. Only after establishing the kinematic

encoding of movements in Purkinje cell responses, can we begin to address the encoding of adaptation, which represents a fundamental change in this encoding.

## 1.3 Significance

Error-sensitivity, or the rate of learning from error, is directly relevant to motor rehabilitation. In this document, we outline the rules that govern modulation of error-sensitivity in humans. Increasing error-sensitivity should produce faster adaptation, affecting the duration of rehabilitation. In addition, we provide evidence in Chapter 2 that error-sensitivity is linked to the phenomenon of savings (i.e., faster re-learning). Understanding this link may provide clue to effectively apply rehabilitation techniques to promote faster re-learning outside the clinic.

In Chapter 3, we begin to identify the neural structures involved in modulation of error-sensitivity. Our results suggest that control of error-sensitivity and its affects on behavior are likely cerebellar dependent. Clinically, patients with cerebellar disease are faced with poor prognoses through conventional rehabilitation. Investigation into the mechanisms by which this error-sensitivity is modulated, either through novel behavioral rehabilitation techniques or brain stimulation, may help translate these findings into novel clinical applications in the future.

Finally, in Chapter 4, we lay the foundation for understanding error-sensitivity via the cellular signals present in the primate cerebellum. These results provide a new framework for understanding the role of the cerebellum in motor adaptation.

## 2 A MEMORY OF ERRORS IN SENSORIMOTOR LEARNING

When we learn to control a novel object, our brain stores a motor memory which we can recall when we are exposed to that object again. The current view of motor memories is that when we are re-exposed to a novel object, the brain recalls the motor commands,  $u$ , that it previously learned. In this view, motor memory is a memory of motor commands, acquired through trial-and-error and/or reinforcement.

In this chapter, we suggest that the brain systematically controls how much it is willing to learn from the current motor error. That is, the brain uses a principled mechanism which determines how much it is willing to learn following an error. Manipulation of this ‘error-sensitivity’ parameter depends on the history of past errors. This suggests that the brain stores a previously unknown form of memory, a memory of errors.

We then provide a mathematical formulation of this foundational idea. In this computational model, error is encoded in the nervous system with a set of basis elements, where each basis element has a preferred error. The model predicts that error-sensitivity for a particular error magnitude would generalize to near-by error magnitudes, similar to how motor output generalizes to nearby target directions. We use this mathematical model to provide insights into a host of previously puzzling experimental data, including savings and meta-learning. Taken together, our results demonstrate that when we are better at a motor task upon re-exposure, it is partly because the brain recognizes the errors it previously experienced.



## 2.1 Introduction

How does the brain alter behavior after experiencing an error? Classic theories assumed that the brain learns some fraction of the error regardless of its history or magnitude (Jordan and Rumelhart, 1992; Kawato et al., 1987). However, recent experiments (Robinson et al., 2003; Soetedjo et al., 2009; Marko et al., 2012) demonstrate that the brain learns relatively more from small errors than large errors, and can modulate its error-sensitivity (Smith and Shadmehr, 2004; Trent and Ahmed, 2013; Gonzalez Castro et al., 2014).

Understanding error-sensitivity is important because it may provide insight into the phenomena of ‘savings’ and ‘meta-learning’. Savings refers to the observation that when a subject adapts to perturbation (A), and then the perturbation is removed (i.e., washout), they exhibit faster readaptation to (A) (Kojima et al., 2004). Remarkably, savings of (A) is present even when washout is followed by adaptation to (-A), a perturbation in the opposite direction (Malone et al., 2011; Sarwary et al., 2013). Current error-dependent models of learning cannot account for these observations (Zarahn et al., 2008; Mawase et al., 2014), nor explain meta-learning, where prior exposure to a random perturbation produces savings (Turnham et al., 2012; Braun et al., 2009).

## 2.2 Materials and Methods

We tested  $n = 113$  naïve human subjects in four experiments. In the first experiment, we used a between-subject design to assess subjects’ error-sensitivity after experiencing one of three different environments. In Experiment 2, we attempted to determine whether changes in error-sensitivity were local to the experienced errors. This experiment features a with-in subject design, which allowed

us to probe error-sensitivity across a range of error magnitudes. Using the results of these first two experiments, we developed a computational model of how the brain modulates error-sensitivity based on the history of previously experienced errors. In Experiment 3, we explicitly tested the predictions of our computational model. Finally, in Experiment 4, we tested whether the behavioral phenomena of savings and meta-learning can be explained by our computational model.

### 2.2.1 Experiment 1

In Exp. 1 we used a between-subject design to test the idea that the nervous system could modulate error-sensitivity. To do so, we used a constant perturbation in probe trials to produce an error in the movement and quantified how much the nervous system learned from this error. Volunteers ( $n = 27$ ,  $23.6 \pm 4.3$  years old, mean  $\pm$  SD, 16 female) were asked to hold the handle of a robotic arm and make rapid out-and-back reaching movements to a target presented 10cm directly in front of them. In some trials, their reach was perturbed by a velocity dependent curl field, where force was related to hand velocity as:

$$\mathbf{f} = \begin{bmatrix} 0 & b \\ -b & 0 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix}. \quad (2.1)$$

In Eq. 2.1,  $\dot{x}$  and  $\dot{y}$  are the horizontal and vertical components of the subject’s hand velocity. The force pushed the subject’s hand perpendicular to the direction of movement on the outward reach, but was turned off on the reach back.

Perturbations were either clockwise ( $b = 13$  N.s/m) or counter-clockwise ( $b = -13$  N.s/m). The subject’s hand was occluded by an opaque horizontal screen located above the plane of the arm. An overhead projector displayed information about hand position and targets on this screen. Continuous feedback of the location of the hand was presented via a cursor (0.3cm in diameter). At the onset of each trial,

subjects were presented with a start circle (1.0cm in diameter). Once the hand was placed inside this circle, following a randomly chosen inter-trial interval [0.25 - 0.75]sec a target circle (1.0cm in diameter) was displayed and an auditory tone was played. The display of the target and the sound of the tone served as the ‘go’ instruction. If the hand passed through the target circle in  $300 \pm 30$ ms, the subject was rewarded with an animation of an explosion, an auditory tone, and a point added to their score. If the hand passed through the circle after 330ms, the target circle turned blue (indicating a movement that was too slow). Otherwise, if the subject’s hand arrived at the target circle in less than 270ms, the circle turned red. The subjects were not required to have the turn-around point of their reach in the target circle; rather, they merely had to pass through the target to obtain reward. In some instances, the reach missed the target entirely. In this case, no reward or timing feedback was provided to the subject. Subjects were instructed to obtain as many points as possible.

Consider an environment in which the perturbations tend to switch slowly, i.e., persist from trial to trial, as compared to one in which the perturbations are rapidly switch. We can define this environment in terms of a Markov chain in which the perturbations can take on one of two states (Fig. 2.1A, top). For example, the perturbation can be + 1 or -1, where +1 refers to a force field that pushes the hand clockwise, and -1 refers to a field that pushes the hand counter-clockwise. The perturbation state can change from one trial to the next, and this change is governed by a transition probability,  $z$ . In the slowly switching environment, the probability of staying in a given perturbation state is high ( $z = 0.9$ ), whereas in the rapidly switching environment, this probability is low ( $z = 0.1$ ). As a result, in the slowly switching environment the perturbations tend to repeat from one trial to the next, whereas in the rapidly switching environment the perturbations tend to change. We hypothesized that the brain would learn more from the error induced by

the perturbation in the slowly switching environment because that perturbation was likely to persist (learning from error in one trial would improve performance on the subsequent trial). However, in the rapidly switching environment the brain would suppress learning from error because the perturbation that produced that error was likely to change (any learning would be detrimental to performance on the subsequent trial).

To produce such environments, we considered a perturbation schedule that was stochastic, as illustrated by the Markov chain in Fig. 2.1A. The perturbation state, labeled by variable  $b$  (indicating the field produced by the robot) was a binomial, taking on one of two values  $[+13, -13]$ Ns./m. For example, suppose that on trial  $n - 1$  the perturbation state is  $b^{(n-1)}$ . Then the perturbation on trial  $n$  is determined by the following probabilities:

$$\begin{aligned} \Pr(b^{(n)} = b^{(n-1)}) &= z \\ \Pr(b^{(n)} = -b^{(n-1)}) &= 1 - z \end{aligned} \tag{2.2}$$

Eq. 2.2 implies that if  $z \approx 1$ , then the perturbation is likely to repeat, i.e., the environment is slowly switching and persistent. However, if  $z \approx 0$ , then the perturbation is likely to change; the environment is rapidly switching. We randomly divided our subjects into three groups and generated a single perturbation schedule for each group using Eq. 2.2. One group ( $n = 9$ ) experienced a slowly switching environment ( $z = 0.9$ ), another group ( $n = 9$ ) experienced a medium switching environment ( $z = 0.5$ ), and a final group ( $n = 9$ ) experienced a rapidly switching environment ( $z = 0.1$ ). All groups began their training in a baseline block (156 trials). In the baseline block there were no perturbations, except for occasional probe trials in which we measured error-sensitivity, described below and illustrated in the inset of Fig. 2.1A. Following the baseline block, subjects experienced 5 blocks of perturbation trials (225 trials each). In each perturbation block there were 5

mini-blocks (learning blocks). Each mini-block included 30 perturbation trials, 10 washout trials, and 5 probe trials. In each mini-block the number of trials with clockwise or counter-clockwise perturbation was equal to 15. In this way, the mean of the perturbations in each block, as well as the mean of the perturbations in each mini-block, was zero for all subjects regardless of the environment. Furthermore, variances of the perturbations were identical across groups. The critical difference was the order of the perturbations. The experiment lasted about an hour. Subjects were allowed a 1-3 minute break between each block of trials.

Our objective was to estimate error-sensitivity during each block and ask whether this quantity changed as the subjects experienced the various environments. We approached the problem by considering a standard model of learning (Mackintosh, 1975; Thoroughman and Shadmehr, 2000; Rescorla and Wagner, 1972; Pearce and Hall, 1980) in which on trial  $n$ , a perturbation  $x$  is imposed on action  $u$  so that the sensory consequences observed by the learner are  $y^{(n)} = u^{(n)} + x^{(n)}$ . On trial  $n$ , the learner predicts the sensory consequences  $\hat{y}^{(n)} = u^{(n)} + \hat{x}^{(n)}$ , and updates its belief about the state of the perturbation from the prediction error  $e^{(n)} = y^{(n)} - \hat{y}^{(n)}$ . Such learning typically depends on a decay factor  $\alpha$ , and error-sensitivity  $\eta^{(n)}$ :

$$\hat{x}^{(n+1)} = \alpha \hat{x}^{(n)} + \eta^{(n)} e^{(n)}. \quad (2.3)$$

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

$$u^{(n)} = u_0 - \hat{x}^{(n)}. \quad (2.4)$$

To measure error-sensitivity, we used probes that consisted of pairs and triplets of

error-clamp trials (Marko et al., 2012). An error-clamp (Scheidt et al., 2001) is a trial in which the robot produces a channel with stiff walls along a line connecting the start position to the target, thereby reducing deviations from a straight line, eliminating error from that trial while allowing one to measure the forces that the subject produces against the channel walls. The error-clamp had the following properties: spring coefficient = 6000 N/m, damping coefficient = 250 N.s/m. On error-clamp trial  $n$ , our proxy for motor output was  $u^{(n)}$ . To find  $u^{(n)}$ , we first regressed the measured force  $f(t)$  that the subject had produced against channel walls onto the ideal force  $f^*(t) = by(t)$  and found the parameters  $k_0$  and  $k_1$  that minimized the quantity  $J = (f(t) - k_1 f^*(t) - k_0)^2$ , and then set  $u^{(n)} = k_1$ . To measure error-sensitivity  $\eta^{(n)}$ , we first used Eq. 2.3 to estimate  $\alpha$  for each subject from all pairs of error-clamp trials that did not have a perturbation (Fig. 2.1A, green probe trials). As the subject did not experience an error in the first error-clamp, the forgetting factor was found by dividing the motor commands in the two trials:  $\alpha = u^{(n+1)}/u^{(n)}$ . Next, we used this estimate of  $\alpha$  to estimate the error-sensitivity,  $\eta^{(n)}$ , from each triplet of error-clamp trials (Fig. 2.1A, lavender probe trials) in which there was a perturbation in the middle trial. The perturbation in this probe was always a counter-clockwise field ( $b = -13$  N.s/m). As a result, we have:

$$\eta^{(n)} = \frac{u^{(n+1)} - \alpha^2 u^{(n-1)}}{e^{(n)}} . \quad (2.5)$$

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

### 2.2.2 Experiment 2

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

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

We enrolled a new group of right-handed volunteers ( $n = 20$ ,  $24.1 \pm 4.5$  years old, mean  $\pm$  SD, including 10 females) who were naïve to the purpose of the experiment. As in Exp. 1, subjects were asked to make rapid out-and-back reaching movements to a target at 10cm. However, unlike Exp. 1, there were no forces to perturb the movement (all movements were in error-clamp). Instead, we perturbed the visual feedback associated with position of the hand. At 100ms after reach onset, we removed the visual feedback, and then re-displayed hand position at the turn-around point of the reach by placing a stationary yellow dot at that location (Fig. 2.2A). In some trials the location of this dot was perturbed by either a 1.1x (magnifying) or a 0.9x (minifying) gain. We restored visual feedback of the hand after this turn-around point, but manipulated the location of the cursor using a gain so that it appeared that the subject had reached to the location indicated by the yellow dot. A trial was considered successful if the yellow dot fell within the target circle. Subjects were rewarded by a visual animation of an explosion, and the addition of a point to their score. They were instructed to maximize the total number of rewarded trials.

The difference between the visual feedback and the target position was our proxy for error. The brain responded to this error by changing the motor commands on the next trial, increasing or decreasing the extent of the reach. Because visual feedback was not available during the outward portion of the reach, the design of

the experiment allowed us to measure error-sensitivity at every trial (the change in the magnitude of the reach divided by the experienced error). Suppose that on trial  $n$ , a perturbation was imposed, resulting in an error  $e^{(n)}$  (defined as the difference between target position and the cursor position displayed to the subject to indicate their hand’s turn-around point). We estimated the forgetting factor in pairs of consecutive error-clamp trials via a technique identical to Exp. 1 (Fig. 2.2A, lavender trials).

Subjects were divided into two groups ( $n = 10$  in each group). Both groups experienced a baseline block (100 trials, no perturbations). Following the baseline block, Group 1 (Fig. 2.2B) experienced three perturbation blocks, each 387 trials, composed of a slowly switching environment ( $z = 0.9$ ), a medium switching environment ( $z = 0.5$ ), and a rapidly switching environment ( $z = 0.1$ ). Group 2 experienced the reverse sequence of environments. All subjects in each group experienced the same perturbation schedule. Each block was composed of 9 mini-blocks (30 perturbation trials, 10 no perturbations, and a probe triplet of trials). The mean of the perturbations within each mini-block, as well as the mean of the perturbations within each block, was zero. We found that a change in perturbation statistics resulted in a change in error-sensitivity (Fig. 2.2C), and that the largest changes occurred where subjects experienced the majority of the errors (Fig. 2.2E and Fig. 2.2F). This suggests that error-sensitivity was a function of the experienced error.

### 2.2.3 Model of error-sensitivity

When participants experience a prediction error, they update their motor command on the next trial to compensate for a fraction of that error. This can be mathematically described by a state space model (Eqs. 2.3 and 2.4), where  $e^{(n)}$  is the error and  $\alpha$  is a retention factor. Eq. 2.3 describes a model in which the learner



uses prediction error to form an estimate of the state of the environment, resulting in a memory of that state. Because our results from Exp. 2 suggested that error-sensitivity was a function of error, we constructed a new set of equations to account for a memory of errors. In this model, the learner has a set of basis elements with which it encodes the error experienced on a given trial. In our simulations, we assumed that  $N$  basis functions with centers located at  $\check{e}_i$  that were uniformly distributed throughout a symmetric error space  $\check{e}_i \in [-P, P], i = 1 \dots N$ . In addition, we assumed that at the beginning of the simulations all weights were equal (i.e. there was a constant error-sensitivity across all error magnitudes),  $\mathbf{w}^{(0)}$ . Therefore, our model that learned to represent the state of the environment had one parameter,  $\alpha$ . Our model that controlled error-sensitivity had two parameters:  $\sigma$ , and  $\beta$ . In total, our model had 3 parameters. In Fig. 2.3B we implemented this model with 10 basis elements, equally spaced between -5 and 5,  $\sigma = 1$ ,  $\alpha = 1$ , and  $\beta = 0.05$ .

In the model, error-sensitivity is a function of error size, and so any change in sensitivity is local to the errors experienced in the recent trials. Suppose that on trial  $n - 1$ , the motor command produces error  $e^{(n-1)} = -1$ . If on the next trial the error  $e^{(n)}$  is of the same sign as  $e^{(n-1)}$ , then  $\text{sign}(e^{(n-1)}e^{(n)}) = 1$  and error-sensitivity is increased around the neighborhood of  $e^{(n-1)}$  (Fig. 2.3A, top). As a result, learning from error is increased about  $e = -1$ , as illustrated by the red line in 2.3B. This means that if this error is ever experienced again, the system will learn more from it than before. On the other hand, if  $e^{(n)}$  is of the opposite sign as  $e^{(n-1)}$ , (Fig. 2.3A, bottom),  $\text{sign}(e^{(n-1)}e^{(n)})$  and sensitivity is decreased about  $e = -1$ , resulting in reduced learning from error around this neighborhood (as illustrated by the blue line in Fig. 2.3B). In Fig. 2.3C we simulated the model in the slow ( $z = 0.9$ ), medium ( $z = 0.5$ ), and rapidly switching ( $z = 0.1$ ) environments (identical parameters as in Fig. 2.3B except  $\beta = 0.001$ ). In all cases, the model learns from error on each trial, as illustrated by the gray line in Fig. 2.3C. However, the errors

in the slowly switching environment tend to repeat, that is  $E[\text{sign}(e^{(n-1)}e^{(n)})] > 0$ , where  $E[\ ]$  is the expected value operator. As a result, in the slowly switching environment error-sensitivity increases, producing an increase in learning from error (Fig. 2.3D, red-line). The errors in the medium switching environment have the following structure:  $E[\text{sign}(e^{(n-1)}e^{(n)})] \approx 0$ . This produces little or no change in sensitivity, resulting in little or no change in learning from error, as illustrated by the green line in Fig. 2.3D. Finally, errors in the rapidly switching environment have the following structure:  $E[\text{sign}(e^{(n-1)}e^{(n)})] < 0$ , that is, error in one trial is usually of the opposite sign of the error in the previous trial. As a result, error-sensitivity decreases (Fig. 2.3D, blue line). The learning from error curves are qualitatively similar to those measured experimentally in Exp. 2 (Fig. 2.3D).

### 2.2.4 Experiment 3

In Exp. 3 we set out to test a critical prediction of the model: that by manipulating the history of errors that were experienced by the subject, we could simultaneously increase error-sensitivity for one range of errors, while decreasing it for another range. We recruited a new group of right-handed volunteers ( $n = 16$ ,  $25.8 \pm 2.6$  years old, mean  $\pm$  SD, including 6 female) who were naïve to the purpose of the experiment. They held a handle attached to a stationary force transducer. The handle was located 20-30cm in front of the subject, such that they could push against it comfortably while seated. The subject’s hand was hidden from view by an opaque horizontal screen. Feedback regarding force generation was provided by an image projected on the screen. The objective of this isometric task was to produce a goal force of 16N.

At the onset of a trial, a start circle and a goal circle (both 0.75N in diameter) appeared. The goal circle was located approximately 15cm from the start circle. The screen was scaled such that a 15cm cursor displacement corresponded to

16N force. A cursor (0.3N diameter) appeared at trial onset. The displacement of the cursor corresponded to the total amount of force that the subject produced, multiplied by a scaling factor:  $s\sqrt{(f_x + f_y)^2}$ , where  $s$  corresponds to a scaling factor that maps units of force into screen displacement (i.e., 15cm/16N).

As the subject began pushing toward the target, we removed visual feedback when the cursor position reached 1/5 of the way to the target ( $>3.20\text{N}$ ), and then placed a yellow dot on the screen (0.5N diameter) in the location corresponding to the maximum force that they produced. In some trials we perturbed the location of this dot by adding an offset,  $x$ . Visual feedback of the cursor was then restored as the force produced by the subject returned to zero. We scaled the position of the cursor continuously during the return so that it appeared that the subject had produced the force signified by the perturbed dot. Once the cursor had returned to the starting circle, the maximum force dot remained on the screen for 0.5s before the cursor, maximum force dot, and target circle disappeared. The subject then waited for an inter-trial-interval to elapse (randomly chosen between [0.25, 0.75]sec) before the next trial began.

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

The perturbation schedule is shown in Fig. 2.3E. The perturbations were designed so that, in theory, subjects would increase their sensitivity to +4N and -4N errors (despite the fact that they never experienced a -4N perturbation), while simultaneously decreasing their sensitivity to +8N and -8N errors. The experiment began with a baseline block (50 trials, no perturbations). Following the baseline

block, we probed sensitivity to +8N and -4N perturbations (labeled as Probe 1 in Fig. 2.3E). We then exposed subjects to 20 repetitions of an alternating [+8N, -8N] environment, followed by 15 repetitions of a stable [0N, +4N] environment. Finally, we again probed sensitivity to +8N and -4N perturbations (labeled as Probe 2 in Fig. 2.3E). The experiment was divided into 6 blocks (105-120 trials each). The experiment took approximately 45 minutes to complete.

### 2.2.5 Experiment 4

According to our model, savings and meta-learning are largely due to a memory of errors. If so, specific manipulations of the history of errors should affect the presence or absence of savings and meta-learning. In Exp. 4, we tested some of these critical predictions. This experiment included  $n = 50$  subjects, 10 in each of the five groups ( $24.3 \pm 5.4$  years old, mean  $\pm$  SD, including 20 females). The subjects were naïve to the purpose of the experiment. They held the handle of the robotic manipulandum. The subject’s hand was hidden from view by an opaque horizontal screen. They were presented with a red circle (1cm diameter), which served as the start and end point for the trial. Subjects were asked to make a rapid shooting movement from the starting circle to a green target circle (0.5cm diameter) located 6cm directly in front of them. They were required to pass through the target within  $150 \pm 50$ ms. Movements that fell outside this range were signaled by the target circle changing color and a low frequency auditory tone. If the subject passed through the target circle within the time range, they were rewarded with an animation of an explosion and a point added to their score.

As subjects began moving towards the target (when the total velocity exceed 0.02m/s), we removed visual feedback for the remainder of the outward motion. When the reach exceeded 6cm eccentricity, a yellow dot (0.5cm diameter) was placed on the screen at that location. Visual feedback was withheld for the duration

of the trial. To aid the subject in returning to the start position, when the hand was within 3cm of the starting circle, the position of the cursor was provided by a white dot (1mm diameter) which blinked at 1Hz (20% duty cycle). In some trials, we manipulated the location of the yellow dot by rotating it relative to the target using the perturbation schedules shown in Fig. 2.5A and Fig. 2.6A.

All groups experienced a baseline block of 90 null (no perturbation) trials. Each group experienced a training condition (90 or 120 trials) followed by 120 trials of washout (null, N), and a final phase in which we tested adaptation to an abruptly imposed 30° counterclockwise (CCW) perturbation over 90 trials (A). The ANA group was trained on a 30° CCW rotation over the course of 90 trials (A) before washout and testing. The BNA group experienced a 30° clockwise rotation (B), followed by washout testing. The BwaitNA group experienced the same perturbation schedule as BNA. However, subjects were asked to wait 1-2 minutes after training of (B). This delay was expected to reduce the adapted motor output (Smith et al., 2006; Criscimagna-Hemminger and Shadmehr, 2008), resulting in reduced after-effects, i.e. smaller errors, when subjects experienced the washout condition. The GNA group experienced a 30° CCW rotation that was gradually imposed over 120 trials. The BGNA group experienced a 30° abruptly imposed CW rotation that was then gradually removed.

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

$$k \exp(-t/\tau) + c . \tag{2.6}$$

In Eq. 2.6 the exponential time constant  $\tau$  has units of trials. Therefore, savings compared to the initial learning of A by the ANA group is represented as a decrease in the value of  $\tau$ .

Indeed we found that the subjects in the ANA and BNA groups learned significantly faster than control in the test of perturbation (A), whereas the GNA,

BwaitNA, and BGNA groups learned at a rate that was no different than control.

### **2.2.6 Data collection and statistical analysis**

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

The standard statistical test used in adaptation studies is repeated-measure ANOVA (RM-ANOVA). In RM-ANOVA, the assumption is that the between subject variance of the measured variable is constant across measurements. However, in motor adaptation studies, the across-subject variance of the measured variable often changes as the experiment progresses, violating a primary assumption of a repeated measures ANOVA. To address this problem, we used the general linear model feature of SPSS (GLM-ANOVA) to test for main effects of trial, group, and group by trial interaction. Our analysis assumed a heterogeneous autoregressive correlation structure of the variance matrix, allowing for between-subject variance to change across repeated measurements (Herzfeld et al., 2014a).

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

## 2.3 Results

Let us begin with a standard model of motor learning (Mackintosh, 1975; Thoroughman and Shadmehr, 2000; Rescorla and Wagner, 1972; Pearce and Hall, 1980) in which on trial  $n$ , a perturbation  $x^{(n)}$  is imposed on action  $u^{(n)}$  so that the sensory consequences are  $y^{(n)}$ . Based on its belief about the environment, the learner predicts the sensory consequences  $\hat{y}^{(n)}$ , and then updates its belief from the prediction error  $e^{(n)} = y^{(n)} - \hat{y}^{(n)}$ . Such learning typically depends on a decay factor  $\alpha$ , and error-sensitivity  $\eta^{(n)}$ , as in Eq. 2.3.

### 2.3.1 History of errors alters error-sensitivity

Consider an environment in which the perturbations persist from trial to trial, and another environment in which the perturbations switch (Fig. 2.1A). In a slowly switching environment, the brain should learn from error because the perturbations are likely to persist (learning from error in one trial will improve performance on the subsequent trial). However, in a rapidly switching environment the brain should suppress learning from error because any learning will be detrimental to performance on subsequent trials.

Three groups of subjects ( $n = 9$  per group) made reaching movements while experiencing force perturbations from either a slow, medium, or rapidly switching environment (Fig. 2.1A). The mean of the perturbations was zero for all blocks (consisting of 30 trials). We measured error in a given trial and then computed the amount that was learned from that error (probe trials, purple bars, Fig. 2.1A). To quantify learning from error on trial  $n$ , we measured the change in force from the trial before to the trial after the perturbation,  $u^{(n+1)} - u^{(n-1)}$  (Fig. 2.1C). In block 1, learning from error was similar in the three groups ( $p > 0.99$ ), and in all probe trials the perturbation produced similar errors (Fig. 2.1B, RM-ANOVA, effect of group

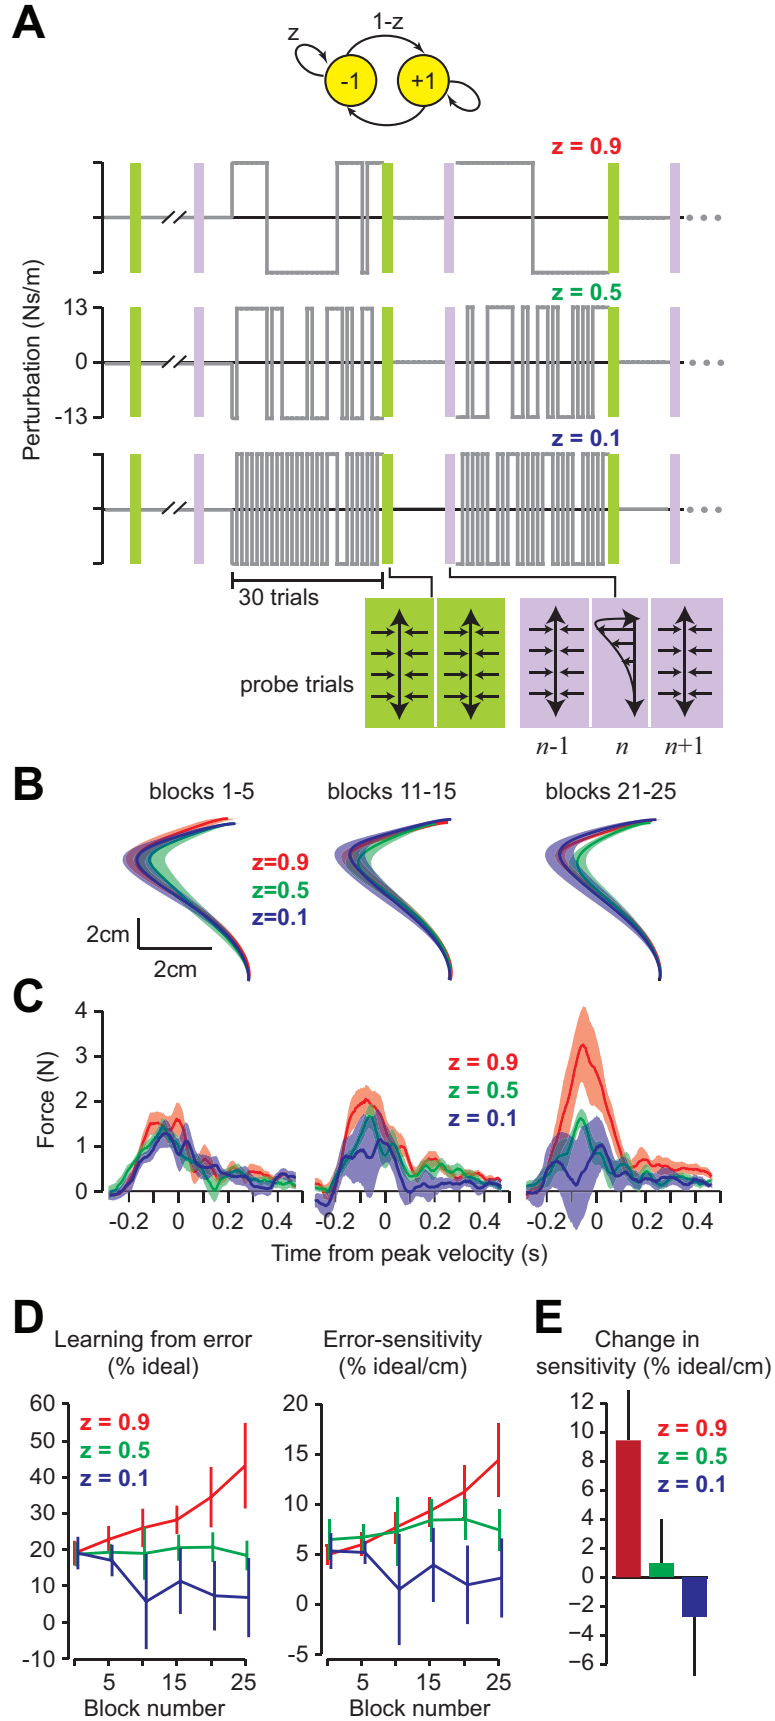




Figure 2.1: History of error alters error-sensitivity. **A.** Reaching paradigm with force field perturbations. The yellow circles note a perturbations state, and indicates probability of remaining in that state. The slow, medium and rapidly switching environments are shown. One group of subjects was trained in each environment. We measured error-sensitivity via probe trials in which subjects experienced a constant perturbation, sandwiched between two error-clamp trials. **B.** Movement trajectories in the perturbation trial of the probe trials. Trajectories were averaged over 5 successive presentations of the probe. The errors in probe trials did not differ between groups. **C.** Learning from error in the probe trials, measured as the change in force from the trial prior to the trial after the perturbation. **D.** Learning from error in the probe trials, plotted as a percentage of the ideal force (left). Error-sensitivity was measured as the trial-to-trial change in the percentage of ideal force divided by error (right). **E.** Change in error-sensitivity between the baseline block and the last 5 error-clamp triplets. Data are mean  $\pm$  SEM.

$p > 0.8$ , interaction,  $p > 0.7$ ). However, subjects that experienced the slowly switching environment increased their learning from error (Fig. 2.1C), whereas those who experienced the rapidly switching environment suppressed this learning.

We measured the force produced on a given trial and computed a coefficient representing percent ideal (Fig. 2.1D). RM-ANOVA indicated a significant block by group interaction ( $p < 0.05$ ), suggesting that the history of perturbations altered the amount of learning from error. Post-hoc tests showed that in the slowly switching environment subjects learned more from error than in the rapidly switching environment ( $p < 0.03$ ). This change in error-sensitivity developed gradually with training (Fig. 2.1D). The slowly switching environment induced an increase in error-sensitivity (Fig. 2.1E, changes in sensitivity from the first half to second half of the experiment, ANOVA,  $p < 0.05$ ).

### 2.3.2 Error-sensitivity is local to the experienced error

Is control of error-sensitivity local to the experienced errors? In Exp. 2, subjects performed rapid out-and-back movements for which no visual feedback was available during the outward part of the reach, with the aim of hitting a target at the

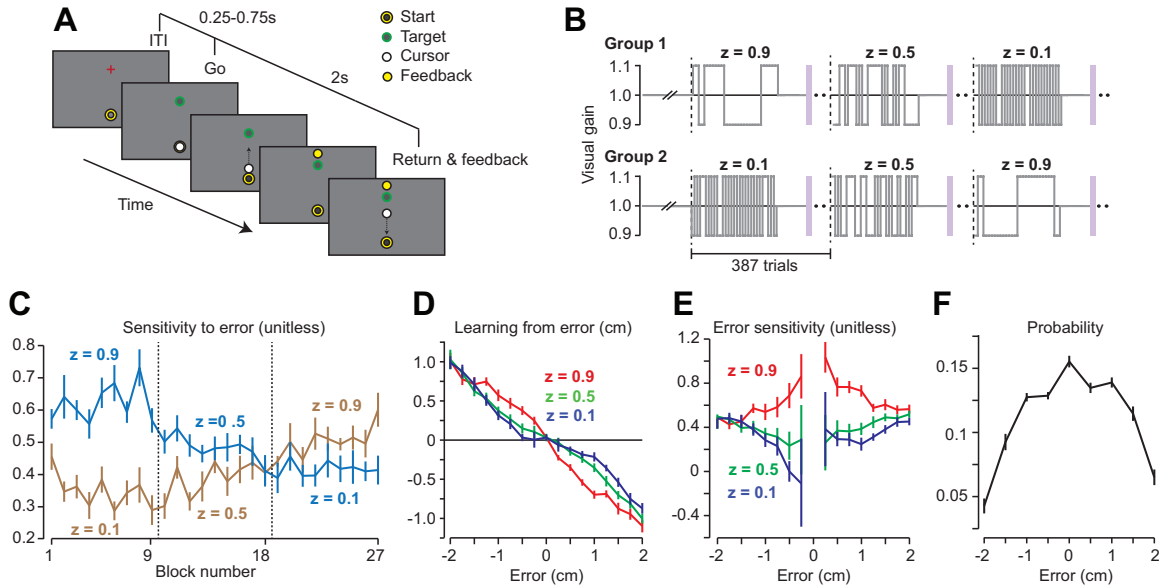


Figure 2.2: Error-sensitivity is local to the experienced errors. **A.** Paradigm with visuomotor gain perturbations. **B.** Perturbation schedule. Dashed lines indicate changes in the statistics of the environment. **C.** Error-sensitivity averaged over all error-sizes measured over each environment block. **D.** Learning from error measured at various error sizes. **E.** Error-sensitivity as a function of error magnitude. **F.** Probability of error.

turn-around point of their movement. An occasional perturbation altered the feedback regarding hand position at the turn-around point (Fig. 2.2A). We measured the relationship between error and learning from error (change in reach extent).

Group 1 ( $n = 10$ ) experienced a perturbation schedule that transitioned from slow, medium, to rapid switching (Fig. 2.2B) whereas Group 2 ( $n = 10$ ) experienced the reverse. In Group 1, error-sensitivity decreased whereas in Group 2 error-sensitivity increased (Fig. 2.2C). We measured the mean error-sensitivity in each environment, resulting in 3 measurements for each subject across the experiment. RM-ANOVA showed a significant main effect of group ( $p < 0.005$ ) and block ( $p < 0.001$ ) and group by block interaction ( $p < 0.001$ ). As the statistics of the perturbation changed, so did the error-sensitivity.

We measured learning from error as a function of error in each environment (Fig. 2.2D). A given error produced greater learning when that error was

experienced in a slow switching environment (Fig. 2.2D, red line) (RM-ANOVA main effects of error size  $p < 10^{-4}$ , and environment  $p < 0.001$ , posthoc between slow versus medium or fast,  $p < 0.001$ ). We quantified error-sensitivity at each error size (Fig. 2.2E) and found that error-sensitivity had not changed globally, but predominantly for smaller error sizes. RM-ANOVA of the absolute sensitivities between 0.25 and 2cm showed a significant main effect of environment ( $p < 10^{-4}$ ), as well as a significant environment by error size interaction ( $p < 0.05$ ). We found a significant difference in error-sensitivity across environments for an error-size of 0.25cm ( $p < 0.05$ ), but no significant difference for an error-size of 2cm ( $p > 0.1$ ). Interestingly, the small error sizes for which the subjects had shown the largest change in error-sensitivity were also the most frequent errors (Fig. 2.2F). This hinted that control of error-sensitivity was a function of error.

