Title
The Cerebellum and Motor Learning

Author
Schlerf, John Edward

Publication Date
2010

Peer reviewed|Thesis/dissertation

The Cerebellum and Motor Learning

by

John Edward Schlerf

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Neuroscience

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Richard B. Ivry, Chair
Professor Martin S. Banks
Professor Jose M. Carmena
Professor Steven L. Lehman

Spring 2010

The Cerebellum and Motor Learning

Abstract

The Cerebellum and Motor Learning

by

John Edward Schlerf
Doctor of Philosophy in Neuroscience

University of California, Berkeley

Professor Richard B. Ivry, Chair

During our daily lives, we make thousands of movements. When we stop and consider that doing something as ordinary as reaching for a glass of juice involves the precise sequential contraction of dozens of muscles simply to move our hand, we appreciate the immense problem that our brains are solving. If we then recognize that both the world and the body are constantly changing, the accuracy with which we move becomes quite staggering. Moving with such proficiency requires the motor system to be continuously learning and adapting. A host of neural structures are important for this behavior. One remarkable part of this system is the cerebellum, or "little brain": a phylogenetically ancient neural structure, containing over half of the neurons in the human central nervous system. Damage to this structure results in a loss of coordination, with marked impairments in the control of eye movements, the timing of simple rhythmic movements, and most intriguingly the ability to adjust well-learned motor skills.

The aim of this dissertation is to explore the processes of motor control and learning, with a special emphasis on the functional contribution of the cerebellum. Following a short introduction (Chapter 1), empirical evidence is provided from two classes of behavior. Chapter 2 deals with the production of rhythmic movements in a population of patients with cerebellar pathology. Chapters 3 through 5 involve the production of goal-directed reaching movements, carefully investigating the representation and correction of errors through the combined use of psychophysics, brain imaging, and patient studies.

In Chapter 2, patients with cerebellar pathology are observed to be impaired when producing rhythmic movements, particularly when the movements contain a

distinct event that can be used to determine the performance error. In Chapter 3, we observe that by reshaping a target region, we can predictably impact the correction of movement errors during reaching movements toward that target. In Chapter 4, we provide physiological evidence of the representation of movement errors within the cerebellum, an effect only observed when appropriate measures are taken to factor out the effects of changes in heart rate. In Chapter 5, we show that patients with cerebellar pathology are impaired in adjusting their movements to counteract a visual perturbation, and furthermore suggest that this impairment is equivalent whether the perturbation is applied suddenly or gradually.

Taken together, this work demonstrates that we learn to make better movements by rapidly evaluating our movements with respect to our goals, and correcting any mistakes with the help of the cerebellum.

To family, friends, and most importantly
the love of my life, my wife Ann.

Contents

List of Figures

List of Tables

Acknowledgments

As science is a collaborative effort, a special thanks is due to all coauthors of these chapters, whose hard work is reproduced here with permission.

Chapter 2: John E. Schlerf, Rebecca M. C. Spencer, Howard N. Zelaznik, and Richard B. Ivry

Chapter 3: John E. Schlerf and Richard B. Ivry

Chapter 4: John E. Schlerf, Richard B. Ivry, and Jörn Diedrichsen

Chapter 5: John E. Schlerf, Jing Xu, Nola M. Klemfuss, Thomas L. Griffiths, and Richard B. Ivry

Financial support was provided by a National Science Foundation Graduate Student Research Fellowship, as well as grants NIH HD060306 and NSF BCS 0726685.

Chapter 1

Introduction

Damage to the cerebellum results in movement coordination deficits (Holmes, 1939; Thach, 1998; Trouillas *et al.*, 1997), though there are also many examples of of cognitive deficits (e.g., Schmahmann & Sherman, 1998; Ravizza *et al.*, 2006), presumably due toconnections with frontal and prefrontal cortices (Kelly & Strick, 2003; Middleton & Strick, 2001). Theorists have long been drawn to the strikingly regular ogranization of the cerebellar cortex, with two independent but nearly identical theories suggesting that this organization is well suited for supervised learning (Albus, 1971; Marr, 1969). Despite a well-characterized description of the clinical outcome of cerebellar dysfunction and relatively thorough neural computational analyses, the function of the cerebellum in remains poorly understood. The current work aims to focus on this issue, exploring the role of the cerebellum in motor learning.

The first experimental chapter of this dissertation (Chapter 2) examines the production of discrete rhythmic movements. Such movements require an accurate representation of time. The cerebellum has been implicated in temporal processing since it was demonstrated to be critical for eyeblink conditioning (Lincoln *et al.*, 1982; McCormick & Thompson, 1984), a behavior that requires learning an adaptive and accurately timed response to an external stimulus. Such conditioning, which proved an important test case of the Marr/Albus theory of cerebellar processing (Ito, 2001), suggests that the accurate representation of time may be a central function of the cerebellar cortex. In support of this hypothesis, participants with cerebellar damage are impaired on a host of tasks in which accurate timing is essential (Ivry & Keele, 1989; Ivry *et al.*, 1988). More recent work has suggested that the utilization of explicit temporal representations may vary for different classes of movements. Cer-

tain movements (for example rhythmic finger tapping) contain within them a distinct event that can be compared to a central clock. However, other movements can be rhythmic without requiring explicit temporal control (involving, instead, "emergent timing"). An example of such a movement is continuous circle drawing (Zelaznik *et al.*, 2002). Surprisingly, participants with cerebellar damage are able to produce rhythmic, continuous circle drawing movements with accurate timing (Spencer *et al.*, 2003).

While the ability to produce explicitly timed movements seems to involve processes specialized for controlling time, the control processes for well-timed continuous movements remain unclear. Recent experimental work suggests that during continuous movement tasks the very first intervals are under the control of the central timing system (Zelaznik *et al.*, 2005). This suggests that an explicit temporal prediction occurs about the timing of that first interval, and typical error correction processes – possibly involving the cerebellar representation (Albus, 1971; Marr, 1969) – are able to adjust the movement before control shifts to other systems. In order to further investigate the contribution of the cerebellum to these movements, we performed a rhythmic movement experiment in which patients with cerebellar ataxia produced either tapping movements or drew continuous circles. The task was designed to focus on variability patterns during the initial movement cycles. In healthy college aged subjects, the variability across trials of these early intervals in continuous circle drawing was much higher than later intervals, while the variability across trials for explicitly timed tasks was constant (Zelaznik *et al.*, 2005). In contrast, both elderly controls and participants with cerebellar ataxia consistently showed higher across-trial variability during the early intervals for both finger tapping and continuous circle drawing. This work was published in the journal Cerebellum, and is reproduced here with permission from *Taylor and Francis* and *Springer*.

Chapter 3 looks at the adaptation of goal-directed reaching movements in order to counteract an external perturbation. In order to produce accurate movements in a dynamic and unstable environment, the motor system needs to be flexible. Previous work on this problem has used two-dimensional tasks (Fine & Thoroughman, 2007; Thoroughman & Taylor, 2005; Thoroughman & Shadmehr, 2000; Lackner & Dizio, 1994; Shadmehr & Mussa-Ivaldi, 1994) in which participants make reaching movements toward a single point. This method makes it difficult to address two important questions. First, there is evidence that when learning a novel skill, participants

exploit task-based redundancies to improve performance (Müller & Sternad, 2004). Point targets are relatively constrained, however, and have no such redundancies. If these redundancies are added, will participants take advantage of them? Second, across many trials movements are planned in order to guarantee the maximal reward, taking into account inherent noise as well as task uncertainty (Trommershäuser *et al.*, 2005, 2003a,b). If we look within a trial, can we see evidence of similar processes, or are they only active in making adjustments between trials? To ask this another way, when participants experience a perturbation, will they attempt to automatically counteract the perturbation, or will their correction be shaped by what they know about the reward contingencies in the environment?

To examine these questions, a novel reaching task was developed. Participants made reaches toward one of two diagonal rectangular targets. While reaching, their arms were perturbed upward by a robotic device that imposed an external force field. By using a third dimension, participants could make corrections along the horizontal direction, an axis that was not contaminated by the applied force field. When the field was unexpectedly turned off, we looked at the online corrections that participants used to get their hands closer to the target. Interestingly, participants integrated the goal into their corrective movements, as the orientation of the target affected the trajectory.

An extension of this task would be to test patients with cerebellar ataxia, asking if damage to this structure disrupts the use of goal-based redundancy. Reach adaptation is typically impaired among participants with cerebellar damage (Martin *et al.*, 1996; Maschke *et al.*, 2004; Smith & Shadmehr, 2005; Tseng *et al.*, 2007; Werner *et al.*, 2009). This is often thought of as further evidence of the Marr/Albus model, given that the learning process here is fundamentally error-driven. This chapter suggests, however, that avoiding errors retroactively is not all the motor system is attempting to do. Online corrections which critically depend on the goal can be observed. Cerebellar patients may thus show a heightened sensitivity to task redundancies in order to compensate for their poor adaptation abilities.

The remaining two chapters of the thesis look directly at the cerebellar contribution to motor learning. We start by investigating error representations themselves before moving to a careful consideration of error-based motor learning. Chapter 4 uses fMRI to examine the response of the human cerebellum to error signals. Despite featuring prominently in the Marr/Albus model and providing a parsimonious ac-

count of learning deficits following cerebellar damage, a signature of error processing in the cerebellum has proven elusive. Studies which show a contribution of the cerebellum to motor learning find that cerebellar activity is strongest early in learning, when errors are largest (e.g., Imamizu *et al.*, 2000; Luauté *et al.*, 2009). However, such studies often fail to account for the increased movement that people produce to correct for their errors. Indeed, after such corrective movements are explicitly accounted for, the data fail to show evidence of a cerebellar response to errors per se (Diedrichsen *et al.*, 2005b; Krakauer *et al.*, 2004).