### 2.3.3 Computational model of error-sensitivity

Current models of sensorimotor learning assume that error-sensitivity,  $\eta^{(n)}$ , is independent of error,  $e^{(n)}$ . This is true for state-space models of learning (Thoroughman and Shadmehr, 2000; Scheidt et al., 2001; Donchin et al., 2003; Smith et al., 2006; Cheng and Sabes, 2006), as well as Kalman filter models of learning (Wei and Körding, 2009; Körding et al., 2007; van Beers, 2009; 2012). However, suppose that sensory prediction errors are encoded in the nervous system with a set of basis elements where each basis element,  $g_i$  has a preferred error  $\check{e}_i$ . Further suppose that error-sensitivity is determined by a population coding:

$$\begin{aligned} \eta(e^{(n)}) &= \sum_{i=1}^N w_i g_i(e^{(n)}) \\ g_i(e^{(n)}) &= \exp \frac{-(e^{(n)} - \check{e}_i)^2}{2\sigma^2} \end{aligned} \quad (2.7)$$

On trial  $n - 1$  the motor command  $u^{(n-1)}$  produces an error  $e^{(n-1)}$ , as

illustrated in the top part of Fig. 2.3A. The nervous system learns from this error and produces motor command  $u^{(n)}$  on the subsequent trial, resulting in  $e^{(n)}$ . In a slowly switching environment (top part of Fig. 2.3A),  $e^{(n)}$  has the same sign as  $e^{(n-1)}$ . In this case, error-sensitivity should increase around  $e^{(n-1)}$  (Fig. 2.3B, red line). On the other hand, in a rapidly switching environment (Fig. 2.3A, bottom),  $e^{(n)}$  has a different sign than  $e^{(n-1)}$ . In this case, error-sensitivity should decrease:

$$\mathbf{w}^{(n+1)} = \mathbf{w}^{(n)} + \beta \text{sign} \left( e^{(n-1)} e^{(n)} \right) \frac{\mathbf{g}(e^{(n-1)})}{\mathbf{g}^T(e^{(n-1)})\mathbf{g}(e^{(n-1)})} \quad (2.8)$$

In Eq. 2.8,  $\mathbf{w} = [w_1, w_2, \dots, w_N]^T$ ,  $\mathbf{g} = [g_1, g_2, \dots, g_N]^T$ , and superscript  $T$  is the transpose operator. This rule is similar to the RPROP algorithm, a heuristic for adjusting the learning rate of machines (Riedmiller and Braun, 1993), but has the unique feature of assuming that error-sensitivity is via population coding of the error space. These equations represent a learner that stores two kinds of memory: a memory of the state of environment ( $\hat{x}$ , Eq. 2.3), and a memory of errors ( $\mathbf{w}$ , Eq. 2.8). We simulated the model (Fig. 2.3C, gray line) and found that in the slow switching environment error-sensitivity increased in the neighborhood of the experienced errors, whereas in the rapidly changing environment error-sensitivity decreased (Fig. 2.3D).

### 2.3.4 Experimental tests of error-sensitivity model predictions

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

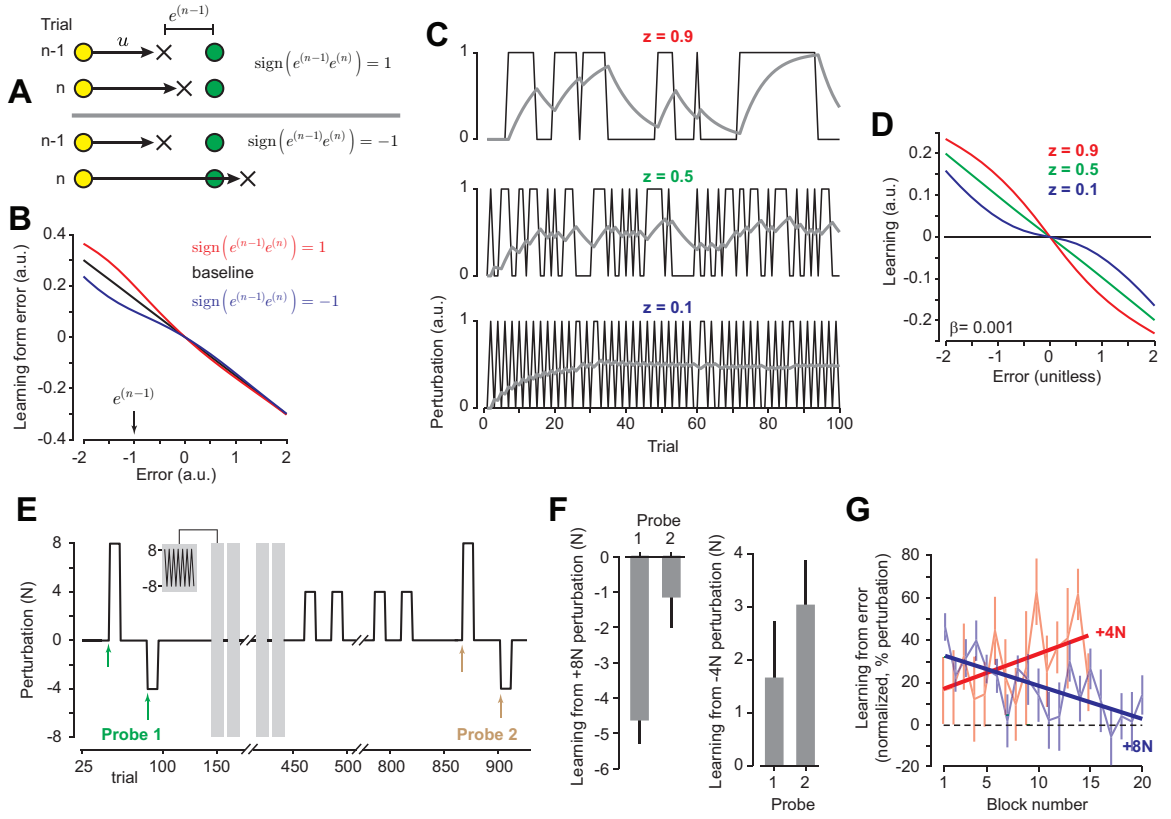


Figure 2.3: Theoretical model and Experiment 3. **A**. On trial  $n - 1$ , the motor command  $u^{(n-1)}$  is generated, resulting in error  $e^{(n-1)} = -1$ . If the error in trial  $n$  is of the same sign as  $e^{(n-1)}$ , then error-sensitivity should increase (top). However, if the error experienced in trial  $n$  has a different sign than  $e^{(n-1)}$  then error-sensitivity should decrease (bottom). **B**. Learning from error following experience of two consecutive errors from A. Error sensitivity around  $e^{(n-1)}$  increases if  $\text{sign}(e^{(n-1)}e^{(n)}) = 1$  and decreases otherwise. **C**. Model performance for slow, medium, and rapidly switching environments (gray line represents  $\hat{x}^{(n)}$ ). However, learning from error, **D**, is increased in the slow switching environment and decreased in the rapidly switching environment. **E**. Experiment 3 perturbation protocol. **F**. Single-trial learning from a  $+8N$  perturbation and a  $-4N$  perturbation in Probes 1 and 2. Learning is increased for the  $-4N$  perturbation, while simultaneously decreased for a  $+8N$  perturbation. **G**. Learning from error normalized by the perturbation magnitude (4 or  $8N$ ) in the first trial of each repetition of the rapid (blue) and slowly switching (red) environments. Learning increased in the slowly switching ( $4N$ ) environment but decreased when the perturbation was rapidly switching ( $8N$ ). Error bars are SEM.

subjects would increase their sensitivity to  $-4N$  errors, while simultaneously decreasing their sensitivity to  $+8N$  errors.

In the baseline block we probed sensitivity to  $+8N$  and  $-4N$  perturbations (Probe 1, Fig. 2.3E). The resulting learning from error is plotted in Fig. 2.3F (Probe 1). At baseline, subjects responded to the  $+8N$  and  $-4N$  perturbations by learning a fraction of each error (Fig. 2.3F). We next produced 20 repetitions of a rapidly switching environment in which the perturbations were  $\hat{\Delta}8N$  (Fig. 2.3E, inset). After a period of washout, we then produced 15 repetitions of a slowly switching environment in which the perturbations were  $0N$  or  $+4N$ . The critical aspect of our design was that the subjects were never exposed to a  $-4N$  perturbation. They nevertheless experienced  $-4N$  errors (because removal of a learned  $+4N$  perturbation results in  $-4N$  error).

The  $8N$  environment induced a decrease in sensitivity to a  $+8N$  error, and subsequent exposure to the  $+4N$  environment resulted in an increase in sensitivity to a  $-4N$  error (Fig. 2.3G, RM-ANOVA showed a significant main effect of perturbation ( $p < 0.03$ ) as well as a perturbation by block interaction ( $p < 0.01$ ). The critical question, however, was whether both of these changes in sensitivity were simultaneously present. Following the slowly switching block of perturbations, we again probed sensitivity to  $+8N$  and  $-4N$  errors (Probe 2, Fig. 2.3E). Compared to the baseline block (Probe 1), learning from a  $+8N$  error had decreased ( $p < 0.005$ ), while simultaneously, learning from a  $-4N$  error had increased ( $p < 0.05$ ) (Fig. 2.3F).

How well could the model explain behavior in this experiment? In Fig. 2.4A we have plotted learning from error as a function of error size for all subjects across the entire data set. The change in error-sensitivity predicted by the model was highly correlated with the change observed in our subjects ( $R^2 = 0.65$ ;  $p < 10^{-8}$ , Fig. 2.4), suggesting that history of error induced changes in error-sensitivity in the region of the experienced errors. Remarkably, we found that the learning from error




Figure 2.4: Comparison of model and experimental results for Exp. 3. **A.** Experimental results. Learning from error as a function of error size, measured across the experiment. Errors were binned across error sizes with a bin width of  $1N$ . Black line represents the best-fit line, corresponding to a constant error-sensitivity across error magnitudes, our estimate of an unbiased learner. Error bars are mean  $\pm$  SEM. **B.** Difference between the measured learning from error and the unbiased learning curve. Predicted curves show the change in error sensitivity as predicted by the model, binned across error-magnitudes. Error bars represent mean  $\pm$  SEM across subjects.

was not monotonic. Rather, subjects learned significantly more from a  $\pm 4N$  error than  $\pm 8N$  error (paired t-test,  $t(15) = 7.76$ ,  $p < 0.001$ ). Fig. 2.4B plots the difference between learning from error and the regression line (our proxy for an unbiased learner). This reflects the change in learning that has been caused by the changes in error-sensitivity. We ran our model on the same sequence of errors that each subject experienced in Exp. 3 and have plotted the predicted change in Fig. 2.4B (scaled by a multiplicative coefficient to convert to units of Newtons). The correlation between the predicted and observed values was  $R^2 = 0.65$ ,  $p < 10^{-8}$ . Therefore, as the model had predicted, we were able to use the history of errors to simultaneously increase learning at one error size, and decrease learning at another.

### 2.3.5 Error-sensitivity model explains savings and meta-learning

This new model of learning provided new insights on a wide range of puzzling experiments, including the phenomena of savings and meta-learning (Fig. 2.7 and Fig. 2.8). It predicted that when one is better at a task than before, it is not because the brain recalled the motor commands, but because it recognized the errors - the errors for which error-sensitivity had been altered. In addition, the model predicted that savings and meta-learning could be blocked by controlling the errors that are experienced during learning.

In Exp. 4 volunteers participated in a visuomotor rotation experiment (Fig. 2.5A and Fig. 2.6,  $n = 10$  per group). We began with a control experiment to establish the basic idea that savings is present despite washout. In the ANA group (Fig. 2.5A), perturbation (A) was imposed ( $+30^\circ$  perturbation), and then following an extended period of washout (N), perturbation (A) was again presented ( $+30^\circ$ ). We expected to observe a faster rate of learning in the second exposure to (A), since this is a protocol that has historically produced savings (Zarahn et al., 2008). According to our model, savings occurs because during the initial exposure to (A) the stable sequence of perturbations increase error-sensitivity, and these errors are re-visited in the subsequent test of (A).

Indeed, we found that subjects in the ANA paradigm experienced a significant amount of savings Fig. 2.5B. If savings occurred because the errors that were experienced during initial learning of (A) were re-visited in the subsequent test of (A), then we should be able to prevent savings by presenting (A) so that the errors that are experienced during initial exposure are unlike the large errors that are experienced during re-exposure to (A). We tested this prediction with the GNA group (Fig. 2.6A). In the GNA group, the perturbation was incremented gradually to (A), as opposed to a sudden presentation. As a result, subjects learned to make



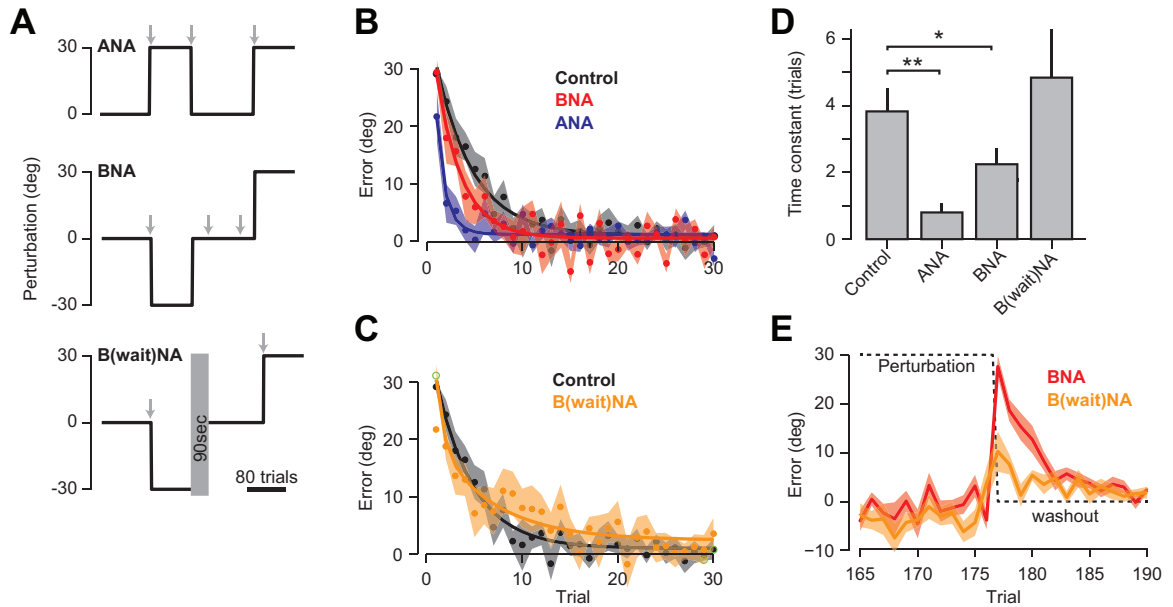


Figure 2.5: Saving occurs only when previously experienced errors are re-visited. **A.** A visuomotor perturbation experiment. Gray arrows indicate 1-2 minute set-breaks. **B.** Performance in the final +30° perturbation. ANA and BNA groups show savings, i.e., faster learning of the perturbation compared to control (naïve). Exponential fits are shown for the group data. **C.** The BwaitNA group does not exhibit savings. **D.** Exponential time constants are compared to controls (\* $p < 0.05$ , \*\* $p < 0.01$ ). A lower time constant indicates faster learning. **E.** Comparison of the errors (i.e., after-effects) experienced by the BNA and BwaitNA groups. The BwaitNA group experienced smaller errors due to the presence of the set-break.

movements typical of learning to (A), but did not experience the same errors.

Furthermore, we should be able to produce savings in a very different way: expose subjects to perturbation (B) and then present sudden washout (Fig. 2.5A, BNA). During washout they are exposed to a sequence of stable errors, which increase error-sensitivity for those errors. Importantly, the washout-induced after-effects are errors that are also experienced during subsequent test of (A). In such a case, the model predicts that the savings occurs because the learner is exposed to errors during sudden washout of (B), i.e., the after-effects. The interesting idea is that the after-effects themselves present a sequence of stable errors, which increase error-sensitivity. Because these after-effect induced errors are re-experienced during subsequent learning of (A), subjects should show meta-learning, despite the fact that they have not previously experienced (A) before.

If the meta-learning in BNA is due to errors that are experienced during washout of (B), we should be able to eliminate it by reducing the washout-induced errors. In BwaitNA, a wait period was inserted between  $-30^\circ$  training and washout. This wait period should reduce the after-effects in the subsequent washout trials (Smith et al., 2006; Criscimagna-Hemminger and Shadmehr, 2008). Therefore, installation of a brief wait period would remove exposure to errors, the same errors that are the part of learning of (A). In this case, the model predicted that we should see no savings (Fig. 2.5E).

We followed this idea with a second group in which we eliminated the after-effects by introducing errors gradually in BGNA group (Fig. 2.5). In BGNA, after exposure to (B) we gradually removed the perturbation so that there would be little or no errors that are similar to those that the subjects would experience during exposure to (A).

In summary, the model predicted savings in ANA but not GNA, meta-learning in BNA but not BwaitNA and BGNA. Our experimental results

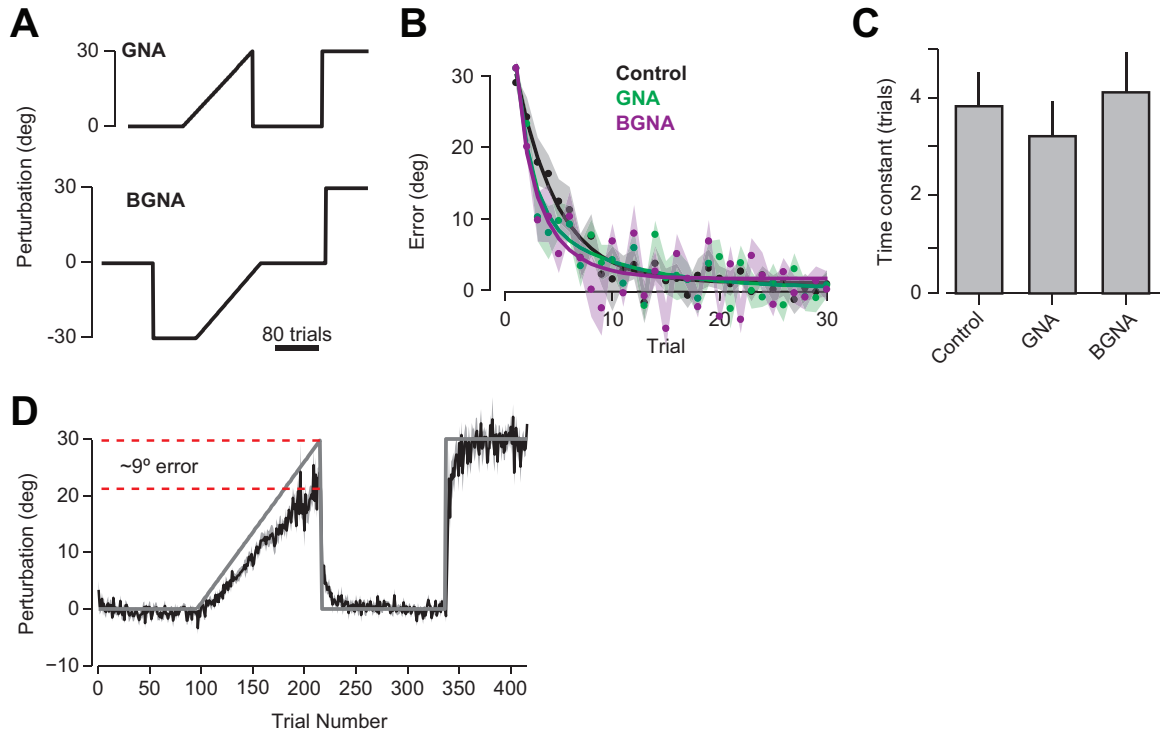


Figure 2.6: Additional groups in Experiment 4: saving and meta-learning occur only when previously experienced errors are re-visited. **A.** Perturbation protocols for groups GNA and BGNA. **B.** Performance during exposure to  $+30^\circ$  perturbation. Exponential fits are shown for the group data. Performances of the GNA and BGNA are not different than control, demonstrating that these perturbation protocols blocked savings. **C.** Exponential time constants are not different than control. A lower time constant indicates faster learning. **D.** Time course of adaptation in the GNA group. Note that the magnitude of the error at the end of gradual learning is significantly smaller than  $30^\circ$ . Data are mean  $\pm$  SEM across subjects.

confirmed these predictions (Fig. 2.5C,D and Fig. 2.6).

## 2.4 Discussion

We have demonstrated that humans store a memory of their previously experienced errors. The history of the experienced errors informs the error-sensitivity used on the current trial. This error-sensitivity is a function of the error magnitude and, therefore, local to the experienced errors. Using these results, we developed a computational model demonstrating a principled technique that can be used to

modulate error sensitivity (Eq. 2.8). In Exp. 3 and 4, we systematically tested the predictions of this model. Here, we demonstrate this this simple computational model can explain a large number of previously collected experimental data.

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

### **2.4.1 Why does learning from error saturate with large perturbations?**

Fine and Thoroughman (2006) examined reaching movements that were perturbed by a force-pulse. The perturbations were drawn from a discrete uniform distribution:  $[\pm 6, \pm 12, \pm 18]N$  and presented in 80% of the trials. They examined learning (change in motor commands from the trial prior to the trial after the perturbation) and noted that this measure did not grow linearly with perturbation size, but saturated (black bars, Fig. 2.7A). We ran our model on a perturbation schedule with the same distribution (50 iterations, 360 trials each, 80% perturbations, discrete perturbations drawn uniformly from  $[\pm 6, \pm 12, \pm 18]N$ ). In our model, error was encoded by 50 Gaussian bases distributed throughout an error space between  $\pm 30N$ , with a standard deviation of  $7N$ . The initial weights of this network were set so that the error-sensitivity was roughly at 20%, a value typical for force field task (Donchin et al., 2003). We allowed the weights to change with  $\beta = 0.005$ . After simulating the motor output for each trial, we estimated the learning following a particular perturbation by measuring the mean trial-to-trial change in motor output for each of the discrete perturbations. The results are shown with the red line in Fig. 2.7A. The correspondence between model and

experimental data is  $R^2 = 0.99 \pm 0.005$  (mean  $\pm$  SD of 50 simulations).

Wei and Körding (2009) measured learning from error in a visual displacement task (Fig. 2.7B, black bars). Subjects were asked to reach in the horizontal direction, and the location of the visual feedback was displaced in the vertical direction. Participants experienced 900 trials in which the perturbation on each trial was drawn from a discrete uniform distribution:  $[0, \pm 1, \pm 2, \pm 4, \pm 8]$ cm, and learning was measured by changes in motor commands in the vertical direction. We simulated the results for 50 randomly generated perturbation schedules with the same characteristics. We distributed 50 Gaussian bases throughout an error space between  $\pm 10$ cm, with a standard deviation of 1cm. The initial sensitivity was set at 5%. We allowed the weights to change with  $\beta = 0.005$  (i.e., unchanged from the above simulations). The model results are shown by the red line in Fig. 2.7B. The correspondence between model and experimental data is  $R^2 = 0.70 \pm 0.17$  (mean  $\pm$  SD of 50 simulations).

Why did the learning saturate for large perturbations? The model explained that this was because the perturbations that were used in these experiments were drawn from a bounded uniform distribution. With such a distribution, error-sensitivity declines (and as a consequence, learning from error saturates) for the large errors produced by the perturbations near the bounds. This is because after experiencing an error from a perturbation near one of the bounds, it is much more likely that the next perturbation will produce a change in the sign of the error than not. To illustrate this, in the Fig. 2.8A (top subplot) we have plotted the quantity  $E [\text{sign}(e^{(n-1)}e^{(n)})]$  as a function of the perturbation size on trial  $n$ . We find that the expected value is positive for small perturbations and negative for large perturbations. As a consequence of this uniform bounded distribution of perturbations, error-sensitivity decreases for large perturbations, as illustrated by the learning from error curve in Fig. 2.8A (right subplot). It is also important to

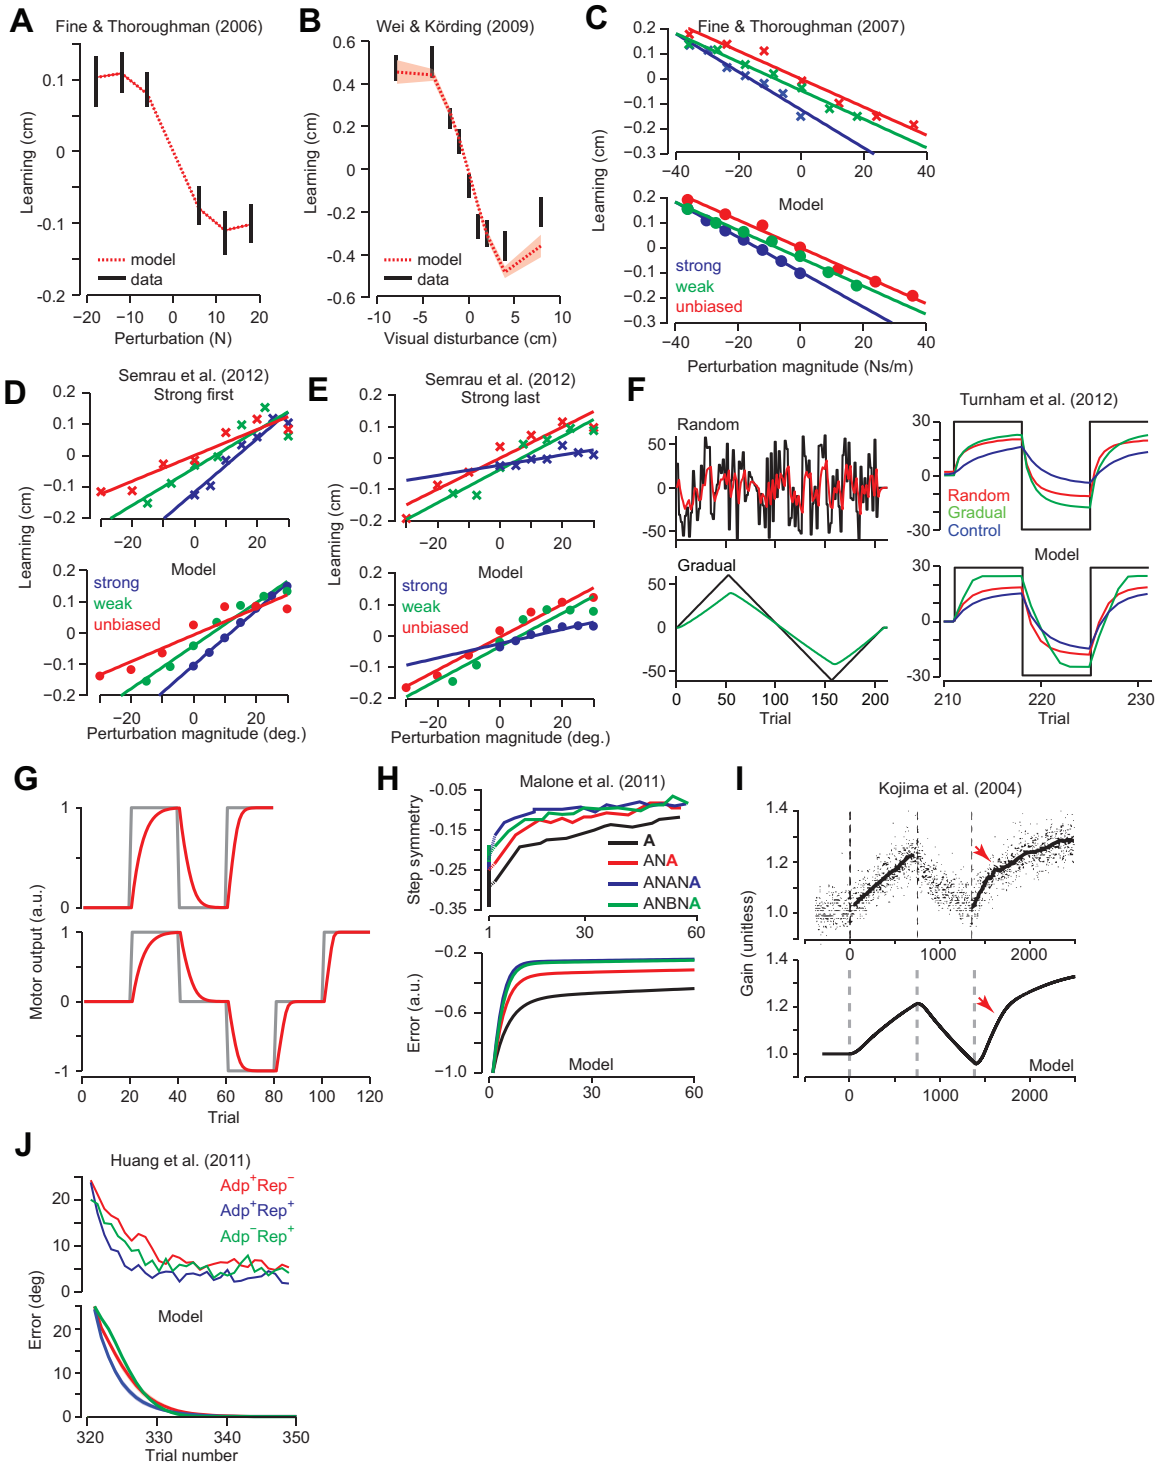


Figure 2.7: Model accounts for a large body of experimental results. **A.** Mean adaptation to a series of discrete force-pulse perturbations. Data reported by Fine and Thoroughman (2006) (black bars, mean  $\pm$  SEM) compared to model (red, mean  $\pm$  SEM). Model SEM is obscured by the line. **B.** Mean adaptation during a visual displacement task reported by Wei and Körding (2009) (black bars) compared to model (red). **C.** Adaptation to discrete force-field perturbations during a reaching task. Volunteers experienced one of three environments in which perturbations were drawn from a distribution that was either unbiased, weakly biased (green x's) or strongly biased (blue x's) (top) (Fine and Thoroughman, 2007). Model results shown in lower panel. **D-E.** Experimental results from a visuomotor rotation task (Semrau et al., 2011). Participants experienced environments with statistics similar to sub-figure C over three separate days (top). Model results for the three environments in the order strong, weak, unbiased are shown in D (bottom); environments in the order unbiased, weak, strong are shown in E. **F.** Example of meta-learning (also termed structural learning). Results from a visuomotor rotation experiment (Turnham et al., 2012). After exposure to a random or gradual environment (left), performance was assessed during learning of a  $[+30, -30, +30]^\circ$  perturbation sequence (right). **G.** Examples of savings. Model results (red) in a perturbation schedule consisting of training (+1), washout (0), re-learning (+1) (grey, top). Model results for a schedule which includes learning of (-1) shows faster learning than control in (-1), as well as faster relearning in (+1) (bottom). **H.** Split-belt gait adaptation results (Malone et al., 2011) in which subjects were exposed to a similar sequence of perturbations as in G. Model results for similar perturbation schedule (bottom). **I.** Saccade adaptation experiment (Kojima et al., 2004) (top) and model (bottom). Individual saccades are shown as black dots; black line shows 150 trial moving average. The red arrow denotes the end of facilitated learning. **J.** Example of data attributed to reinforced repetition. Model results for a paradigm similar to (Huang et al., 2011). Savings is present in the Adp+Rep+ group. In all plots, shaded error regions are model mean  $\pm$  SEM, across randomly generated perturbation sequences.

note that because the mean of the perturbation distribution is zero, the positive expected values around the mean perturbation produce little or no changes in learning from error function (Fig. 2.8A, right subplot). That is, even though sensitivity increases for small errors ( $e \approx 0$ ), learning from error is small for these errors. The dominant effect is the reduced error-sensitivity for errors other than those produced by the mean of the perturbation.

### 2.4.2 Why does error-sensitivity depend on the mean of the perturbation distribution?

Fine and Thoroughman (2007) performed a force-field adaptation study in which the perturbations were similar to their previous study (Fine and Thoroughman, 2006). However, subjects practiced in one of three environments: unbiased, in which the magnitude of the force-field was drawn from a zero mean discrete uniform distribution [+36, +24, +12, -12, -24, -36]Ns/m, weakly biased [+18, +9, -9, -18, -27, -36]Ns/m, or strongly biased [-6, -12, -18, -24, -36]Ns/m (Fig. 2.7C). If error-sensitivity is independent of error history, the three groups of data points should have the same slope. However, the authors found that the slope of the learning curve versus perturbation magnitude was greater (more steep) for the strongly biased distribution, and smaller for the unbiased distribution. We ran our model on the same distributions (50 simulation runs per distribution) and have plotted the results in Fig. 2.7C. We used the same choice of model parameters for this simulation as we did for Fine and Thoroughman (2006), except that the error space was enlarged to accommodate the larger perturbations ( $N = 50$ ,  $P = 50N$ ,  $\sigma = 7N$ ,  $\beta = 0.005$ ). The model results are shown in Fig. 2.7C. The correspondence between the model and experimental data is  $R^2 = 0.97 \pm 0.008$  (mean  $\pm$  SD). Note that the model SEM bars are very small in Fig. 2.7C and are obscured by the circle data markers.




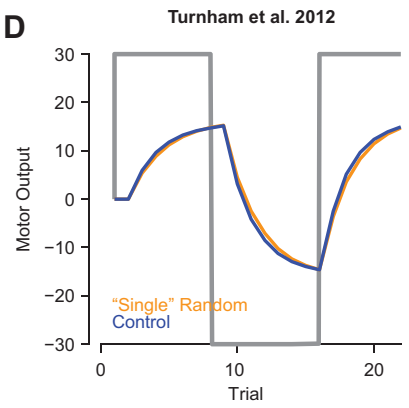
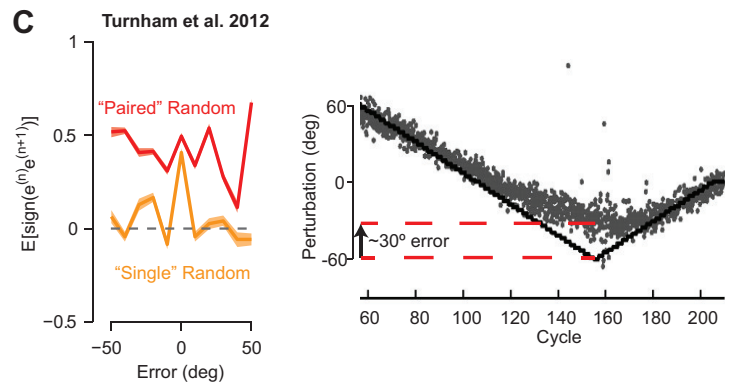


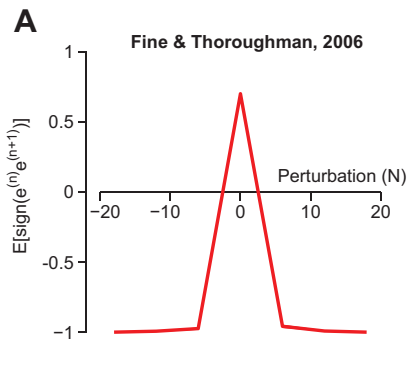


Figure 2.8: The model was simulated on each perturbation distribution in the various experiments and then the quantity  $E[\text{sign}(e^{(n-1)}e^{(n)})]$  was plotted as a function of perturbation magnitude in that distribution. **A.** Simulation results for the perturbation distribution in Fine and Thoroughman (2006). In the right subplot the baseline learning curve is shown by the black line and the final learning curve is shown by the red line. **B.** Simulation results for the perturbation distributions in Fine and Thoroughman (2007). **C.** Simulation results for the random perturbation sequence in Turnham et al. (2012) labeled as “paired” random. Additional results for a schedule in which perturbations are random on each trial rather than each pair is labeled as “single” random. Mean errors from the gradual training in Turnham et al. (2012), showing sustained errors of approximately  $30^\circ$ , similar to those tested at the end of the experiment. **D.** Simulation results for the “single” random perturbation. The “single” random schedule does not show savings whereas the “paired” random schedule from Turnham et al. (2012) shows savings (Fig. 2.8F). **F.** Learning from error as a function of error magnitudes for all groups from the simulation of Turnham et al. (2012). Learning from error increases for the gradual and paired-random perturbations, but not single-random perturbations. In all plots, shaded error regions are mean  $\pm$  SEM, across randomly generated perturbation sequences.

In Fig. 2.8B we have plotted the quantity  $E[\text{sign}(e^{(n-1)}e^{(n)})]$  as a function of perturbation size for this experiment. Each function has a peak near the mean of the corresponding distribution, but there is a clear asymmetry in the function associated with the strongly biased distribution, whereas the function is symmetric for the unbiased distribution. When the distribution of the perturbation is unbiased, the function  $E[\text{sign}(e^{(n-1)}e^{(n)})]$  is symmetric and generally negative for non-zero errors, suppressing learning from errors that arise from both positive and negative perturbations. In contrast, when the distribution is strongly biased, the expected value is asymmetric, only suppressing learning from errors that arise from large negative perturbations. In both cases, sensitivity is reduced to perturbation near the bounds of the uniform distribution, as in (Wei and Körding, 2009; Fine and Thoroughman, 2006). However, because of the distribution of errors in the strongly biased group, the expected change in sensitivity is asymmetric. As a consequence, learning from perturbation (Fig. 2.7C) is shallow for the unbiased distribution since sensitivity is reduced for most perturbation magnitudes, but steep for the strongly

biased distribution.

### 2.4.3 Why does error-sensitivity depend on the sequential order of the perturbation distributions?

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

Semrau et al. (2011) performed a visuomotor rotation experiment. Similar to Fine and Thoroughman (2007), perturbations were presented in 80% of the trials. However, subjects were exposed to a sequence of three different perturbation distributions: unbiased  $[\pm 30, \pm 20, \pm 10, 0]^\circ$ , weakly biased  $[+30, +22.5, +15, +7.5, 0, -7.5, -15]^\circ$ , or strongly biased  $[30, 25, 20, 15, 10, 5, 0]^\circ$ . The order of the environments was counterbalanced across the two groups of participants. The authors found that there was an order effect: learning in a given environment depended on the specific order in which the environments were experienced. Subjects that experienced the strongly bias distribution first showed a steep learning function (Fig. 2.7D, top subplot, blue line), whereas participants that experienced the strongly biased perturbation last (Fig. 2.7E, top subplot, blue line) showed a shallow learning function. We constructed a perturbation schedule which copied this design and simulated our model with the same order of perturbation distributions. As before, the error space was encoded by 50 Gaussian bases throughout a  $40^\circ$  error space with an initial sensitivity of 20% ( $\sigma = 10^\circ$ ,  $\beta = 0.005$ ,  $\alpha = 0.8$ ). The model's results are shown in Fig. 2.7D and Fig. 2.7E. The correspondence between model and experimental data is  $R^2 = 0.86$  (Fig. 2.7D,  $p < 10^{-8}$ ) and  $R^2 = 0.93$  (Fig. 2.7E,  $p < 10^{-10}$ ).

The group that experienced the strongly biased distribution last had already experienced the weakly biased and unbiased distributions. Each of these prior

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

#### **2.4.4 Model explains data attributed to structural learning**

Let us now consider a remarkable observation termed structural learning. A key experiment is that of Turnham et al. (2012), in which subjects perform a visuomotor rotation task. The participants in the random condition trained in an environment in which the perturbation on the odd trials was drawn randomly and without replacement from a discrete sequence uniformly spanning the range  $[-60^\circ$  to  $+60^\circ]$ . Importantly, the even trials had a perturbation that was always the same as in the previous odd trial. That is, perturbations repeated twice in a row. [In the experiment, subjects were tested in 8 directions, termed a ‘cycle’. Here, we considered a simpler version of this experiment in which there is only one direction and a trial represents a cycle.] An example of such a random distribution of perturbations is shown in Fig. 2.7F (left sub-figure, black line). Following this training, the subjects were tested in a series of constant visuomotor perturbations of  $[+30, -30, +30]^\circ$ . The authors found that as compared to a control group, subjects that had this prior training in the random perturbations showed faster learning in the subsequent constant perturbations (Fig. 2.7F, top-right). That is, the experience of the random perturbations appeared to facilitate learning of a constant perturbation, a phenomenon that the authors interpreted as structural learning. A second group completed a gradually imposed perturbation which spanned the range  $[-60^\circ$  to  $+60^\circ]$ . Importantly, subjects in this experiment did not fully compensate for

the errors when the gradually imposed perturbations became large (c.f. Fig. 3 of (Turnham et al., 2012)), resulting in persistent and repeating errors (see Fig. 2.8C, right subplot).

We constructed similar perturbation schedules as the authors of this experiment. We used simulation parameters that were identical to those used in the visuomotor rotation experiment described above for Semrau and colleagues except that we increased the size of the error range to accommodate the larger perturbations. As before, we used 50 Gaussian bases to span the error range  $[-150, 150]^\circ$ , with an initial sensitivity of 20%. We assumed that each of these bases had a standard deviation of  $10^\circ$ . Finally, we set  $\alpha = 0.8$  in Eq. 2.3 and  $\beta = 0.005$ . Our model reproduced the basic observation that following training in the random sequence, learning was faster than control in the series of constant perturbation of  $[+30, -30, +30]^\circ$  (Fig. 2.7F, bottom-right). In addition, the model reproduced the result that learning was faster than control following the gradual perturbation. Finally, the model reproduced the observation that the gradual group performed slightly better than the random group in both the test of +30 and test of -30 conditions.

According to our model, the key fact in these experiments was the repetition of errors in the initial training, and then re-experiencing of these errors in the subsequent testing. In the random group, the repetition came about because every even trial had the same perturbation as the previous odd trial, resulting in sequence of errors that were likely to have the same sign, up-regulating error-sensitivity for the error experienced in the odd trial. In the gradual group, the repetition came about because the subjects could not adapt as fast as the gradually imposed perturbation, resulting in errors that accumulated and repeated near the end of the gradual perturbation (see Fig. 2.8C, right subplot for data from the original experiment, and Fig. 2.7F, bottom subplot for simulation data).

The model explained that in the random group, the faster than control learning that was observed was not due to a memory of perturbations (as the mean of the perturbation sequence was zero), but due to the accumulation of memory of errors. The  $+30^\circ$  constant perturbation produced a  $+30^\circ$  error, for which error-sensitivity had increased due to exposure to the previous 'random' perturbation. This is shown by the quantity  $E[\text{sign}(e^{(n-1)}e^{(n)})]$ , plotted for various error sizes in Fig. 2.8C (left subplot, labeled 'paired random'). Note that the  $E[\text{sign}(e^{(n-1)}e^{(n)})] > 0$  for all errors, resulting increased error-sensitivity for all errors. The model explained that this increase in error-sensitivity would not have occurred if the perturbations on the even trials were unrelated to the perturbations on the previous odd trial (Fig. 2.8C, left subplot, labeled 'single random') where for all errors. As a result, we found that error-sensitivity remained at near baseline, and the model did not show savings, as illustrated in Fig. 2.8D.

The model explained that in the gradual group, the faster learning that was observed (with respect to control) for  $+30^\circ$  perturbation was because at the end of gradual training subjects experienced repeated exposure to  $+30^\circ$  errors (Fig. 2.8C, right sub-plot). The repeated  $+30^\circ$  errors resulted in increased error-sensitivity for this error, which accounted for the fact that when they were tested with a  $+30^\circ$  perturbation, the resulting  $+30^\circ$  errors produced faster learning than in control.

In Fig. 2.8F we have summarized the results of the various simulations. Learning from error increased in the paired-random perturbations, as well as in the gradual condition. As a result, the model suggested that data attributed to structural learning could be explained by memory of errors. The random and gradual conditions had resulted in a memory of errors for which error-sensitivity had increased, and during testing the subjects experienced errors similar to those they had experienced before. It was the repetition of errors, followed by subsequent revisiting of these errors that resulted in faster learning than control.

### 2.4.5 Model explains savings following washout

A number of studies have considered the phenomenon of savings by conducting experiments in which subjects are trained with a constant perturbation, and then the perturbation is removed for a long duration, producing washout. Intriguingly, following this period of washout the re-learning of the perturbation is faster than control. Current state-space models in which the only memory is one of perturbations cannot account for such data (Zarahn et al., 2008; Mawase et al., 2014). Instead, some models have proposed that savings arises from a hypothetical ability of the brain to recognize a context and protect the memory from washout (Lee and Schweighofer, 2009). In addition, other models have proposed that savings arises because during training certain motor commands are reinforced by repetition and reward (Huang et al., 2011). However, neither of these two hypotheses can explain the meta-learning results that we highlighted in Fig. 2.7F. Let us show that the same memory of errors that accounted for meta-learning also readily accounts for these savings experiments. Consider a typical scenario termed A, Null, A (ANA). In this simulation we exposed the model to 20 trials of a +1 (a.u.) perturbation, 20 trials of washout (0 perturbation) followed by 20 trials of relearning of the +1 perturbation (Fig. 2.7G, top). We assumed an error region of  $\pm 3$  consisting of 20 bases whose initial sensitivity was 20% ( $\sigma = 0.25$ ,  $\beta = 0.01$ ,  $\alpha = 1$ ). We found that despite washout, and the fact that memory of perturbation had returned to zero, the model exhibited savings, i.e., faster rate of learning in the second exposure to the +1 perturbation. The reason for this savings, the model explained, was the fact that the previous errors had been experienced in a stable environment, enhancing error-sensitivity for the errors that were again experienced in the second exposure to +1. In a second simulation, we exposed the model to A, Null, B, Null, A (ANBNA), as shown in Fig. 2.7G, lower subplot. The idea in this

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

#### **2.4.6 Model explains savings in a gait-adaptation experiment**

Malone et al. (2011) considered an experiment with a sequence of perturbations similar to the one that we simulated in Fig. 2.7G. In their split-belt gait adaptation task, subjects in the ANA group were asked to adapt to a perturbation in which the belt under the non-dominant leg was moving twice as fast as the dominant leg for 15 minutes. Subjects in this group then returned the next day and were exposed to a tied belt-condition (washout), followed by an additional re-learning period in the split belt condition. Subjects in the ANANA group were exposed to an additional cycle of 15 minutes of tied belt followed by 15 minutes of split-belt before leaving at the end of day 1. In the ANBNA group, subjects were exposed to 15 minutes of tied belt followed by 15 minutes of adaptation to a split-belt condition in which the non-dominant leg was moving half as fast as the dominant leg (opposite of the ‘A’ perturbation) before leaving after day 1. The results of this study are presented in Fig. 2.7H (top). We simulated a similar perturbation schedule in which each of the 15 minute perturbation/null sessions was approximated by 100 trials of a perturbation with magnitude  $\pm 1$  (adaptation) or 0 (washout). We distributed 25 bases throughout an error space of  $\pm 5$ . Initial sensitivity was set at 1% ( $\sigma = 0.5$ ,  $\alpha = 0.9$ ,  $\beta = 0.0005$ ) (Fig. 2.7H, bottom). Consistent with the experimental data,



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

#### **2.4.7 Model explains the limited range of savings**

Kojima et al. (2004) performed a saccade adaptation experiment in monkeys to quantify savings. In this experiment, a standard intra-saccade step paradigm (McLaughlin, 1967) was used to produce errors that resulted in adaptation. After collecting 400-800 saccades during a ‘gain-up’ adaptation period (shown in Fig. 2.7I, top), the direction of the intra-saccade step was reversed until the animal was making saccades with an approximate gain of 1.0. The duration of this period of counter-adaptation was approximately the same as in the adaptation period. Finally, the monkey was exposed to a period of re-adaptation on the gain-up perturbation, but this block contained a larger number of trials than did the initial learning block (Fig. 2.7I, top). The behavior showed clear evidence of savings, but the important observation was that the faster re-learning was present only in the first 100 or so trials, after which the learning curve returned to a rate similar to initial adaptation (red arrow, Fig. 2.7I, top). Why was the faster learning present for only a limited number of trials?

To simulate this experiment, we constructed a perturbation schedule consisting of 300 trials of baseline movements, followed by adaptation to a  $3.5^\circ$

intra-saccade step over 750 trials. The learning of this perturbation was then washed out using the counter perturbation over 650 trials, before being re-exposed 1100 trials of the gain-up perturbation (Fig. 2.7I, bottom). We distributed 50 bases in an error-region spanning  $\pm 6^\circ$  with initial sensitivity of 0% ( $\sigma = 0.25^\circ$ ,  $\beta = 0.00001$ ,  $\alpha = 1.0$ ). Just as in the experimental data, our model also showed the fast initial re-learning, and then a return to slow learning following the first 100 or so trials. The model explained this limited range of savings by noting that the inflection point occurred near the limit of the previously exposed errors. That is, saving was present only up until the errors that were previously experienced - the errors for which error-sensitivity had been up-regulated.

#### **2.4.8 Model explains savings that was attributed to reinforced repetition**

A current hypothesis posits that in some conditions, savings may be the result of reinforcement of motor actions during the adaptation period. In their experiment, Huang et al. (2011) constructed 4 different perturbation schedules for a visuomotor rotation task. In the first group, Adp-Rep-, subjects made movements with veridical feedback to targets drawn from a uniform distribution between  $70^\circ$  and  $110^\circ$ . This group then learned a constant  $25^\circ$  perturbation to a target located at  $95^\circ$ . In another group, Adp+Rep-, subjects moved to random targets between  $70^\circ$  and  $110^\circ$  (identical to the Adp-Rep- group), except that the perturbations were randomly selected for each target from the uniform bounded distribution  $[0, 40]^\circ$ . Therefore, this group adapted to the mean of the perturbation ( $20^\circ$ ), but did not repeat actions to a particular target. In the Adp+Rep+ group, subjects were presented targets from a uniform distribution, but perturbations were chosen so that the correct solution to all rotations would be to move in the  $70^\circ$  direction. The final group, Adp-Rep+, made repeated movements to the  $70^\circ$  target in the absence of a

perturbation. Each group experienced 80 trials of null movements, followed by 160 training trials, 80 trials of washout, and then a final test phase to the  $25^\circ$  perturbation for an additional 80 trials.

We simulated the same experiment for 50 random perturbation schedules for all groups except Adp-Rep-, which the authors found was not significantly different than the repetition control group, Adp-Rep+. We incorporated a model of movement generalization into our model which simulated the effects of generalization of motor commands to nearby targets (identical to (Tanaka et al., 2009)). We distributed 50 basis elements throughout an error-region of  $\pm 50^\circ$  with an initial sensitivity of 10% ( $\sigma = 10^\circ$ ,  $\beta = 0.05$ ,  $\alpha = 1.0$ ). The authors found that only the Adp+Rep+ group showed savings after a washout block. Our model reproduced this result (Fig. 2.7J), and explained that the reason was that the Adp+Rep+ group experienced errors that up-regulated error-sensitivity which were then re-visited during the test of savings. We made two assumptions in this simulation. First, that motor commands generalize to nearby targets according to the model described by Tanaka et al. (2009). Second, that the weights of the error-sensitivity bases are not target specific. That is, change in error-sensitivity is a function of error, and therefore generalizes fully to other targets. While we have not explicitly tested either of these predictions, we note that repetition of a rewarded action alone cannot account for the meta-learning results we have noted above.

#### **2.4.9 Why do gradual perturbations sometimes produce savings, and sometimes not?**

It is puzzling that in certain examples of a gradual perturbation there can be evidence of savings (Fig. 2.7F), whereas in other examples of a gradual perturbation savings is precluded (Fig. 2.6A, GNA group). Why? The model explains that the critical factor is the history of errors during learning. In the experiment shown in

Fig. 2.7F in which savings is present, the gradual perturbation produced large errors, as shown in Figure Fig. 2.8C (right subplot). These residual errors at the end of the gradual perturbation were the same errors that were experienced by the subjects when they were tested in a  $+30^\circ$  perturbation. In contrast, in the GNA group (Fig. 2.6A), the errors at the end of the gradual perturbation were much smaller than the errors that are experienced at the onset of the  $+30^\circ$  perturbation. As a consequence of these differing history of errors, one form of gradual training results in savings (Fig. 2.7F), while another does not (Fig. 2.6C).

## 2.5 Conclusion

We found that during learning, the brain controlled error-sensitivity in a principled way: learning more from error when perturbations were likely to persist, and less when perturbations were likely to change. Error-sensitivity modulation was specific to the experienced errors, suggesting that training produced a memory of errors. This idea accounted for a host of puzzling observations, including saturation of error-sensitivity (Wei and Körding, 2009; Marko et al., 2012; Fine and Thoroughman, 2006), the phenomenon of meta-learning (Turnham et al., 2012), examples of savings (Kojima et al., 2004; Malone et al., 2011; Sarwary et al., 2013), and reinforced repetition(15).

The model predicted that meta-learning vanishes when a small delay or gradual washout alters the history of errors (Fig. ??A), demonstrating that savings depends crucially on the memory of errors that is accumulated during training. This memory of errors likely exists in parallel with the two traditional forms of motor memory, memory of perturbations (Smith et al., 2006) and memory of actions (Huang et al., 2011).

In our model, we chose to describe the learner as a process with a single

time-scale. However, data suggest that learning from error depends on a fast and a slow process with different error-sensitivities (Smith et al., 2006; Kording et al., 2007; Herzfeld et al., 2014a). We speculate that the memory of errors exerts its influence through the error-sensitivity of the fast process, and its manipulation through history of errors may be a useful strategy to speed recovery during rehabilitation (Patton et al., 2013).

### 3 CONTRIBUTIONS OF THE CEREBELLUM TO ACQUISITION OF MOTOR MEMORIES

In Chapter 2, we suggested that, as we learn a novel motor task, the brain stores a previously unknown form or motor memory: a memory of errors. This memory can then be used to modulate error-sensitivity for the experienced errors. What brain region is responsible for maintaining this memory of errors? Previous motor control studies have suggested that the primary motor cortex (M1) and the cerebellum are two brain regions required for motor learning.

In this chapter, we investigate the contributions of the cerebellum and motor cortex to acquisition and retention of human motor memories during a force field reaching task. We show that anodal transcranial direct current stimulation (tDCS) of the cerebellum, a technique that is thought to increase neuronal excitability, increased the ability to learn from error and form an internal model of the field, while cathodal cerebellar stimulation reduced this error-dependent learning. In addition, cathodal cerebellar stimulation disrupted the ability to respond to error within a reaching movement, reducing the gain of the sensory-motor feedback loop. By contrast, anodal M1 stimulation had no significant effects on these variables.

During sham stimulation, early in training the acquired motor memory exhibited rapid decay in error-clamp trials. With further training the rate of decay decreased, suggesting that with training the motor memory was transformed from a labile to a more stable state. Surprisingly, neither cerebellar nor M1 stimulation altered these decay patterns. Participants returned 24 hours later and were re-tested in error-clamp trials without stimulation. The cerebellar group that had

learned the task with cathodal stimulation exhibited significantly impaired retention, and retention was not improved by M1 anodal stimulation.

In summary, non-invasive cerebellar stimulation resulted in polarity-dependent up- or down-regulation of error-dependent motor learning, indicating a modulation of error-sensitivity. In addition, cathodal cerebellar stimulation during acquisition impaired the ability to retain the motor memory overnight. Thus, in the force field task we found a critical role for the cerebellum in both formation of motor memory and its retention. Taken together, our results suggest that the cerebellum may be a prime candidate for storage of previously experienced errors and control of error-sensitivity, as measured via the rate of learning.

### **3.1 Introduction**

When we interact with a novel object, we learn through trial and error to control that object, producing a motor memory that can be recalled the next time the object is encountered. Force field learning has been used as an experimental paradigm to uncover some of the processes that the brain relies on to accomplish this feat. In a typical experiment, the participant holds the handle of a robotic arm and makes a reaching movement, experiencing novel forces that displace the hand, resulting in error. This error engages short- and long-latency feedback pathways, producing a within-movement motor response to the error. In the subsequent reach the brain predicts some of the novel forces from the onset of the movement, resulting in partial compensation for the robot-induced forces. This trial-to-trial change in the motor commands has a specific pattern: the within-movement error feedback response is shifted earlier in time to produce a predictive response (Thoroughman and Shadmehr, 1999). With training, some of the

modifications to the motor commands become a motor memory, as exemplified by the observation that the memory is disengaged when the robot handle is released (Kluzik et al., 2008), and is recalled day (Criscimagna-Hemminger and Shadmehr, 2008; Joiner and Smith, 2008) or months (Shadmehr and Brashers-Krug, 1997) later when the robot handle is grasped.

Formation of this motor memory appears independent of human medial temporal lobe structures (Shadmehr et al., 1998), but dependent on the integrity of the cerebellum (Smith and Shadmehr, 2005; Criscimagna-Hemminger et al., 2010; Donchin et al., 2012) and the motor cortex (Li et al., 2001; Richardson et al., 2006; Arce et al., 2010; de Xivry et al., 2011b;a). In particular, a study in humans demonstrated that reversible disruption of the thalamic projections of the cerebellum to the cortex produced within-subject impairments in the ability to learn the force field task (Chen et al., 2006). Therefore, the current evidence points to the cerebellum as one of the structures that plays a critical role in the acquisition of this motor memory. Here, we used transcranial direct current stimulation (tDCS) to alter function of the cerebellum and quantified the effect of this disruption on the ability to learn the force field task. tDCS of the cerebellum is thought to affect the excitability of Purkinje cell (Galea et al., 2009). Anodal cerebellar stimulation, which is thought to elevate the excitability of Purkinje cells, has been shown to increase rates of adaptation in visuomotor (Galea et al., 2009; Block and Celnik, 2013) and gait (Jayaram et al., 2012) tasks, whereas cathodal cerebellar stimulation, which is thought to reduce Purkinje cell excitability, has been shown to decrease rates of gait adaptation (Jayaram et al., 2012). By contrast, anodal stimulation of the motor cortex (M1) had no effect on the rate of visuomotor adaptation, the size of after-effects, or the rate of de-adaptation upon removal of the perturbation (Galea et al., 2011). However, immediately after adaptation and removal of anodal M1 tDCS, those in the stimulation group showed a reduced rate at which the resulting



memory decayed in the absence of visual feedback (Galea et al., 2011). These findings led Galea et al. (2011) to propose that whereas the cerebellum may be critical for learning from error, the motor cortex plays a role in retention of the resulting memory. By contrast with the findings of Galea et al. (2011), Hunter et al. (2009) applied anodal stimulation to the motor cortex in a force field task and observed a larger reduction in signed kinematic errors during adaptation than in a sham tDCS condition, suggesting that motor cortical stimulation increased learning from error. Therefore, whereas current evidence suggests that stimulation of the human cerebellum can affect learning from error, it is unclear whether stimulation of the motor cortex affects learning from error and/or retention.

Here, we compared the effects of cerebellar and M1 stimulation on the process of acquisition and retention of motor memories in a force field paradigm. Given previous observations in other motor learning paradigms, we expected that M1 stimulation would not affect the rate of learning from error, whereas anodal cerebellar stimulation would increase this rate and cathodal cerebellar stimulation would decrease the rate of learning. In addition, to specifically test the hypothesis that anodal stimulation of M1 enhances retention of motor memories (Galea et al., 2011), we tested the effects of M1 anodal stimulation on both short-term retention (via blocks of error-clamp trials during the training blocks), and long-term retention (at 24 hours following completion of training).

## **3.2 Materials and Methods**

Fifty healthy self-reported right-handed volunteers (21 females; mean age  $\pm$  STD of  $24 \pm 4.7$  years, range 18-38 years) with no known neurological or psychiatric disorders participated in our study. All participants were naïve to the purpose of the experiment and gave written informed consent. The study was approved by the

Johns Hopkins School of Medicine Institutional Review Board. Participants were screened prior to enrollment in the study to ensure that they did not have conditions that would exclude them from a brain stimulation study (cardiac pacemakers, history of seizure, or aneurysm clips). Participants were also screened to ensure that they were not taking any neurological drugs.

### **3.2.1 Experiment 1: cerebellar stimulation**

We recruited  $n = 37$  participants for this experiment. They were divided into three groups: sham ( $n = 12$ ), anodal cerebellar ( $n = 15$ ), and cathodal cerebellar ( $n = 10$ ) stimulation. During analysis of the data we noted that one participant in the cerebellar cathodal group exhibited large errors during field trials and failed to compensate for the forces over the course of the experiment. Although it is possible that this is related to the stimulation (as we will see, cathodal stimulation impaired the ability to learn), to err on the side of caution, the data from this participant were not included in our report.

tDCS (2mA, 25 minutes) was delivered by a Phoresor II device (model PM850, IOMED) through two 5x5 cm saline-soaked sponge electrodes (Ferrucci et al., 2008; Galea et al., 2011; 2009) The current density was approximately 0.08 mA/cm<sup>2</sup>. For the anodal tDCS group, the anode was centered on the right cerebellar cortex, 3 cm lateral to theinion (Galea et al., 2009; Ugawa et al., 1995) with the cathode positioned on the right buccinator muscle (i.e. on the cheek) (Galea et al., 2009; 2011). For the cathodal group the electrode polarity was reversed such that the cathode was placed over the right cerebellar cortex. The procedures for the sham group were identical to the other groups. Anode and cathode positions were counterbalanced between cerebellum and buccinators. The current was increased over a period of 30 sec and then decreased back to zero. With this procedure, participants are unable to reliably distinguish real from sham

stimulation (Gandiga et al., 2006; Kaski et al., 2012).

Both the experimenter and the participant were blind to the type of stimulation, as a third person uninvolved in the experiment controlled the tDCS settings. As illustrated in Fig. 3.1A, stimulation began with block n2 and concluded with block g2, lasting no more than 25 minutes. Brain stimulation was applied on Day 1 only. On Day 2, all participants performed block b1. Additionally, block b2 was performed by a subset of participants:  $n = 12/12$  anodal cerebellar,  $n = 10/12$  sham, and  $n = 8/10$  cathodal.

### **3.2.2 Experiment 2: motor cortex stimulation**

To determine whether the effects observed with anodal stimulation of the cerebellum were unique to this structure, or could also be elicited via anodal stimulation of the motor cortex,  $n = 14$  additional participants were recruited. They performed the identical experiment during anodal tDCS of left M1 (2mA, 25 min., 5x5 cm electrodes, induced current density of 0.08 mA/cm<sup>2</sup>). The anode was positioned on the scalp overlying the 'motor hotspot' of the right first dorsal interosseus (FDI) muscle, that is, the optimal position at which a consistent motor evoked potential, as recorded via EMG, could be elicited using minimal intensity transcranial magnetic stimulation (70mm coil coupled with a Magstim 200). We used FDI (rather than biceps) muscle to localize M1, primarily because it is more easily activated via TMS. The size of the tDCS electrode (25cm<sup>2</sup>) makes it likely that coverage included both muscle representations. The other electrode was positioned on the skin overlying the contralateral supraorbital region.

### **3.2.3 Behavioral procedures**

All volunteers participated in a standard force field task (Shadmehr and Mussa-Ivaldi, 1994). Using the right hand, each participant held the handle of a

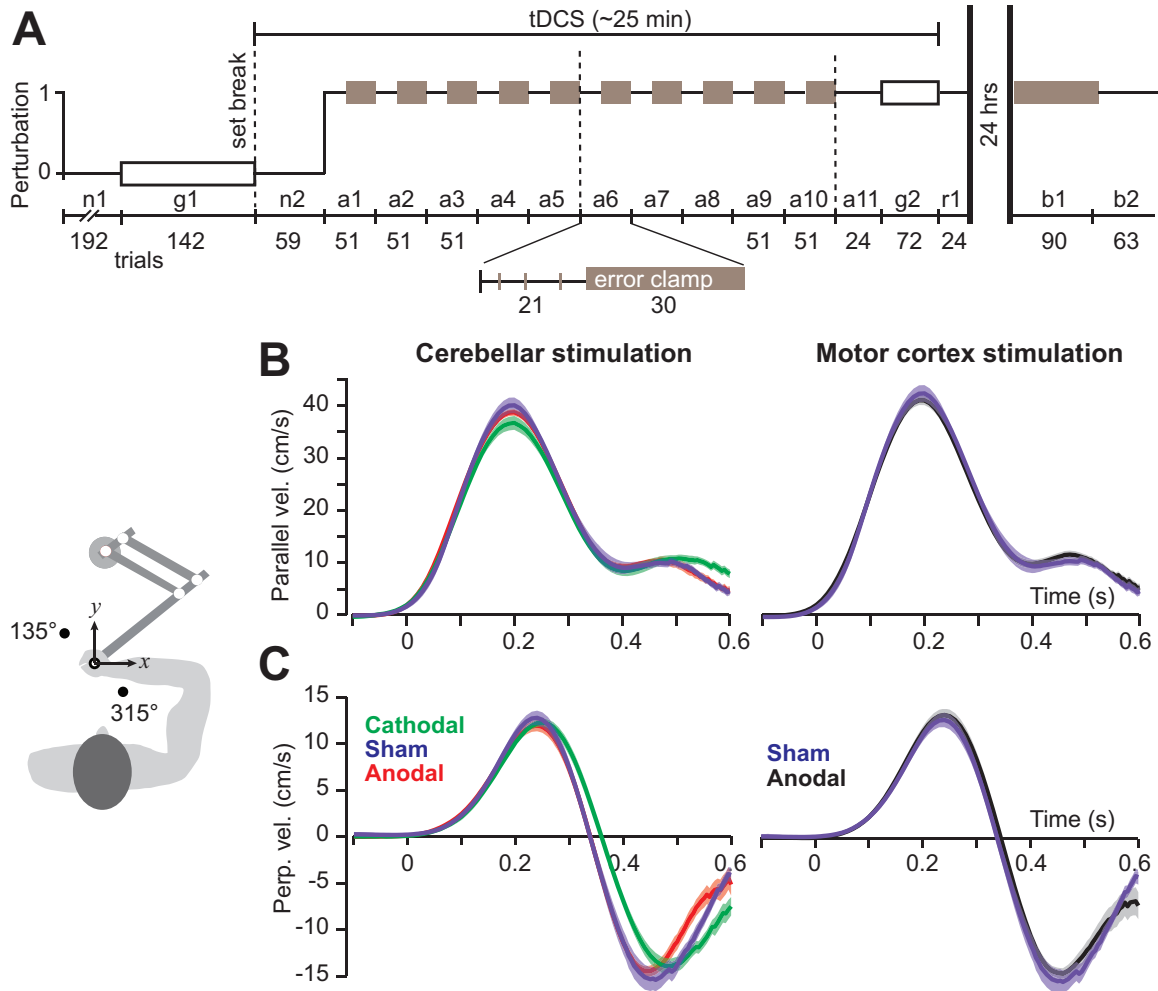


Figure 3.1: Experiment protocol and effects of stimulation on feedback control. **A.** Volunteers were instructed to hold the handle of a manipulandum and reach to one of two targets that appeared at 10cm. After a period of null field training (no perturbation, blocks n1 and n2), a clockwise curl force field was introduced. During training, short blocks of field trials were followed by short blocks of error-clamp trials (blocks a1-a11). Blocks g1 and g2 refer to trials in which generalization of learning was assayed at nearby targets. Block r1 provided relearning after block g1. The right cerebellum or the left motor cortex was stimulated during Day 1. Retention was assessed on Day 2 by means of a block of error-clamp trials (b1) followed by a block of re-exposure to the field (b2). Dashed lines are set breaks (around 1min). **B.** Hand velocity parallel to the direction of target during stimulation of the right cerebellum or the left motor cortex, first 10 trials of block a1. **C.** Hand velocity perpendicular to the direction of target, first 10 trials of block a1. Cathodal stimulation of the cerebellum slowed the error-feedback response, indicated by the later time at which the perpendicular velocity trace crosses zero. Trajectories during anodal stimulation of the motor cortex or cerebellum were indistinguishable from the sham group. Data are mean  $\pm$  SEM.

manipulandum and made center-out movements to a target (1cm diameter, Fig. 3.1). The reach was perturbed by a velocity dependent clockwise curl force field that pushed the hand perpendicular to the direction of motion: where is force on the hand,  $N.s/m$ , and is hand velocity. In the starting posture, the hand was positioned such that the shoulder and elbow were at  $45^\circ$  and  $90^\circ$  respectively (Fig. 3.1A). Participants were unable to see their hand, which was occluded by an opaque horizontal screen. Instead, visual feedback regarding hand position was provided by a cursor (0.5cm diameter) that was continuously projected onto the horizontal screen.

On each trial (except generalization trials, see below), one of the two targets appeared on the screen (pseudo-randomized with equal probability). Targets 1 (T1) and 2 (T2) were positioned at 10 cm at  $135^\circ$  and  $315^\circ$  (Fig. 3.1). The trial was successful if the hand arrived at the target within 400-500ms after movement onset, with success indicated by an ‘explosion’ of the target (an animation). Feedback regarding movements that were too fast or too slow was indicated via changes in target color. After completion of the trial, the robot brought the hand back to the start position. Participants were instructed to maximize the number of successful trials. In some trials, an ‘error-clamp’ was applied (Scheidt et al., 2001). In these trials, the force field was turned off. Normally, removal of the field produces an after-effect. However, in error-clamp trials the hand path was constrained to a straight line to the target via stiff walls (spring coefficient 2000 N/m, damping coefficient 25 N.s/m). The stiff walls allowed us to measure the forces that the participant produced, serving as a proxy for the motor output that the brain generated in order to compensate for the force field expected from the robot.

The experiment was conducted over two consecutive days (Fig. 3.1A). On Day 1, the session began with two blocks of training in the null field without brain stimulation. Block n1 consisted of 192 trials to targets T1 and T2, including 48

interspersed error-clamp trials. Block g1 consisted of 142 trials to targets at  $\pm 45^\circ$ ,  $\pm 90^\circ$ ,  $112.5^\circ$ ,  $\pm 135^\circ$ ,  $157.5^\circ$ ,  $180^\circ$ , and  $225^\circ$ . Brain stimulation was started at the onset of block n2. This was followed by another block of null field training (59 trials, including 15 error-clamp) to targets T1 and T2 (block n2). Participants then experienced alternating field and error-clamp blocks (labeled a1-a11). As illustrated in Fig. 3.1A, each of these blocks consisted of 21 field trials with 3 randomly inserted error-clamp, followed by 30 trials of error-clamp. Block a11 consisted of 24 field trials (including 5 error-clamp).

During blocks a1-a11, participants alternated between short blocks of field and error-clamp trials. This enabled measurement of two distinct properties of learning: 1) in field trials we assayed error-dependent learning by quantifying how the motor output changed from one trial to the next as a function of error, and 2) in error-clamp trials we assayed the stability of the developing memory by quantifying how the motor output decayed within blocks in the absence of error (Smith et al., 2006; Criscimagna-Hemminger et al., 2010).

Training on Day 1 concluded with 72 generalization trials (block g2, including 36 error-clamp) in which we quantified motor output to locations near the trained targets. The generalization targets were at  $\pm 22.5^\circ$ ,  $\pm 45^\circ$ , and  $\pm 90^\circ$  degrees with respect to the training target T1. The reaches to the generalization targets were always in error-clamp. The generalization block consisted of cycles in which there was one movement to T1, followed by error-clamp movements to successive generalization targets chosen randomly so that every cycle included one of each of the target positions. Following the generalization trials, we concluded Day 1 training with 24 trials in a re-learning block to targets T1 and T2 (block r1, including 5 error-clamp). At the end of Day 1 participants ranked their level of attention (1: least attentive, 7: most attentive), fatigue (1: least fatigued, 7: most fatigued), and perceived head discomfort (1: no discomfort, 7: extreme

discomfort/pain) using a visual scale. Retention of the motor memory was assessed on Day 2 by means of an error-clamp block (b1, 90 trials), followed by re-exposure to the field (b2, 63 trials, including 9 error-clamp). Stimulation was not applied on Day 2. All procedures were identical to between the two experiments, with the exception that on Day 2, all M1 participants were tested in block b1, but not b2.

### 3.2.4 Data collection and statistical analysis

A force transducer measured the forces applied by the participant at the robot handle and optical encoders measured position of the robot. The sensors and transducers were sampled at a rate of 200 Hz. Movement onset was defined as the time when the reach exceeded 10% of the maximum velocity in the direction of the target. Data from aborted trials, trials in which participants moved in the wrong direction (exceeding 0.02 m from a line connecting the starting position and the target), and trials in which hand velocity did not exceed 0.08 m/s were excluded (< 4% of all trials). Using these criteria, the following percentage of trials was removed prior to analysis: 3.8% (cerebellar anodal), 4.7% (cerebellar cathodal), 3.7% (M1 anodal) and 3.7% (sham). All other trials were included in the analysis. For each participant, the force profile measured during error-clamp trials in baseline block n2 was subtracted from error-clamp trials during adaptation. To quantify how well the forces that participants produced matched the perturbation forces, we computed a force index: the force  $f(t)$  produced by the participant in an error-clamp trial was compared to the ideal force  $f^*(t) = B\dot{\mathbf{x}}$  (field strength times the hand velocity) by finding the coefficient that minimized the following:

$$\int_{t=0}^T (af^*(t) - f(t))^2 \quad (3.1)$$

In Eq. 3.1,  $T$  is time at end of the reach. We will refer to the  $a$  that minimizes Eq. 3.1 as the force index.

Statistical analyses were performed using R (R-project, Vienna, Austria). Motor adaptation studies often show changes in the across-participant variance of the learned parameter as the experiment progresses. This implies that the covariance between two pairs of samples changes over the course of the experiment, violating the compound symmetry assumption of repeated measures ANOVA. Therefore, we used the generalized linear model (GLM) feature of R (`glms`) to test for fixed effects of block, stimulation type, and block by stimulation interactions Laird and Ware, 1982. We constructed three models with different covariance structures, including compound symmetry, similar to the statistical model used for repeated measures ANOVA, autoregressive, and unstructured correlations. We compared the fit of these models using Akaike’s information criteria (Akaike, 1998) and noted that an autoregressive structure provided the best fit in all tested cases. This autoregressive correlation structure assumed that consecutive measurements had a correlation given by the product of the measured variance and the discounting parameter,  $r$ , where  $r \approx 1$ . Therefore, the correlation between any two within participant measurements decreased as the temporal distance between the measurements increased:

$$S = \begin{bmatrix} s^2 & s^2 r & \dots & s^2 r^n \\ s^2 r & s^2 & \dots & s^2 r^{n-1} \\ \vdots & \vdots & \ddots & \vdots \\ s^2 r^n & s^2 r^{n-1} & \dots & s^2 \end{bmatrix}. \quad (3.2)$$

This approach is in contrast to standard repeated measures ANOVA, which assumes that the correlation between any two measurements is constant. In addition, use of a GLM accounts for unbalanced designs, in which the number of subjects per group



is not equal (an unbalanced design may violate the assumption of orthogonal interaction effects when using a repeated measure ANOVA). All estimation was performed by the linear mixed-effects procedure built into R. In cases where we used a GLM, we represented each participant's response as a single point per block, typically by using the mean value of the outcome variable for each participant within a block. Estimates of the unknown parameters were found using maximum likelihood. We report the adjusted type III error in all cases, which accounts for an unequal number of observations between groups.

When possible, we included the data from Experiment 2 (M1 stimulation) in the statistical tests for Experiment 1 (cerebellar stimulation). Using a single GLM to test for the effect of tDCS across the cerebellum and cortex reduced the total number of statistical tests, thereby reducing spurious multiple comparison effects.

In cases where we found a significant main effect of stimulation, or stimulation by block interaction, we performed post hoc tests on the simple effect of stimulation to determine which groups were significantly different from sham. To guard against false positives that can arise from multiple comparisons, we used Dunnett's t-test for this post hoc comparison. Dunnett's t-test is a multiple comparison corrected approach that is used when a single control group (the sham group) is compared to other groups. All figures show mean  $\pm$  SEM, unless otherwise specified.

### **3.3 Results**

In our experiment, short blocks of field trials alternated with short blocks of error-clamp trials (Fig. 3.1A). The two day experiment enabled us to measure three separate components of learning: 1) in field trials of Day 1 we assayed error-dependent learning by quantifying how the motor output improved from one

trial to the next, 2) in error-clamp trials of Day 1 we assayed the stability of the developing memory by quantifying how the motor output decayed within blocks in the absence of error, and 3) in error-clamp trials of Day 2 we assayed how much of the acquired memory was retained over a 24 hour period. Our principal question was with regard to effects of stimulation of the cerebellum and the motor cortex on these three components.

Our study included four stimulation groups: sham, anodal cerebellar, cathodal cerebellar and anodal M1. Because the same protocol was used for all groups, a single statistical model (GLM) assessed effects across all stimulation groups. However, for clarity of presentation we first report the effects of cerebellar stimulation on a particular set of variables, and then present effects of M1 stimulation on the same set of variables. After completion of the adaptation blocks on Day 1, subjects ranked their level of attention, fatigue, and perceived head discomfort using a visual scale. Self-reported ratings of attention, fatigue, and perceived pain did not differ with stimulation (all  $p > 0.05$ ).

### **3.3.1 Effects of tDCS in the null field**

To test whether brain stimulation affected basic characteristics of movement such as reaction time and peak velocity, we compared performance in a null field condition in which there was no stimulation (block n1, last 50 trials), to a null field condition with stimulation (block n2). We analyzed peak velocity of the reaching movements for each group (sham, anodal cerebellar, cathodal cerebellar, and anodal M1) and found there was no effect of stimulation type ( $F(3, 46) = 0.29, p > 0.8$ ) nor a tDCS by block interaction ( $F(3, 46) = 2.4, p > 0.05$ ). We considered other kinematic measures such as perpendicular displacement or velocity at various times into the movement (100ms and 200ms) and found no significant effect of stimulation type, nor any interaction.