A recent physiological study motivated us to reexamine this question. Horn *et al.* (2004) noted that the inferior olive, the source of the "teaching signal" in the cerebellar cortex under the Marr/Albus model, was only activated by the presence of an unexpected sensory signal, rather than the absence of an expected sensory signal (which also signals an error). Thus, we looked for asymmetric error representations within the cerebellar cortex in humans performing reaching movements under conditions in which online corrections were not possible. Contrary to the hypothesis of Horn *et al.* (2004), we did not observe asymmetries in activation patterns to the two types of errors. In fact, the results showed a very strong negative response to errors, opposite what would be expected based on theories of error-based learning and the cerebellum. However, we also observed that this decline in the hemodynamic response was correlated with a decrease in heart rate. Task-free fMRI studies have recently demonstrated that heart rate has a measurable affect on the BOLD signal measured with fMRI (Chang *et al.*, 2009; Shmueli *et al.*, 2007). Indeed, when heart rate is regressed out of the data, an error representation became visible within the arm area of the cerebellar cortex. This study provides a strong warning about the relevance of physiological covariates in fMRI studies. Failure to account for the heart rate changes would have led to the erroneous rejection of the hypothesis that errors are represented in the cerebellum. Moreover, the results reconcile a troubling discrepancy between clinical studies that emphasize the importance of the cerebellum for error-based learning and neuroimaging studies that have failed to identify cerebellar responses to error. Further investigations into the nature of error signals are possible, for example whether the region of cerebellar cortex processing the error is sensitive to the sensory modality of the error signal itself.

The final experimental chapter considers the cerebellar contribution to motor adaptation directly, examining participants with spinocerebellar ataxia during per-

formance of a visuomotor adaptation task. In the spirit of the research on continuous rhythmic movements (Spencer *et al.*, 2003), we attempted to create a situation where adaptation was predicted to not involve the cerebellum, contrasting this with a condition in which we expected adaptation to be severely disrupted in patients with cerebellar pathology. This was done by comparing the learning of visuomotor perturbations introduced gradually to perturbations introduced all at once. When a rotation is introduced in small steps, participants have been observed to perform accurately with a 90 degree rotation (where moving to the left moves a cursor straight up) with little awareness that any manipulation has occurred (Kagerer *et al.*, 1997). For a variety of reasons, we expected that small errors might be corrected solely by cerebral mechanisms, particularly as awareness of the gradual onset rotations is so poor.

However, when we examined performance on this task, introducing the perturbation gradually had little effect on the amount of learning. Neither elderly controls nor ataxics demostrated stronger learning of gradual onset rotations. A detailed analysis was performed with a Kalman Filter model of the visuomotor performance, allowing us to test quantitative predictions about learning in the presence of noisy information. This model was used to ask whether the impaired learning could be described as a direct consequence of the motor deficits in the patients. Despite revealing some subtle differences between the sudden-onset and gradual-onset rotations, ataxics remained impaired beyond what should be optimal given their noisy motor systems.

Taken together, these chapters provide a number of novel insights into the function of the cerebellum in motor control and motor learning.

Chapter 2

Timing of rhythmic movements in patients with cerebellar degeneration[1]

Abstract

A distinction in temporal performance has been identified between two classes of rhythmic movements: those requiring explicit timing of salient events marking successive cycles, i.e., event timing, and continuous movements in which timing is hypothesized to be emergent. Converging evidence in support of this distinction is reviewed, including neuropsychological studies showing that individuals with cerebellar damage are selectively impaired on tasks requiring event timing (e.g., tapping). Recent behavioral evidence in neurologically healthy individuals suggests that for continuous movements (e.g., circle drawing), the initial cycle is marked by a transformation from event to emergent timing, allowing the participant to match their movement rate to an externally defined cycle duration. We report a new experiment in which individuals with cerebellar ataxia produced rhythmic tapping or circle drawing movements. Participants were either paced by a metronome or unpaced. Ataxics showed a disproportionate increase in temporal variability during tapping compared to circle drawing, although they were more variable than controls on both tasks. However, two predictions of the transformation hypothesis were not confirmed. First, the

[1]Previously published in *The Cerebellum*, Volume 6, Number 3; September, 2007. doi:10.1080/14734220701370643 All figures and text reproduced with permission from *Springer*.

ataxics did not show a selective impairment on circle drawing during the initial cycles, a phase when we hypothesized event timing would be required to establish the movement rate. Second, the metronome did not increase variability of the performance of the ataxics. Taken together, these results provide further evidence that the integrity of the cerebellum is especially important for event timing, although our attempt to specify the relationship between event and emergent timing was not successful.

2.1 Introduction

Many models of cerebellar function emphasize a critical role for this structure in the representation of precise temporal information (Zelaznik *et al.*, 2005; Ivry & Keele, 1989; Perrett *et al.*, 1993; Medina *et al.*, 2000). These models seek to account for a range of task domains associated with the cerebellum, including the production of well-timed movements, certain types of sensorimotor learning, and various perceptual tasks in which precise timing is essential.

The production of rhythmic movements has been one of the most widely-employed tasks for studying timing. Rhythm production tasks are appealing for a variety of reasons. Analytically, formal models have been developed to differentiate between sources of variability (Wing & Kristofferson, 1973; Collier & Ogden, 2004). Methodologically, the required movements are sufficiently simple that they can be performed by participants with neurological disorders (Ivry *et al.*, 1988; O'Boyle *et al.*, 1996; Spencer & Ivry, 2005) as well as in the constrained environments used in neuroimaging research (e.g., Desmond *et al.*, 1997; Rao *et al.*, 1997). Studies using these models and methods have provided some of the foundational evidence in support of the hypothesis that the cerebellum operates as an internal timing system.

Our recent behavioral work with rhythmic movement tasks has led to a further refinement of the cerebellar timing hypothesis. In the original behavioral work, healthy college-aged participants did not show correlated individual differences on timing precision between tapping, a discrete task, and circle drawing, a smooth, continuous task (Robertson *et al.*, 1999; Zelaznik *et al.*, 2000, 2002). This dissociation led us to propose a distinction between event and emergent timing (Ivry *et al.*, 2002). In event timing, salient events such as the initiation point for each cycle or the point of contact with a surface (as in table tapping) define the temporal interval: the participant controls the timing of these events to match the task goal (e.g., maintain

a movement rate defined by a metronome). In emergent timing, this goal may be achieved indirectly by exploiting the dynamics of a biomechanical system. For example, in drawing rhythmic circles, movement rate can be maintained by adopting the appropriate dynamics of an oscillatory system.

Although asserting the presence of two timing control processes may seem to lack parsimony, the distinction between event and emergent timing has been supported by a number of studies designed to test non-intuitive, neurophysiologically-based predictions. Most relevant for this paper is the work of Spencer *et al.* (2003) showing that individuals with cerebellar pathology are selectively impaired on tasks theorized to require event timing. First, participants with unilateral cerebellar lesions were impaired in a timed tapping task, as well as an intermittent circle drawing task when each of these tasks were performed with the ipsilesional hand (the 'impaired' hand) compared to their contralesional hand (the 'unimpaired' hand). Of greater interest was the result that on the continuous circle drawing task, hypothesized to use emergent timing, performance of the ipsilesional and contralesional hands was virtually identical. In an additional experiment, only finger movements were performed, either produced intermittently (tapping or pausing prior to each cycle) or continuously. Again, an increase in variability associated with cerebellar pathology was restricted to the intermittent conditions, consistent with the predictions of the event timing hypothesis.

The initial work on the event/emergent dichotomy focused on the idea that distinct psychological representations and their associated neural systems underlie control of superficially similar tasks. One question that was ignored in this earlier work concerns the interaction between internal control processes and externally-defined task goals. In particular, how do participants establish the initial movement rate? For event timing, this is relatively straightforward: a target rate, indicated by a metronome, defines the target interval for the output of an internal timing system. This mapping process is consistent with claims that the cerebellum provides a common computation for perception and action. The metronome beats define salient temporal events and the cerebellum is important in the representation of these intervals, either for tracking the metronome, producing responses timed to be synchronized with these events (paced tapping), or for producing responses that mark similar events once the metronome is terminated (unpaced tapping).

But how is an externally defined temporal goal achieved when timing is emergent?

People are generally quite good at producing rhythmic circles at a rate that closely matches a metronome-defined goal. Zelaznik *et al.* (2005) proposed a transformation hypothesis in which the metronome guides an initial event-based representation. The participant adopts a movement rate based on previous experience and adjustments are instituted following a comparison of the produced and desired rate. Thus, timing during this initial phase is hypothesized to be event-based. The transformation to emergent timing for continuous movements occurs once the produced rate approximates the desired rate. At this point, the control system would no longer need to refer to an internal temporal signal. Rather, the dynamics could be sustained to keep the movement cycle constant at the metronome-defined rate. For example, in circle drawing, continuous variation of a spatiotemporal oscillator would suffice to maintain consistent timing by keeping angular velocity constant. Such a control scheme would not work for movements with discontinuities.

Zelaznik *et al.* (2005) reported evidence in support of the transformation hypothesis. Using an individual difference approach with neurologically healthy individuals, they found that temporal variability was positively correlated between tapping and circle drawing on the first cycle of a repetitive movement. For subsequent intervals, the correlation was not significant. Thus, the initial cycle of circle drawing appears to rely on a timing process common to finger tapping. By the transformation hypothesis, this common process would be event-based timing.

Given our assumption that the cerebellum is critical for event timing, two predictions can be derived from the transformation hypothesis with respect to the performance of individuals with damage to this structure. First, if movement timing is initially event-based, then participants with cerebellar damage should show elevated temporal variability when drawing continuous circles on the first interval (or intervals). We test this prediction in the current experiment, calculating variability for each interval (interval 1 to interval 26) across trials. Note that this procedure is quite different than that adopted in most studies where variability is calculated within a trial (based on the values for each interval) and then averaged across trials.

Second, we speculate that the presence of a metronome signal may induce event-based timing, given that the tones define singular events for each cycle. While Spencer *et al.* (2003) used a metronome in some of their experiments, these data were not analyzed. In the present study, participants performed trials in which the metronome was present for all 26 intervals (paced trials) or in which the metronome was only

present for the first 7 intervals (partial-paced trials). If our speculation is correct, individuals with cerebellar damage should be impaired during the metronome-based intervals during both movement conditions. By including both paced and partial-paced trials, we are able to directly compare temporal variability at similar time points within a trial rather than confound paced and unpaced phases with the initial and latter phases of a trial.

2.2 Methods

2.2.1 Participants

Nine individuals with bilateral cerebellar degeneration (ataxics) were recruited for this study (Table 2.1). For five of these participants, genetic testing confirmed a variant of spinocerebellar ataxia (SCA3: confirmed one participant, familiar history indicates diagnosis likely in another; SCA6: three participants). For the remaining four, the etiology of their cerebellar degeneration was unknown. The ataxics were all evaluated with the International Cooperative Ataxia Rating Scale (ICARS, see Trouillas *et al.*, 1997). The mean score (see Table 2.1) was 40.7, with a range of 17.5 (mild to moderate) to 60 (severe).