### 3.3.2 Effects of tDCS on reaction time

We quantified reaction times during the null field and force field parts of the experiment. To check whether the stimulation itself produced a change in reaction times, we compared reaction times before stimulation (block n1) to reaction times during stimulation (block n2). We performed a GLM with factors of group, block, and a group by block interaction. We found a main effect of group ( $F(3, 46) = 2.7$ ,  $p < 0.05$ ), but no group by block interaction. Post hoc tests indicated that the sham group in general reached with slightly shorter reaction times (around 20ms, the effect reached significance in comparison of sham versus cerebellar anodal). However this difference in reaction times was not due to the onset of the stimulation, as it was present even before stimulation onset. Therefore, for unknown reasons the sham group reached with slightly shorter reaction times than other groups.

We quantified the reaction time in the early and late phases of training (first 5 and last 5 blocks). In the early phase of training we found that all stimulation groups had longer reaction times than sham (group effect  $F(3, 46) = 5.86$ ,  $p = 0.0007$ ; post hoc testing revealed a significant difference between all stimulation groups and sham,  $p < 0.001$  in each case). However, by the late phase of training this difference had disappeared (no group effect  $F(3, 46) = 1.8$ ,  $p = 0.14$ , no effect of block,  $F(5, 230) = 1.8$ ,  $p = 0.12$ , and no group by block interactions,  $F(15, 230) = 1.2$ ,  $p = 0.30$ ). Finally, we checked to see if there was a difference in reaction times between the various groups that received stimulation. We found no effect of stimulation type ( $F(2, 35) = 0.82$ ,  $p > 0.4$ ), and no stimulation type by block interaction ( $F(18, 315) = 0.93$ ,  $p > 0.5$ ). That is, stimulation modality did not alter reaction time. In summary, there were no significant differences in reaction times between the various tDCS groups. However, the reaction times throughout the experiment were shorter (by about 20ms) for the sham group than other groups.

This was not because of brain stimulation, as the differences existed even in the first null block in which there was no stimulation. The differences in reaction time between the tDCS and sham groups disappeared by late phase of training (last 5 blocks, a5-a10), during which all groups exhibited comparable reaction times.

Effect of cerebellar stimulation on feedback control It has been hypothesized that the motor response to error during a movement can act as a teaching signal, driving corrective changes in the motor commands generated in the subsequent movement (Kawato, 1996). To test whether cerebellar stimulation affected the motor response to error during the reach, we assessed the effect of tDCS on hand velocity perpendicular to the direction of target. We focused our analysis on the first 10 trials of block a1, that is, during the earliest period of exposure to the field, before significant learning had occurred.

Hand velocities parallel and perpendicular to the target were computed separately for every trial and every participant, and then averaged across the first 10 trials (shown in Fig. 3.1B and Fig. 3.1C). Parallel velocity appeared indistinguishable between the groups (Fig. 3.1B). Analysis of the peak parallel velocity confirmed that there was no effect of stimulation in the parallel direction (one-way ANOVA, main effect of tDCS,  $F(3, 46) = 1.9, p > 0.1$ ). Furthermore, the magnitude of the peak perpendicular velocity, which is a proxy for the early motor response to the perturbation, was not affected by stimulation (one-way ANOVA, main effect of tDCS,  $F(3, 46) = 0.45, p > 0.7$ ). However, a closer examination of the perpendicular velocity trace suggested that the feedback response to the perturbation appeared to be delayed in the cathodal cerebellar group, separating from the other groups approximately 350ms into the movement (Fig. 3.1C). To quantify this potential delay in the feedback response, we considered the time at which the perpendicular velocity crossed zero. This quantity represents the time at which participants had compensated for the cumulative effects of the field, hence

allowing us to assess the time within a trial when participants in each tDCS condition compensated for the field. For the cathodal cerebellar group this time was later than for the sham and anodal cerebellar groups (one-way ANOVA, main effect of tDCS,  $F(3, 46) = 3.4$ ,  $p < 0.05$ ; post-hoc Dunnett's t-test, cathodal versus sham,  $p < 0.05$ ). By contrast, anodal M1 stimulation had no discernible effects on the ability to respond to sensory feedback: the time when the perpendicular velocity crossed zero was not significantly different between anodal cerebellar stimulation, anodal M1 stimulation and the sham group (peak perpendicular velocity: no main effect of tDCS; zero crossing: post-hoc test, M1 versus sham,  $p > 0.1$ ).

In summary, we found that cathodal cerebellar stimulation impaired the ability of participants to respond to error feedback during the reach. This delay was not due to a general slowness in visual processing, as reaction times were comparable between various groups that received tDCS.

### **3.3.3 Effect of cerebellar stimulation on learning from error**

To quantify learning from error, we focused on reach kinematics in field trials.

Fig. 3.2A shows average reach trajectories in representative blocks of the experiment for the cerebellar tDCS groups. In the null block (n2) the trajectories appeared indistinguishable. When the perturbation was introduced (block a1), the hand was displaced from its nominal trajectory, and with training the trajectories converged to an 'S' shaped path that over-compensated for the perturbation early in the movement and under-compensated late in the movement. In healthy individuals, over-compensation is a characteristic of learning in curl force fields (Thoroughman and Shadmehr, 2000; Izawa et al., 2008). However, this characteristic of force field learning is reduced or missing in people with cerebellar damage (Criscimagna-Hemminger et al., 2010). Here, we found that anodal cerebellar stimulation enhanced over-compensation, whereas cathodal stimulation

reduced it (Fig. 3.2A). This is further illustrated in the perpendicular velocity traces, as shown in Fig. 3.2B. In the early adaptation block (a1), at the onset of the movement the perpendicular velocity was in the positive direction, reflecting the effect of the perturbation. However, with further training (blocks a6 and a10) the perpendicular velocity progressively shifted in the negative direction, reflecting over-compensation. Over-compensation was evident by block a6 in the anodal condition, but appeared to develop more slowly in the cathodal condition.

To quantify these patterns, we focused on a measure early in the movement, hand velocity perpendicular to the direction of the target at 100ms after reach onset, and a measure relatively late in the movement, maximum perpendicular displacement. Analysis of other trajectory measures (e.g. perpendicular displacement at 50ms or 200ms) confirmed the same pattern of results.

We first considered a measure early in the movement (100ms) (Fig. 3.2C). In block a1, the groups showed comparable performance. There was no significant difference between the mean perpendicular velocity at 100ms in block a1 (one-way ANOVA, main-effect of tDCS,  $F(3, 46) = 0.2, p > 0.8$ ). However, as training progressed, performance of the three groups diverged. In particular, over-compensation emerged fastest in the anodal group and slowest in the cathodal group (when the data values fall below zero, the motor commands exhibited over-compensation). A GLM with factors of block (a1 to a11) and tDCS found a main effect of stimulation type ( $F(3, 46) = 2.8, p < 0.04$ ) and block ( $F(10, 460) = 23.1, p < 10^{-3}$ ). Post hoc comparisons indicated that the cathodal cerebellar group exhibited slower learning, resulting in an increase in the overall perpendicular velocity compared to the sham group (Dunnett's t-test,  $p < 0.01$ ). In contrast, the anodal group showed faster learning compared to sham (Dunnett's t-test,  $p < 0.05$ ). We next considered a measure that focused on the late part of the movement (peak displacement from a straight line). Fig. 3.2D plots maximum

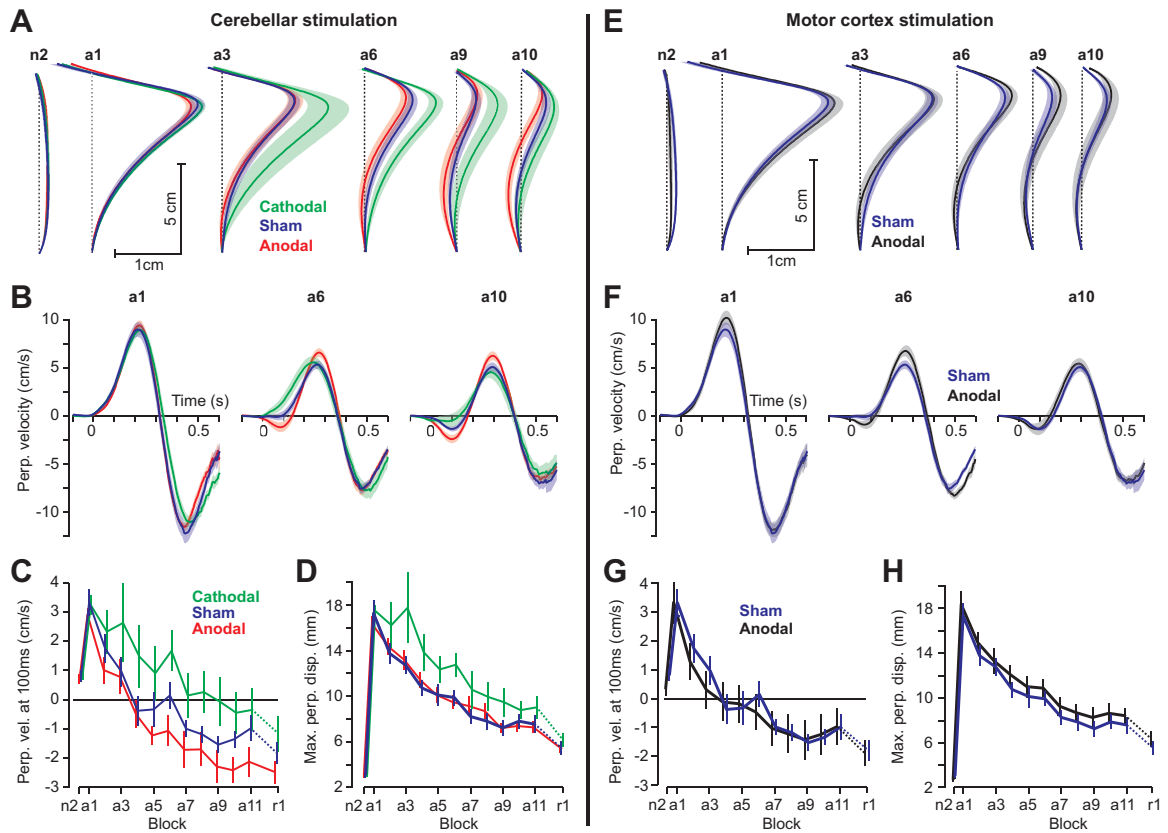


Figure 3.2: Reach kinematics and measures of error-dependent learning during cerebellar (left column) and motor cortex (right column) stimulation. **A.** Hand paths during cerebellar stimulation. The reach starts at the bottom and ends at the top. Figures show across participants mean  $\pm$  SEM of hand position for each tDCS group during blocks of training labeled at top of the figure. With training, both the anodal and sham tDCS groups exhibited ‘over-compensation’ early in the reach; this effect appeared larger in the anodal group and smaller in the cathodal group. **B.** Hand velocity (cm/s) perpendicular to the direction of target. Over-compensation gradually emerges across blocks, as reflected in the negative hand velocities in the early period after movement onset (0 - 100ms). **C.** Perpendicular velocity at 100ms after reach onset during selected blocks of training. Positive values represent clockwise deviation of the hand, and negative values represent over-compensation. **D.** Maximum displacement of the hand perpendicular to the direction of the target. **E-H.** The same as parts A-D, but for anodal stimulation of the motor cortex.

displacement caused by the perturbation, averaged for each block. The maximum displacement curves of the cerebellar anodal and sham groups appeared indistinguishable, whereas the cathodal group exhibited larger maximum displacement, indicating reduced compensation for the force field. GLM analysis identified a significant main effect of tDCS ( $F(3, 46) = 2.6, p = 0.05$ ) and block ( $F(10, 460) = 28.9, p < 10^{-3}$ ), and post-hoc analysis confirmed that the cathodal group exhibited significantly larger maximum displacement than the sham group (Dunnett's t-test,  $p < 0.001$ ) whereas the anodal group showed no significant difference compared to sham (Dunnett's t-test,  $p > 0.9$ ). This reduced compensation in the cathodal group relative to sham may be due to stimulation-induced impairments in learning, or due to impairment of the feedback response reported earlier.

In summary, kinematic measures during training illustrated that anodal cerebellar stimulation increased the learning rate, whereas cathodal stimulation reduced this rate.

### **3.3.4 Robustness of statistical results**

We had  $n = 9$  subjects in the cerebellar cathodal group and a larger number of subjects in the sham and cerebellar anodal groups. To what extent could this imbalance in the study population size have affected our conclusion regarding impairment of learning in the cerebellar cathodal group?

In both GLM and standard one-way ANOVA, the smaller number of subjects in a group increases in the estimate of the between-subject variability. Therefore, the cathodal group is at a statistical disadvantage in terms of the likelihood of finding significant results when it is compared to other groups. Despite this, we found that the altered rate of learning during cathodal cerebellar stimulation was the strongest effect in the dataset, showing the highest levels of significance (even



compared to anodal cerebellar ( $n = 15$ ) versus sham ( $n = 12$ )).

A reasonable way to deal with unequal sample size is to use an autoregressive structure of the GLM in which each group has a different variance value (diagonal elements on the variance-covariance matrix). This ensures that when performing post-hoc contrasts, smaller groups, with measured higher variance values, are at a statistical disadvantage. Despite using this conservative approach, we found consistent effects of cerebellar cathodal stimulation.

To directly test the robustness of our inference we performed a bootstrap analysis in which we randomly sampled  $n = 9$  subjects from the sham group (without replacement). For each of 100 iterations, we performed a GLM and a post-hoc comparison between the resampled sham ( $n = 9$ ) and cerebellar cathodal groups ( $n = 9$ ) using the metric of perpendicular displacement at 100ms. We found that the mean p-value for this corrected comparison was  $0.013 \pm 0.018$  (mean  $\pm$  SD), indicating that the cathodal group learned significantly slower than sham, even when the group sizes were equalized.

### **3.3.5 Effects of motor cortex stimulation on learning from error**

Were the changes in rates of learning specific to the stimulation of the cerebellum? Fig. 3.2E and Fig. 3.2F display reach trajectories and perpendicular velocities for anodal M1 and sham groups during various stages of training. Anodal M1 stimulation did not appear to induce significant changes in reach kinematics. For example, the over-compensation early in the reach and the ‘S’ shape of the hand path appeared unaffected by M1 stimulation. These patterns were quantified via perpendicular velocity at 100ms (Fig. 3.2G) and maximum perpendicular velocity (Fig. 3.2H). Following a GLM analysis (reported in the cerebellar section above), a post-hoc comparison did not find a significant difference between M1 and sham

groups (Dunnett’s t-test, 100ms:  $p > 0.9$ , max displacement:  $p > 0.1$ ).

In summary, we observed an increased rate of learning when the cerebellum received anodal stimulation, but not when anodal stimulation was applied to the motor cortex.

### 3.3.6 Effects of cerebellar stimulation on stability of the motor memory

To assess stability of the acquired motor memory, we focused on the force patterns that the participants produced in error-clamp trials. Fig. 3.3A shows examples of these forces. In general, cerebellar stimulation did not alter the shape of the force profiles. Rather, participants who received cathodal cerebellar stimulation tended to produce smaller forces. We quantified force traces in error-clamp trials by comparing them to the ideal force, as defined in Eq. 3.1, and computed a force index (reflecting the fraction of compensation). This measure is shown in Fig. 3.3B. Two features stand out: 1) block after block, the force index increases, compensating for a greater amount of the perturbation, and 2) within each error-clamp block the force index decreases, reflecting decay of motor output and de-instantiation of the motor memory in the absence of error (Vaswani and Shadmehr, 2013).

To quantify the between-block change in the force index, independent of the within-block decay, we focused on the first five trials of each block. For each participant we computed the average force index across these five trials in each block. A GLM revealed a significant main effect of block ( $F(9, 414) = 10.0$ ,  $p < 10^{-3}$ ) as well as a significant effect of tDCS ( $F(3, 46) = 6.2$ ,  $p < 0.001$ ). A post-hoc test revealed that the cathodal cerebellar group produced significantly smaller forces in error-clamp trials compared to sham (Dunnett’s t-test,  $p < 0.001$ ) (Fig. 3.3B). Reach kinematics had shown that anodal stimulation of the cerebellum led to a larger amount of over-compensation than cathodal or sham stimulation

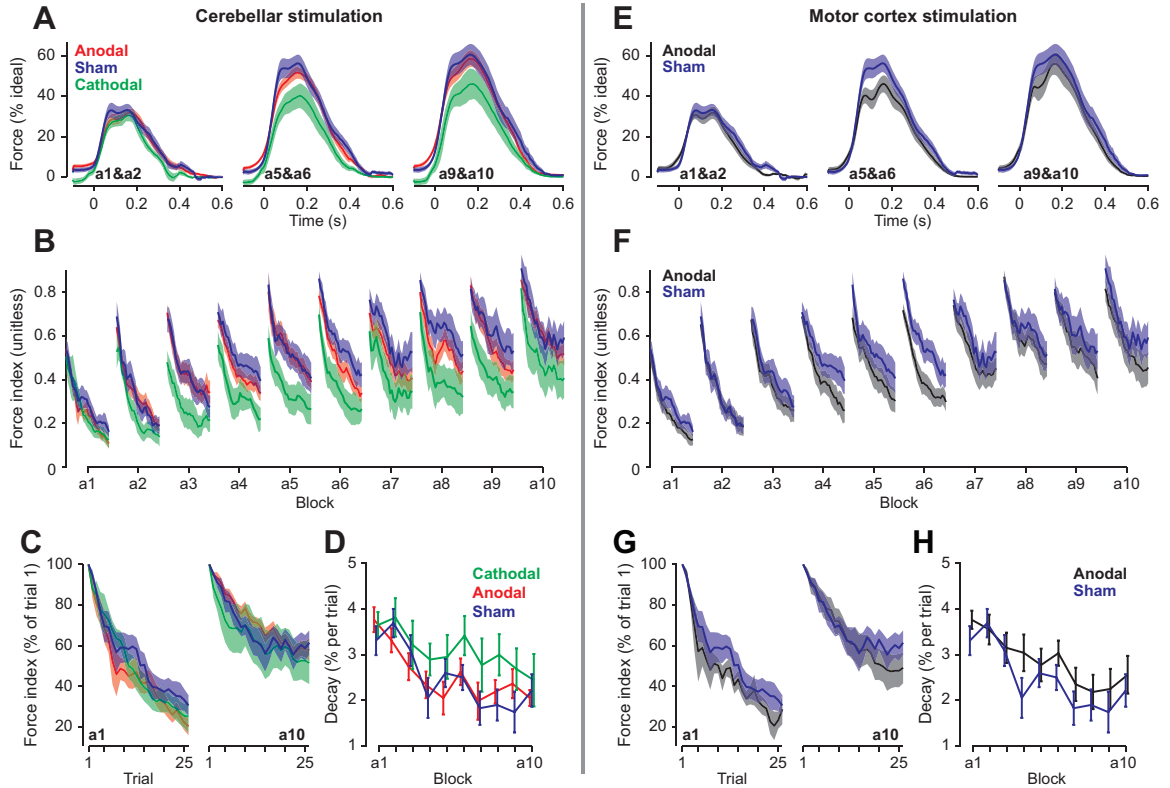


Figure 3.3: Force in error-clamp blocks and measures of decay in motor output during cerebellar or motor cortex stimulation. **A.** Force in error-clamp trials in various blocks as a percentage of ideal force in the cerebellar stimulation group. The ideal force was computed at time of peak velocity during the reach by multiplying velocity by field strength. **B.** Force index (Eq. 3.1, a unitless variable) as computed in error-clamp blocks. Data were smoothed using a sliding window with a bin width of 5 trials. **C.** In error-clamp blocks forces decay. However, with training the memory becomes resistant to decay. Forces were normalized to the first trial of the error-clamp block. The traces represent data from blocks a1 and a10. Data were smoothed using a sliding window with a bin width of 5 trials. **D.** Decay per trial in each block was estimated by fitting a line to the data shown in part B. The slope of the regression line represents the rate of change in units of percent force index per trial. **E-H.** The same as parts A-D but for anodal stimulation of the motor cortex. Data are mean  $\pm$  SEM.

(Fig. 3.2C). However, the force measurements in error-clamp trials did not suggest a difference between the anodal and the sham groups (Dunnett's t-test,  $p > 0.4$ , and as shown in Fig. 3.3B). We hypothesized that the reason for this may be that our force measure (force index) had quantified the entire trajectory, rather than focusing on the early component of the movement (when over-compensation occurs). Therefore, we performed further analysis of the force, but now focused on movement onset. We defined movement onset as the time when the reach exceeded 10% of maximum velocity in the direction of the target. This corresponds to time zero on the force traces shown in Fig. 3.3A. Focusing on force at movement onset, and for the first five trials of each error-clamp block, GLM showed a significant main effect of block ( $F(9, 414) = 3.2, p < 10^{-3}$ ) and a significant main effect of stimulation type ( $F(3, 46) = 5.4, p < 0.01$ ). Post-hoc tests revealed that the anodal group produced significantly larger forces compared to sham (Dunnett's one-sided test,  $p < 0.05$ ) but the cathodal group continued to produce significantly smaller forces than sham (Dunnett's one-sided test,  $p < 0.001$ ). In summary, force measurements early in the reach at start of error-clamp blocks confirmed kinematic measurements in field trials, demonstrating increased learning with anodal cerebellar stimulation and decreased learning with cathodal cerebellar stimulation. A critical question was whether cerebellar stimulation affected the force decay patterns in error-clamp trials. To assess the within error-clamp block change in the force index, we computed the rate at which this index decayed. For example, in a1 the force index was around 0.55 at the start of the error-clamp block (Fig. 3.3B). This implies that in block a1, in the field trials that had preceded the start of the error-clamp block, participants learned about 55% of the ideal force. As the error-clamp block in a1 ended the forces had decayed to approximately 20%. Therefore, the small number of field trials in a1 produced a great deal of learning (55% of force was learned), but the resulting memory exhibited decay in the absence of error (loss of around 63%).

In a10 the force index was around 0.85 at start of the error-clamp block and decayed to around 0.55 by the end of the block, exhibiting about 35% loss. A useful way to visualize these patterns is to normalize the force measure with respect to the first trial of each error-clamp block, illustrated in Fig. 3.3C. At the start of training the memory could be described as ‘fast,’ exhibiting rapid decay (Smith et al., 2006). With training, the memory decayed less in the error-clamp block, becoming ‘slow’.

To quantify the decay patterns we fitted the data in Fig. 3.3B to a single line for each block and each participant and measured the slope of that line. The results are shown in Fig. 3.3D, represented as percent decay per trial. GLM analysis revealed a significant main effect of block ( $F(9, 411) = 6.9, p < 0.001$ ), but no significant effect of tDCS ( $F(3, 46) = 1.6, p > 0.1$ ), and no interactions ( $F(27, 411) = 0.8, p > 0.7$ ). Therefore, cerebellar stimulation did not significantly alter the rate of decay in error-clamp blocks. Similar results were obtained with other measures of performance, such as force at peak velocity.

In summary, analysis of forces in error-clamp trials demonstrated that cathodal cerebellar stimulation slowed the rate at which the brain learned to predict and compensate for the perturbation. In blocks of error-clamp trials, these forces decayed. Early in training the decay per trial was large, but with further training decay per trial became smaller, suggesting that with training the memory gained stability. Cerebellar stimulation did not significantly alter these decay patterns. Therefore, cerebellar stimulation affected the rate of learning, but not the rate of decay of the resulting memory as assayed in error-clamp trials.

### **3.3.7 Effects of motor cortex stimulation on stability of the motor memory**

Fig. 3.3E displays the force traces produced by the group that received anodal stimulation of the motor cortex, and Fig. 3.3F summarizes these results using the

force index. To quantify the between-block change in the force index, independent of the within-block decay, we focused on the first five trials of each block. For each participant we computed the average force index across these five trials in each block. As reported in the above results, a GLM had revealed a significant effect of block, and significant effect of tDCS. However, a post-hoc test revealed no significant effects of anodal M1 stimulation (Dunnett's t-test,  $p > 0.1$ ). Fig. 3.3G and Fig. 3.3H display the decay properties of the force index, demonstrating that with training the decay rates were reduced (statistics reported above). Nevertheless, anodal M1 stimulation did not produce any significant changes in these decay patterns (as demonstrated by the lack of stimulation effect in the GLM, reported above). Therefore, in contrast to what was observed in a visuomotor rotation experiment (Galea et al., 2011), in this force field task anodal M1 stimulation did not alter the rate of decay of the motor memory in the absence of error.

### 3.3.8 Effect of stimulation on generalization

After completion of block a11, participants were tested in a generalization block (g2), to assess transfer of performance from trained targets to nearby untrained targets. Reaches to the generalization targets were in error-clamp. A fraction of the reaches to the trained targets were in field trials (to prevent decay of the motor output), and the remaining reaches were in error-clamp trials (to assess the force index). The force index for the generalization trials was expressed as a fraction of the average index for the two trained targets (Fig. 3.4A). Generalization was tested using a compound-symmetric GLM with factors of direction and tDCS group.

There was a significant effect of direction ( $F(6, 272) = 10.5, p < 0.001$ ), reflecting transfer of training to untrained target locations, but no effect of tDCS ( $F(3, 46) = 0.3, p > 0.8$ ) and no interaction effect ( $F(18, 272) = 1.2, p > 0.2$ ).

Therefore, stimulation of the motor cortex or the cerebellum did not significantly

alter generalization patterns.

### 3.3.9 Effect of stimulation on over-night retention

Training on Day 1 ended with a final block of field trials (r1), which significantly improved performance with respect to block a11 (Fig. 3.2D, perpendicular velocity at 100ms, main effect of block  $F(1, 46) = 21.9$ ,  $p < 0.001$ , with no effects of tDCS, and no interaction). Subsequently, participants left the experiment room and returned a day later. Testing on Day 2 began with a block of error-clamp trials (block b1), which precluded re-exposure to the previously trained force field. As before, we quantified the forces that subjects produced on each error-clamp trial of Day 2 using a force index, and the results are plotted in Fig. ??B and Fig. ??C. A one-way ANOVA on the force index averaged across block b1 revealed a significant effect of stimulation type ( $F(3, 46) = 3.1$ ,  $p < 0.05$ ). Post-hoc tests revealed a significant difference between the cathodal and sham groups only (Dunnett's t-test,  $p < 0.05$ ). Therefore, participants who had received cerebellar cathodal stimulation on Day 1 produced smaller forces on Day 2.

However, learning on Day 1 in the cerebellar cathodal group had been impaired by tDCS. Indeed, the cathodal cerebellar participants did not attain the same level of task performance as the other groups. Hence, between-group differences on Day 2 do not simply reflect differences in retention. To address this issue, we used an approach based on previous work. Joiner and Smith (2008) trained groups of volunteers in a force field task for various durations, yielding different levels of task performance. They then tested each group on Day 2 in error-clamp trials. The authors found that final performance on Day 1 was not a good predictor of forces exerted on Day 2. Rather, a specific component of performance on Day 1 was a good predictor of Day 2: the component attributed to 'the slow process', that which in the absence of error shows little decay. The authors showed that, despite

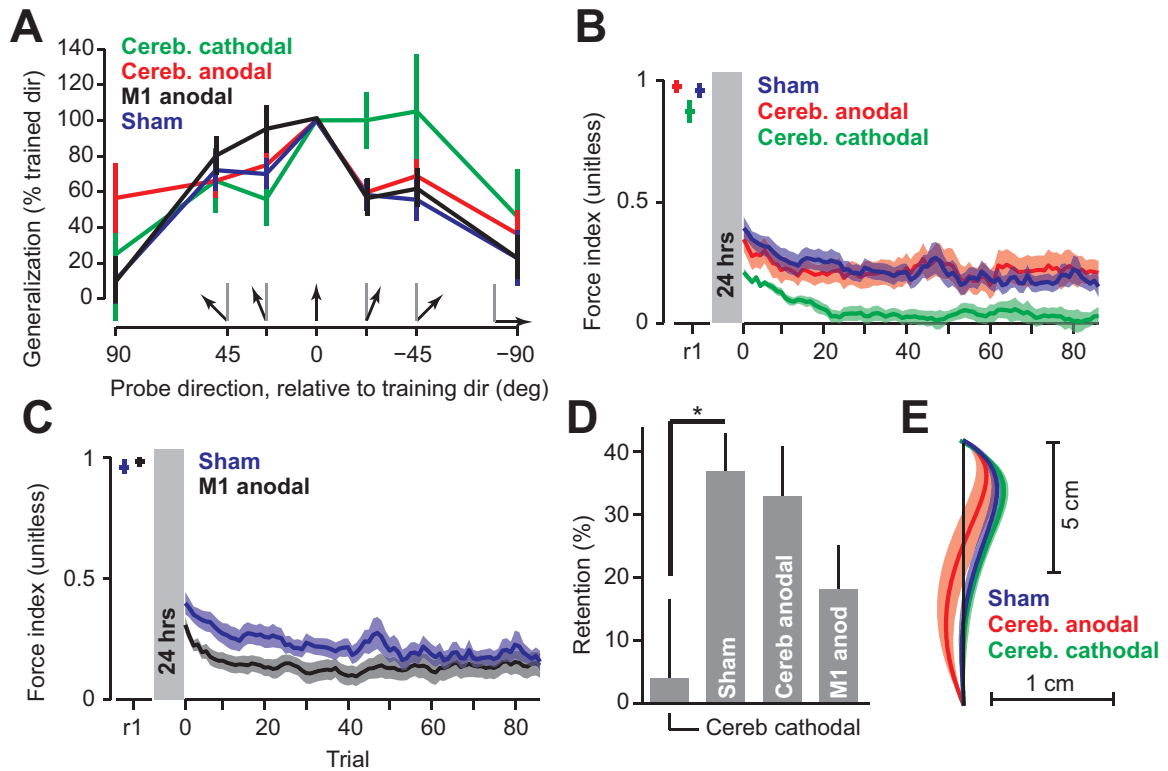


Figure 3.4: Effects of tDCS on generalization and retention. **A.** Generalization was assayed at end of training in Day 1 (block g2). Force index for each participant at the various probe targets was normalized to each participant's own force index in the trained targets (direction 0). The arrow indicates direction of probe target and the gray line indicates direction of trained target. Data are mean  $\pm$  SEM. **B.** Retention at 24 hours following completion of training. Force index at the end of training on Day 1 (block r1), and during testing on Day 2 (block b1) for cerebellar stimulation. The data in block b1 were smoothed with a sliding window using a bin width of 5 trials. **C.** The same as in part B, but for anodal stimulation of the motor cortex. **D.** Overnight retention, measured as average force index in block b1 as a percentage of force index in the last 10 trials of block a10. \* indicates  $p < 0.05$ . **E.** Hand paths in block b2. Data are mean  $\pm$  SEM.



different final levels of task performance on Day 1, the amount of force participants produced on Day 2 was a constant fraction of this slow-process component of the forces produced on Day 1. In our experiment, the error-clamp blocks on Day 1 were 30 trials in duration, long enough to be dominated by the slow process, as the fast process has a time constant in which forces decay by 95% by the 8th trial (Smith et al., 2006). Therefore, to compute retention, we averaged the force index during the final two error-clamp blocks on Day 1 (blocks a9 and a10), and then compared this for each participant to the average force index during the error-clamp block on Day 2 (block b1). The results are plotted in Fig. 3.4D. ANOVA indicated a significant effect of stimulation type ( $F(3, 46) = 3.19, p < 0.05$ ), and a post-hoc test revealed impaired retention in the cathodal group (Dunnett's t-test,  $p < 0.05$ ).

Following the error-clamp block b1, re-learning was assessed in a block of field trials (block b2 was examined in a subset of participants from each cerebellar tDCS group,  $n = 12/12$  anodal cerebellar,  $n = 10/12$  sham, and  $n = 8/10$  cathodal). Fig. 3.4E shows hand trajectories during block b2. All groups exhibited faster re-learning (i.e., 'savings'), showing a maximum perpendicular displacement in the first 10 trials of b2 that was within 95% of the final value of block r1 in the previous day. In addition, the anodal group exhibit greater over-compensation than the cathodal group (one-way ANOVA on the perpendicular velocity at 100ms identified a main effect of stimulation type,  $F(2, 28) = 3.7, p < 0.05$ , and Dunnett's post-hoc t-test showed significantly larger over-compensation in the anodal versus sham groups,  $p < 0.05$ ).

In summary, we found that retention, as measured by the ratio of force produced in error-clamp trials on Day 2 with respect to end of Day 1, was impaired in the cerebellar cathodal group, but unaffected by anodal M1 or anodal cerebellar stimulation.

## 3.4 Discussion

We performed a two day experiment to measure effects of non-invasive brain stimulation on the ability to learn to reach in a force field. We found that increasing the excitability of the cerebellum via anodal tDCS increased the rate of learning, while decreasing cerebellar excitability via cathodal tDCS impaired the ability to respond to sensory feedback and decreased the rate of learning. On Day 1, training resulted in a motor output that decayed in the absence of error. This decay was fast in the early part of training, but with further training the decay slowed, suggesting that with training the motor memory gained stability. Stimulation of the cerebellum or the motor cortex did not alter these decay patterns. On Day 2, when re-exposed to the same learning context, participants reproduced some of the motor commands that they had learned the previous day. Participants who had acquired the task while receiving cathodal cerebellar stimulation exhibited impaired retention, whereas anodal stimulation of the motor cortex or the cerebellum did not alter overnight retention.

### 3.4.1 Feedback control

When the nervous system detects an error during a reach, motor commands that correct the error and bring the hand to the target originate in the spinal cord, the motor cortex, and the cerebellum. If the cerebellar deep nuclei are cooled, the early component of the error-feedback response (associated with a response in the agonist muscle) is generally unaltered, but the later component (associated with a response in the antagonist muscle) is delayed (Vilis and Hore, 1980). Here, we observed that cathodal cerebellar stimulation reduced the feedback gain of the arm, resulting in motor commands that were slower than normal in correcting for the perturbation. How might cathodal stimulation of the cerebellum affect the error-feedback

response? Results from TMS experiments (Ugawa et al., 1995) suggest that cathodal tDCS decreases the resting membrane potential of cerebellar neuron (Galea et al., 2009), apparently decreasing the proportion of cells that respond to input. Because Purkinje cell activity is modulated by unexpected sensory feedback in the context of a self-generated movement (Gilbert and Thach, 1977; Brooks and Cullen, 2013), a reduced sensitivity to mossy fiber input may underlie the impairment in error-feedback response.

Given the extensive evidence regarding the role of the motor cortex in feedback control (Evarts and Tanji, 1976; Kimura and Gomi, 2009) it seems likely that disruption of M1 via cathodal stimulation, something that we did not attempt, would also affect the ability of the brain to respond to a perturbation. An interesting future experiment would be to compare the effects of cathodal stimulation of the cerebellum with M1.

### **3.4.2 Learning from error**

We used reach kinematics to assay learning from error and found that anodal cerebellar stimulation enhanced error-dependent learning, whereas cathodal cerebellar stimulation impaired it. This is consistent with results reported in a visuomotor rotation task, in which anodal cerebellar stimulation enhanced learning (Galea et al., 2011), and in a walking task, where anodal cerebellar stimulation enhanced learning while cathodal stimulation impaired it (Jayaram et al., 2012). Together, the results suggest that the cerebellum is a unique structure that supports the general process of error-dependent motor learning.

In the force field task, the result of learning is not a return to the null, unperturbed trajectory (Izawa et al., 2008). Rather, movements exhibit over-compensation early in the reach and under-compensation late in the reach, resulting in an S-shaped path to the target. Why do the motor commands exhibit

over-compensation? The motor commands that are produced in response to the perturbation during the reach may act as a teacher for the brain (Kawato, 1996), driving the change in motor commands that are generated in the subsequent movement (O’Shea et al., 2014). Recordings from muscles show a gradual and orderly transition of the motor commands from one that responds to the perturbation force during the reach (early in training), to one that predicts it near the onset of the reach (late in training) (Thoroughman and Shadmehr, 1999). The fact that cathodal cerebellar stimulation impaired both functions suggests that the later function may benefit from the former.

In contrast to the effects of cerebellar stimulation, we did not observe any effect of anodal stimulation of M1 on learning from error (2mA, 25min, 25cm<sup>2</sup> electrodes). This is consistent with an earlier work in which we found no effect of anodal or cathodal stimulation (1mA, 20min, 25cm<sup>2</sup> electrodes) of M1 in a similar force field task (de Xivry et al., 2011b). Similarly, in a visuomotor rotation task, Galea et al. (2011) used anodal stimulation of M1 (2mA, 15min, 25cm<sup>2</sup> electrodes) and found no effects of stimulation during training in the presence of the perturbation and no differences in the subsequent after-effects when the perturbation was removed. By contrast, Hunter et al. (2009) reported that anodal tDCS of M1 (1mA, 17min, 35cm<sup>2</sup> electrodes) produced a larger reduction in kinematic errors from the first to the 4th block of force-field training than did sham stimulation (their ‘signed-error’ measure). Following training in the field, these subjects were exposed to a null field condition, in which they exhibited after-effects. However, using the same ‘signed-error’ measure Hunter et al. (2009) did not find an effect of tDCS on the resulting after-effect. Hence, the results of tDCS studies do not, at present, paint a consistent picture of the function of M1 during motor learning. Most of the studies to date, however, have found that learning from error is not affected by anodal stimulation of M1.

### 3.4.3 Functional stages of motor memory

While there are many factors that can affect kinematic performance in field trials, including changes in muscle co-contraction (Thoroughman and Shadmehr, 1999), and changes in the gain of the long-latency sensory feedback pathways (Kimura and Gomi, 2009; Ahmadi-Pajouh et al., 2012) in error-clamp trials these factors are eliminated. Forces that participants produce in error-clamp trials are a proxy for a model that the brain constructs, associating state of the limb to expected perturbation forces (Hwang and Shadmehr, 2005; Sing et al., 2009). In Smith et al. (2006), we predicted that early in training, motor memory was ‘fast’, decaying rapidly in the absence of error, but that with further training, the memory was transformed to ‘slow’, showing gradual decay. Here we found direct evidence for this prediction: we observed that early in training the decay rates of motor output in error-clamp trials were high, but with further training the decay rates declined by about 50% (Fig. 3.3C). Therefore, with increased practice motor memory gained stability, as reflected in its decay properties in the absence of error. What was the neural basis of this transformation?