MRI or CT scans were reviewed for all of the participants and confirmed that there was evidence of significant atrophy in the cerebellum. The extent and distribution varied considerably across the group. There was no evidence of significant atrophy in extra-cerebellar regions, although MRI and post-mortem studies of patients with certain SCA subtypes (e.g., SCA3) have reported pathology in the basal ganglia and brainstem (Klockgether *et al.*, 1998; Rüb *et al.*, 2005). While we expect that some of our participants may have extra-cerebellar pathology, the pathology is most evident in the cerebellum in accord with their clinical presentation.

A control group, on average matched in age and education to the ataxics, was also recruited. These individuals reported no history of neurological or psychiatric problems.

All participants provided informed consent, and were compensated for their time. The protocol was approved by the Committee for the Protection of Human Subjects at the University of California, Berkeley.

Table 2.1: Population Data

Subject	Diagnosis	Gender	Age	Years since Diagnosis	Education	ICARS Total	Ataxia, Upper	Ataxia, Lower
P1	Atrophy	M	77	32	12	45	12.75	5.5
P2	SCA6	F	57	1	16	17	4	2
P3	Atrophy	M	68	15	17	42.75	13.5	8.5
P4	Atrophy	M	67	8	20	17.75	5	3
P5	SCA3	M	51	11	18	49.75	16.5	8.5
P6	Atrophy	M	54	10	14	34.5	4	2
P7	SCA6	M	52	6	16	21.5	5	0
P8	SCA6	F	46	12	16	54	16	4
P9	Atrophy/SCA3	M	72	8	12	60	18	3
C1	Healthy	F	75		19			
C2	Healthy	M	66		16			
C3	Healthy	F	65		18			
C4	Healthy	M	70		21			
C5	Healthy	F	61		16			
C6	Healthy	M	53		14			
C7	Healthy	M	52		12			
C8	Healthy	F	46		14			
C9	Healthy	M	72		12			
Patient	**Average:**		60.4, SD 10.7	9, SD 8,7	15.7, SD 2.6	40.7, SD 15.1		
Control	**Average:**		62.2, SD 10.0		15.8, SD 3.1			

2.2.2 Task and Procedure

Participants produced rhythmic finger tapping and continuous circle drawing movements. For finger tapping, movements involved flexion-extension of the index finger with the finger contacting a table surface at the end of the flexion phase. For continuous circle drawing, movements primarily involved rotation about the shoulder and elbow, similar to that employed in our previous studies (Spencer *et al.*, 2003). The participant held a cylindrical manipulandum (1 cm diameter, 3 cm tall) between the thumb, index and middle fingers, with the instructions to keep the tip of this object in contact with the table surface at all times. A 7 cm diameter circle was taped to the table surface to serve as a template for this condition. As in previous studies, the instructions emphasized that this template was to provide a guide for movement amplitude and that it was not necessary to carefully trace along the circumference.

The target cycle duration for all movements was 900 ms and a complete trial consisted of 26 movement cycles. The target interval was specified by an auditory metronome. This metronome was played at the start of each trial and the participant began to move when the rate was internalized (usually within 2 or 3 tones). We manipulated the number of cycles for which the metronome was presented following movement onset. In the paced condition, the metronome was present throughout the trial. In the partial-paced condition, the metronome was present until the participant produced seven movement cycles (seven circles or eight taps) and was then terminated for the remaining cycles. Thus, there were a total of four conditions, created by the factorial combination of movement type (tapping or circle drawing) and metronome status (paced or partial-paced).

The test session began with practice trials for two of the four conditions: finger tapping in the paced condition and continuous circle drawing in the partial-paced condition. Participants received practice trials until they understood both the required movements as well as the metronome conditions. Then each participant completed four blocks of trials (one per condition), with each block consisting of 10 trials. The order of the movement types (finger tapping, circle drawing) was counterbalanced across participants. Within a movement type, the two metronome variants were always tested in the order: paced then partial-paced. Participants were also tested with a third type of movement in which the circle drawing was performed in a constrained space (a grooved circular ring), with the idea that this apparatus might improve performance by reducing the demands on trajectory control. However, this device led to

an increase in variability for all participants, perhaps due to the effects of friction or corrections generated when the hand contacting the edge of the groove. As such, we will not report these data.

Movements were recorded by a 3-dimensional kinematic tracking system (Ascension mini-Bird, http://www.ascension-tech.com). A magnetic transmitter was located below the table, and a lightweight sensor was secured with tape to the back of the index finger of the dominant hand, with the wire from the sensor secured loosely to the arm. The position of the sensor was sampled at a mean rate of 145 Hz. This sampling rate fluctuates slightly across trials, although not within trials.

2.2.3 Data analysis

Movement trajectories for each trial were analyzed using custom programs in MATLAB. First, the trajectories were smoothed using a 30-Hz Butterworth low-pass filter. For tapping, the duration of each cycle was then defined as the interval between successive contact points with the table surface. These contact points were identified with a velocity criterion, set as the first sample in which velocity fell below 3% of the maximum downward velocity for that flexion phase. For circle drawing, cycle duration was defined as the interval between successive crossings of the point on the circle most distant from the participant. In order to accommodate the impaired spatial trajectories of individuals with cerebellar degeneration, this point was identified on each cycle by a velocity criterion, and taken as the first point at which the velocity in the y-dimension fell below 3% of the maximum velocity along that axis. Using these criteria, the trajectories were marked into individual cycles using a semi-automated method with manual oversight (a procedure similar to that used in Zelaznik *et al.* 2005).

The data from each trial were divided into two phases, early and late. The early phase consisted of the first seven intervals. The metronome was present during these early intervals for both the paced and partial-paced conditions. Because preliminary analyses revealed no differences between the two conditions, we combine the early phase data across the two conditions to increase statistical power in the analyses. The late phase consisted of the remaining intervals (approximately 19 intervals). We report the late phase data separately for the paced condition (with metronome) and the partial-paced condition (without metronome).

For each trial, the mean and standard deviation of the cycle durations were cal-

culated separately for the early and late phases. Note that for the paced condition, the division between early and late is arbitrary; for the partial-paced condition, this division corresponds to presence or absence of the metronome. To adjust for variation in the mean duration, the coefficient of variation (CV, standard deviation divided by the mean) was used as the measure of temporal variability. The CV was calculated independently for each trial and then averaged across the trials for a given condition. Intervals exceeding +/-50% of the mean value for a given trial were excluded.

As a further test of the transformation hypothesis, we also calculated the mean and CV scores for each cycle defined by ordinal position (e.g., mean and CV for cycle 1, cycle 2, etc.; see Zelaznik *et al.* 2005). For this calculation, the data are tabulated across trials rather than within a trial. Variability can be considerably inflated here if the overall rate varies significantly from trial to trial. In order to reduce effects caused by this form of trial-by-trial variation, cycle durations were normalized by dividing each interval by the mean cycle duration for that trial. The transformation test can be applied at two points in the partial-paced condition: at the start of the trial and when the metronome is turned off. As the trials for the paced and partial-paced conditions begin identically, we again combined the data across the paced and partial-paced conditions in the individual interval analysis of the early phase. This doubled the number of possible observations per participant for each interval (maximum of 20).

2.3 Results

Trials were excluded from analysis due to computer error (0.5% of trials for control participants, 1.3% of trials for ataxics) or the presence of extended pauses between cycles (stopping for more than 50 ms during continuous circle drawing; 0.5% of trials for control participants, 3% of trials for ataxics). Of the remaining trials, individual cycles were excluded if the duration was greater than 150% or less than 50% of the mean cycle duration. This criterion caused the exclusion of 45 cycles for the ataxic group (approximately 0.5% of the total number of cycles; 6 during circle drawing, 39 during tapping) and two cycles for control participants (both tapping). After applying these various screening criteria, all trials had at least 22 cycles for the subsequent analyses.