In our study, we found no effect of cerebellar or M1 tDCS on the rate of decay of the motor memory. By contrast, Galea et al. (2011) in a visuomotor rotation task found that anodal M1 stimulation reduced the decay rate of the learned motor output (assayed after learning/tDCS had finished, specifically when no visual feedback was provided). In our experiment, we repeatedly measured the decay rate of the evolving motor memory in the absence of error (error-clamp trials). Despite repeated measurements, we found no effects of M1 anodal stimulation. Of course, a null result does not constitute evidence of no effect. Nevertheless, our null effect observations are consistent with another work on force field learning (de Xivry et al., 2011b), in which anodal or cathodal stimulation to M1, or anodal

stimulation to the posterior parietal cortex (1mA, 20min), did not change the decay rates. What could explain this difference with respect to Galea et al.'s findings?

First, learning in force fields and visuomotor rotations engage distinct areas of the cerebellum (Donchin et al., 2012), and the cerebral cortex (Diedrichsen et al., 2005). This difference in functional anatomy may underlie the reported differences in the effects of M1 stimulation in force field and visuomotor tasks. For example, our earlier work on visuomotor rotation (Hadipour-Niktarash et al., 2007) found that M1 TMS during exposure did not affect the ability to learn from error, but resulted in a motor memory that was fragile, exhibiting rapid decay. Hence, for visuomotor rotation, our earlier results and those of Galea et al. (2011) are consistent, both finding no evidence of a functional role for M1 in learning from error, and both suggesting a role for M1 in the decay of the resulting motor memory. By contrast with these results for visuomotor rotation, in a force field task rTMS of M1 did not induce a deficit in retention, as assayed immediately after learning/stimulation (Baraduc et al., 2004), an effect that appears inconsistent with the predictions of Galea et al. (2011). The existing brain stimulation data suggest that different functional substrates mediate learning in visuomotor rotation and force fields.

Second, in our experimental design we included periodic error-clamp blocks, interleaved amongst blocks of learning in the field. Because the error-clamp blocks induce decay, they may reduce the overall amount of learning achieved during the task, and also reduce the rate of repetition of the motor commands, a natural component of most motor learning paradigms. Repetition is thought to produce a form of memory that is distinct from the memory that is produced from error-dependent learning (Diedrichsen et al., 2010; Huang et al., 2011). Importantly, repetition may produce a memory that depends on the cerebral cortex (de Xivry et al., 2011a). Hence, it is possible that if we had included a greater degree of repetition of motor commands in our training protocol, anodal M1 stimulation may

have slowed memory decay in error-clamp trials. Future work could test this hypothesis.

Finally, the way in which decay is assayed may change the effects of anodal M1 stimulation. Here, we measured decay using error-clamp trials in which the proprioceptive and visual error components of each movement were artificially constrained. In contrast, Galea et al. (2011) measured decay in trials in which visual feedback was withheld. Continuous feedback versus no-feedback trials have been shown to elicit differences in the rate of adaptation, and by extension, this difference likely impacts on the decay of acquired motor memories (Kitago et al., 2013). In particular, Galea et al. (2011) showed that when subjects were exposed to washout after learning, using a full-visual feedback condition, which requires a combination of learning from error as well as extinction of the acquired memories, there was no difference in the rate of decay with M1 anodal tDCS versus sham.

In summary, whereas our study of force field learning found that anodal M1 stimulation did not change the decay properties of the motor memory during acquisition, as assayed using error-clamp trials, the same stimulation in visuomotor rotation has been reported to reduce the decay rate of the acquired motor memory measured after learning. Hence, in combination, brain stimulation evidence to date suggests a role for M1 in stabilizing the motor memory that results from visuomotor rotation but not force field learning.

#### **3.4.4 Retention**

When participants returned on Day 2, they held the robot handle and reached in error-clamp trials. They produced forces that were correlated with those that they had learned on Day 1, demonstrating retention. These forces were significantly smaller in the cathodal cerebellar group than other groups. However, the critical question was whether this effect was a reflection of the fact that they had learned to

a lesser degree on Day 1, or whether the performance on Day 2 was evidence for reduced retention over and beyond the basic effect associated with acquisition. Therefore, to measure retention, we faced the issue that learning had been impaired in Day 1 in the cerebellar cathodal group: they had not reached the same levels of performance as other groups. To solve this problem, we used the analytic approach developed by Joiner and Smith (2008), which showed that retention of force field learning on Day 2, as assayed in error-clamp trials, was a constant fraction of the slow-component of forces produced on Day 1. We found that cathodal cerebellar stimulation showed significantly impaired retention. Anodal cerebellar or M1 stimulation had no effect on overnight retention.

Our results on the potential role of the cerebellum in retention are intriguing because of other results from the force field learning literature. Imaging studies of the cerebellum in the force field task suggest that during multi-week training activity in the anterior cerebellar cortex decreases while activity in the deep nuclei increases (Nezafat et al., 2001). In other motor tasks (e.g. VOR or optokinetic reflex), there is also evidence for this interplay between the cerebellar cortex and nuclei during acquisition and retention (Kassardjian et al., 2005; Okamoto et al., 2011a;b).

The fact that we did not observe an effect of anodal M1 stimulation contrasts with the results of Reis et al. (2009), who examined a skill learning task and found that M1 anodal stimulation produced greater over-night learning gains than sham. This highlights a potential difference between error-dependent learning, which appears to rely predominantly on the cerebellum, and skill learning, which has been proposed to rely more on the cerebral cortex.



### 3.4.5 Generalization

Generalization can be viewed as a signature of the tuning properties of the cells that participate in learning (Shadmehr, 2004). Force field learning produces narrow generalization to neighboring directions of movements, and broad generalization to neighboring positions of the arm, consistent with a neural coding that relies on proprioception (Hwang and Shadmehr, 2005). We have previously found that stimulation of M1 altered spatial generalization patterns, producing greater generalization in joint coordinates of the arm (de Xivry et al., 2011b). Here, we found that stimulation of the cerebellum or M1 did not affect directional generalization patterns, i.e., learning declined as a function of distance to the trained target. An important future experiment is to compare the effects of cerebellar and M1 stimulation on spatial generalization.

### 3.4.6 Limitations

Given the size of the tDCS electrodes ( $25\text{cm}^2$ ), and the dipole nature of a direct current stimulation montage, it seems likely that stimulation was not confined solely to the cerebellum or M1. For instance, it is well-established that M1 tDCS alters the excitability of the motor cortico-spinal tract (Nitsche and Paulus, 2000), and also changes functional brain activity in distal inter-connected brain regions, with the pattern of spread varying with cognitive state (Lang et al., 2005; Stagg et al., 2011). The functional consequence of these distal changes is unclear. With cerebellar stimulation, physiological evidence (MEPs) indicates that the tDCS-induced changes in measures of cerebellar-brain inhibition do not arise from local spread of current to the adjacent brainstem or visual cortex (Galea et al., 2011; 2009). Nevertheless, It is possible that cerebellar tDCS affects processing in M1 and thalamus by changing tonic neural activity in the cerebello-thalamo-cortical

pathway. Hamada et al. (2012) aimed to test this physiologically, by assessing sensory evoked potentials in M1 before and after anodal cerebellar tDCS, but they found no change in the excitability of these pathways. Hence, while future work is required to characterize the spatial distribution of tDCS-induced changes in functional brain activity, the available evidence, though not conclusive, does suggest that the current induced by the stimulation protocols used here probably affected mainly the cerebellum or the motor cortex.

We measured feedback response during the early phase of learning, and not in a situation where the perturbations were random. This potentially confounds the ability to learn from error (trial-to-trial change in motor commands), with the ability to correct for error (within trial change in motor commands). However, we think that we can dissociate these two factors: the main effect of learning from error was to produce changes very early in the movement, reflected in the perpendicular displacement at near movement onset (Fig. 3.2B), whereas the stimulation induced differences that we attributed to feedback control occurred late in the movement (Fig. 3.1B). Regardless, we envision a future experiment that includes continuous measurements of muscle activity in the context of feedback responses during cerebellar or M1 stimulation.

We found that anodal M1 stimulation produced no significant enhancement of learning or retention. Given the substantial neurophysiological evidence for involvement of M1 in the force field task (Li et al., 2001; Arce et al., 2010), and the fact that rTMS of M1 impairs overnight retention (Richardson et al., 2006) and the ability to switch from learning of one field to another (Cothros et al., 2006), an important next experiment is to compare the effects of cathodal M1 stimulation with cathodal cerebellar stimulation.

### **3.5 Conclusion**

In summary, we demonstrated that anodal stimulation of the cerebellum enhanced the error-dependent learning process, whereas cathodal stimulation impaired it. We demonstrated that with training, the motor memory was transformed from a process that decayed rapidly in the absence of error, to one that decayed slowly. Neither cerebellar nor motor cortical stimulation affected this transformation. Finally, we found that cathodal stimulation of the cerebellum during acquisition resulted in impaired retention as measured in 24 hours. Overall, we found a critical role for the human cerebellum in the ability to correct for error during a movement, the ability to learn from that error, and the ability to retain the resulting motor memory.

## 4 ENCODING OF ACTION BY THE PURKINJE CELLS OF THE CEREBELLUM

We showed in Chapter 3 that anodal transcranial direct current stimulation of the cerebellum, a stimulation technique that is thought to increase neuronal excitability, leads to faster learning of a novel motor task. In addition, we found that error-sensitivity is down-regulated when cathodal cerebellar stimulation is used. Our results, therefore, suggest a role for the cerebellum in manipulation of error-sensitivity.

However, to understand the mechanisms involved in modulation of error-sensitivity and its role in adaptation, it is useful to go beyond behavior and identify the neural substrates responsible for motor learning in the cerebellum. However, the storage of motor memories in the cerebellar circuitry has remained elusive since a clear encoding of movement kinematics has not been demonstrated.

Previous results has conclusively demonstrated that execution of accurate movements depend critically on the cerebellum. In this chapter, we focus on the simplest of all voluntary movements: saccades. These rapid eye movements have been shown to depend on the presence of an intact cerebellum, suggesting that the primary output cells in the cerebellum, Purkinje cells (P-cells), likely predict motion of the eye. Yet, this encoding has remained a long-standing puzzle: P-cells show little consistent modulation with respect to saccade amplitude, or direction, and critically, their discharge lasts longer than duration of a saccade.

In this chapter, we analyze P-cell discharge data from the oculomotor vermis of behaving monkeys, and find that individual neurons increase or decrease their activity during saccades. We then estimate the combined effect of these two

populations via their projections on the caudal fastigial nucleus (cFN) and uncover a simple-spike population response that precisely predicts the real-time motion of the eye. Our results suggest that it is not the response of an individual Purkinje cell that is responsible for encoding of saccade kinematics, but rather the combined response of P-cells across a population.

We then take our results a crucial step further: when we organize the P-cells according to each cell's complex-spike directional tuning, the simple-spike population response predicted both the real-time speed and direction of the saccade multiplicatively via a gain-field. This suggests that the cerebellum predicts the real-time motion of the eye during saccades via the combined inputs of P-cells onto individual nucleus neurons. A gain-field encoding of simple spikes emerges if the P-cells that project onto a nucleus neuron are not selected at random, but share a common complex-spike property. Only using this hypothesized organization of the cerebellum can we begin to understand the role of P-cell responses in motor learning tasks.

## 4.1 Introduction

Previous studies have focused on bursting activity of Purkinje-cells (P-cells) during saccades (Thier et al., 2000; Catz et al., 2005; 2008) and found no consistent modulation with saccade amplitude (Ohtsuka and Noda, 1995; Helmchen and Büttner, 1995) speed (Helmchen and Büttner, 1995; Thier et al., 2000; Kase et al., 1980), or direction (Ohtsuka and Noda, 1995). A recent simulation (Gad and Anastasio, 2010) suggested that P-cells that pause during saccades may be important in understanding the responses observed in the deep cerebellar nucleus neurons. The main puzzle that we wished to tackle was how the P-cells encoded the real-time motion of the eye.

## 4.2 Materials and Methods

We analyzed data from  $n = 72$  Purkinje cells in the oculomotor vermis of 5 rhesus monkeys as they made saccades to visual targets.

### 4.2.1 The dataset

The data set included Purkinje-cell (P-cell) discharge from the oculomotor vermis (OMV) in 5 rhesus monkeys (*Macaca mulatta*; males; 5.0-7.4kg; monkeys B, F, W, K, and KO). These data were collected during two previous studies (Kojima et al., 2010; Soetedjo et al., 2008). A scleral search coil was surgically implanted into one eye of each monkey, allowing measurement of eye kinematics via standard techniques (Fuchs and Robinson, 1966). Following recovery from surgery, the monkeys were trained to make saccades to visual targets (less than  $0.4^\circ$  in diameter) in a dimly lit room. The targets appeared within  $25^\circ$  of center. Monkeys were rewarded with applesauce for keeping their eyes within a virtual window which extended  $\pm 3^\circ$  in both the horizontal and vertical directions about the target. After the monkeys were trained to saccade to a single target, they were trained to make saccades between successively presented targets.

Once the monkey could reliably track the targets for an extended period of time, a recording chamber was implanted on the midline of the cerebellum (14.5mm posterior of the interaural axis and directed straight down), providing access to the OMV. Single-unit activity was recorded with homemade tungsten electrodes with an iron-particle coating ( $100\text{k}\Omega$  impedance at 1kHz). The position of the electrode for each recording was measured with respect to the center of the recording chamber, providing approximate coordinates of each P-cell within OMV. We recognized OMV by observing saccade-related changes in background activity. An isolated unit was classified as a P-cell if it produced a complex-spike (CS), which

was identified online as a positive action potential with multiple wavelets. We focused our recordings on units which showed a saccade-related change in simple-spike activity in at least one direction (a burst, pause, or a combination of the two). Saccade activity was assessed while the monkey made  $15^\circ$  saccades from the central fixation point to one of eight targets spaced at  $45^\circ$  intervals.

Neurophysiology data was sampled at 50kHz by a Power 1401 digitizer (Cambridge Electronic Design, Cambridge, UK) and subsequently band-pass filtered between 30Hz and 10kHz. The location of the eye, as measured by the scleral search coil, was sampled at 1kHz. Data were displayed in real-time on a computer monitor running Spike2 and saved for offline analysis (Soetedjo and Fuchs, 2006).

We performed spike-sorting to isolate the simple-spike activity of each P-cell. The timing of each simple-spike was identified and subsequently down-sampled to 1kHz to coincide with the timing of the behavioral recordings. We identified saccades via an absolute velocity threshold, marking onset of the saccade as the time when the speed of the eye exceeded  $20^\circ/\text{s}$ . The end of the saccade was similarly defined as the time when speed fell below  $20^\circ/\text{s}$ . Trials in which the monkey moved its eyes in the wrong direction ( $> 90^\circ$  with respect to the target), or trials in which the error at the endpoint exceeded  $15^\circ$ , were removed from the analysis (4.6% of all trials). The peak speed for each saccade was determined as the maximum magnitude of the velocity vector in the direction of the presented target.

To convert the simple-spikes of the P-cells into firing rates, we computed the inter-spike interval between two consecutive spikes, and then replaced the period between these two spikes with a box-car of magnitude equal to the reciprocal of the interval. We smoothed the resulting time-sequence with a normalized Gaussian of 2.5ms standard deviation (integral one), and then used this rate function to compute the mean and peak firing rate of the cell during the saccade period. The peak firing rate is the mathematical maximum of the firing rates during the saccade

period, for both pause and burst cells. We use the term “firing rate” to refer to this instantaneous rate, and quantify it using units of Hertz (Hz).

### 4.2.2 Statistical analysis

Statistical analyses were conducted in R (R-project, Vienna, Austria). To assess the relationship between two variables we report Pearson’s  $R^2$  coefficient as well as the results of a repeated measures analysis of variance (RM-ANOVA). If the data used for an RM-ANOVA failed the test for sphericity, we report the Greenhouse-Geisser corrected statistics. Paired sample t-tests were used to assess differences in response characteristics as a function of direction (CS-on/off). In cases where we used independent samples t-tests, we also assumed unequal variances between the two groups. All tests were two-sided with a significance level of 0.05, unless otherwise noted.

### 4.2.3 Complex-spikes

After we had identified a P-cell that exhibited phasic changes in simple-spike activity during a saccade, we identified the cell’s preferred complex-spike direction, termed CS-on. To do so we induced errors at the end of a saccade and recorded the resulting complex spikes in the P-cell. To induce errors, we used the intra-saccadic step paradigm (McLaughlin, 1967), blanking the original target during execution of the saccade and replacing it with another target so that at saccade termination the eyes appeared to miss the target by approximately  $5^\circ$ , as illustrated in Fig. 4.3A. The monkey began by fixating the central point. The central point disappeared and a target appeared at either  $12^\circ$  or  $15^\circ$  in one of 8 directions (spaced at  $45^\circ$  intervals). During the execution of the saccade to the target (when the eye velocity exceed  $70^\circ$ ), the target was back-stepped by  $5^\circ$  relative to the original target. As a result the saccade over-shot the target, resulting in an error of approximately  $5^\circ$ . A



second presentation of this intra-saccadic step paradigm in the opposite direction brought the monkey's eyes back to the central fixation point. For each monkey, we collected more than 30 trials in each direction. We counted the number of complex-spikes in the 50-200ms period following termination of the first (primary) saccade. The back-step direction (error direction) which elicited the highest probability of complex-spikes was classified as the CS-on direction (Fig. 4.3B).

The paradigm that we used for identifying the error vector that produced the highest probability of complex-spikes (i.e., CS-on direction) suffered from the short-coming that saccade direction and error direction were  $180^\circ$  apart. Therefore, we examined whether probability of CS was dependent only on the direction of error, or whether it was also affected by the saccade direction that preceded that error (Soetedjo and Fuchs, 2006). To dissociate these two variables, for  $n = 39$  cells we systematically varied both error and saccade directions. An example of this is shown for a cell in Fig. 4.3C. This cell (N1) had a high probability of complex-spikes when the error vector was  $-45^\circ$ , regardless of whether the primary saccade was at direction of error+ $0^\circ$ , or direction of error+ $180^\circ$ . Similarly, the cell had a low probability of complex-spikes when the error vector was  $+135^\circ$ , regardless of whether the primary saccade was at direction of error+ $0^\circ$ , or direction of error+ $180^\circ$ .

Across the population of  $n = 39$  cells probability of complex-spikes was modulated by direction of the error vector, but not the direction of the saccade (Fig. 4.3D): two-way RM-ANOVA showed a main effect of error direction ( $p < 10^{-4}$ ) but no effect of saccade direction ( $p > 0.5$ ).

For  $n = 39$  cells we were able to maintain excellent isolation of the P-cell throughout the recording period. For these cells, we manually identified a subset of complex-spike waveforms which served as a template for matching against all other complex-spikes during the recording. Using a Gaussian mixture model, we determined the presence/absence of a complex-spike in 30ms overlapping windows,

providing us with the onset of a complex-spike with millisecond resolution. Looking at the same period (50-200ms following the primary saccade), we were able to construct an estimate of the probability of observing a complex-spike as a function of the angle between the endpoint of the primary saccade and the back-stepped target with better than  $45^\circ$  resolution (Fig. 4.6D, brown). For the remaining  $n = 33$  cells, we used the manually identified CS-on direction for all analyses.

The data sets that we analyzed were collected under two slightly different experimental protocols (Soetedjo et al., 2008; Kojima et al., 2010). In all cases, we focused our analysis on the simple-spike related activity of the primary saccade rather than changes in activity due to adaptation or complex-spikes. For  $n = 16$  units, we presented primary targets in either the CS-on or CS-off direction and then pseudo-randomly modified the magnitude of the intra-saccadic step (ranging from  $-9^\circ$  to  $+9^\circ$ ). Because the direction of intra-saccadic step (backwards, forwards, or no back-step) was randomized, the saccade to the presented target did not undergo adaptation. For the remaining cells, we asked the monkey to make primary saccades whose magnitude was either approximately  $15^\circ$  or  $25^\circ$ . Approximately 20 trials in both the CS-on and CS-off directions were presented without an intra-saccadic step ( $0^\circ$ ). After this block, we induced saccade adaptation in both the CS-on and CS-off directions by consistently stepping the target inwards ( $5-11^\circ$ ). We did not observe a qualitative difference between the population response for those saccades during the adaptation period and those during the stationary period (changes in eye velocity during adaptation coincided with changes in the magnitude of the population response). We therefore elected to use data from all saccades.

#### 4.2.4 Population response

About 50 P-cells synapse onto a single neuron in the caudal fastigial nucleus (Person and Raman, 2012). The P-cells that project onto a single cFN

neuron may be organized by their inputs from the inferior olive (De Zeeuw et al., 2011; Kojima et al., 2011). In this scenario, the olive projections divide the P-cells into clusters where each cluster of P-cells projects onto one cFN neuron. We therefore made the assumption that P-cells that projected onto a cFN neuron shared a critical feature: they had the same CS-on direction (Fig. 4.5A). We computed what a typical cFN neuron would receive from this population of P-cells by estimating a population response.

A recent study had shown that each simple-spike induced by the presynaptic firing of a P-cell influenced the post-synaptic cell in the deep cerebellar nucleus by producing an inhibition that had a 2.5ms time constant (Person and Raman, 2012). We therefore convolved the simple-spike train of each recorded P-cell with a normalized (i.e., integral of one) Gaussian of 2.5ms standard deviation. In all figures, we dissociate this quantity from the measurement of firing rate computed via the inverse of the inter-spike interval by using units of Hz to quantify firing rates, and units of spikes/second to quantify population response.

We calculated the population response by a bootstrapping procedure in which we sampled at random  $n = 50$  P-cells (with replacement) from our population of recorded neurons. In our recorded sample, we had roughly equal numbers of bursting ( $n = 39$ ) and pausing ( $n = 33$ ) cells. Therefore, in our bootstrapped population of 50 cells, we also had roughly equal numbers of bursting and pausing cells ( $27.1 \pm 1.9$  bursting cells, mean  $\pm$  SD). As the bursting and pausing populations had different saccade-related response profiles, standard error of the mean would result in a significant over-estimation of the variance of the population response that would be experienced at a nuclear neuron. We therefore elected to bootstrap 50 different populations of 50 randomly chosen P-cells. The standard error bars in these plots represent mean  $\pm$  SEM across bootstrapped samples.

## 4.3 Results

We analyzed simple-spike activity of 72 oculomotor vermis (OMV, cerebellar lobules VI and VII) P-cells from five monkeys during saccades. Our population included cells that exhibited increased activity (bursting;  $n = 39$ , Fig. 4.1A) or decreased activity (pausing;  $n = 33$ , Fig. 4.1B). Consistent with previous reports (Helmchen and Büttner, 1995; Thier et al., 2000; Catz et al., 2008), the majority of the neurons were poorly modulated by saccade amplitude (Fig. 4.1C, and Fig. 4.2). However, the mean firing rate of burst cells (but not pause cells) increased significantly with saccade peak speed (Fig. 4.1D,  $p < 10^{-10}$ ). Previous work had demonstrated that the population response encoded additional saccade-related information that was not reliably present in the responses of individual neurons (Thier et al., 2000; Dash et al., 2013; Prsa et al., 2009). To examine the population response, we measured change in firing rates (from baseline) for the bursting and pausing cells during slow ( $400^\circ/\text{s}$ ) and fast ( $650^\circ/\text{s}$ ) saccades (Fig. 4.1E), pooled across all directions. The onset of change in firing rates in both populations generally led saccade onset by more than 50ms. The termination of activity was also significantly later than the saccade: a  $650^\circ/\text{s}$  saccade was  $38 \pm 1.2\text{ms}$  in duration (mean  $\pm$  SEM), whereas activity of burst and pause cells persisted for a more than 100ms. Given that the cerebellum is thought to play a critical role in termination of ipsiversive saccades (Fuchs et al., 1993; Robinson et al., 1993), how can P-cells be involved in controlling the eye if their activity persists so much longer than the saccade?

### 4.3.1 P-cell population response predicts saccade speed

P-cells project to the caudal fastigial nucleus (cFN), where about 50 P-cells converge onto a cFN neuron (Person and Raman, 2012). For each P-cell we computed the probability of a simple-spike in 1ms time-bins during saccades of a

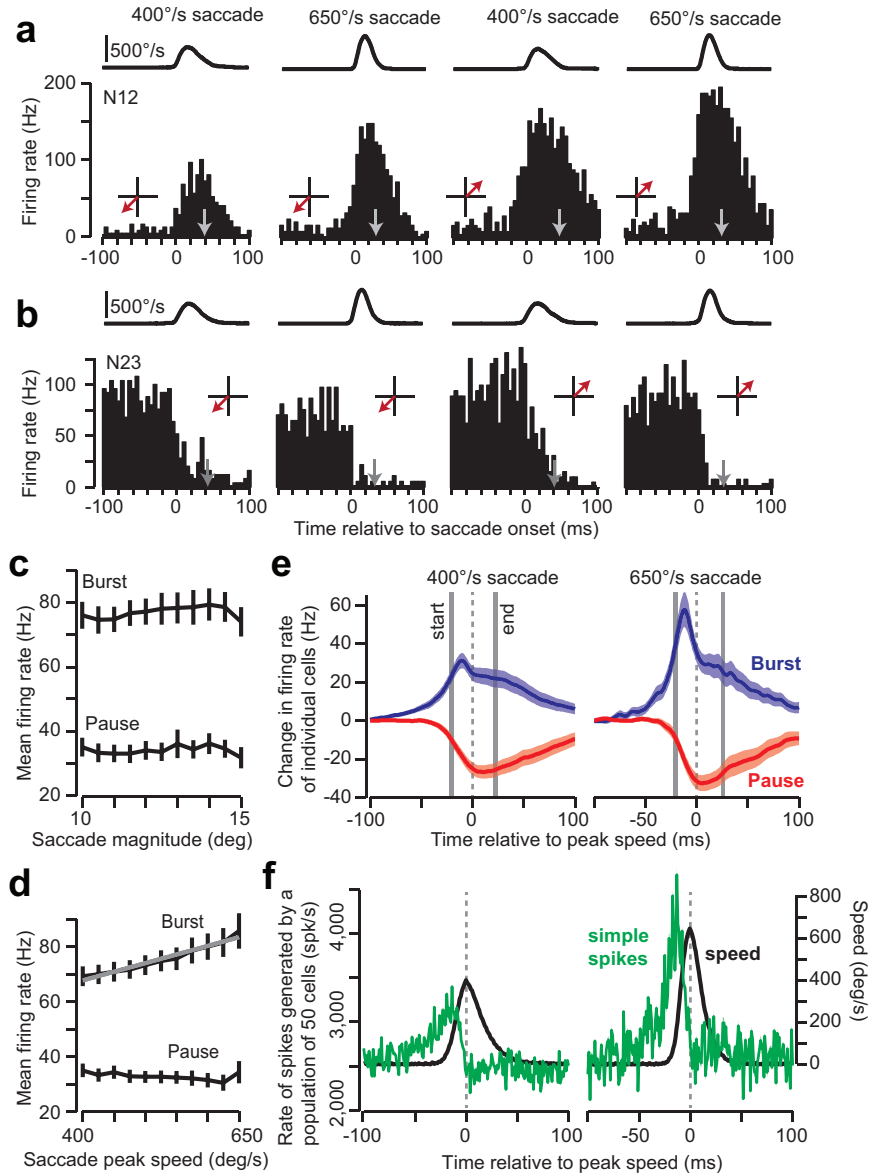


Figure 4.1: A population of burst and pause P-cells together predict eye speed in real-time. Perisaccade histograms for a bursting **A** and pausing **B** P-cell during saccades of various speeds and directions (red arrow). The trace on the top row is saccade speed. The gray arrow indicates saccade end. **C** and **D**. Mean firing rates over the duration of saccade computed across all directions. Changes in speed produced an increase in the firing rate of the burst cells but not the pause cells. **E**. Change in firing rates (with respect to baseline) of the bursting and pausing P-cells for two saccade speeds. Gray bars are onset and termination of the saccade (width is SEM). **F**. The total rate of simple-spikes produced by a random selection of 50 P-cells.

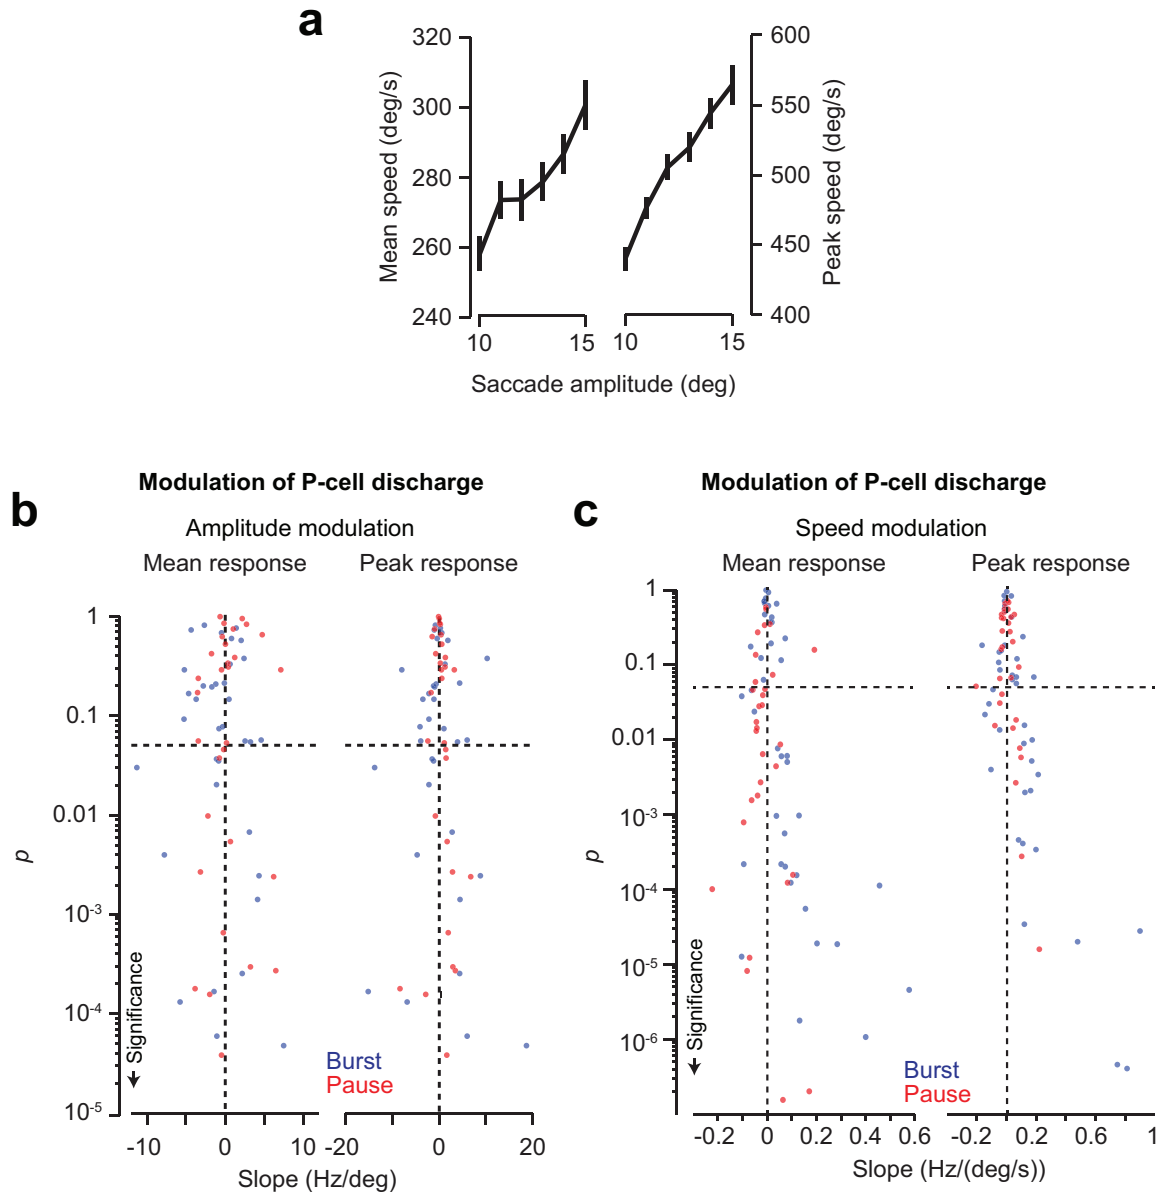


Figure 4.2: Firing rates of individual P-cells as a function of saccade amplitude and peak speed. **A.** Increase in saccade amplitude produced robust increases in mean and peak saccade speed (mean:  $R^2 = 0.86$ ,  $p < 10^{-4}$ ; peak:  $R^2 = 0.99$ ,  $p < 10^{-9}$ ). **B.** For each neuron, we correlated the average firing rate and the peak firing rate (computed over the saccade duration and averaged over all directions) with saccade amplitude. Some neurons increased their firing rates with increasing saccade amplitude (positive slope) and some neurons decreased their responses (negative slope). However, mean and peak firing rates of a majority of neurons (47/72) were not significantly modulated with saccade amplitude. As a result, activity of neither the burst nor the pause cells showed a significant modulation with saccade amplitude (Fig. 4.1C). **C.** Majority of neurons (45/72) had a significant linear relationship between firing rates and peak saccade speed. In particular, mean and peak response of burst cells showed a significant increase with peak speed (Fig. 4.1D).

given peak speed, averaged across all directions. We then chose 50 P-cells at random and computed the total number of simple-spikes generated by the population at each millisecond of time, resulting in an estimate of the rate of presynaptic spikes converging onto a cFN cell. The results (Fig. 4.1F) revealed a real-time encoding of the speed of the eye: the peak of the activity preceded peak speed, increased in magnitude when speed increased, and returned to baseline just before saccade termination ( $R^2$  at the optimal delay, 400 °/s:  $R^2 = 0.52$ ,  $p < 10^{-22}$ ; 650 °/s:  $R^2 = 0.62$ ,  $p < 10^{-43}$ ). It appeared that the simple-spikes of the pause and burst cells combined together to predict motion of the eye.

Let us hypothesize that the P-cells that project to a nucleus neuron are not selected randomly, but are organized by their inputs from the inferior olive (De Zeeuw et al., 2011). That is, suppose that the olive projections divide the P-cells into clusters where each cluster of P-cells projects onto a single nucleus neuron. The input from the olive produces complex-spikes (CS) in the P-cells. We found that if we organized the simple-spikes of the P-cells based on each cell's CS properties, additional features of the population activity were unmasked.

We measured CS properties of each P-cell by inducing a post-saccadic error through displacement of the target during the saccade, and then measured the probability of CS as a function of the direction of this error (Fig. 4.3, also Section 4.2.3). For each P-cell, the direction of error that produced the largest probability of a CS during the 50-200ms post-saccade period was labeled as CS-on, and the opposite direction was labeled as CS-off9 (Fig. 4.4). We then made the assumption that the P-cells that projected onto a nucleus neuron all had the same CS-on direction (Fig. 4.5A). Under this assumption, we computed the rate of presynaptic simple-spikes that a nucleus neuron would receive from the cluster of P-cells (see Section 4.2.4). We did this by convolving each P-cell simple-spike train with a 2.5ms standard-deviation normalized Gaussian, approximating the temporal

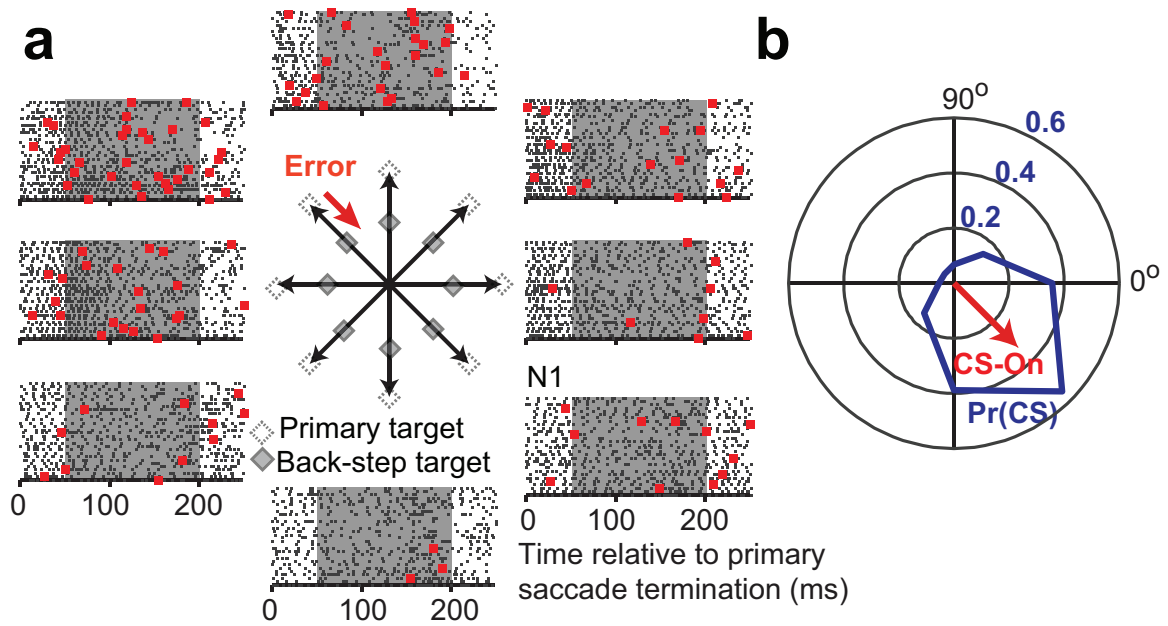


Figure 4.3: Determination of complex-spike (CS) properties of P-cells. **A.** Response of a P-cell during the 250ms period after completion of a saccade (simple-spikes are gray, complex-spikes are red). CS-on was determined via a back-step paradigm in which the target was jumped (unfilled target to filled target) during saccade execution. Black arrow indicates saccade vector, red arrow indicates error vector. We computed the probability of CS in the 50-200ms period following saccade termination. **B.** The probability of a CS as a function of the direction of the error vector. For this neuron, the highest probability (CS-on) occurred when the error vector was in direction  $-45^\circ$ . Direction of CS-off for this cell was  $135^\circ$ .

characteristics of the inhibition produced in the nucleus neuron due to a simple-spike in the P-cell (Person and Raman, 2012; Telgkamp et al., 2004).