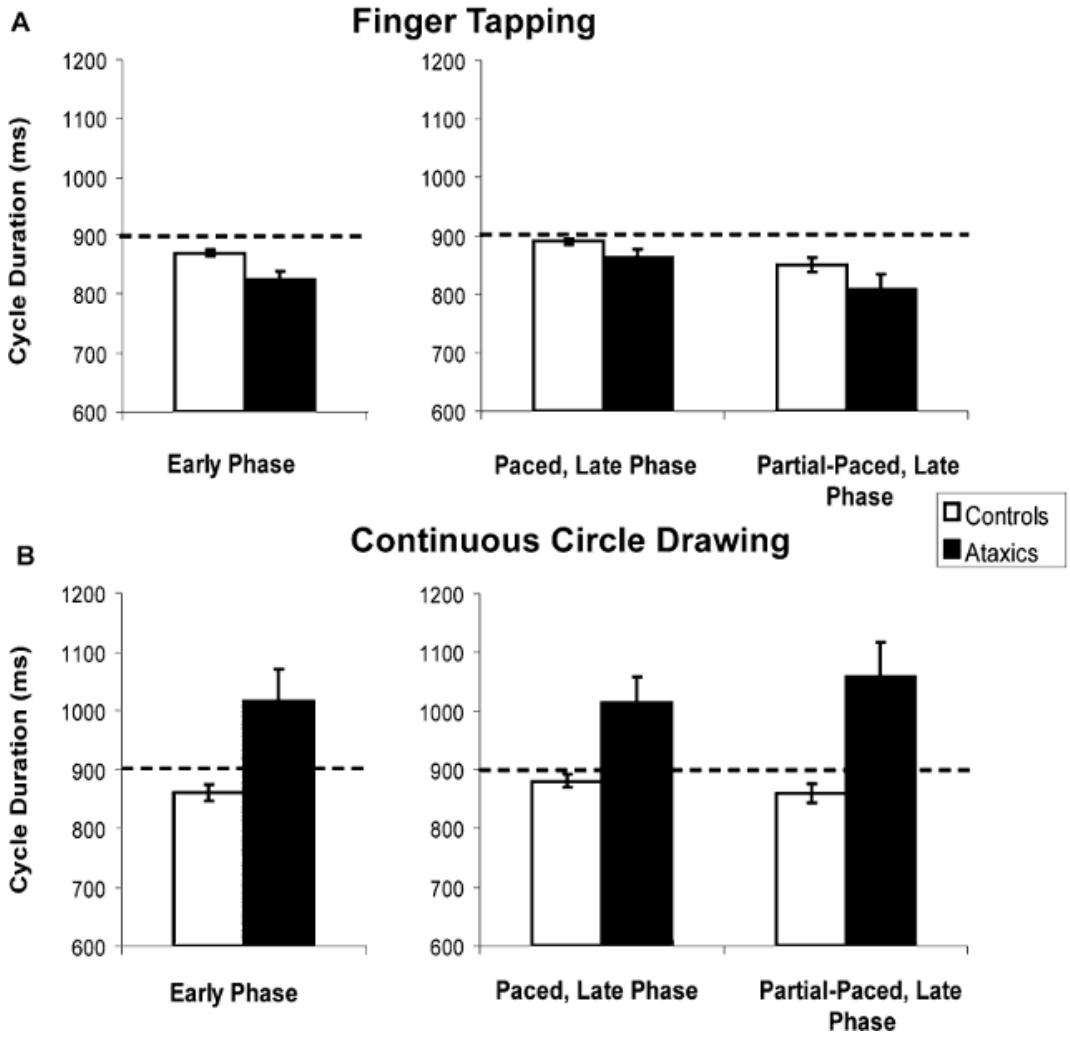

Figure 2.1: Cycle duration. Plotted are the average cycle durations, in milliseconds, across participants in the control group (white bars) or ataxic group (black bars). Error bars reflect standard errors. (A) Finger tapping. Ataxics are consistently faster than controls both during the early phase (first seven intervals) and the late phase (intervals 8 through 26) of both conditions. (B) Circle drawing. Ataxics are consistently slower than controls in all phases.

2.3.1 Cycle duration

Figure 2.1 shows the mean cycle durations. Control participants generally produced mean cycle durations that were well-matched to the target rate of 900 ms. The mean produced duration for the controls is slightly faster than the target rate during the early phase of the trials. On some trials, the movements were initiated at a faster rate and subsequently slowed to the target rate over the first few cycles. The ataxics exhibited considerably more variation with some systematic deviations. During finger tapping, the ataxics moved more quickly than the target rate, both when the metronome was present ($t[16]=2.522$, $p=0.023$) and when it was absent ($t[16]=2.099$, $p=0.052$). However, when drawing circles, they moved considerably slower than the target rate (paced condition: $t[16]=3$, $p=0.008$; partial-paced condition: $t[16]=3.342$, $p=0.004$). It is possible that the target rate was faster than the ataxics were capable of producing; however, we do not expect that this provides a complete account of these data. In pilot work, we have observed a similar lack of correspondence between the target and produced rates even at slower and self-chosen rates.

Given that the early phase consisted of just seven cycles and there were strategic differences in how participants initiated the trials, our statistical analysis of cycle duration is restricted to the late phase. Note that the data from this phase come from the same ordinal positions for the paced and partial-paced conditions. A 2x2x2 ANOVA was conducted with within-subject factors of movement type (tapping vs. circle drawing) and metronome (present vs. absent) and the between-subject factor of group (control vs. ataxic). Significant main effects were observed for group ($F[1,16]=6.32$, $p=0.023$) and movement type ($F[1,16]=13.86$, $p=0.002$), while the effect of the metronome was marginally reliable ($F[1,16]=4.435$, $p=0.051$). Movement type interacted with group ($F[1,16]=13.89$, $p=0.002$) and metronome ($F[1,16]=13.65$, $p=0.002$) and the three-way interaction was also significant ($F[1,16]=6.03$, $p=0.023$). This interaction reflects the fact that the mean cycle duration for the patients was closer to the target rate when the metronome was present: during tapping this led to a slowing down of movement rate, and during circle drawing this led to a speeding up of movement rate. Thus, while the ataxics were not able to match their movement rate to the metronome as well as the control participants, their performance was influenced by the metronome. In the late phase of the trials, the mean cycle duration for the ataxics on both tasks was closer to the target rate when the metronome was present.

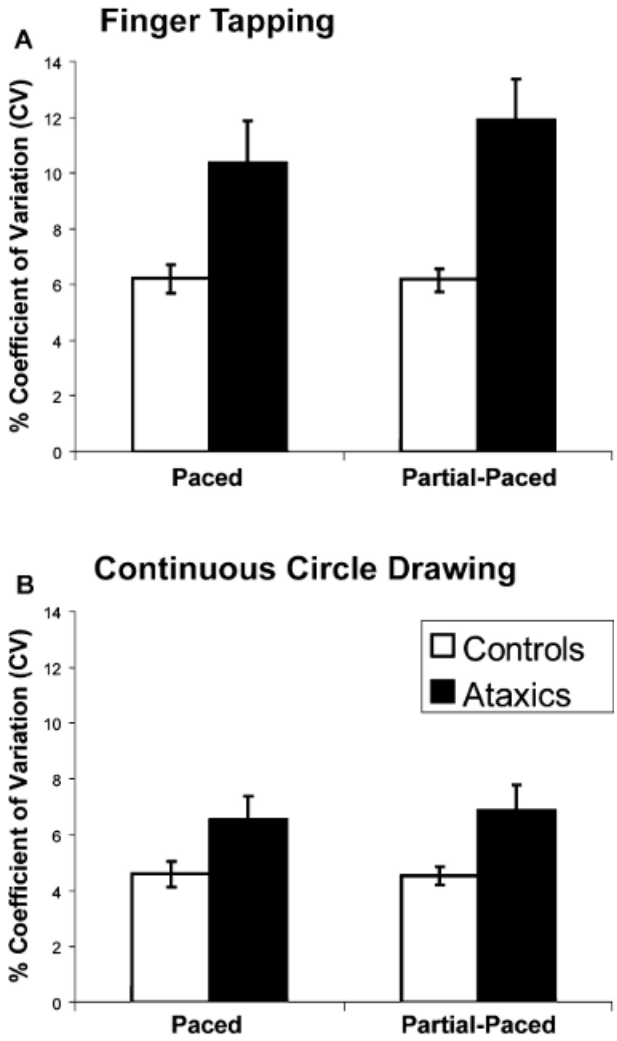

Figure 2.2: Trial variability during the late phase. Plotted are the average coefficients of variation (CV), calculated for each trial, for controls (white bars) and ataxics (black bars). Error bars reflect standard error. Ataxics are consistently more impaired during both finger tapping (A) and continuous circle drawing (B), though the impairment is more severe during finger tapping than circle drawing. The presence of the metronome during the paced condition lowered variability during finger tapping.

2.3.2 Event vs. Emergent Timing: Within-Trial Variability

Figure 2.2 presents the variability data calculated across cycles within a trial and then averaged over trials. We present two panels for each movement type: late phase of the paced condition, and late phase of the partial-paced condition. The late phase data allow a comparison of temporal variability over similar cycles (intervals 8-26) as a function of whether or not the metronome is present. To reduce the contribution of global drift in performance (e.g., speeding up or slowing down), the data were detrended by fitting a regression line to the data and measuring variability relative to the predicted value of this regression line.

The current design contains two of the key conditions reported in Spencer et al. (Spencer *et al.*, 2003) that provided neuropsychological evidence in support of the event/emergent distinction: unpaced rhythmic tapping and circle drawing following synchronization with a metronome. As can be seen in the figure, the results provide a partial replication. Consistent with Spencer et al., the ataxics exhibit a greater impairment than controls during finger tapping compared to circle drawing, in accord with the hypothesis that the former relies on event timing. However, unlike Spencer et al., the ataxics were also significantly more variable than controls during continuous circle drawing. The same pattern is found during the late phase of the paced conditions. Here, too, participants in the ataxic group were more variable than control participants and this difference was larger (both in absolute percentage and proportionally) during tapping than circle drawing.

To statistically evaluate these effects, we performed the same 3-way ANOVA as for the cycle duration data, using the factors movement type, metronome, and group. The main effects of task (tapping vs. circle drawing) ($F[1,16]=565.0$, $p<0.001$) and group ($F[1,16]=10.6$, $p<0.005$) were reliable. More important, the interaction of these two factors was significant ($F[1,16]=511.9$, $p<50.003$); while the ataxics were more variable than the control participants on both tasks, their increase in temporal variability was significantly greater during tapping.

The controls were unaffected by the presence or absence of the metronome. In contrast, ataxics were less variable when the movements were paced by the metronome, resulting in a reliable group by metronome interaction, ($F[1,16]=6.84$, $p=0.019$). Although the improvement during the paced portion appears to be restricted to tapping, the three-way interaction was not significant ($F[1,16]<1$). The fact that the performance of the ataxic group improved in the presence of the metronome is at odds with

one prediction derived from the transformation hypothesis. If the metronome induced eventbased timing, we expected the opposite pattern: a disproportionate increase in variability when drawing circles with a metronome.

2.3.3 Event vs. emergent timing: Between-trial variability

A stronger test of the predictions of the transformation hypothesis requires a microanalysis of the variability of individual intervals. As shown by Zelaznik *et al.* (2005), only variability from the first interval during circle drawing was significantly correlated with tapping variability (on all intervals). These results were interpreted to indicate that an event-based temporal representation guides performance initially, allowing the performer to match the target interval of the metronome, before transitioning to control in which timing is emergent. Following this logic, we expected that individuals with cerebellar pathology would exhibit increased temporal variability when compared with control participants on the first interval in both tapping and circle drawing. The impairment during circle drawing should become attenuated over successive trials; in contrast, it should remain constant during tapping.

To assess this prediction, we calculated variability across trials on an interval-by-interval basis for the initial seven paced cycles (again, combining across the paced and partial-paced conditions to improve statistical power). Figure 2.3 shows these data for tapping (a) and circle drawing (b). For both conditions and for both groups, variability is highest on the first interval and then decreases to asymptotic levels by about the third interval. Contrary to the prediction of the transformation hypothesis, there is no indication that the ataxics were most impaired on the first interval during circle drawing. Indeed, in terms of mean values, the difference between the ataxic and control groups was actually smallest for this interval. When these data were statistically evaluated, a 2 (group)x2 (movement type)x7 (interval) ANOVA revealed a significant effect of interval ($F[1,16]=34.56$, $p<0.001$), but this factor did not interact with group, movement type, or the predicted 3-way interaction. Indeed, the only reliable interaction was between group and movement type ($F[1,16]=12.90$, $p=0.002$), reflecting once again that the ataxics were more variable during tapping.