Fig. 4.5B shows the change in population response from baseline when a saccade was made in the same direction as CS-off. The response rose above baseline before saccade onset, peaked prior to peak speed, and then returned to near baseline. The peak response scaled robustly with saccade amplitude (Fig. 4.5C,  $R^2 = 0.93$ ,  $p < 10^{-5}$ ). We observed a strong correspondence between the real-time population response and the real-time speed (lower plot of Fig. 4.5D, Fig. 4.7). The population response led eye speed by an average of  $21.2 \pm 0.4$ ms (correlation analysis in the CS-off direction, mean  $\pm$  SEM). Peak population response precisely



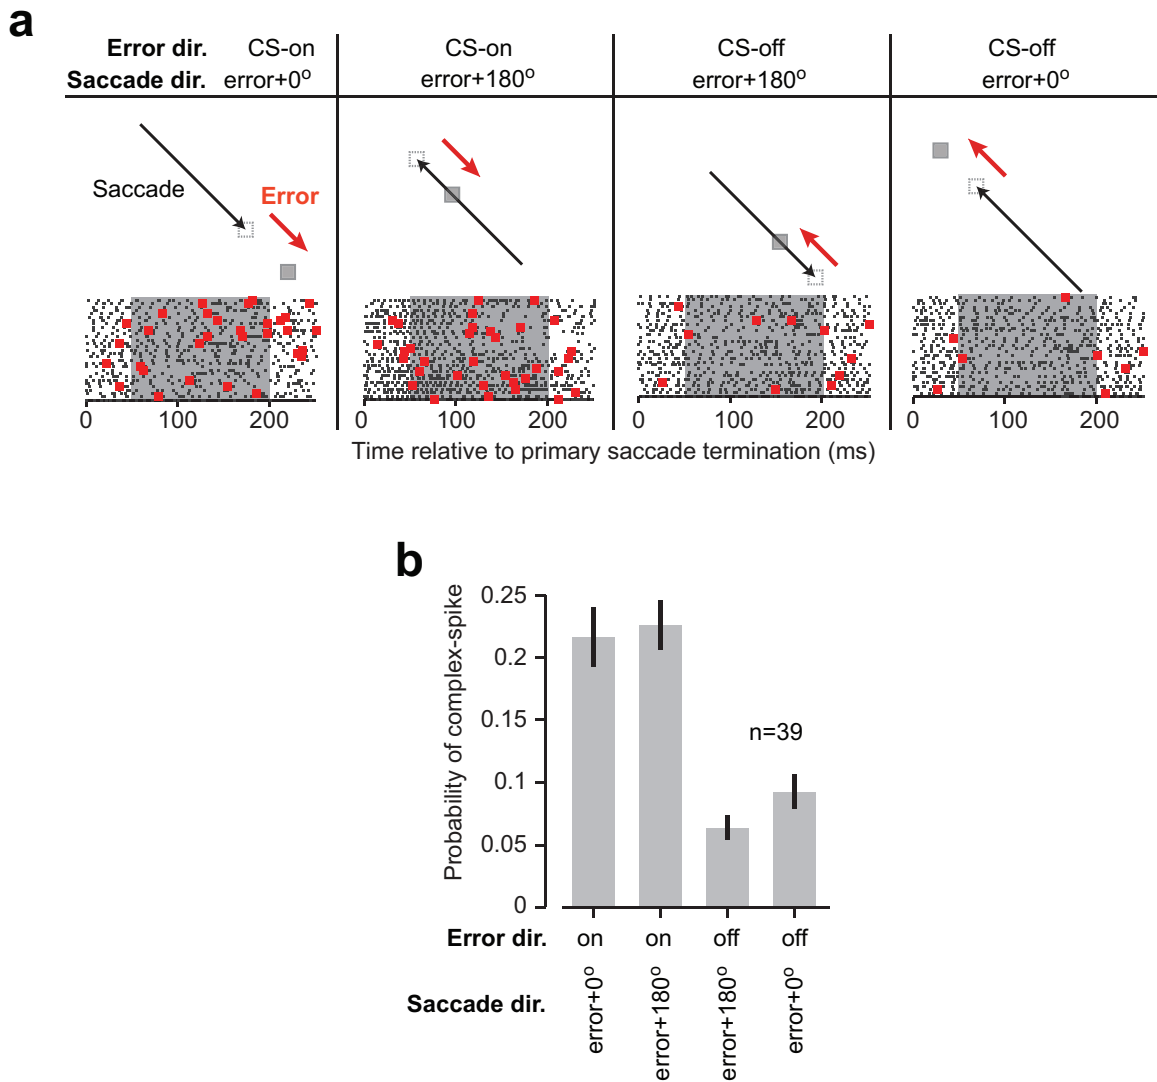
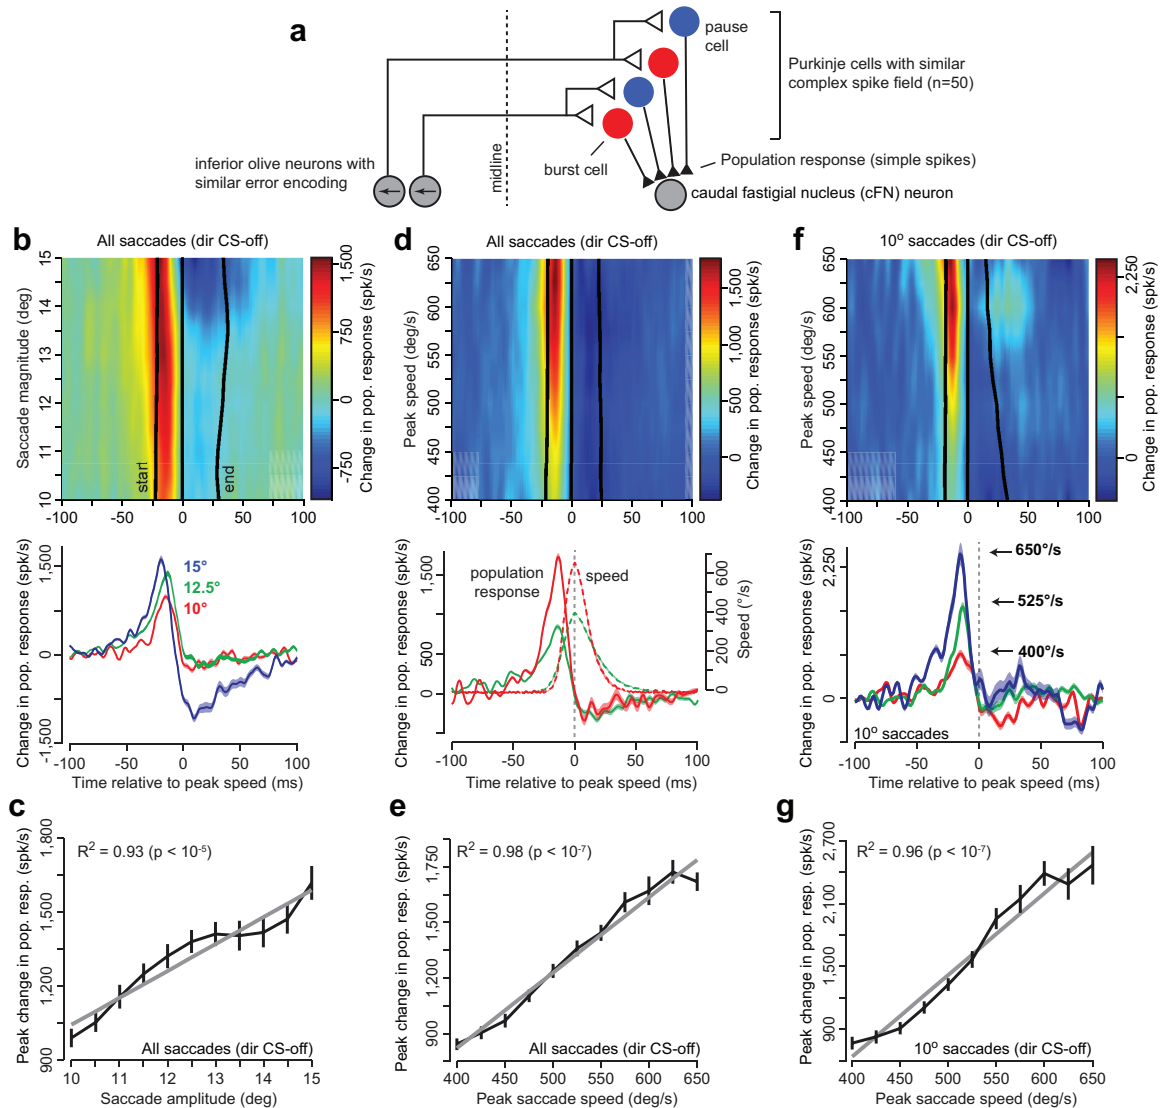


Figure 4.4: Complex spikes encode direction of error, not direction of the saccade that preceded that error. **A**. The response of the same cell shown in Fig. 4.3 as a function of direction of saccade and direction of error. The probability of CS is high when the direction of error is at  $-45^\circ$ , despite the fact that saccade direction may be at  $-45^\circ$  or  $+135^\circ$ . **B**. Population statistics from  $n = 39$  P-cells in which the probability of CS was quantified as a function of direction of the error vector and direction of the saccade that preceded that error. Probability of CS depended on direction of error, not direction of saccade.



predicted peak speed (Fig. 4.5E,  $R^2 = 0.98$ ,  $p < 10^{-7}$ ).

We took advantage of natural variability in saccades to further test the relationship between the population response and speed. We sorted all 10° saccades (direction CS-off) according to peak speed (Fig. 4.5F) and found that despite the constant amplitude, the population response precisely predict the actual peak speed of the eye (Fig. 4.5G,  $R^2 = 0.96$ ,  $p < 10^{-7}$ ).

Figure 4.5: A cluster of P-cells, organized by their complex-spikes, produced a population response that predicted in real-time the motion of the eye. **A.** Hypothesized organization of the oculomotor vermis. To compute a population response, we measured the simple-spikes of each P-cell as a function of saccade direction with respect to the CS-on direction of that cell. For the P-cells shown here, the CS-on is an error vector to the left (arrow). **B.** Change in population response (with respect to baseline) as a function of saccade amplitude in  $0.5^\circ$  bins, for saccades in the CS-off direction. Data in the amplitude axis were smoothed by a first-order Savitzky-Golay filter with a width of 3 bins (Thier et al., 2000). Bottom plot shows the population response for three representative amplitudes. **C.** Peak population response increased linearly with saccade amplitude. P-values indicates significant linear correlation. **D.** Population response as a function of saccade peak-speed. Bottom plot shows representative responses with their corresponding speed traces. **E.** Peak population response increased linearly with saccade peak-speed. **F.** Population response for  $10^\circ$  saccades ( $\pm 1^\circ$ ), as a function of saccade peak-speed. Bottom plot shows the population response for slow, medium, and fast saccades of  $10^\circ$  amplitude. **G.** Peak population response increased linearly as a function of peak-speed even for a fixed magnitude saccade. Error bars are SEM.

### 4.3.2 P-cell responses, organized by error, predict saccade kinematics as a gain-field

Therefore, when the simple-spikes were organized according to each cell’s CS directional preference, the population response for saccades of constant direction predicted nearly all of the variability in saccade peak speed. No previous work, to our knowledge, had revealed how direction of a saccade is encoded in the activity of P-cells. For the burst and pause cells, the mean and peak firing rates during the saccade did not vary as a function of direction (Fig. 4.8). However, organizing the population response according to each P-cell’s CS tuning preference revealed a clear encoding of direction: the peak response was greater if the saccades were in the same direction as CS-off as compared to CS-on (Fig. 4.6A, t-test  $p < 10^{-16}$ ). Indeed, the peak population response rose linearly as a function of peak speed in both directions, but with a larger gain when the saccade direction was congruent with CS-off (Fig. 4.6B, RM-ANOVA with main effects of peak-speed,  $p < 10^{-15}$ ;

CS-direction  $p < 10^{-7}$ ; and a speed by CS-direction interaction,  $p < 10^{-15}$ ).

To more closely examine the effects of saccade direction, we plotted the population response across saccade directions with respect to CS-on (Fig. 4.6C). We found that the population response was highest for saccades made in the CS-off direction, with an encoding of direction that resembled a cosine function (Fig. 4.6D).

That is, when we organized the P-cells into clusters in which each cluster was composed of bursting and pausing cells (54% burst cells), all with a common CS-on direction (Fig. 4.5A), and then computed the population response, i.e., an estimate of the rate of simple spikes that converged onto a single cFN neuron during a saccadic eye movement. We found that when the saccade was in direction CS-off, the peak population response correlated near perfectly with peak saccade speed at  $R^2 = 0.98$  (Fig. 4.5E). When the saccade was in direction CS-on, the peak population response also correlated near perfectly with peak saccade speed at  $R^2 = 0.99$  ( $p < 10^{-8}$ ) (Fig. 4.6B). A change in saccade direction altered the gain relating the population response to saccade speed (Fig. 4.6D): the gain was lowest when the saccade was in direction CS-on, and highest when the saccade was in direction CS-off. These results suggested that the P-cells, clustered by their CS-on direction, as a population produced simple spikes that were related to the real-time motion of the eye via a gain-field:

$$s(t) = |\dot{\mathbf{x}}(t + \Delta)| [a \cos(\theta - \theta_{CS}) + b] + c \quad (4.1)$$

In this equation,  $s(t)$  is the rate of simple spikes,  $\dot{\mathbf{x}}$  is the magnitude of the eye velocity vector at time  $t + \Delta$ ,  $a$  is a scaling factor,  $\theta$  is saccade direction,  $\theta_{CS}$  is the direction of CS-off for that cluster of P-cells, and  $b$  and  $c$  are baseline offsets. The fit of this equation to the measured data (real-time motion of the eye vs. real-time population response,  $t = -100\text{ms}$  to  $+150\text{ms}$  with respect to peak saccade speed, for

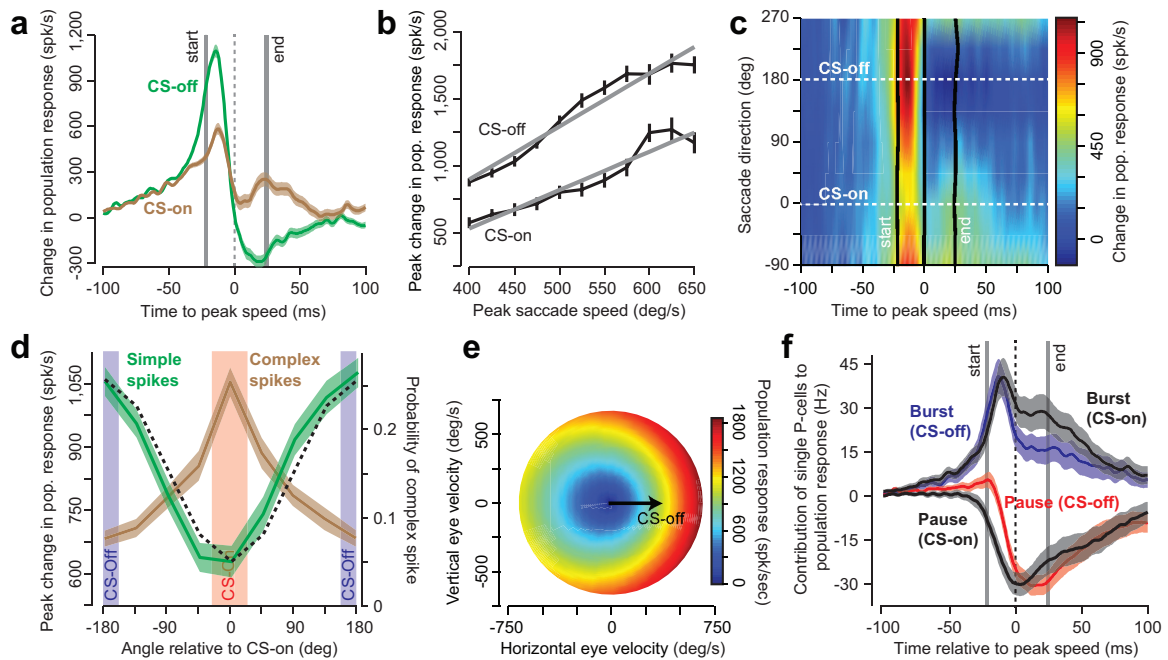


Figure 4.6: Population response of P-cells predicted saccade speed and direction in real-time as a gain-field. **A.** Population response for saccades in direction CS-on and CS-off. The population response is larger when the saccade is in the CS-off direction. **B.** Peak population response grew linearly with saccade speed, but had a higher gain for saccades in CS-off direction. **C.** Real-time population response as a function of saccade direction relative to CS-on. Data smoothed as in Fig. 4.5B. **D.** Peak population response (labeled as simple-spikes) as a function of saccade direction with respect to CS-on. Brown curve shows probability of observing a complex-spike as a function of the angle relative to each neuron's CS-on. Black curve indicates cosine fit of probability of CS. **E.** Gain-field encoding by a cluster of P-cells whose CS-off direction is to the right (Eq. 4.1). **F.** Contribution of single P-cells to the population response. A change in direction coincides with a shift in timing of the pause cells. Error bars are SEM.

saccades in directions CS-on and CS-off, and peak speeds of 400-650 deg/s, binned by 25 deg/s) was highly significant ( $R^2 = 0.80$ ,  $p < 10^{-5}$ ), resulting in the following parameter values:  $a = 0.37$  spk/deg,  $b = 1.88$  spk/deg,  $c = 2588$  spk/s with a time lead of  $\Delta = 19$  ms. The resulting gain-field encoding of eye motion is depicted in Fig. 4.6E.

### 4.3.3 Individual Purkinje cells do not predict saccade kinematics

In contrast to organizing the P-cells into clusters and computing a population response, we also attempted to relate activity of single cells to motion of the eye. With this approach we found that the mean activity of the individual pause cells did not vary with saccade speed (Fig. 4.1D,  $p > 0.20$ ), and mean activity of the individual pause cells did not vary with saccade direction (Fig. 4.7A,  $p > 0.4$ ). Mean activity of the burst cells increased with saccade speed (Fig. 4.1D,  $p < 10^{-10}$ ), but mean activity of the burst cells did not vary with direction (Fig. 4.7A,  $p > 0.4$ ). Therefore, consistent with previous reports (Ohtsuka and Noda, 1995; Kase et al., 1980), we could not detect an encoding of direction in the peak or mean activity of individual cells. If the mean or peak response of the individual cells did not vary significantly with direction, how did the population response produce an encoding of direction (i.e., the cosine tuning)? We found that a change in saccade direction produced a subtle shift in the timing of the discharge in the pause cells (Fig. 4.6F). The pause in these cells occurred earlier when the saccade was in the CS-on direction, and later when the saccade was in the CS-off direction. In contrast, the burst timing was not dependent on saccade direction. As a result, the population response exhibited a smaller gain when the saccade was in direction CS-on, and a larger gain when the saccade was in direction CS-off (Fig. 4.6B). We next asked how well the population of P-cells, clustered by their complex spike properties, predicted

the real-time speed of the eye. According to Eq. 4.1, if we consider the data across all directions, using the measured population response we should be able to predict the actual motion of the eye:

$$|\hat{\mathbf{x}}(t + \Delta)| = \frac{1}{b} [s(t) - c] \quad (4.2)$$

In the above equation,  $\hat{\mathbf{x}}(t + \Delta)$  is the predicted real-time speed of the eye. The parameters  $b$  and  $c$  were identical to those determined in Eq. 4.1. Using Eq. 4.2 we plotted the predicted real-time motion of the eye (Fig. 4.8A, saccades with peak speeds of 400, 525, and 650 deg/s). The predicted motion led the actual motion by +19ms, and was highly correlated with the actual motion (data for all saccades ranging from 400-650 deg/s, binned by 25 deg/s,  $R^2 = 0.74$ ,  $p < 10^{-5}$ ). Of particular importance was the fact that the predicted speed not only rose above baseline before saccade onset, it returned to below baseline before the end of the saccade. That is, the population response predicted in real-time the motion of the eye, and therefore could play a central role in controlling that motion, particularly in terminating the saccade.

We next asked whether individual neurons could predict the real-time speed of the eye. For each recorded cell, we fitted the parameters in Eq. 4.2 for all saccades from 400-650 deg/s in 25 deg/s bins, leading to three parameter estimates for each neuron ( $b$ ,  $c$ , and  $\Delta$ ). We found that after finding the best fit for each neuron, the predicted eye speed did not return to baseline until long after saccade termination (Fig. 4.8B). Across the neurons, the average delay was not significantly different than zero (two-sided t-test,  $p > 0.5$ ), indicating that when we used individual cells, rather than the population, the predicted motion did not lead or lag the real-time actual motion of the eye. The mean squared error (MSE) for individual neurons was 245% ( $\pm 10\%$ ) of the MSE achieved by the population

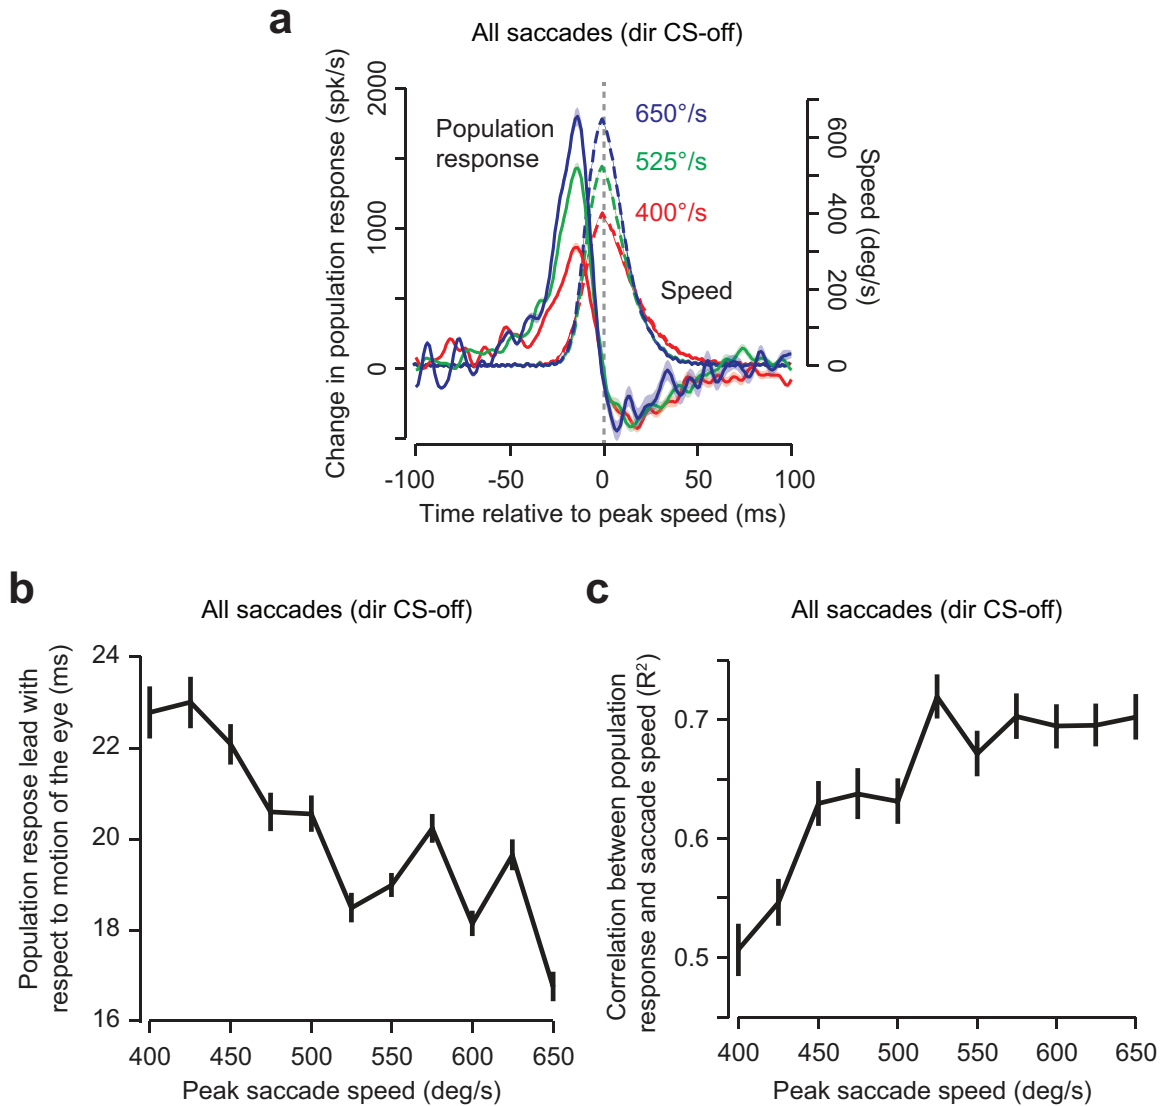


Figure 4.7: The simple-spike population response of P-cells, organized by their complex-spike properties (Fig. 4.5A), correlated with motion of the eye in real-time. **A.** Population response for saccades in direction CS-off for three different peak speeds. **B.** Temporal lead of the population response with respect to saccade speed as computed by finding the temporal shift that maximized the cross-correlation. **C.** Correlation between the population response and the temporally shifted eye speed trace (measured as  $R^2$ ).



response. That is, the real-time motion of the eye predicted by individual neurons was much poorer than the population estimate.

We next considered whether a population composed entirely of either burst or pause cells could predict the real-time motion of the eye. We fitted Eq. 4.2 to the mean activities of the burst and pause cells, and plotted the results in Fig. 4.8C and Fig. 4.8D. For the population composed entirely of burst cells, the predicted speed led the actual speed by +11ms, but critically did not return to baseline until long after saccade termination (Fig. 4.8C). The mean squared error for this bursting population was 146% of the MSE achieved by the population response composed of both bursting and pausing cells. A population consisting solely of pausing cells predicted that speed followed the actual speed by 9ms (that is, the best fit was a lag, not a lead). Similar to the exclusively bursting population, the predicted speed did not return to baseline until long after saccade termination (Fig. 4.8D). The mean squared error for pausing population was 162% of the MSE achieved by the population response which included all cells.

Therefore, our results suggest that the combined activity of burst and pause cells, but not the activity of either population individually (Fig. 4.8), aligned to CS-off, produced a population response that exhibited gain-field encoding: the magnitude of the population response increased linearly with speed, and was cosine tuned in direction, with a multiplicative interaction between speed and direction.

How did the activity of individual cells within the burst and pause clusters produce this directional encoding in the population response? The main contributors were the pause cells, which started their pause approximately 10ms earlier when the saccade was in the CS-on direction (Fig. 4.6F), a change which was independent of saccade speed (Fig. 4.10). This subtle shift in the timing of spikes produced an increase of the population response when saccade direction changed from CS-on to CS-off (Fig. 4.6A).

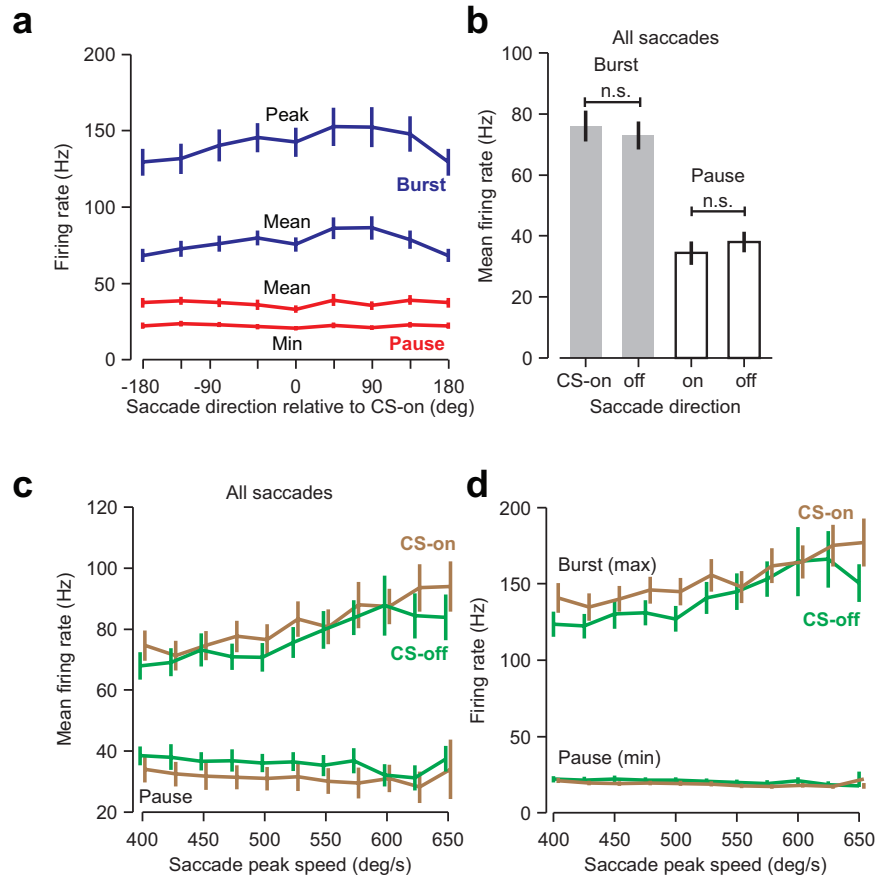


Figure 4.8: Mean and peak/trough firing rates of the burst and pause cells were poorly modulated by saccade direction. **A.** Maximum, minimum, and mean firing rates averaged across burst or pause cells with respect to saccade direction, relative to CS-on direction of each cell. **B.** Mean firing rates of the burst and pause cells, as measured across all saccades, were not significantly different for saccades in the CS-on vs. CS-off direction (burst  $p > 0.10$ , pause  $p > 0.05$ ). **C.** Mean firing rates of the burst and pause cells as a function of saccade speed, for saccades in the CS-on versus CS-off direction. Saccade speed modulated mean firing rates of the burst cells, but there were no significant interaction between saccade direction and speed ( $p > 0.6$ ), nor a significant effect of saccade direction ( $p > 0.7$ ). **D.** Peak (maximum) firing rates of the burst cells and the minimum firing rate of the pause cells as a function of saccade speed, for saccades in the CS-off and CS-on directions. We asked whether the maximum response of the burst cells or the minimum response of the pause cells was significantly modulated by direction. Separate RM-ANOVAs showed that for the burst cells, peak activity increased as a function of saccade peak speed ( $p < 0.001$ ), but this relationship was unaffected by saccade direction ( $p > 0.4$ ). For the pause cells, the response was not affected by saccade speed ( $p > 0.6$ ), and this relationship was not modulated by saccade direction ( $p > 0.4$ ). In summary, we found that saccade direction did not significantly alter the encoding of peak speed in either the mean or minimum/maximum activity of P-cells.

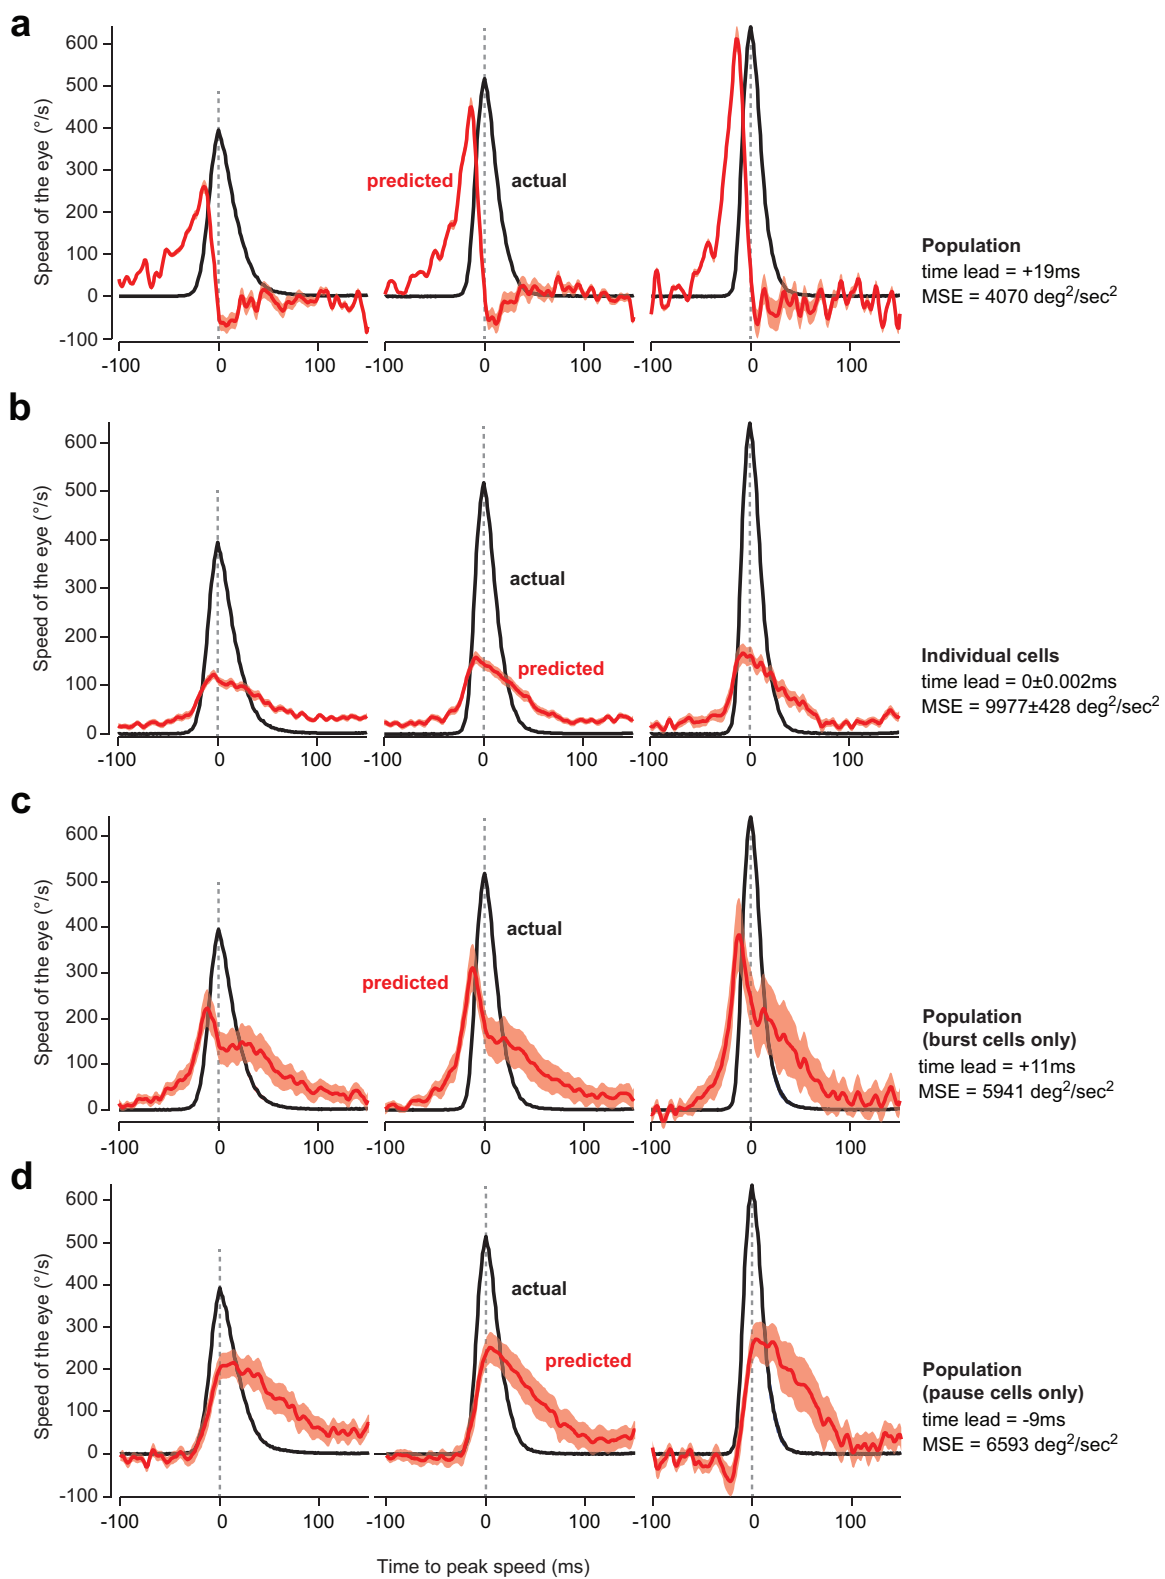


Figure 4.9: Population of P-cells, organized by their complex spike properties, predicted the real-time speed of the eye better than activity of individual cells. **A.** We employed Eq. 4.2 and used the measured population response of P-cells to predict the real-time speed of the eye. The plot shows the predicted speed for saccades of 400, 525, and 650 deg/sec. The predicted speed led the actual speed by 19ms. MSE is the mean squared error between the predicted and actual eye trajectory at the optimal value of  $\Delta$ . **B.** The result of fitting Eq. 4.2 to the response of individual neurons. **C.** The result of fitting Eq. 4.2 to the discharge of a population composed exclusively of burst cells. **D.** The result of fitting Eq. 4.2 to the discharge of a population composed exclusively of pause cells.

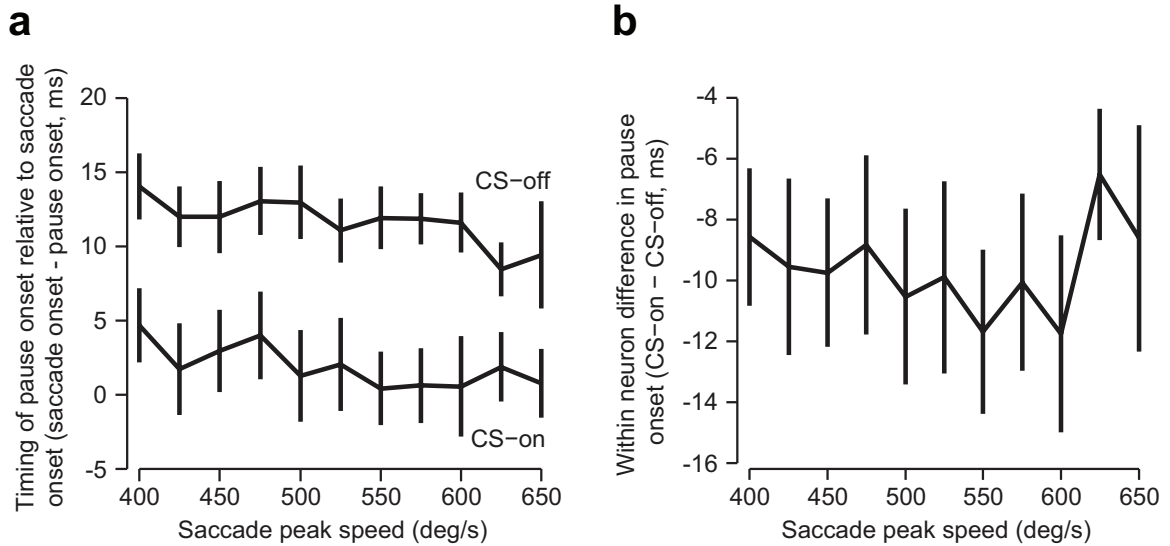


Figure 4.10: Change in saccade direction was associated with a change in the timing of the reduction of discharge in the pause cells, i.e., pause onset (see Fig. 4.6F). **A.** Timing of pause onset with respect to saccade onset for saccades of various speeds and directions. We computed the pause onset as the time when the neuron's response reached 20% of its minimum response. Positive numbers indicate that the pause onset occurred before saccade onset. **B.** Within neuron measure of pause onset for saccade in direction CS-on, minus onset from saccades in direction CS-off. Negative numbers indicate that the pause onset occurred earlier for saccades in the CS-on direction.

### 4.3.4 Anatomical distribution of Purkinje cells supports population hypothesis

We found that the anatomical distribution of P-cells, as labeled by their CS-off direction, was not random, but lateralized (Soetedjo et al., 2008) (Fig. 4.11), confirming previous anatomical studies suggesting that olivary projections are contralateral (Yamada and Noda, 1987; Kralj-Hans et al., 2006). P-cells with rightward CS-off were more likely to be on the right side of the cerebellum (t-test,  $p < 10^{-4}$ ). This indicates that saccades made in the same direction as CS-off were typically ipsiversive whereas saccades congruent with CS-on were contraversive. In contrast, pause and burst cells were uniformly distributed across the cerebellum ( $p > 0.4$ ).

Our results rely critically on our hypothesis that P-cells organize into clusters with roughly equal number of pause and burst cells, all with a common complex-spike tuning preference (Fig. 4.5A). If, contrary to our hypothesis, pause and burst cells organized into separate clusters, the population response would not predict the real-time motion of the eye (Fig. 4.1E). Similarly, if each cluster was not composed of roughly equal number of pause and burst cells, the population response could not predict the real-time speed of the eye (Fig. 4.12, Section 4.4.1). The fact that burst and pause cells were distributed uniformly across the recording locations, and not lateralized as we found with the CS tuning properties, suggests that a cluster is composed of both burst and pause. Finally, if we ignored the CS properties of the P-cells, and made the typical assumption that simple-spikes were sufficient to uncover the coordinate system of encoding motion, then the gain-field representation of speed and direction would disappear (Fig. 4.13, Section 4.4.2).

In summary, organizing the P-cell into clusters where all the cells shared a common complex-spike property resulted in simple spikes that encoded speed and

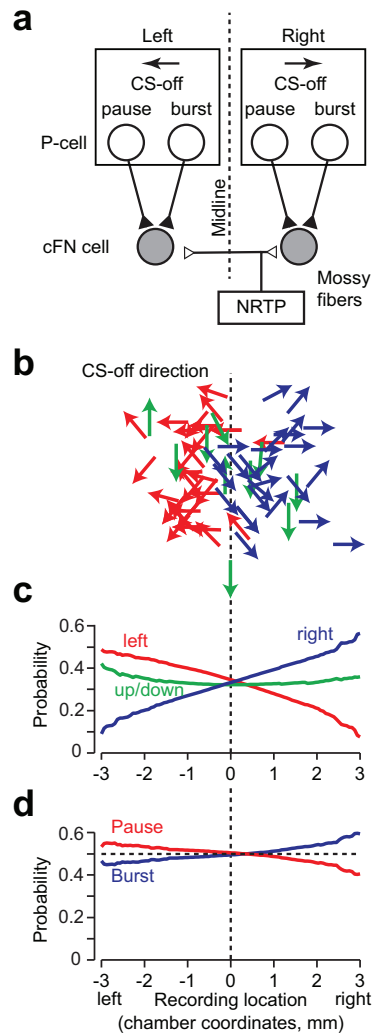


Figure 4.11: CS-dependent organization of the P-cells. **A.** Hypothesized anatomical organization of the oculomotor vermis (OMV). Bursting and pausing P-cells are organized into clusters, with the cells in each cluster sharing a common CS direction. Neurons on the right side of the OMV project to right cFN neurons and have CS-off directions to the right. **B.** Distribution of the CS-off directions from recorded neurons in chamber coordinates. Vertical dotted line shows the line which best separates rightwards CS-off direction neurons (blue) from leftwards CS-off direction neurons (red). **C.** Probability of having a rightwards (blue), up/down (green) or leftwards CS-off direction as a function of chamber coordinates. P-cells with CS-off to the left were more probable on the left side of the cerebellum. P-cells with CS-off to the right were more probable on the right-side of the cerebellum. **D.** Pause (red) and burst (blue) P-cells were equally likely at all recorded locations.

direction in real-time via a gain-field.

## 4.4 Discussion

In summary, the population of P-cells, clustered by their CS-on direction, produced simple spikes that predicted the real-time motion of the eye. In this encoding, the speed and direction of the motion of the eye were multiplicatively encoded via a gain field, with the gain being largest for when the eye moved toward CS-off, and lowest for when the eye moved toward CS-on. Neither individual neurons, nor populations that exclusively included burst or pause cells, predicted the real-time speed of the eye with the accuracy of a population which combined these two cell types.

Our results have broad implications regarding function and organization of the cerebellum. During saccades, the transformation of efference copy (via mossy fibers) into prediction of kinematic state, a concept termed a forward model, does not occur in the individual P-cells, but via combined activity of burst and pause P-cells that produce inhibition at the deep cerebellar nucleus. It is this inhibition produced by the combined activity of both groups of P-cells that predicts the motion of the eye during a saccade via a gain field, multiplicatively encoding speed and direction of movement. Therefore, our results demonstrate that the forward model, a theoretical concept central to conductance views of motor control, and often hypothesized to depend on the cerebellum, (Xu-Wilson et al., 2009; Izawa et al., 2012; Miall and Wolpert, 1996), is represented during a saccade not in the activity of individual P-cells, but in the population activity that converges onto the cells in the deep cerebellar nuclei.

This encoding of movements is present only if there is a specific anatomical organization in the cerebellum: the projections from P-cells to nucleus neurons are not random, but likely organized by the complex-spike properties of the P-cells.

That is, a nucleus cell receives projections from P-cells that share the same complex-spike field.

#### 4.4.1 The importance of bursting and pausing cells in the population

The proportion of bursting and pausing cells in our data set were approximately equal ( $n = 39$  bursting,  $n = 33$  pausing), resulting in 54% burst cells. However, two previous reports suggested that there may be a higher concentration of bursting cells in the OMV (Ohtsuka and Noda, 1995; Helmchen and Büttner, 1995). Both reports found that bursting neurons comprised approximately 70% of all cells in OMV. Therefore we asked if changes in the ratio of bursting to pausing P-cells would significantly alter our primary result that the population activity of OMV predicted eye velocity.

We performed the same analysis as in Fig. 4.5D, in which we estimated the population response, but rather than choosing 50 neurons randomly from our pool of 72 neurons, we always chose 35 bursting cells and 15 pausing cells (corresponding to 70% bursting in the population). We found that our primary result was robust to this modest change in ratio of bursting to pausing cells. That is, when 70% of the population was chosen to be bursting neurons, the peak response scaled linearly with the peak velocity of the saccade ( $R^2 = 0.94$ ,  $p < 10^{-6}$ ). Importantly, the timing of the population response waveform remained tightly coupled with saccade kinematics in that the response returned to near zero at saccade termination (Fig. 4.12).

However, this tight coupling of the population response and saccade speed disappeared if the population of P-cells was composed of a super majority of bursting or pausing cells. To illustrate this, we estimated what the population response by simulating populations that had differing proportions of burst and pause cells. The resulting population responses under the assumption that the






Figure 4.12: The population response was sensitive to the fraction of pause and burst cells that composed a cluster of P-cells. In our data set, 54% of the population was composed of burst cells. We computed the population response under the assumption that the membership of a cluster was 54% burst cells. Here, we tested how sensitive the population response was to this membership ratio. The vertical lines indicate saccade onset and offset for all saccades pooled across direction and speed. As the percentage of burst cells in the cluster becomes larger than 70%, or smaller than 50%, the population response no longer returns to baseline at saccade offset.

population is composed entirely of bursting cells (100%), or 90% bursting, or 50% bursting, or 0% bursting (i.e., entirely pausing) are plotted in Fig. 4.12. We found that as the membership within the population became highly skewed toward bursting or pausing, the population response no longer returned to zero at saccade termination. Therefore, the ability of the population to predict in real-time the speed of the eye was present only if the composition of the P-cells that converged onto a single cFN neuron (i.e., a P-cell cluster) included roughly equal number of pausing and bursting cells.

#### 4.4.2 The importance of organizing the P-cells via their complex-spike properties

How critical is our assumption that each cluster is composed of P-cells that all share a similar complex-spike field? The traditional approach to measure population activity is to organize the cells based on their simple-spike activity, where for each cell the “preferred direction” is computed as the direction of action for which the cell shows the greatest simple-spike firing rate. This kind of analysis effectively assumes that the P-cells organize into clusters in which the cells share the same preferred direction of saccade, where preferred direction is measured via the rate of simple-spikes (Fig. 4.13A). Let us consider the consequences of this hypothesis.

We estimated the preferred direction of P-cells in two ways. First, we defined the preferred direction of each neuron as the direction which elicited the largest mean simple-spike firing rate during the saccade (max response). We then computed the population response for saccades of various speed (Fig. 4.13B). We found that the population response was no longer modulated by saccade speed ( $p > 0.3$ ).

Next, we calculated the preferred direction as the saccade direction in which there was the largest change in the magnitude of the simple-spike firing rate (max modulation). For bursting cells, this was the direction which elicited the largest mean simple-spike firing rate during the saccade. For pausing cells, this was the direction which featured the lowest simple-spike firing rate response (i.e., the direction of the largest pause). We then computed the population response for saccades of various speeds in this preferred direction (Fig. 4.13C). Again we found that the population response was not modulated by the speed of the saccade ( $p > 0.7$ ). Therefore, when we organized clusters of P-cells based on their simple-spike responses rather than their complex-spike responses, the population response no longer encoded the speed of the saccade. An important limitation is

that our cells were recorded one at a time, yet in computing the population response we analyzed the data as if the cells were recorded simultaneously. Future work is needed to confirm our observations with simultaneous multi-cellular recordings.

## 4.5 Conclusion

Together, our results suggest three principles of cerebellar function during control of saccadic eye movements. First, the cerebellum predicts real-time motion not in the time-course of individual P-cell simple-spikes, nor in the individual activities of the bursting or pausing populations, but in the combined activities of these two populations via the simple-spikes that converge onto cells in the deep cerebellar nucleus. A similar population coding has been suggested during smooth pursuit (Krauzlis, 2000). Second, this population input to each nucleus neuron encodes direction and speed via a gain-field. Because a similar encoding has been shown in the posterior parietal cortex during saccades (Andersen et al., 1985), as well as in the motor cortex during reaching (Paninski et al., 2004), our observation in the cerebellum suggests a common principle of encoding in disparate regions of the motor system. Finally, the gain-field encoding was present if we assumed a specific anatomical organization: a cluster of P-cells that projected onto a single nucleus neuron was composed of approximately equal numbers of bursting and pausing P-cells, all sharing a common complex-spike property. Because the complex-spikes of a P-cell are due to input from the inferior olive, the gain-field encoding predicts that the oculomotor vermis is organized into clusters of P-cells that share similar climbing fiber projections from the inferior olive. This in turn suggests that motor memories are anatomically clustered in the cerebellum by the errors that were experienced during movements (Herzfeld et al., 2014b).

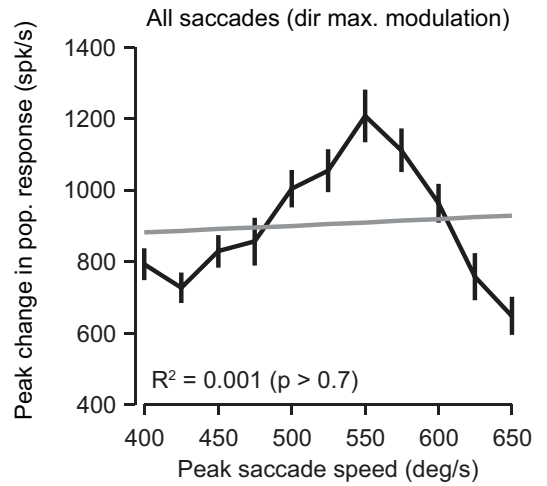
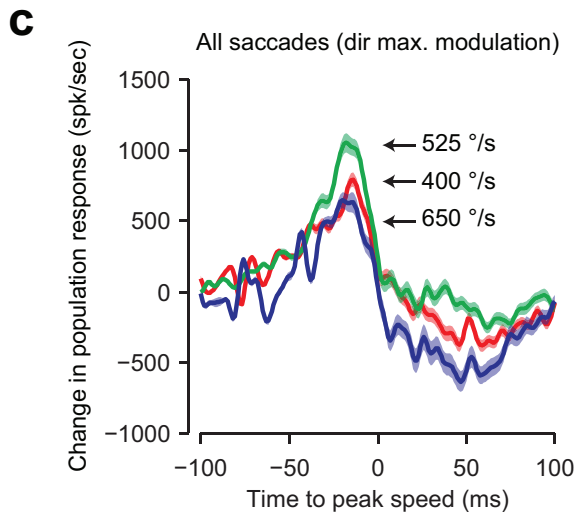

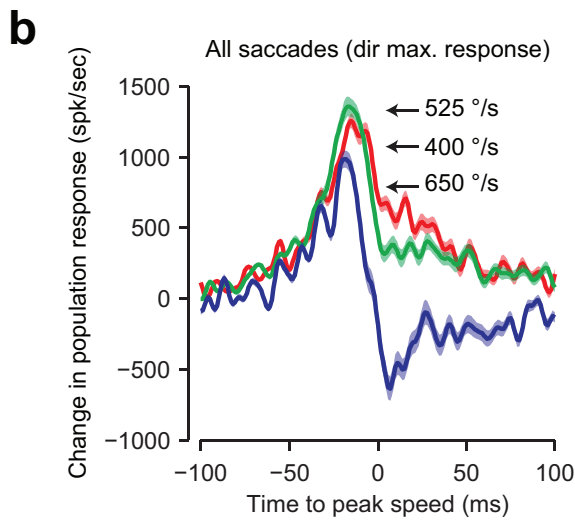



Figure 4.13: Gain-field encoding of saccade kinematics in the population response of the P-cells disappeared if the P-cells were organized by their simple-spike activity. **A.** In this analysis we assumed that a collection of 50 P-cells projected onto a single cFN neuron, with the property that all the P-cells shared a similar simple-spike preferred direction. Therefore, the cluster was organized based on the simple-spike properties of the P-cells, not their complex-spike properties. **B.** The population response for saccades made in the direction for which each P-cell showed the largest mean firing rate (simple-spikes), for various saccade peak speeds. The peak population response was not modulated with saccade speed. Error bars are boot-strap estimated SEM. **C.** The population response for saccades made in the direction of maximal modulation. For burst cells, this was the direction for which the P-cell showed the largest mean firing rate, whereas for pause cells, this was the direction associated with the minimum activity (largest pause). The peak population response was not modulated with saccade speed when clusters were organized based on the direction of maximal simple-spike modulation.

## 5 CONCLUSIONS & FUTURE DIRECTIONS

Motor learning seeks to answer a fundamental question: how do we learn from error? When a subject experiences an error during a motor task, they adapt their next movement to compensate. Even when perturbations are random, trial-to-trial changes in the motor output are still present (Marko et al., 2012; Izawa et al., 2008; Donchin et al., 2003), demonstrating that on every movement the brain learns from error.

Previous theoretical models of motor learning suggest that error-sensitivity, which corresponds to the fraction of the error that the subject learns on the next trial, is fixed both as a function of trial and as a function of error. That is, regardless of the magnitude of the error that a subject experiences, that subject will learn the same fractional amount of that error on the next trial. However, previous results as well as our own experimental evidence suggests that this is not the case (Marko et al., 2012; Fine and Thoroughman, 2006; Wei and Körding, 2009).

In Chapter 2, we suggest that the brain stores a previously unknown form of motor memory: a memory of errors. Imagine that the brain makes a movement in trial  $n$ , resulting in an error,  $e^{(n)}$ . The subject then updates their next motor command, also resulting in an error,  $e^{(n+1)}$ . Given these two error signals, the brain can evaluate whether the amount of adaptation (i.e., the level of error-sensitivity) is appropriate for the error experienced in trial  $n$ . For instance, imagine that the movement  $u^{(n)}$  resulted in an undershoot of the target. If an updated movement,  $u^{(n+1)}$  also results in an undershoot, then the error-sensitivity associated with  $e^{(n)}$  needs to be larger. However, if  $u^{(n+1)}$  resulted in an overshoot, then error-sensitivity is too large and needs to be smaller.

We show that human subjects use the history of prior errors to modulate how much they are willing to learn from error. The simple computational rule that governs manipulation of this error-sensitivity explains a host of previous motor control results. In particular, our error-sensitivity hypothesis describes why subjects learn faster following initial learning, savings. The model suggests that savings occurs not because subjects recall the previous motor commands they produced during initial learning, but rather recall the errors that they experienced, resulting in higher error-sensitivity upon re-exposure. In addition, this model explains how subjects are able to learn a perturbation faster than naïve even if they have never experienced the motor commands required to compensate for the perturbation. That is, the phenomenon of meta-learning can also be explained by a memory of errors.

What region of the brain is responsible for maintaining this memory of errors? In Chapter 3 we used non-invasive brain stimulation to manipulate the process of learning a novel perturbation. We showed that positive (anodal) stimulation of the cerebellum resulted in significantly faster motor learning of a force-field perturbation. That is, when we increased the excitability of neurons in the cerebellum, subjects tended to increase their error-sensitivity and learn the perturbation faster. Conversely, when we applied negative (cathodal) stimulation of the cerebellum, error-sensitivity decreased. Since cathodal stimulation has been associated with reduced neuronal excitability, our results suggest that decreasing the activity of neurons in the cerebellum results in lower error-sensitivity and therefore slower motor learning.

We contrast these results with anodal and cathodal tDCS stimulation of the primary motor cortex. Application of either anodal or cathodal stimulation did not substantially alter the learning rate of subjects. Taken together, our results suggest that the cerebellum is crucially responsible for determining error-sensitivity, and

that the activity of cells within the cerebellum likely corresponds directly to how much a subject is willing to learn from an error.

Our results, along with numerous previous results, provides a critical role for the cerebellum in motor adaptation. Why has the neural basis of cerebellar-dependent adaptation remained elusive? Much of the problem centers on the role of Purkinje cells in the control of movement. Lesions and inactivations of the cerebellum affect not only motor learning, but also accurate execution of movements. These results lead to the hypothesis that the primary neurons of the cerebellum should be related to movement parameters. However, decades of recordings in the cerebellum do not provide compelling evidence that kinematics are encoded in Purkinje cell firing rates. How can we begin to address the question of motor adaptation, which entails a change in the encoding of movement kinematics, without first identifying how Purkinje cells are affecting movements?

In Chapter 4, we provided a resolution to this problem: encoding of kinematics is not present in the response of individual Purkinje cells, rather it is the combined activity of populations of Purkinje cells via their projections to the deep cerebellar nuclei that encodes movement parameters. Further, we suggest that the subpopulation (micro-cluster) of neurons that project to an individual neuron in the deep cerebellar nuclei are not selected at random, but rather share a common preference for error via climbing fiber projections from the inferior olive. When the cerebellum is organized in this fashion, a beautiful encoding of kinematics emerges. Both the speed and direction of movement are encoded via a gain-field, which appears to be a common type of encoding across multiple movement centers in the primate brain. Using these results, we can begin to determine the neural substrates of adaptation in the cerebellum.



## 5.1 Future directions

The cerebellum is the key neural substrate for adaptation (Smith and Shadmehr, 2005; Criscimagna-Hemminger et al., 2010; Donchin et al., 2012) suggesting that the responses of the major output neurons of the cerebellum, Purkinje cells, should reflect the changes in movement that occur following an error. Classical theories of cerebellar learning are based on three fundamental principles (Medina and Lisberger, 2008). First, when a movement occurs and the result is an error, that error is transmitted to the cerebellum via climbing fiber projections from neurons in the inferior olive (Gilbert and Thach, 1977), resulting in complex spikes (CSs) in P-cells. Second, the presence of a CS results in plasticity in the cerebellum, which affects the generation of P-cell simple spikes (Ito, 2001). Finally, changes in the responses of these cells affects neurons in the primary output structure of the cerebellum, the deep cerebellar nuclei, resulting in a modification of motor behavior. That is, there is a direct link between P-cell responses and movement. This relationship is then modified following an error.

This hypothesis of motor learning, sometimes termed the Marr-Albus-Ito hypothesis, has guided cerebellar research for the last half century. While this hypothesis has many testable elements, the role of the cerebellum in motor adaptation has remained unclear. For instance, this classical view of the cerebellar learning suggestions that the role of complex spikes is to convey a prediction error signal (Marr, 1969; Albus et al., 1971). Recent neurophysiological studies, however, do not lend credence to this limited view of CS as merely an encoder of prediction error. For instance, in a recent study, monkeys adapted their saccades (a rapid movement of the eyes) to a constant visual perturbation (Catz et al., 2005). As the monkey adapted, performance improved and the magnitude of the errors declined. Yet, the probability of CSs in some P-cells increased, opposite of what was expected

by a strict interpretation of the error-encoding hypothesis. Another study noted that as saccade adaptation proceeded, errors decreased, yet there were no corresponding changes in the average probability of CSs (Soetedjo and Fuchs, 2006). Therefore, the role of CSs in adaptation remains more complicated than would be predicted by the Marr-Albus-Ito hypothesis.

In addition, the classical theory of cerebellar learning implies a link between the activity of P-cells in the cerebellum and behavior. That is, the activity of P-cells directly affects the subject's motor output and the relationship between P-cell responses and behavior changes during learning. However, evidence linking P-cell response and behavior has been lacking. For instance, accurate execution of saccades depends critically on the oculomotor vermis (OMV) (Barash et al., 1999) as well as its target in the caudal fastigial nucleus (cFN) (Goffart et al., 2004; Kurzan et al., 1993; Pélisson et al., 2003). This leads to the prediction that P-cell responses should be well correlated with the kinematic measures of the eye movement. However, this encoding has historically remained unclear: P-cells show little consistent modulation with respect to kinematic parameters such as saccade amplitude (Ohtsuka and Noda, 1995; Helmchen and Büttner, 1995), speed (Helmchen and Büttner, 1995; Kase et al., 1980; Thier et al., 2000), or direction (Ohtsuka and Noda, 1995). Without a clear understanding of the encoding of movement, the problem of learning from error, which entails a change in this encoding, has been even more difficult to solve. However, our results in Chapter 4 now provide a framework that we can use to understand cerebellar adaptation.

### **5.1.1 Cerebellar organization provides clues for determining the neurophysiological basis of adaptation**

When we experience a movement error, this error is transmitted to the cerebellum from the inferior olive via climbing fibers, resulting in complex spikes in Purkinje

cells (Gilbert and Thach, 1977). A population of approximately 50 P-cells project to a single neuron in the caudal fastial nucleus (Person and Raman, 2012). In Chapter 4, we show evidence for a new model of cerebellar organization: P-cells that project to a single nucleus neuron are not selected at random, but rather share a common preference for error (see Fig. 4.5A). In this model, the encoding of movement is not in the response of an individual P-cell, but rather in the combined activity in this micro-cluster of 50 P-cells. Using this hypothesized anatomy, we show that decoding of movement kinematics from micro-clusters of P-cells is now possible (Herzfeld et al., 2015).

We suggest that adaptation may also be understood when P-cells are organized using this hypothesized anatomy. When a movement error occurs, complex spikes would occur in the set of micro-clusters which prefer that particular error. Presence of a CS results in plasticity and affects the output of the micro-cluster, and in turn behavior. That is, while changes in the simple-spike responses of individual P-cells may be heterogeneous, we would expect to observe consistent changes at the micro-cluster level. In summary, we may only be able to understand execution of movements and adaptation when P-cell responses are organized by their preference for error.

### **Cerebellar organization protects previously learned memories**

Using behavioral psychophysics, in Chapter 2 we demonstrate that the brain stores a previously unknown form of memory: a memory of errors. That is, the history of the errors a subject experiences affects how much they are willing to learn from an error in the future. This result provides an interesting interpretation of the phenomenon of savings. Savings occurs because the subject experiences an set of errors which are then re-experienced upon the second exposure, resulting in faster learning than naïve. Importantly, washout of the perturbation following initial

learning preserves this memory of errors. How can the cerebellum accomplish this task? Using the hypothesized organization of the cerebellum described in Chapter 4, we can begin to understand these results.

We have suggested that the cerebellum is organized into micro-clusters wherein P-cells share a common preference for error. When a subject experiences a perturbation (A), the errors they encounter tend to be in the same direction, leading to complex spikes in a micro-cluster whose preference for error aligns with those errors. During washout of (A), the errors experienced by the subject are opposite the errors that were experienced during learning of (A). That is, the errors during washout change direction. This change in direction will result in CSs in a different micro-cluster of P-cells, and the original micro-cluster of P-cells will no longer receive CSs. Therefore, the memory of perturbation (A) is functionally distinct from the motor memory of washout. When the subject is re-exposed to perturbation (A), the original micro-cluster is reengaged, leading to an improvement in performance, or savings. Our hypothesized organization of the cerebellum shows how the anatomy results in protection of previously learned motor memories, providing a novel view of savings.

### **Linking complex spikes to adaptation**

Previous studies have failed to provide a consistent link between CSs and error. For instance, as monkeys adapt to a saccadic perturbation, the magnitude of the error decreases. However, the probability of complex spikes as reported by previous studies does not follow this decrease in error (Catz et al., 2005). How can we rectify this result with our previous observations that complex spikes signal foveal error (e.g., Fig. 4.3)?

Imagine that P-cells are organized into micro-clusters based on their preference for foveal error, as demonstrated by the results in Chapter 4. P-cells are

active for saccades in all directions (see Fig. 4.6). However, micro-cluster responses are larger when the saccade is made in direction CS-off than CS-on. If we imagine that the cerebellum is divided into many micro-clusters, each with a different error preference, when the subject makes a saccade at least two micro-clusters are active: one whose preference for error is close to the saccade direction and one whose preferred error direction is opposite the saccade direction. How do these two populations interact? We found anatomical evidence for a complex spike dependent organization of P-cells (Fig. 4.11): an error vector to the right produces complex spikes for P-cells on the left of the vermis, but an error vector to the left produces complex spikes on the right of the vermis. Cells on the right of the vermis project to cells in the right cFN whereas cells on the left project to left cFN. Activity in the cFN is related to acceleration of the contralateral eye and deceleration of the ipsilateral eye. Therefore, it is likely that the motion of the eye results from a functional subtraction of activities of P-cell micro-clusters in their preferred and anti-preferred directions.

When the subject makes a saccade to the left, the resulting saccade kinematics are determined by the difference in the activities of the two micro-clusters on the right (R) and left (L) sides of the cerebellum (e.g.,  $R - L$ ). We can understand gain-down saccade adaptation as a decrease in the population activity of (R), an increase in (L), or both. Conversely, inducing an increase in the gain of the saccade could result from an increase in the activity of (R), a decrease in the response (L), or a combination of the two.

Further imagine that the monkey is performing a saccadic task in which the target is consistently stepped inwards during leftwards saccades. When the monkey finishes their primary saccade, there is a foveal error pointing towards the right, and the magnitude of this error decreases over the course of adaptation. We hypothesize that, at the start of adaptation, when the error is large, the probability of CSs in

the (R) micro-cluster will also be large (since the error is in the preferred CS direction of the (R) micro-cluster). Conversely, the probability of CSs in the (L) population will be substantially less than the baseline probability of CSs since the error here resides in the (L) micro-cluster's anti-preferred direction. As the monkey adapts, the probability of CSs in the (R) micro-cluster will decrease whereas the probability of CSs in the (L) micro-cluster will increase. If we were to average the probability of CSs across multiple P-cells without first characterizing their preferred error direction, then the probability of CSs would not precisely track the magnitude of the saccadic error. However, by organizing P-cells by their preference for error, a differential encoding of error emerges.

A large increase in the probability of CSs in population (R) would result in a decrease in the simple-spike activity of this population. Conversely, the lack of a baseline level of CSs in the (L) micro-cluster would result in an overall increase in the simple-spike responses. Therefore, it is likely that adaptation is a combination of changes in simple-spike activity whose direction is dependent on the preferred error direction of each P-cells.

In this way, we can rectify the results of previous studies of saccadic adaptation with the Marr-Albus-Ito hypothesis. Complex spikes *do* signal error, however, this error is defined along a preferred direction. That is, the presence of a complex spike represents the presence of an error along the cell's preferred direction whereas lack of complex spike (below baseline) signals an error in the anti-preferred direction. Future studies can test these intriguing hypotheses.

## BIBLIOGRAPHY

- Mohammad Ali Ahmadi-Pajouh, Farzad Towhidkhah, and Reza Shadmehr. Preparing to Reach: Selecting an Adaptive Long-Latency Feedback Controller. *The Journal of Neuroscience*, 32(28):9537–9545, November 2012. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.4275-11.2012.
- Hirotougu Akaike. Information Theory and an Extension of the Maximum Likelihood Principle. In Emanuel Parzen, Kunio Tanabe, and Genshiro Kitagawa, editors, *Selected Papers of Hirotugu Akaike*, Springer Series in Statistics, pages 199–213. Springer New York, 1998. ISBN 978-1-4612-7248-9 978-1-4612-1694-0. doi: 10.1007/978-1-4612-1694-0\_15.
- James S. Albus, Datu Techltiques Branch, Communicated Donald, and H. Perkel. *A Theory of Cerebellar Function*. 1971.
- R. A. Andersen, G. K. Essick, and R. M. Siegel. Encoding of spatial location by posterior parietal neurons. *Science (New York, N.Y.)*, 230(4724):456–458, October 1985. ISSN 0036-8075.
- Fritzie Arce, Itai Novick, Yael Mandelblat-Cerf, and Eilon Vaadia. Neuronal Correlates of Memory Formation in Motor Cortex after Adaptation to Force Field. *The Journal of Neuroscience*, 30(27):9189–9198, July 2010. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.1603-10.2010.
- Pierre Baraduc, Nicolas Lang, John C. Rothwell, and Daniel M. Wolpert. Consolidation of Dynamic Motor Learning Is Not Disrupted by rTMS of Primary Motor Cortex. *Current Biology*, 14(3):252–256, February 2004. ISSN 0960-9822. doi: 10.1016/j.cub.2004.01.033.
- Shabtai Barash, Armenuhi Melikyan, Alexey Sivakov, Mingsha Zhang, Mitchell Glickstein, and Peter Thier. Saccadic Dysmetria and Adaptation after Lesions of the Cerebellar Cortex. *The Journal of Neuroscience*, 19(24):10931–10939, December 1999. ISSN 0270-6474, 1529-2401.
- Hannah Block and Pablo Celnik. Stimulating the Cerebellum Affects Visuomotor Adaptation but not Intermanual Transfer of Learning. *The Cerebellum*, 12(6):

781–793, April 2013. ISSN 1473-4222, 1473-4230. doi:  
10.1007/s12311-013-0486-7.

Daniel A. Braun, Ad Aertsen, Daniel M. Wolpert, and Carsten Mehring. Motor Task Variation Induces Structural Learning. *Current Biology*, 19(4):352–357, February 2009. ISSN 0960-9822. doi: 10.1016/j.cub.2009.01.036.

Jessica X. Brooks and Kathleen E. Cullen. The Primate Cerebellum Selectively Encodes Unexpected Self-Motion. *Current Biology*, 23(11):947–955, June 2013. ISSN 0960-9822. doi: 10.1016/j.cub.2013.04.029.

Nicolas Catz, Peter W. Dicke, and Peter Thier. Cerebellar Complex Spike Firing Is Suitable to Induce as Well as to Stabilize Motor Learning. *Current Biology*, 15(24):2179–2189, December 2005. ISSN 0960-9822. doi: 10.1016/j.cub.2005.11.037.

Nicolas Catz, Peter W. Dicke, and Peter Thier. Cerebellar-dependent motor learning is based on pruning a Purkinje cell population response. *Proceedings of the National Academy of Sciences*, 105(20):7309–7314, May 2008. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.0706032105.

Haiyin Chen, Sherwin E. Hua, Maurice A. Smith, Frederick A. Lenz, and Reza Shadmehr. Effects of Human Cerebellar Thalamus Disruption on Adaptive Control of Reaching. *Cerebral Cortex*, 16(10):1462–1473, January 2006. ISSN 1047-3211, 1460-2199. doi: 10.1093/cercor/bhj087.

Sen Cheng and Philip N. Sabes. Modeling Sensorimotor Learning with Linear Dynamical Systems. *Neural Computation*, 18(4):760–793, March 2006. ISSN 0899-7667. doi: 10.1162/neco.2006.18.4.760.

Nicholas Cothros, Stefan Köhler, Erin W. Dickie, Seyed M. Mirsattari, and Paul L. Gribble. Proactive Interference as a Result of Persisting Neural Representations of Previously Learned Motor Skills in Primary Motor Cortex. *Journal of Cognitive Neuroscience*, 18(12):2167–2176, November 2006. ISSN 0898-929X. doi: 10.1162/jocn.2006.18.12.2167.

Sarah E. Criscimagna-Hemminger and Reza Shadmehr. Consolidation Patterns of Human Motor Memory. *The Journal of Neuroscience*, 28(39):9610–9618,



September 2008. ISSN 0270-6474, 1529-2401. doi:  
10.1523/JNEUROSCI.3071-08.2008.

Sarah E. Criscimagna-Hemminger, Amy J. Bastian, and Reza Shadmehr. Size of Error Affects Cerebellar Contributions to Motor Learning. *Journal of Neurophysiology*, 103(4):2275–2284, April 2010. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00822.2009.

Suryadeep Dash, Peter W. Dicke, and Peter Thier. A vermal Purkinje cell simple spike population response encodes the changes in eye movement kinematics due to smooth pursuit adaptation. *Frontiers in Systems Neuroscience*, 7:3, 2013. doi: 10.3389/fnsys.2013.00003.

Jean-Jacques Orban de Xivry, Sarah E. Criscimagna-Hemminger, and Reza Shadmehr. Contributions of the Motor Cortex to Adaptive Control of Reaching Depend on the Perturbation Schedule. *Cerebral Cortex*, 21(7):1475–1484, January 2011a. ISSN 1047-3211, 1460-2199. doi: 10.1093/cercor/bhq192.

Jean-Jacques Orban de Xivry, Mollie K. Marko, Sarah E. Pekny, Damien Pastor, Jun Izawa, Pablo Celnik, and Reza Shadmehr. Stimulation of the Human Motor Cortex Alters Generalization Patterns of Motor Learning. *The Journal of Neuroscience*, 31(19):7102–7110, November 2011b. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.0273-11.2011.

Chris I. De Zeeuw, Freek E. Hoebeek, Laurens W. J. Bosman, Martijn Schonewille, Laurens Witter, and Sebastiaan K. Koekkoek. Spatiotemporal firing patterns in the cerebellum. *Nature Reviews Neuroscience*, 12(6):327–344, June 2011. ISSN 1471-003X. doi: 10.1038/nrn3011.

Jörn Diedrichsen, Yasmin Hashambhoy, Tushar Rane, and Reza Shadmehr. Neural Correlates of Reach Errors. *The Journal of Neuroscience*, 25(43):9919–9931, October 2005. ISSN 0270-6474, 1529-2401. doi:  
10.1523/JNEUROSCI.1874-05.2005.

Jörn Diedrichsen, Olivier White, Darren Newman, and Níall Lally. Use-Dependent and Error-Based Learning of Motor Behaviors. *The Journal of Neuroscience*, 30(15):5159–5166, April 2010. ISSN 0270-6474, 1529-2401. doi:  
10.1523/JNEUROSCI.5406-09.2010.

- Opher Donchin, Joseph T. Francis, and Reza Shadmehr. Quantifying Generalization from Trial-by-Trial Behavior of Adaptive Systems that Learn with Basis Functions: Theory and Experiments in Human Motor Control. *The Journal of Neuroscience*, 23(27):9032–9045, August 2003. ISSN 0270-6474, 1529-2401.
- Opher Donchin, Kasja Rabe, Jörn Diedrichsen, Níall Lally, Beate Schoch, Elke Ruth Gizewski, and Dagmar Timmann. Cerebellar regions involved in adaptation to force field and visuomotor perturbation. *Journal of Neurophysiology*, 107(1):134–147, January 2012. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00007.2011.
- E. V. Evarts and J. Tanji. Reflex and intended responses in motor cortex pyramidal tract neurons of monkey. *Journal of Neurophysiology*, 39(5):1069–1080, September 1976. ISSN 0022-3077, 1522-1598.
- R. Ferrucci, S. Marceglia, M. Vergari, F. Cogiamanian, S. Mrakic-Sposta, F. Mameli, S. Zago, S. Barbieri, and A. Priori. Cerebellar Transcranial Direct Current Stimulation Impairs the Practice-dependent Proficiency Increase in Working Memory. *Journal of Cognitive Neuroscience*, 20(9):1687–1697, March 2008. ISSN 0898-929X. doi: 10.1162/jocn.2008.20112.
- Michael S. Fine and Kurt A. Thoroughman. Motor Adaptation to Single Force Pulses: Sensitive to Direction but Insensitive to Within-Movement Pulse Placement and Magnitude. *Journal of Neurophysiology*, 96(2):710–720, August 2006. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00215.2006.
- Michael S. Fine and Kurt A. Thoroughman. Trial-by-Trial Transformation of Error Into Sensorimotor Adaptation Changes With Environmental Dynamics. *Journal of Neurophysiology*, 98(3):1392–1404, September 2007. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00196.2007.
- A. F. Fuchs and D. A. Robinson. A method for measuring horizontal and vertical eye movement chronically in the monkey. *Journal of Applied Physiology*, 21(3):1068–1070, May 1966. ISSN 8750-7587, 1522-1601.
- A. F. Fuchs, F. R. Robinson, and A. Straube. Role of the caudal fastigial nucleus in saccade generation. I. Neuronal discharge pattern. *Journal of Neurophysiology*, 70(5):1723–1740, November 1993. ISSN 0022-3077, 1522-1598.