Given that the CV was not equal during tapping and circle drawing (e.g., 19), we also compared the groups in terms of the percent increase in variability on an interval-by-interval basis. On this measure, the transformation hypothesis would predict an interaction, where the ataxics become less impaired relative to controls over the course

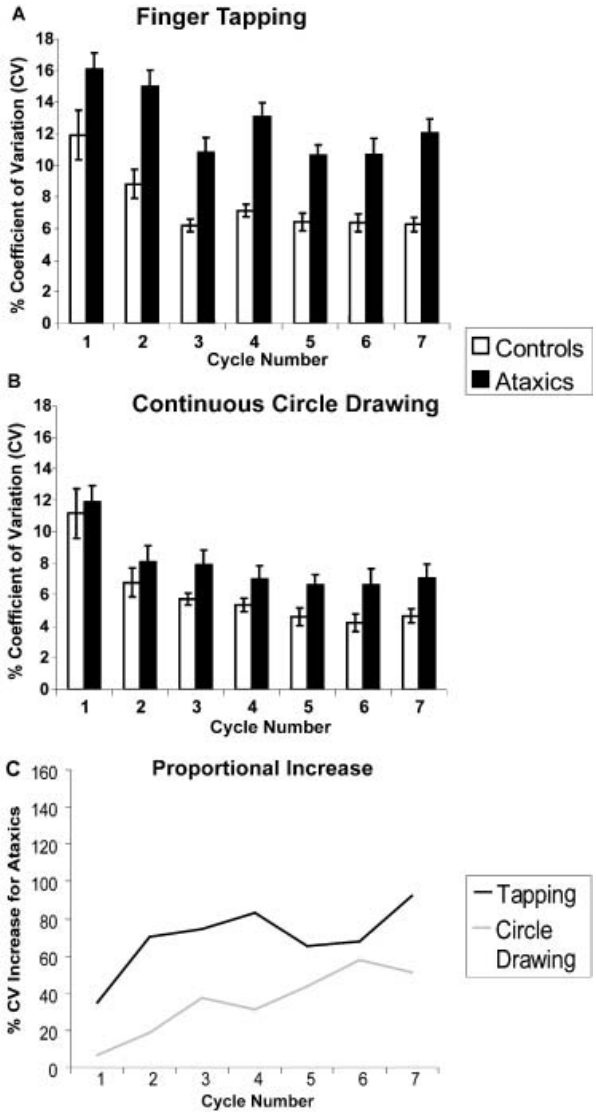

Figure 2.3: Interval variability during the early phase. CVs are calculated for each interval across trials following procedures detailed in the text. Paced and partial-paced conditions are combined, as the metronome is present for both conditions. Error bars reflect standard error of the mean. (A) Ataxics (black bars) are consistently impaired at all intervals during finger tapping. (B) Variability for circle drawing decreases for both groups after the first cycle, but ataxics remain impaired. (C) Plotted is the proportional impairment for the ataxic group relative to the control group. Unexpectedly, finger tapping (black lines) shows an increasing impairment over subsequent intervals. Contrary to the prediction of the transformation hypothesis, the proportional impairment on circle drawing (gray line) is lowest for the first interval.

of a circle drawing trial. As can be seen in Figure 2.3c, the average divergence between the performance of ataxic and control participants actually increased over the course of these initial seven intervals. For circle drawing, the initial increase in variability was only 7%; by the seventh interval the increase was up to 51%. Thus, the ataxics became proportionately more variable over successive intervals, opposite of the prediction of the transformation hypothesis. A similar trend was observed for tapping, with an initial 35% deficit rising to 92% by the seventh interval.

Finally, the transition from paced to unpaced movements in the partial-paced condition may offer another look at the transformation from event to emergent timing. Figure 2.4 plots the percent increase in CV for the ataxic group during the late phase, either with or without the metronome. We combined the data across pairs of cycles here since there are, at most, only 10 trials for each cycle. Again, there is no indication of a pronounced rise in variability during circle drawing when the metronome is turned off (Figure 2.4b). Rather, the timing deficit for these participants is relatively constant across the trial. The difference between the two types of movements was less pronounced in the late phase of the paced condition, in large part because the performance on tapping improved when the metronome was present.

2.4 Discussion

The cerebellar timing hypothesis was proposed as a system-level functional characterization of this subcortical structure. In its original formulation, the emphasis was on describing the functional domain of the cerebellum (e.g., Ivry & Keele, 1989). The ability to control the fine timing between successive gestures is a fundamental prerequisite for skilled movement. Similarly, precise timing is essential for certain forms of sensorimotor learning. While the timing hypothesis has offered a parsimonious account of cerebellar function over a range of task domains, there remains much to be understood about both the underlying mechanisms of internal timing and the specificity of the cerebellum in timing. Indeed, the modular perspective explicitly motivating the timing hypothesis, that of a specialized system that is accessible to disparate neural systems requiring this form of representation, has been challenged by recent work suggesting that temporal coding may be a more local process (Brody *et al.*, 2003; Leon & Shadlen, 2003; Roux *et al.*, 2003). By this latter hypothesis, temporal representation may be a ubiquitous feature of neural processing, emergent

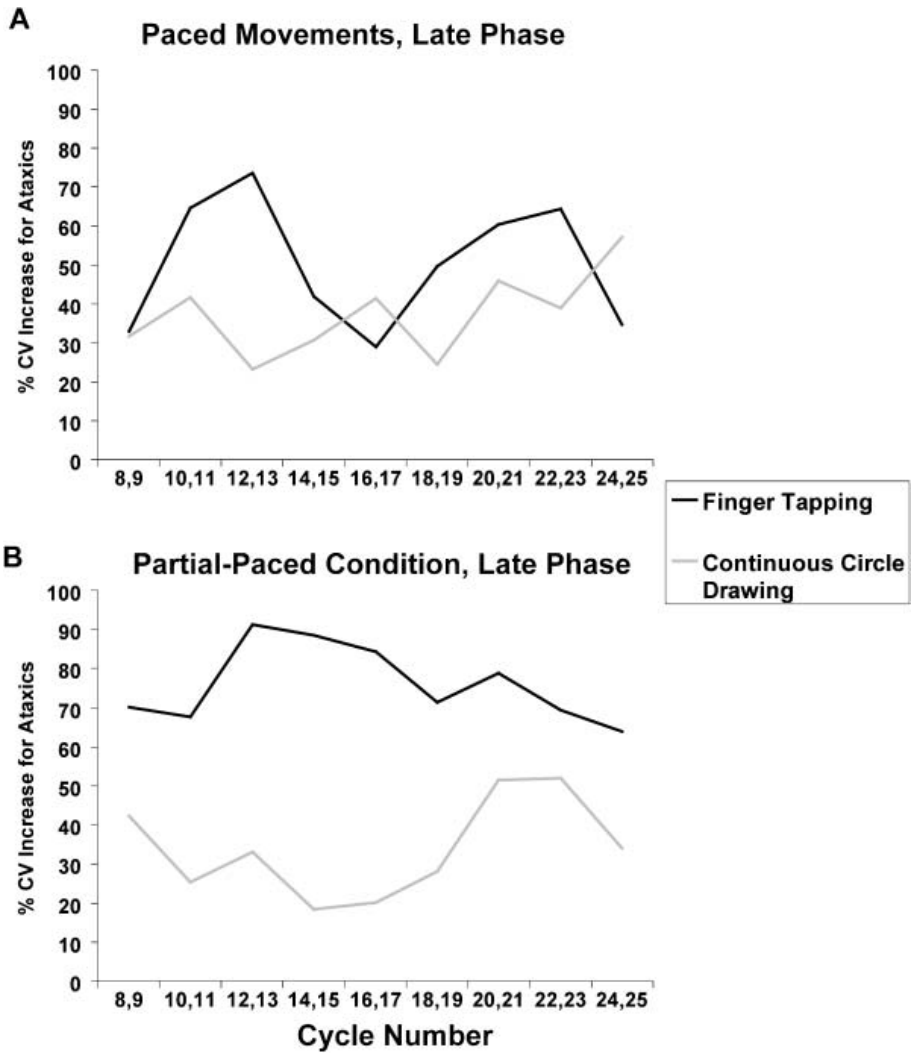

Figure 2.4: Interval variability during the late phase. Plotted are the proportional impairments for the ataxic group relative to the control group for intervals 8 through 25. To account for the slim data (10 trials per condition), CV was averaged across pairs of intervals. (A) During paced movements, ataxics are more impaired during finger tapping for most intervals, although there is some overlap in the curves. (B) During the partial-paced condition, in which the metronome is not present, patients are more impaired during finger tapping than during circle drawing at all intervals.

in domain-specific representations. For example, the duration of a visual stimulus might be incorporated in the representation of the stimulus itself.

2.4.1 Event Timing and the Cerebellum

We have recently proposed a hybrid account of timing in which the functional domain of the cerebellum is limited to tasks in which an explicit representation of time is essential to the task goal. This event timing model was motivated by correlational work indicating that distinct processes were associated with the control of different classes of rhythmic movements (Zelaznik *et al.*, 2005; Robertson *et al.*, 1999; Zelaznik *et al.*, 2000, 2002). These findings suggested that, for movements marked by salient features, control might involve a process that encodes the requisite time between successive events. In contrast, when such features are absent, timing may be emergent, secondary to a more continuous form of control. Spencer et al. (Spencer *et al.*, 2003) reported that damage to the cerebellum selectively impaired performance on tasks linked to event timing, providing neuropsychological evidence for the role of the cerebellum in event timing.

Although secondary to the main purposes of the current study, we included tasks that were similar to those used in Spencer *et al.* (2003). The results provide a partial replication. The participants with cerebellar degeneration exhibited a disproportionately greater impairment during finger tapping compared to circle drawing. The current results further extend our understanding by showing that this effect is not limited to movements performed in the absence of a metronome. Compared to controls, the relative impairment during finger tapping was evident both with and without a metronome.

Contrary to Spencer *et al.* (2003), the individuals with ataxia were also more variable than control participants during continuous circle drawing (see also, Bo *et al.*, 2005). This result seems especially puzzling given that many of the ataxic participants (6 of 9) in the currnt study had also participated in Spencer et al. The disease process has advanced in some of these individuals and this may contribute to their greater variability in the current study. However, upon closer examination, the pattern of results is similar, despite the different statistical outcomes. In Spencer et al., the ataxic group showed a 13% increase in variability on the circle drawing task and a 56% increase during tapping compared to controls. This interaction was significant and a post-hoc contrast restricted to circle drawing was not reliable. In the current study,

the percent increase in variability is 52% and 92% for circle drawing and tapping, respectively, with the interaction again significant as is the paired contrast for each task. Thus, the results from both studies emphasize a disproportionate impairment during tapping, consistent with the event timing hypothesis.