- Yash P. Gad and Thomas J. Anastasio. Simulating the shaping of the fastigial deep nuclear saccade command by cerebellar Purkinje cells. *Neural Networks*, 23(7): 789–804, September 2010. ISSN 0893-6080. doi: 10.1016/j.neunet.2010.05.007.
- Joseph M. Galea, Gowri Jayaram, Loni Ajagbe, and Pablo Celnik. Modulation of Cerebellar Excitability by Polarity-Specific Noninvasive Direct Current Stimulation. *The Journal of Neuroscience*, 29(28):9115–9122, July 2009. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.2184-09.2009.
- Joseph M. Galea, Alejandro Vazquez, Neel Pasricha, Jean-Jacques Orban de Xivry, and Pablo Celnik. Dissociating the Roles of the Cerebellum and Motor Cortex during Adaptive Learning: The Motor Cortex Retains What the Cerebellum Learns. *Cerebral Cortex*, 21(8):1761–1770, January 2011. ISSN 1047-3211, 1460-2199. doi: 10.1093/cercor/bhq246.
- Prateek C. Gandiga, Friedhelm C. Hummel, and Leonardo G. Cohen. Transcranial DC stimulation (tDCS): A tool for double-blind sham-controlled clinical studies in brain stimulation. *Clinical Neurophysiology*, 117(4):845–850, April 2006. ISSN 1388-2457. doi: 10.1016/j.clinph.2005.12.003.
- P. F. C. Gilbert and W. T. Thach. Purkinje cell activity during motor learning. *Brain Research*, 128(2):309–328, June 1977. ISSN 0006-8993. doi: 10.1016/0006-8993(77)90997-0.
- Laurent Goffart, Longtang L. Chen, and David L. Sparks. Deficits in Saccades and Fixation During Muscimol Inactivation of the Caudal Fastigial Nucleus in the Rhesus Monkey. *Journal of Neurophysiology*, 92(6):3351–3367, December 2004. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.01199.2003.
- Luis Nicolas Gonzalez Castro, Alkis M. Hadjiosif, Matthew A. Hemphill, and Maurice A. Smith. Environmental Consistency Determines the Rate of Motor Adaptation. *Current Biology*, 24(10):1050–1061, May 2014. ISSN 0960-9822. doi: 10.1016/j.cub.2014.03.049.
- Arash Hadipour-Niktarash, Christine K. Lee, John E. Desmond, and Reza Shadmehr. Impairment of Retention But Not Acquisition of a Visuomotor Skill Through Time-Dependent Disruption of Primary Motor Cortex. *The Journal of Neuroscience*, 27(49):13413–13419, May 2007. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.2570-07.2007.

- Masashi Hamada, Gionata Strigaro, Nagako Murase, Anna Sadnicka, Joseph M. Galea, Mark J. Edwards, and John C. Rothwell. Cerebellar modulation of human associative plasticity. *The Journal of Physiology*, 590(10):2365–2374, May 2012. ISSN 1469-7793. doi: 10.1113/jphysiol.2012.230540.
- C. Helmchen and U. Büttner. Saccade-related Purkinje cell activity in the oculomotor vermis during spontaneous eye movements in light and darkness. *Experimental Brain Research*, 103(2):198–208, March 1995. ISSN 0014-4819, 1432-1106. doi: 10.1007/BF00231706.
- David J. Herzfeld, Damien Pastor, Adrian M. Haith, Yves Rossetti, Reza Shadmehr, and Jacinta O’Shea. Contributions of the cerebellum and the motor cortex to acquisition and retention of motor memories. *NeuroImage*, 98:147–158, September 2014a. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2014.04.076.
- David J. Herzfeld, Pavan A. Vaswani, Mollie K. Marko, and Reza Shadmehr. A memory of errors in sensorimotor learning. *Science*, 345(6202):1349–1353, December 2014b. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1253138.
- David J. Herzfeld, Yoshiko Kojima, Robijanto Soetedjo, and Reza Shadmehr. Encoding of action by the Purkinje cells of the cerebellum. *Nature*, 526(7573): 439–442, October 2015. ISSN 0028-0836. doi: 10.1038/nature15693.
- Vincent S. Huang, Adrian Haith, Pietro Mazzoni, and John W. Krakauer. Rethinking Motor Learning and Savings in Adaptation Paradigms: Model-Free Memory for Successful Actions Combines with Internal Models. *Neuron*, 70(4): 787–801, May 2011. ISSN 0896-6273. doi: 10.1016/j.neuron.2011.04.012.
- Timothy Hunter, Paul Sacco, Michael A. Nitsche, and Duncan L. Turner. Modulation of internal model formation during force field-induced motor learning by anodal transcranial direct current stimulation of primary motor cortex. *The Journal of Physiology*, 587(12):2949–2961, June 2009. ISSN 1469-7793. doi: 10.1113/jphysiol.2009.169284.
- Eun Jung Hwang and Reza Shadmehr. Internal models of limb dynamics and the encoding of limb state. *Journal of Neural Engineering*, 2(3):S266, 2005. ISSN 1741-2552. doi: 10.1088/1741-2560/2/3/S09.

- Masao Ito. Cerebellar Long-Term Depression: Characterization, Signal Transduction, and Functional Roles. *Physiological Reviews*, 81(3):1143–1195, July 2001. ISSN 0031-9333, 1522-1210.
- Jun Izawa, Tushar Rane, Opher Donchin, and Reza Shadmehr. Motor Adaptation as a Process of Reoptimization. *The Journal of Neuroscience*, 28(11): 2883–2891, December 2008. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.5359-07.2008.
- Jun Izawa, Sarah E. Criscimagna-Hemminger, and Reza Shadmehr. Cerebellar Contributions to Reach Adaptation and Learning Sensory Consequences of Action. *The Journal of Neuroscience*, 32(12):4230–4239, March 2012. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.6353-11.2012.
- Gowri Jayaram, Byron Tang, Rani Pallegadda, Erin V. L. Vasudevan, Pablo Celnik, and Amy Bastian. Modulating locomotor adaptation with cerebellar stimulation. *Journal of Neurophysiology*, 107(11):2950–2957, June 2012. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00645.2011.
- Wilsaan M. Joiner and Maurice A. Smith. Long-Term Retention Explained by a Model of Short-Term Learning in the Adaptive Control of Reaching. *Journal of Neurophysiology*, 100(5):2948–2955, November 2008. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.90706.2008.
- Michael I. Jordan and David E. Rumelhart. Forward Models: Supervised Learning with a Distal Teacher. *Cognitive Science*, 16(3):307–354, July 1992. ISSN 1551-6709. doi: 10.1207/s15516709cog1603\_1.
- Manabu Kase, David C. Miller, and Hiroharu Noda. Discharges of Purkinje cells and mossy fibres in the cerebellar vermis of the monkey during saccadic eye movements and fixation. *The Journal of Physiology*, 300(1):539–555, March 1980. ISSN 1469-7793. doi: 10.1113/jphysiol.1980.sp013178.
- D. Kaski, S. Quadir, M. Patel, N. Yousif, and A. M. Bronstein. Enhanced locomotor adaptation aftereffect in the “broken escalator” phenomenon using anodal tDCS. *Journal of Neurophysiology*, 107(9):2493–2505, May 2012. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00223.2011.

- Charles D. Kassardjian, Yao-Fang Tan, Ji-Yeon J. Chung, Raquel Heskin, Michael J. Peterson, and Dianne M. Broussard. The Site of a Motor Memory Shifts with Consolidation. *The Journal of Neuroscience*, 25(35):7979–7985, August 2005. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.2215-05.2005.
- M. Kawato. Learning internal models of the motor apparatus. *The acquisition of motor behavior in vertebrates*, page 409, 1996.
- M. Kawato, K. Furukawa, and R. Suzuki. A hierarchical neural-network model for control and learning of voluntary movement. *Biological Cybernetics*, 57(3): 169–185, 1987. ISSN 0340-1200.
- Toshitaka Kimura and Hiroaki Gomi. Temporal Development of Anticipatory Reflex Modulation to Dynamical Interactions During Arm Movement. *Journal of Neurophysiology*, 102(4):2220–2231, October 2009. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.90907.2008.
- Tomoko Kitago, Sophia L. Ryan, Pietro Mazzoni, John W. Krakauer, and Adrian M. Haith. 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, 2013. doi: 10.3389/fnhum.2013.00307.
- JoAnn Kluzik, Jörn Diedrichsen, Reza Shadmehr, and Amy J. Bastian. Reach Adaptation: What Determines Whether We Learn an Internal Model of the Tool or Adapt the Model of Our Arm? *Journal of Neurophysiology*, 100(3): 1455–1464, September 2008. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.90334.2008.
- Yoshiko Kojima, Yoshiki Iwamoto, and Kaoru Yoshida. Memory of Learning Facilitates Saccadic Adaptation in the Monkey. *The Journal of Neuroscience*, 24(34):7531–7539, August 2004. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.1741-04.2004.
- Yoshiko Kojima, Robijanto Soetedjo, and Albert F. Fuchs. Changes in Simple Spike Activity of Some Purkinje Cells in the Oculomotor Vermis during Saccade Adaptation Are Appropriate to Participate in Motor Learning. *The Journal of Neuroscience*, 30(10):3715–3727, October 2010. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.4953-09.2010.

- Yoshiko Kojima, Robijanto Soetedjo, and Albert F. Fuchs. Effect of inactivation and disinhibition of the oculomotor vermis on saccade adaptation. *Brain Research*, 1401:30–39, July 2011. ISSN 0006-8993. doi: 10.1016/j.brainres.2011.05.027.
- Konrad P. Kording, Joshua B. Tenenbaum, and Reza Shadmehr. The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience*, 10(6):779–786, June 2007. ISSN 1097-6256. doi: 10.1038/nn1901.
- Ines Kralj-Hans, Joan S. Baizer, Catherine Swales, and Mitchell Glickstein. Independent roles for the dorsal paraflocculus and vermal lobule VII of the cerebellum in visuomotor coordination. *Experimental Brain Research*, 177(2): 209–222, September 2006. ISSN 0014-4819, 1432-1106. doi: 10.1007/s00221-006-0661-x.
- R. J. Krauzlis. Population coding of movement dynamics by cerebellar Purkinje cells. *Neuroreport*, 11(5):1045–1050, April 2000. ISSN 0959-4965.
- R. Kurzan, A. Straube, and U. Büttner. The effect of muscimol micro-injections into the fastigial nucleus on the optokinetic response and the vestibulo-ocular reflex in the alert monkey. *Experimental Brain Research*, 94(2):252–260, June 1993. ISSN 0014-4819, 1432-1106. doi: 10.1007/BF00230293.
- Nicolas Lang, Hartwig R. Siebner, Nick S. Ward, Lucy Lee, Michael A. Nitsche, Walter Paulus, John C. Rothwell, Roger N. Lemon, and Richard S. Frackowiak. How does transcranial DC stimulation of the primary motor cortex alter regional neuronal activity in the human brain? *European Journal of Neuroscience*, 22(2):495–504, July 2005. ISSN 1460-9568. doi: 10.1111/j.1460-9568.2005.04233.x.
- Jeong-Yoon Lee and Nicolas Schweighofer. Dual Adaptation Supports a Parallel Architecture of Motor Memory. *The Journal of Neuroscience*, 29(33): 10396–10404, August 2009. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.1294-09.2009.
- Chiang-Shan Ray Li, Camillo Padoa-Schioppa, and Emilio Bizzi. Neuronal Correlates of Motor Performance and Motor Learning in the Primary Motor Cortex of Monkeys Adapting to an External Force Field. *Neuron*, 30(2): 593–607, May 2001. ISSN 0896-6273. doi: 10.1016/S0896-6273(01)00301-4.

- N. J. Mackintosh. A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82(4):276–298, 1975. ISSN 1939-1471(Electronic);0033-295X(Print). doi: 10.1037/h0076778.
- Laura A. Malone, Erin V. L. Vasudevan, and Amy J. Bastian. Motor Adaptation Training for Faster Relearning. *The Journal of Neuroscience*, 31(42):15136–15143, October 2011. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.1367-11.2011.
- Mollie K. Marko, Adrian M. Haith, Michelle D. Harran, and Reza Shadmehr. Sensitivity to prediction error in reach adaptation. *Journal of Neurophysiology*, 108(6):1752–1763, September 2012. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00177.2012.
- David Marr. A theory of cerebellar cortex. *The Journal of Physiology*, 202(2):437–470.1, June 1969. ISSN 0022-3751.
- Firas Mawase, Lior Shmuelof, Simona Bar-Haim, and Amir Karniel. Savings in locomotor adaptation explained by changes in learning parameters following initial adaptation. *Journal of Neurophysiology*, 111(7):1444–1454, April 2014. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00734.2013.
- Samuel C. McLaughlin. Parametric adjustment in saccadic eye movements. *Perception & Psychophysics*, 2(8):359–362, August 1967. ISSN 0031-5117, 1532-5962. doi: 10.3758/BF03210071.
- Javier F. Medina and Stephen G. Lisberger. Links from complex spikes to local plasticity and motor learning in the cerebellum of awake-behaving monkeys. *Nature Neuroscience*, 11(10):1185–1192, October 2008. ISSN 1097-6256. doi: 10.1038/nn.2197.
- R. C. Miall and D. M. Wolpert. Forward Models for Physiological Motor Control. *Neural Networks*, 9(8):1265–1279, November 1996. ISSN 0893-6080. doi: 10.1016/S0893-6080(96)00035-4.
- Reza Nezafat, Reza Shadmehr, and Henry H. Holcomb. Long-term adaptation to dynamics of reaching movements: a PET study. *Experimental Brain Research*, 140(1):66–76, September 2001. ISSN 0014-4819, 1432-1106. doi:



10.1007/s002210100787.

- M. A. Nitsche and W. Paulus. Excitability changes induced in the human motor cortex by weak transcranial direct current stimulation. *The Journal of Physiology*, 527(3):633–639, September 2000. ISSN 1469-7793. doi: 10.1111/j.1469-7793.2000.t01-1-00633.x.
- K. Ohtsuka and H. Noda. Discharge properties of Purkinje cells in the oculomotor vermis during visually guided saccades in the macaque monkey. *Journal of Neurophysiology*, 74(5):1828–1840, November 1995. ISSN 0022-3077, 1522-1598.
- Takehito Okamoto, Shogo Endo, Tomoaki Shirao, and Soichi Nagao. Role of Cerebellar Cortical Protein Synthesis in Transfer of Memory Trace of Cerebellum-Dependent Motor Learning. *The Journal of Neuroscience*, 31(24): 8958–8966, June 2011a. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.1151-11.2011.
- Takehito Okamoto, Tomoaki Shirao, Fumihiro Shutoh, Toshinori Suzuki, and Soichi Nagao. Post-training cerebellar cortical activity plays an important role for consolidation of memory of cerebellum-dependent motor learning. *Neuroscience Letters*, 504(1):53–56, October 2011b. ISSN 0304-3940. doi: 10.1016/j.neulet.2011.08.056.
- Jacinta O’Shea, Valérie Gaveau, Matthieu Kandel, Kazuo Koga, Kenji Susami, Claude Prablanc, and Yves Rossetti. Kinematic markers dissociate error correction from sensorimotor realignment during prism adaptation. *Neuropsychologia*, 55:15–24, March 2014. ISSN 0028-3932. doi: 10.1016/j.neuropsychologia.2013.09.021.
- Liam Paninski, Shy Shoham, Matthew R. Fellows, Nicholas G. Hatsopoulos, and John P. Donoghue. Superlinear Population Encoding of Dynamic Hand Trajectory in Primary Motor Cortex. *The Journal of Neuroscience*, 24(39): 8551–8561, September 2004. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.0919-04.2004.
- James L. Patton, Yejun John Wei, Preeti Bajaj, and Robert A. Scheidt. Visuomotor Learning Enhanced by Augmenting Instantaneous Trajectory Error Feedback during Reaching. *PLOS ONE*, 8(1):e46466, January 2013. ISSN 1932-6203.

doi: 10.1371/journal.pone.0046466.

John M. Pearce and Geoffrey Hall. A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87(6):532–552, 1980. ISSN 1939-1471(Electronic);0033-295X(Print). doi: 10.1037/0033-295X.87.6.532.

Denis Pélisson, Laurent Goffart, Alain Guillaume, and Julie Quinet. Visuo-motor deficits induced by fastigial nucleus inactivation. *The Cerebellum*, 2(1):71–76, March 2003. ISSN 1473-4222, 1473-4230. doi: 10.1080/14734220310015629.

Abigail L. Person and Indira M. Raman. Purkinje neuron synchrony elicits time-locked spiking in the cerebellar nuclei. *Nature*, 481(7382):502–505, January 2012. ISSN 0028-0836. doi: 10.1038/nature10732.

Mario Prsa, Suryadeep Dash, Nicolas Catz, Peter W. Dicke, and Peter Thier. Characteristics of Responses of Golgi Cells and Mossy Fibers to Eye Saccades and Saccadic Adaptation Recorded from the Posterior Vermis of the Cerebellum. *The Journal of Neuroscience*, 29(1):250–262, July 2009. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.4791-08.2009.

Janine Reis, Heidi M. Schambra, Leonardo G. Cohen, Ethan R. Buch, Brita Fritsch, Eric Zarahn, Pablo A. Celnik, and John W. Krakauer. Noninvasive cortical stimulation enhances motor skill acquisition over multiple days through an effect on consolidation. *Proceedings of the National Academy of Sciences*, 106(5):1590–1595, March 2009. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.0805413106.

RA Rescorla and AW Wagner. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In AH Black and WF Prokasy, editors, *Classical Conditioning II: Current Research and Theory*, pages 64–99. Appleton-Century-Crofts, 1972.

Andrew G. Richardson, Simon A. Overduin, Antoni Valero-Cabré, Camillo Padoa-Schioppa, Alvaro Pascual-Leone, Emilio Bizzi, and Daniel Z. Press. Disruption of Primary Motor Cortex before Learning Impairs Memory of Movement Dynamics. *The Journal of Neuroscience*, 26(48):12466–12470, November 2006. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.1139-06.2006.

- M. Riedmiller and H. Braun. A direct adaptive method for faster backpropagation learning: the RPROP algorithm. In , *IEEE International Conference on Neural Networks, 1993*, pages 586–591 vol.1, 1993. doi: 10.1109/ICNN.1993.298623.
- F. R. Robinson, A. Straube, and A. F. Fuchs. Role of the caudal fastigial nucleus in saccade generation. II. Effects of muscimol inactivation. *Journal of Neurophysiology*, 70(5):1741–1758, November 1993. ISSN 0022-3077, 1522-1598.
- Farrel R. Robinson, Christopher T. Noto, and Scott E. Bevans. Effect of Visual Error Size on Saccade Adaptation in Monkey. *Journal of Neurophysiology*, 90(2):1235–1244, August 2003. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00656.2002.
- A. M. E. Sarwary, L. P. J. Selen, and W. P. Medendorp. Vestibular benefits to task savings in motor adaptation. *Journal of Neurophysiology*, 110(6):1269–1277, September 2013. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00914.2012.
- Robert A. Scheidt, Jonathan B. Dingwell, and Ferdinando A. Mussa-Ivaldi. Learning to Move Amid Uncertainty. *Journal of Neurophysiology*, 86(2): 971–985, August 2001. ISSN 0022-3077, 1522-1598.
- Jennifer A. Semrau, Amy L. Daitch, and Kurt A. Thoroughman. Environmental experience within and across testing days determines the strength of human visuomotor adaptation. *Experimental Brain Research*, 216(3):409–418, December 2011. ISSN 0014-4819, 1432-1106. doi: 10.1007/s00221-011-2945-z.
- R. Shadmehr and T. Brashers-Krug. Functional stages in the formation of human long-term motor memory. *Journal of Neuroscience*, 17(1):409–419, 1997. ISSN 0270-6474.
- R. Shadmehr and F.A. Mussa-Ivaldi. Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience*, 14(5 II):3208–3224, 1994. ISSN 0270-6474.
- R. Shadmehr, J. Brandt, and S. Corkin. Time-dependent motor memory processes in amnesic subjects. *Journal of Neurophysiology*, 80(3):1590–1597, 1998. ISSN

0022-3077.

- Reza Shadmehr. Generalization as a behavioral window to the neural mechanisms of learning internal models. *Human Movement Science*, 23(5):543–568, November 2004. ISSN 0167-9457. doi: 10.1016/j.humov.2004.04.003.
- Gary C. Sing, Wilsaan M. Joiner, Thrishantha Nanayakkara, Jordan B. Brayanov, and Maurice A. Smith. Primitives for motor adaptation reflect correlated neural tuning to position and velocity. *Neuron*, 64(4):575–589, November 2009. ISSN 1097-4199. doi: 10.1016/j.neuron.2009.10.001.
- Maurice A Smith and Reza Shadmehr. Modulation of the rate of error-dependent learning by the statistical properties of the task. *Advances in Computational Motor Control*, 3(2004), 2004.
- Maurice A. Smith and Reza Shadmehr. Intact Ability to Learn Internal Models of Arm Dynamics in Huntington’s Disease But Not Cerebellar Degeneration. *Journal of Neurophysiology*, 93(5):2809–2821, May 2005. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00943.2004.
- Maurice A Smith, Ali Ghazizadeh, and Reza Shadmehr. Interacting Adaptive Processes with Different Timescales Underlie Short-Term Motor Learning. *PLoS Biol*, 4(6):e179, May 2006. doi: 10.1371/journal.pbio.0040179.
- Robijanto Soetedjo and Albert F. Fuchs. Complex Spike Activity of Purkinje Cells in the Oculomotor Vermis during Behavioral Adaptation of Monkey Saccades. *The Journal of Neuroscience*, 26(29):7741–7755, July 2006. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.4658-05.2006.
- Robijanto Soetedjo, Yoshiko Kojima, and Albert F. Fuchs. Complex Spike Activity in the Oculomotor Vermis of the Cerebellum: A Vectorial Error Signal for Saccade Motor Learning? *Journal of Neurophysiology*, 100(4):1949–1966, October 2008. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.90526.2008.
- Robijanto Soetedjo, Albert F. Fuchs, and Yoshiko Kojima. Subthreshold Activation of the Superior Colliculus Drives Saccade Motor Learning. *The Journal of Neuroscience*, 29(48):15213–15222, February 2009. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.4296-09.2009.

- C. J. Stagg, G. Jayaram, D. Pastor, Z. T. Kincses, P. M. Matthews, and H. Johansen-Berg. Polarity and timing-dependent effects of transcranial direct current stimulation in explicit motor learning. *Neuropsychologia*, 49(5): 800–804, April 2011. ISSN 0028-3932. doi: 10.1016/j.neuropsychologia.2011.02.009.
- Hirokazu Tanaka, Terrence J. Sejnowski, and John W. Krakauer. Adaptation to Visuomotor Rotation Through Interaction Between Posterior Parietal and Motor Cortical Areas. *Journal of Neurophysiology*, 102(5):2921–2932, November 2009. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.90834.2008.
- Petra Telgkamp, Daniel E Padgett, Veronica A Ledoux, Catherine S Woolley, and Indira M Raman. Maintenance of High-Frequency Transmission at Purkinje to Cerebellar Nuclear Synapses by Spillover from Boutons with Multiple Release Sites. *Neuron*, 41(1):113–126, January 2004. ISSN 0896-6273. doi: 10.1016/S0896-6273(03)00802-X.
- Peter Thier, Peter W. Dicke, Roman Haas, and Shabtai Barash. Encoding of movement time by populations of cerebellar Purkinje cells. *Nature*, 405(6782): 72–76, May 2000. ISSN 0028-0836. doi: 10.1038/35011062.
- K. A. Thoroughman and R. Shadmehr. Electromyographic correlates of learning an internal model of reaching movements. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 19(19):8573–8588, October 1999. ISSN 1529-2401.
- K. A. Thoroughman and R. Shadmehr. Learning of action through adaptive combination of motor primitives. *Nature*, 407(6805):742–747, October 2000. ISSN 0028-0836. doi: 10.1038/35037588.
- Michael C. III Trent and Alaa A. Ahmed. Learning from the value of your mistakes: evidence for a risk-sensitive process in movement adaptation. *Frontiers in Computational Neuroscience*, 7:118, 2013. doi: 10.3389/fncom.2013.00118.
- Edward J. A. Turnham, Daniel A. Braun, and Daniel M. Wolpert. Facilitation of learning induced by both random and gradual visuomotor task variation. *Journal of Neurophysiology*, 107(4):1111–1122, February 2012. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00635.2011.

- Yoshikazu Ugawa, Yoshikazu Uesaka, Yasuo Terao, Ritsuko Hanajima, and Ichiro Kanazawa. Magnetic stimulation over the cerebellum in humans. *Annals of Neurology*, 37(6):703–713, June 1995. ISSN 1531-8249. doi: 10.1002/ana.410370603.
- Robert J. van Beers. Motor Learning Is Optimally Tuned to the Properties of Motor Noise. *Neuron*, 63(3):406–417, August 2009. ISSN 0896-6273. doi: 10.1016/j.neuron.2009.06.025.
- Robert J. van Beers. How Does Our Motor System Determine Its Learning Rate? *PLoS ONE*, 7(11):e49373, November 2012. doi: 10.1371/journal.pone.0049373.
- Pavan A. Vaswani and Reza Shadmehr. Decay of Motor Memories in the Absence of Error. *The Journal of Neuroscience*, 33(18):7700–7709, January 2013. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.0124-13.2013.
- T. Vilis and J. Hore. Central neural mechanisms contributing to cerebellar tremor produced by limb perturbations. *Journal of Neurophysiology*, 43(2):279–291, 1980. ISSN 0022-3077.
- Kunlin Wei and Konrad Körding. Relevance of Error: What Drives Motor Adaptation? *Journal of Neurophysiology*, 101(2):655–664, February 2009. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.90545.2008.
- Minnan Xu-Wilson, Haiyin Chen-Harris, David S. Zee, and Reza Shadmehr. Cerebellar Contributions to Adaptive Control of Saccades in Humans. *The Journal of Neuroscience*, 29(41):12930–12939, October 2009. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.3115-09.2009.
- Jinzo Yamada and Hiroharu Noda. Afferent and efferent connections of the oculomotor cerebellar vermis in the macaque monkey. *The Journal of Comparative Neurology*, 265(2):224–241, November 1987. ISSN 1096-9861. doi: 10.1002/cne.902650207.
- Eric Zarahn, Gregory D. Weston, Johnny Liang, Pietro Mazzoni, and John W. Krakauer. Explaining Savings for Visuomotor Adaptation: Linear Time-Invariant State-Space Models Are Not Sufficient. *Journal of Neurophysiology*, 100(5):2537–2548, November 2008. ISSN 0022-3077,

1522-1598. doi: 10.1152/jn.90529.2008.

**David J. Herzfeld, B.S., M.S.**  
**Born on March 23, 1988 in Milwaukee, WI**

---

**Contact Information**

David J. Herzfeld, B.S., M.S.  
416 Traylor, 720 Rutland Ave.  
Baltimore, MD 21205 USA  
www.davidherzfeld.com

Department of Biomedical Engineering  
Johns Hopkins University School of Medicine  
*Cell:* (414) 690-8965  
*E-mail:* dherzfe1@jhmi.edu

**Education**

Ph.D., Johns Hopkins University School of Medicine  
Department of Biomedical Engineering, June 2016

- Advisor: Reza Shadmehr, Ph.D.
- Area of Study: Motor Control & Motor Learning

M.S. (5 Year Master's Degree), Marquette University  
Department of Biomedical Engineering, July 2011

- Thesis Topic: *Modeling and Computational Framework for the Specification and Simulation of Large-scale Spiking Neural Networks*
- Advisor: Scott A. Beardsley, Ph.D.
- Area of Study: Bioinstrumentation & Biocomputation

B.S., Marquette University  
Department of Biomedical Engineering, May 2010

- *Summa cum laude*, Top Scholar in Curriculum
- Biocomputing specialization
- Emphasis on computational modeling using parallel and distributed systems

**Research Experience**

Ph.D. Student

January 2012 to June 2016

- Department of Biomedical Engineering, Johns Hopkins University School of Medicine
- Computation Motor Control Laboratory
- Laboratory Director: Reza Shadmehr, Ph.D.
- Laboratory topics include psychophysical motor learning experiments, functional imaging experimental design and analysis, and neurophysiology analysis.

Master's Student (5 Year Master's Degree)

May 2008 to July 2011

- Department of Biomedical Engineering, Marquette University
- Integrative Neural Systems Laboratory
- Laboratory Director: Scott A. Beardsley, Ph.D.



- Laboratory topics include computational neural modeling, implementation and analysis of psychophysics tasks, and EEG/fMRI analysis.

Research Assistant

May 2009 to May 2010

- Department of Math, Statistics, and Computer Science, Marquette University
- Bistro Lab (Bioinformatics & Biomathematical Modeling)
- Laboratory Director: Craig A. Struble, Ph.D.
- Implemented a series of software applications to deploy campus grids of virtualized Linux systems in Windows environments.

### Teaching Experience

- Learning Theory, Teaching Assistant & Guest Lecturer, Johns Hopkins University (Spring 2014)
- Intro. to Embedded Microcontrollers, Instructor, Johns Hopkins University (Winter Intersession 2014)
- Biocomputer Design Lab II, Teaching Assistant, Marquette University (Spring 2011)
- Embedded Biomedical Instrumentation, Teaching Assistant, Marquette University (Spring 2011)
- Computing for Biomedical Engineers, Teaching Assistant & Guest Lecturer, Marquette University (Fall 2010)
- Biocomputer Design Lab I, Teaching Assistant & Guest Lecturer, Marquette University (Fall 2010)

### Professional Experience

Consultant, Ardian, Inc., Medtronic

May 2010 to Present

- *Project:* Renal ablation radio frequency generator used in surgical procedure to permanently reduce/eliminate chronic hypertension.
- Developed algorithms and research framework for evaluating treatment outcomes and a reduction in treatment error conditions.
- Work resulted in U.S. and international patents

### Research Interests

- Motor learning & motor control
- Cerebellar contributions to learning
- Computational neural modeling
- Neuroengineering
- Neural coding/decoding

### Honors and Awards

- Martin & Carol Macht Award, Johns Hopkins Young Investigators' Day, 2016

- Siebel Scholar, Class of 2016. Awarded annually for academic excellence and demonstrated leadership to 93 top students from the world's leading graduate schools
- Mette Strand Award, Johns Hopkins Young Investigators' Day, 2015
- Society for the Neural Control of Movement Fellowship, 2015
- NIH Ruth L. Kirschstein National Research Service Award (NRSA F31), 2014 to 2016
- Neuroengineering Training Initiative (NETI) Scholar, 2011 to 2012
- Richard W. Jobling Fellowship, Marquette University, September 2010
- Top Scholar In Curriculum, Marquette University, April 2010
- Anthony J. and Rose E. Bagoszzi Medical Research Fellowship, Marquette University, September 2009
- National Dean's List, Marquette University, 2006 to 2010
- Allen Edmonds Scholarship, Marquette University, 2006 to 2010

### **Peer Reviewed Publications**

- Herzfeld, D.J.**, Kojima, Y., Soetedjo, R., and R. Shadmehr. Encoding of action in the Purkinje cells of the cerebellum. *Nature*, 526:439-442.
- Herzfeld, D.J.**, Kojima, Y., Soetedjo, R., and R. Shadmehr. Encoding of action in the Purkinje cells of the cerebellum. *TCMC: Translational and Computational Motor Control*, 2014.
- Herzfeld, D.J.**, Vaswani, P.A., Marko, M.K., and R. Shadmehr. A memory of errors in sensorimotor learning. *Science*, 345:1349-1352, 2014.
- Herzfeld, D.J.\***, Pastor, D.\*, Haith, A.M., Rossetti, Y., Shadmehr, R., and J. O'Shea. Contributions of the cerebellum and the motor cortex to acquisition and retention of motor memories. *NeuroImage*, 98:147-158, 2014.
- Herzfeld, D.J.**, Vaswani, P.A., Marko, M.K. and R. Shadmehr. Sensitivity of motor adaptation depends on the history of experienced errors. *TCMC: Translational and Computational Motor Control*, 2013.
- Herzfeld, D.J.**, and S.A. Beardsley. Synaptic Weighting for Physiological Responses in Recurrent Spiking Neural Networks. *IEEE Engineering in Medicine and Biology*. 2011:4187-90, 2011.
- Herzfeld, D.J.**, and S.A. Beardsley. Improved multi-unit decoding at the brain-machine interface using population temporal linear filtering. *J Neural Eng.*, 7(4):046012, 2010.
- Herzfeld, D.J.**, Olson, L.E. and C.A. Struble. Pools of virtual boxes: building campus grids with virtual machines. *HPDC: Proceedings of the ACM*. 667-675, 2010.
- Herzfeld, D.J.**, Kojima, Y., Soetedjo, R., and R. Shadmehr. Cerebellar output predicts movement kinematics as a forward model, *In preparation*.
- Herzfeld, D.J.**, Kojima, Y., Soetedjo, R., and R. Shadmehr. Neural correlates of learning in the vermis of the cerebellum, *In preparation*.

### **Invited Reviews**

- Herzfeld, D.J.** and R. Shadmehr. Motor variability is not noise, but grist for the learning mill. *Nat Neurosci*, 17(2):149-50, 2014.
- Herzfeld, D.J.** and R. Shadmehr. Cerebellum estimates the sensory state of the body. *Trends in Cognitive Sciences*, 18(2)66-7, 2013.

### **Selected Abstracts**

- Herzfeld, D.J.**, Kojima, Y., Soetedjo, R., and R. Shadmehr. Sensory prediction errors during saccade adaptation drive cerebellar complex spikes and learning. *Soc. Neurosci*, 2016. San Diego, CA.
- Herzfeld, D.J.**, Kojima, Y., Soetedjo, R., and R. Shadmehr. Encoding of action in the Purkinje cells of the cerebellum. *Soc. Neurosci*, 2015. Chicago, IL.
- Herzfeld, D.J.**, Kojima, Y., Soetedjo, R. and R. Shadmehr. Encoding of action in the Purkinje cells of the cerebellum. *Gordon Cerebellar Conference*, 2015. Lewiston, ME.
- Shadmehr, R. and **D.J. Herzfeld**. Changes in error-sensitivity account for sensorimotor savings. *Soc. Neurosci*, 2014. Washington, D.C.
- Herzfeld, D.J.**, Kojima, Y., Soetedjo, R., and R. Shadmehr. Encoding of prediction error by complex spikes of the cerebellum. *Soc. Neurosci*, 2014. Washington, D.C.
- Herzfeld, D.J.**, Vaswani, P.A., and R. Shadmehr. Sensitivity of motor adaptation depends on the history of experienced errors. *Soc. Neurosci*, 2013. San Diego, CA.
- Herzfeld, D.J.**, Vaswani, P.A., Marko, M.K., Kojima, Y., Soetedjo, R. and R. Shadmehr. Sensitivity of motor adaptation depends on the history of experienced errors. *Gordon Cerebellar Conference*, 2013. New London, NH.
- Herzfeld, D.J.**, and R. Shadmehr. Sensitivity to error depends on perturbation statistics. *Soc. Neurosci*, 2012. New Orleans, LA.
- Herzfeld, D.J.**, and S.A. Beardsley. Localization of synaptic changes using simulated hemodynamic responses. *Soc. Neurosci*, 2011. Washington, D.C.
- Herzfeld, D.J.**, and S.A. Beardsley. Improved multi-unit decoding at the brain-machine interface using population temporal linear filtering. *Soc. Neurosci*, 2009. Chicago, IL.

### **Selected Invited Talks**

- Encoding of action in the Purkinje cells of the cerebellum. Gordon Cerebellar Conference, Lewiston, ME, August 2015.
- A memory of errors in sensorimotor learning, Johns Hopkins Young Investigator's Day, April 2015.
- Encoding of action in the Purkinje cells of the cerebellum, Neural Control of Movement, Raleigh, N.C., April 2015.
- Encoding of action in the Purkinje cells of the cerebellum, Biomedical Engineering Seminar Series, Johns Hopkins University, January 2015.
- A memory of errors in sensorimotor learning, Brotz Seminar, Marquette University, September 2014.

- A memory of errors in sensorimotor learning, Workshop on Neural Population Dynamics Underlying Sensorimotor Integration, Janelia Farms (Howard Hughes Medical Institutes), Ashburn VA, June 2014.
- Learning from error: history of past errors dictates sensitivity to error, TCMC: Translational and Computational Motor Control, San Diego, November 2013.
- Deploying Condor using VirtualBox on Windows Hosts, Condor Week, University of Wisconsin - Madison, April 2009.

### **Patents**

**Herzfeld, D.J.**, Ballakur S., Beitel R.J., Friedrichs P., Wu A., Zarins D., and Leung M.S. Devices, systems and methods for evaluation and feedback of neuromodulation treatment. Patent filed October, 2011. US and International Patents (11888774, PCT/US2011/057740, 13281269).

**Herzfeld, D.J.**, Beitel R.J., Friedrichs P., and Wu A. Evaluation and feedback for a radiofrequency treatment. Provisional patent filed 2010.

### **Funding**

**F31NS090860** (Role: PI)

01/01/15 - Current

#### **A memory of errors in motor adaptation**

NIH Ruth L. Kirschstein National Research Service Award which focuses on the role of error-sensitivity to the processes of savings and meta-learning.

Understanding how to increase error-sensitivity may speed motor recovery during rehabilitation training following neurotrauma or disease.

### **Professional Service**

- **Reviewer:** PLOS Computational Biology, Experimental Brain Research, Neural Networks
- **Assistant Reviewer:** Nature, Science, Nature Neuroscience, Brain, Journal of Neuroscience, Journal of Neurophysiology