In Spencer *et al.* (2003), the more dramatic dissociation of event and emergent timing was obtained in the within-subject comparison involving patients with unilateral lesions. The absence of any indication of a difference between the impaired and unimpaired limb during circle drawing led to the strong claim that the cerebellum was not involved in the control of continuous rhythmic movements. The current results suggest that this claim should be reconsidered. Alternative characterizations of cerebellar function, such as hypotheses suggesting putative roles for the cerebellum in forward modeling (e.g., Blakemore *et al.*, 2001; Diedrichsen *et al.*, 2005a; Kawato *et al.*, 2003; Wolpert *et al.*, 1998), state estimation (e.g., Paulin, 2005; Scott, 2004), and error correction (e.g., Diedrichsen *et al.*, 2005b), may suggest ways in which the cerebellum contributes to the control of both discrete and continuous rhythmic movements. However, it remains incumbent to explain why the cerebellum should be more heavily taxed during tapping compared to circle drawing. The emphasis on an essential role for the cerebellum in event timing offers one hypothesis (Spencer *et al.*, 2003; Diedrichsen *et al.*, 2005b).

2.4.2 The Transition to Emergent Timing: Tests of the Transformation Hypothesis

We still do not understand how a rate is established in tasks which do not require the event timer. While consistent timing can be achieved during circle drawing by maintaining a constant angular velocity, the adopted velocity must produce cycles matched to the duration specified by a metronome. Zelaznik *et al.* (2005) proposed that there is an initial transformation between event-based and emergent timing. Within the context of rhythmic movements, the initial movement cycle(s) may be compared to representations of an explicit temporal goal. Once an acceptable match is achieved, observably accurate timing could emerge from a secondary control parameter, at least in the absence of a metronome.

The primary goal of the current study was to test the transformation hypothesis, given the assumption that the cerebellum contributes specifically to event timing.

As such, we expected that during circle drawing, the ataxics would exhibit increased variability during the early cycles of a trial compared to the latter cycles. We failed to obtain support for this prediction. The ataxic group did not show dispro-portionately elevated variability on the first interval; in fact, their impairment increased over successive cycles, similar to what was observed during tapping.

We do note that the best test of the transformation hypothesis comes from the first interval. However, obtaining a reliable estimate of temporal variability for this interval is problematic. While we instructed the participants to begin moving when they had internalized the target duration, there are considerable fluctuations in how participants initiated their movements from trial to trial. Moreover, fluctuations in the internal representation of the target rate make it more difficult to obtain reliable estimates in a between-trial analysis compared to a within-trial analysis. In the latter, the mean of the central representation is assumed to be fixed for a given trial and noise about this mean reflects one source of temporal variability. Systematic changes within a trial (e.g., drift) can be discounted by simple analytic tools (e.g., linear detrending). These tools do not appear to be appropriate for a between-trials analysis, even though there are likely shifts in the mean between trials. Thus, it is more difficult to estimate between-trial variability because the observed intervals might reflect variation about a constant mean, shifts in that mean between trials, or some combination of the two. We attempted to correct for this by normalizing the mean duration on a trial-by-trial basis. Nonetheless, our data set is considerably smaller, and likely less reliable, than that used in Zelaznik *et al.* (2005).

We also failed to obtain support for the transformation hypothesis in a second, within-trial analysis. If we assume that the metronome induces an eventbased representation, then we expected that the ataxic group would become more variable when performing the circle drawing task with the metronome compared to when the metronome was absent. Contrary to this prediction, variability for the ataxics decreased when the metronome was present for both tapping and circle drawing. It is interesting to note that control participants do not typically show a reduction in temporal variability during paced tapping compared to unpaced tapping. In fact, healthy individuals usually become more variable due to the operation of an error correction process (Pressing, 1998; Vorberg & Wing, 1996; Helmuth & Ivry, 1996), an effect that was present, although not reliably so, in the current study. Previous work has shown that error correction processes that operate during rhythmic movements are

unaffected by cerebellar pathology Diedrichsen *et al.* (2005a); Molinari *et al.* (2005). While a metronome leads to an increase in variability in healthy individuals, the observed improvement for the ataxic group during paced movements suggests that intact error correction and/or other processes influenced by the metronome may help offset increases in variability that arise from cerebellar pathology.

2.5 Conclusion

In summary, the predictions of the transformation hypothesis were not supported. At least two possibilities exist. First, the data set may not have been sufficiently robust to allow for a cycle-by-cycle analysis, the strongest test of the transformation hypothesis. Second, contrary to the assumption of Zelaznik *et al.* (2005), the correlations observed between tapping and the initial cycle during circle drawing may not have reflected the use of an event-based representation in both tasks. Rather, the correlations may reflect the operation of a different, shared process that is especially relevant at movement onset. By this view, the current data indicate that the cerebellum is not involved with these initiation processes (Spencer *et al.*, 2005).

Note

In Zelaznik *et al.* (2005), the initial cycle was produced in the absence of a metronome; in the current study, this interval was produced with a pacing metronome. We have conducted a study in which the metronome was terminated prior to the first movement cycle. Under those conditions, we also failed to observe a selective impairment on the initial cycles during circle drawing and between-trial variability was considerably larger (Schlerf *et al.*, 2005).

Chapter 3

Task goals influence online corrections and adaptation when reaching in an unstable environment

Abstract

Everyday movements often have multiple solutions. Many of these solutions arise from biomechanical redundancies. Often, however, the goal does not require a unique movement. To examine how people exploit task-related redundancy, participants produced three-dimensional reaching movements, moving to one of two rectangular targets that were diagonally oriented in the frontal (X,Y) plane. On most trials, the movement was perturbed by a velocity-dependent force in the vertical (Y) direction. Since participants were free to move in 3D space, online corrections could involve movement along the perturbed Y dimension, as well as the non-perturbed lateral (X) direction. If the motor system exploits task redundancies, then corrections along the lateral dimension should depend on the orientation of the target. Consistent with this prediction, participants modified both the X and Y coordinates of the trajectory over the course of learning, and the lateral component was sensitive to the orientation of the target. Furthermore, participants produced online corrections with a lateral component that brought the hand closer to the target. These results suggest that we

not only correct for mismatches between expected and experienced forces, but also exploit task-specific redundancies to efficiently improve performance.

3.1 Introduction

Reaching to grasp an object can usually be accomplished by a large set of kinematic patterns and final joint configurations, the so-called degrees of freedom problem. Despite the potential selection problem created by an excess of degrees of freedom, the motor system effortlessly chooses a motor command which achieves the desired end position. Optimal control theory has provided a formal framework for understanding the constraints underlying this process (Todorov & Jordan, 2002). For example, certain combinations of changes about the muscles and/or joints are less costly require less energy expenditure or increase end-state comfort than others. We thus perform optimally when we select an action with the lowest cost.

The principles of optimal control can also be used to understand how people make adjustments to ongoing movements (Diedrichsen *et al.*, 2010). Assuming that there is a control cost to such corrections, on-line adjustments should be most evident when they help ensure the desired task outcome. Conversely, adjustments should be reduced for deviations that are irrelevant to the task outcome. Consistent with these expectations, greater variability is observed along a dimension that is irrelevant for task outcome (a redundant dimension) than along a dimension which is crucial for task outcome (Todorov & Jordan, 2002). This idea is also encompassed in the notion of an "uncontrolled manifold," which proposes that variability is allowed among the set of coordinates in task space (or manifold) for which the task outcome is equivalent (Cusumano & Cesari, 2006; Scholz & Schöner, 1999).

Redundancy does not solely arise because of the excessive degrees of freedom in the motor system. Many tasks entail goals that afford redundancies unrelated to the biomechanics of our limbs. When closing a door, for example, the same force can be applied anywhere along the vertical axis of the door with equivalent results. Similarly, to increase stability, we can grasp a stairway railing at multiple locations. In such situations, we should expect that an optimal planning system would allow greater variability along redundant dimensions defined in task-space. Indeed, when learning a new skill, people have been shown to exploit such redundancies, producing greater variability along task-irrelevant dimensions compared to task-relevant dimen-

sions (Cusumano & Cesari, 2006; Müller & Sternad, 2004).

While we know task-based redundancy affects learning over many trials, the role of this form of redundancy in online control of well-learned skills has received less attention. This issue is important, not only when considering movement execution, but also for planning and learning. The stability of motor performance is sensitive to many factors. Fatigue, injury, clothing, gravity, and the current posture all affect how a limb responds to a neural command. Despite these multiple sources of variability, we manage to move with comparable proficiency across an extensive range of conditions. This robust performance requires adjustment of the neural command to fit the context, indicating that the motor system is highly adaptive. While adaptation has been the focus of a substantial body of literature, such studies have generally been limited to conditions in which the task goal is defined as a single point (with some tolerance) in a two-dimensional workspace. For example, in studies involving force field perturbations (Lackner & Dizio, 1994; Shadmehr & Mussa-Ivaldi, 1994; Taylor & Thoroughman, 2007; Thoroughman & Shadmehr, 2000) or visuomotor transformations (Fishbach & Mussa-Ivaldi, 2008; Mazzoni & Krakauer, 2006; Sober & Sabes, 2003; Tseng *et al.*, 2007), the task goal is defined by a target location in 2-d space. In these contexts, adaptation requires the adjustment of an internal model to counteract the effect of the perturbation such that the movement terminates in the vicinity of the target location. While optimal control models have provided elegant accounts of learning under such conditions (Fishbach & Mussa-Ivaldi, 2008), the focal nature of the targets in such studies preclude the analysis of whether adaptive processes exploit task-based redundancies. The goal of the current study was to address this problem.

Dimensional redundancy is not present in the typical force field study such as when participants make planar movements in a viscous curl field. The force field involves two dimensions, and thus there is no irrelevant dimension. To introduce redundancy, we had participants reach in a 3-dimensional workspace (see Figure 3.1). We presented the target as a rectangle, oriented diagonally on a virtual surface in the frontal (X,Y) plane. Contact at any point within the rectangular region was considered a successful reach. Using this target instead of a single point made the extraneous third dimension relevant to task performance. To evaluate whether on-line feedback and adaptive processes incorporate information regarding task-based redundancies, we introduced a consistent force perturbation during the movement. This perturbation was restricted to the vertical dimension, thus displacing the hand from the target at an oblique

angle. Any lateral component of the correction is irrelevant to force adaptation (e.g., does not cancel out the perturbation), but nonetheless, remains relevant to task performance.

We focused on how participants learned to respond to this perturbation. If learning involves generating an accurate model of the environmental perturbation, then the participants' behavior should be independent of the orientation of the target. That is, we would expect to observe an anticipatory trajectory that counteracts the perturbing effects of the force field. Alternatively, learning may incorporate task-based redundancy related to the rectangular targets. This hypothesis predicts that participants will not only adjust their trajectories to counteract the effects of the vertical perturbation, but will also show systematic deviations along the X axis that bring the hand closer to the target.

3.2 Materials and Methods

3.2.1 Participants

21 right handed, college aged individuals (10 male, 11 female, mean age 19.5 +/- 1.8) with normal or corrected-to-normal vision participated in this experiment. All volunteers provided informed consent, and were compensated for their time in accordance with the Committee for the Protection of Human Subjects at UC Berkeley.

3.2.2 Apparatus

Participants grasped the handle of a robotic manipulandum (PHANToM 3.0L, http://www.sensable.com) capable of recording position and generating force along any of the three Cartesian axes. The robot was controlled by custom software written in Visual C++, using the OpenHaptics library. Control signals to the robot were updated at 1000 Hz and the output of the device was subsampled at 200 Hz for offline analysis. As a safety measure, the force output was capped at 9.0 N. The handle was allowed to rotate freely around any axis (roll, pitch, and yaw), but the angle of handle rotation was not recorded. Participants viewed the environment through a mirror. While the mirror precluded vision of the participant's arm, the cursor indicating hand position was presented to appear near the actual location of the hand, facilitating the subjective feeling of immersion in a 3D environment (see Figure 3.1a).

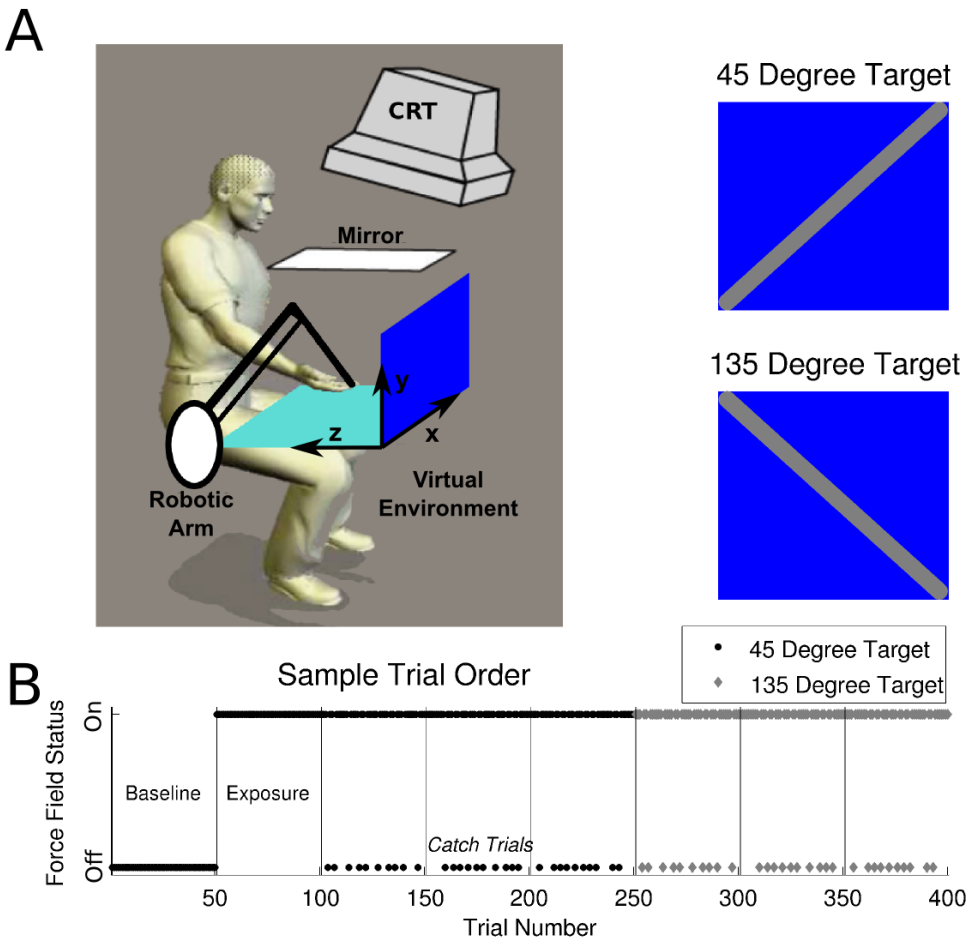

Figure 3.1: Overview of the task. (A) Participants controlled a cursor in a 3D virtual environment using a manipulandum. The goal was to aim for one of two oblique targets. (B) Participants move without the field in a baseline block, then move with the field in the first exposure block. Following are three test blocks which contain catch trials without the force on 20% of trials. Finally, the target is rotated by 90 degrees for three additional test blocks.

3.2.3 Task

By moving the manipulandum, the participant controlled a 6 mm, white spherical cursor that moved in the three-dimensional workspace. The cursor remained visible at all times during the experiment. At the onset of a trial, the cursor was pulled gently by the robot to the start position, an 8 mm sphere located at the participants' vertical midline. In virtual space, the simulated start position was approximately 10 cm in front of the eyes and 25 cm below eye level. After the cursor was maintained within this sphere for 1000 ms, one of two rectangular targets was presented along the back wall of the environment, 12 cm from the start location. The target was 12 cm long and 1 cm wide, with the long axis oriented at either 45 degrees or 135 degrees from horizontal (angles increase counterclockwise). The participant was required to make a single reaching movement, attempting to land within the target. At the termination of the reach (contact with the virtual back wall), an 8 mm sphere (the "feedback sphere") was presented on the surface to provide additional feedback of the movement endpoint as well as movement speed. If the movement duration was less than 275 ms, the feedback sphere was red, instructing the participant to slow down. If the movement duration was greater than 325 ms, the feedback sphere was green, instructing the participant to speed up. Movement durations between 275 and 325 ms fell within the desired speed criterion; on these trials, the feedback sphere was white.

To further motivate the participants, a running point tally was presented on the screen after each trial. If the cursor landed in the target within the appropriate movement duration window, five points were awarded. If the endpoint location was outside the target region, the score was decreased by 1 point. If the endpoint location was accurate but the movement duration was outside the desired range, the score remained unchanged.

Participants first practiced the task until they were comfortable moving within the virtual environment and could readily interpret the feedback. This typically involved 10-20 reaches. The main experiment began with a training block of 50 reaches in a null field (no perturbing forces). The orientation of the target was fixed for the entire block (45 deg or 135 deg, counterbalanced across participants). The force field was then introduced in a second block of 50 trials, with the target orientation the same as in the training block. In this exposure block, the movements were perturbed by a viscous curl field in which an upward vertical force was generated as a function of

velocity into the workspace (z-axis). While the perturbation was velocity-dependent, the viscosity term was position-dependent to ensure that all trials were identical during the initial phase of the movement. Viscosity was zero for the initial 30 mm of movement, and was then quickly ramped up to 7.5 Ns/m over the next 30 mm of movement.

Following the learning block, the participant completed six test blocks of 50 trials each (Fig 1B). For the first three test blocks, the orientation of the target surface was always the same as in the training and learning blocks. For the last three test blocks, the target was rotated by 90 degrees. Within each test block, the force field was present for 40 (80%) of the trials and turned off on 10 of the trials (20% catch trials). The catch trials were randomly determined with the constraint that a catch trial was always preceded by at least two force field trials. A short break was provided between each block.

3.2.4 Data analysis

Data analysis was performed in Matlab (http://www.mathworks.com). Velocity profiles were manually reviewed to identify and remove trials in which participants stopped moving forward before hitting the wall or took a wind-up by initially moving the cursor away from the wall before moving forward (criterion of 20 mm from start position). Approximately 5% of trials were discarded based on these criteria.

For the remaining trials, trajectories were standardized such that they had one value per mm along the Z axis (by binning existing values and interpolating missing values). This procedure simplified the analysis along the lateral (X) and vertical (Y) dimensions.

Analysis of the vertical dimension assesses the response along the axis of the force perturbation. Since the force field had no horizontal component, analysis of the lateral dimension is largely free from the effects of the perturbation. As such, analysis of the response along the X axis provides a more direct test of the effect of target orientation on adaptation.

3.3 Results

Participants found the task quite challenging. During the baseline block, participants landed in the target region on an average of 14.8 (sd = 4.1) trials out of the final

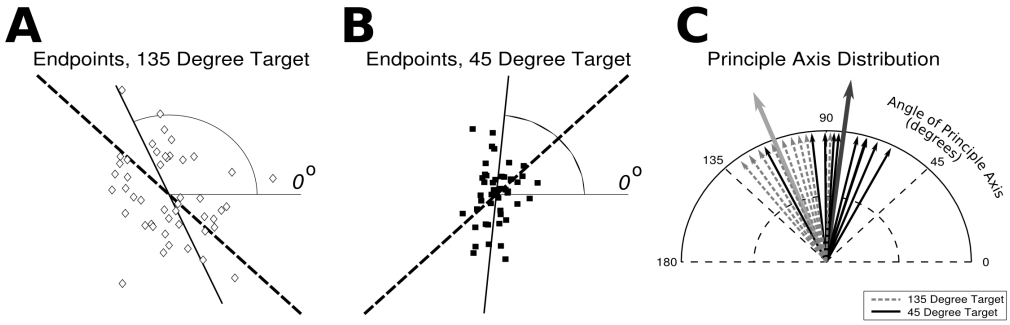

Figure 3.2: Endpoint distribution during the baseline block. Shown are sample endpoint distributions for the 135 degree target (A, white diamonds) and the 45 degree target (B, black squares). The dashed lines represent the target, the solid line represents the principal component. The angle between this line and the horizontal axis is plotted in (C) for the two targets, with solid lines indicating the means of the distribution.

25. Reach endpoints were tightly grouped near the center of the target (Figure 3.2a). We used principal component analysis to examine the endpoint distribution in greater detail. We then converted the unit vector representing the first principal component to an angle using the arctangent (Figure 3.2). When reaching to the 45 degree target, the principal axis was oriented on average at 82.5 deg (sd=15.3); when reaching to the 135 degree target, the principal axis was oriented at 112.3 deg (sd=12.2). This indicates that the distribution of endpoints is greater along the vertical dimension than along the lateral. Nonetheless, the angles of the principal axes were influenced by the orientation of the target, an effect confirmed by a Wilcoxon rank-sum test (ranksum 63, $p<0.002$).

We begin with a descriptive summary of the results before turning to a series of statistical tests. When the force field was turned on during the learning block, the movements were initially perturbed upward, with a marked upward deviation in the distribution of endpoint locations along the vertical axis. Participants rapidly learned to compensate for this perturbation (Fig 3.3A), with the endpoint value dropping by 2.5 cm from the first five exposure trials to the final 5 trials during the exposure block ($t(20)=4.7$, $p<0.001$). Interestingly, the perturbation was not fully corrected along the vertical axis: asymptotic values of the endpoint distribution along the vertical

axis, reached after about 10-15 trials, were approximately 2 cm higher than baseline ($t(20)=4.1$, $p< 0.001$). Over the last 25 trials of the learning block, participants landed in the target region an average of only 6.1 (sd = 3.4) times.

The pattern for the lateral component is more complex. Note that the force field perturbation, if not fully compensated along the vertical axis, will result in an endpoint that is to the left of the 45 degree target and to the right of the 135 degree target. With the introduction of the force field at the start of the learning block, participants exhibited an immediate shift to the left. This deviation is likely a biomechanical consequence of an upward perturbation of the right arm. Importantly, a target-specific effect on the distribution of the X coordinate of endpoint locations becomes evident over the course of learning, with the deviations resulting in endpoints that are brought closer to the target. For the 45 degree target, the endpoints shifted gradually in the rightward direction (black lines in Figure 3.3b). In contrast, for the 135 degree target, the endpoints remained shifted to the left (gray lines in Figure 3.3b). The divergence of the two functions along the X axis increased the likelihood that the endpoint location would fall within the target surface. This profile is consistent with the hypothesis that the participants' response to the perturbations incorporated properties of the target orientation. Comparing the change in the lateral endpoint between the first five trials in the exposure block and the last five trials in the exposure block failed to reveal a significant difference ($t(19) = 1.64$, $p = 0.12$). However, the difference between baseline performance and the final five endpoints in the exposure block was influenced by the target angle ($t(19)=2.17$, $p<0.05$).

Performance remained relatively stable over the subsequent three test blocks with the same target orientation. The test blocks included catch trials, trials in which the perturbing vertical force was not presented. These catch trials provide a probe of the participants' underlying internal model of the task environment. For comparison, we used the trials that immediately preceded the catch trials (PreCatch). Note that the force field was presented on the precatch trials. Assuming the planning process anticipates the upward perturbation, we expected that the average endpoints on catch trials would fall below the endpoints on precatch trials.

By examining the lateral component on these trials, we can assess whether an on-line correction process incorporates target information. Specifically, will an on-line correction in the absence of a vertical perturbation include a lateral component that increases the likelihood that the hand will end within the target surface?

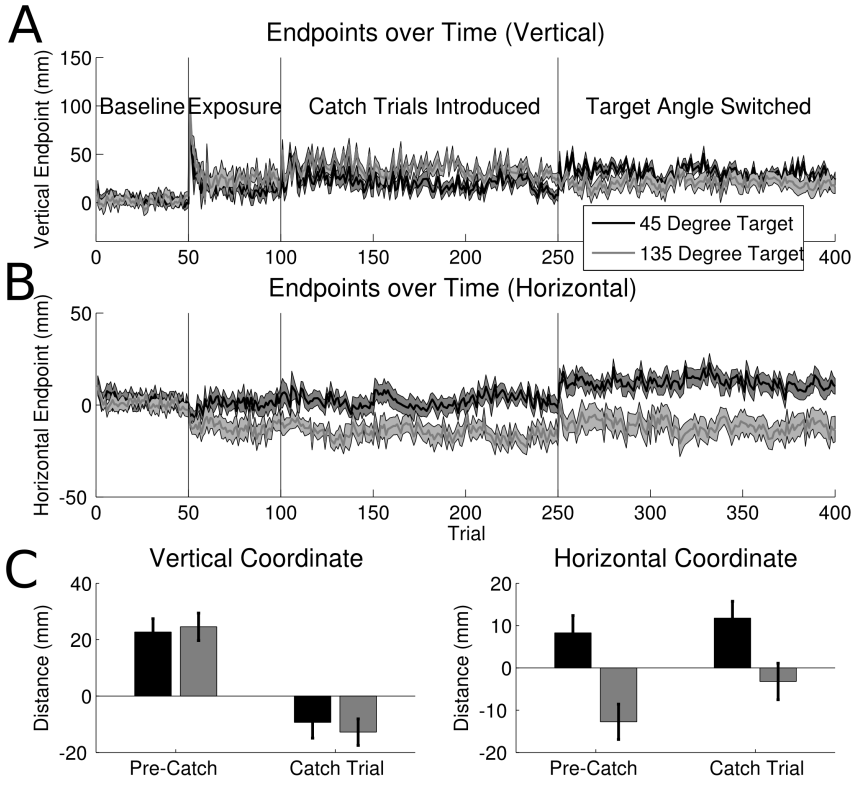

Figure 3.3: Endpoint performance during the experiment. Shown are the evolution of reach endpoints over the course of the experiment when reaching for the 45 degree target (black) or the 135 degree target (gray). The vertical coordinate, which is increased and rapidly drops when the force field is applied, is shown in A. The X coordinate, which diverges for the two targets over learning, is shown in B (negative numbers mean to the left of center). The average endpoints during pre-catch and catch trials are shown in C and D for the vertical and lateral coordinates, respectively.

Examination of the vertical dimension showed that the presence of the force field (precatch trials vs. catch trials) had a strong effect on the endpoint locations (Figure 3.3c). Overall, the average endpoint location was 23.7 mm (sd 15.0) above the center of the target during precatch trials, and 10.9 mm (sd 15.6) below the center of the target during catch trials. The vertical endpoint was minimally affected by the target orientation.

In terms of the lateral dimension (Figure 3.3d), participants maintained the overall bias that they exhibited on the force field trials (to the right for the 45 degree target and to the left for the 135 degree target, consistent with aiming slightly above center). This effect adds further support to the hypothesis that the planned trajectory incorporates features of both the expected perturbation and target orientation. During catch trials the trajectories for both targets moved in a rightward direction (relative to precatch trials). However, this effect was larger for the 135 degree target. Note that for the 135 degree target, a lower trajectory will cause the participant to be too far to the left. Thus, the larger shift to the right increases the likelihood that the movement will terminate within the target.

To statistically analyze performance during the test blocks, we performed two 2 x 2 repeated measures ANOVAs, with within-subject factors Target Angle (45deg / 135 deg) and Trial Type (Precatch/Catch). For the vertical dimension, this analysis revealed a main effect of Trial Type ($F(1,20)=126.4$, $p<0.001$), confirming that the vertical endpoints were lower on catch trials compared to force field trials. There was no effect of target angle ($F(1,20) = 0.11$, $p=0.77$). Interestingly, the interaction approached significance ($F(1,20) = 3.8$, $p = 0.06$). For the lateral dimension, both main effects were reliable (Target Angle: $F(1,20)=11.6$, $p<0.003$; Trial Type: ($F(1,21)=47.3$, $p<0.001$). Moreover, the interaction term was highly significant ($F(1,20)=34.2$, $p<0.001$). The main effect is consistent with the hypothesis that the participants anticipate the force field and incorporate a lateral deviation that is expected to move the hand closer to the target (rightward for 45 deg, leftward for 135 deg). The interaction, in which a rightward shift on catch trials is even more pronounced for the 135 deg. target, indicates that the target information is also incorporated into on-line corrections.

The internal model of the expected force perturbation is violated on the catch trials. This produces an error signal that should produce a change in behavior on the subsequent trial (Thoroughman Shadmehr, 2000), although the effects of these

changes would be small since catch trials are infrequent and solitary. As such, post-catch trials provide a second look at adaptation. We would expect adaptation here to be evident in the movement heading before the force field is applied. To examine such trial-by-trial changes, we compared the trajectories on catch and post-catch trials. For this analysis, we reduced the effects of overall drift in performance by subtracting a running baseline (defined as the pre-catch trial) from each trajectory. Since catch trials are not predictable, we expect the average catch trial heading before the expected force onset to be identical to the precatch trial, equivalent to expecting a value of 0 after adjustment. Since the expected vertical force is then absent, participants end up reaching to a location significantly lower than the midpoint of the target. Assuming this error influences their planning for the subsequent trial, we would expect the initial trajectory on the post-catch trial to be higher. Consistent with such a prediction, catch trials were unaffected during the initial 30 mm (Figure 3.4a), while an upward shift in the trajectory was evident on the post-catch trials (Figure 3.4b).

Considering the horizontal dimension, our hypothesis predicts that following a catch trial we should see a lateral component to the post-catch trial adjustment. Since the endpoint analysis suggests that, when reaching to the 135 degree target, the online correction is more strongly rightward compared to when reaching to the 45 degree target, we should see a similar effect on the post-catch trial. Looking first at the catch trial trajectories (Figure 3.4c), the lateral displacement is simailar to that observed on the pre-catch trial (i.e., has a difference of 0) before the expected force onset. Immediately following this point, trajectories toward both targets deviated to the right. Near the end of the reach on catch trials, we observed a strengthening of the rightward deviation for the 135 degree target, while for the 45 degree target, participants were making adjustments to reduce the rightward deviation. Turning to the post-catch trials, prior to force onset we see a strong overall leftward deviation when reaching toward the 45 degree target, with much less deviation, though still slightly leftward, when reaching toward the 135 degree target. This remains consistent with the directions of the online corrections.

To quantify these effects statistically, we compared the adjusted position at 30 mm into the reach, the first point at which the force field, if present, would be nonzero, for both catch and post-catch trials. For this analysis, we used two 2x2 repeated measures ANOVAs, with within-subjects factor of Target Angle and Trial

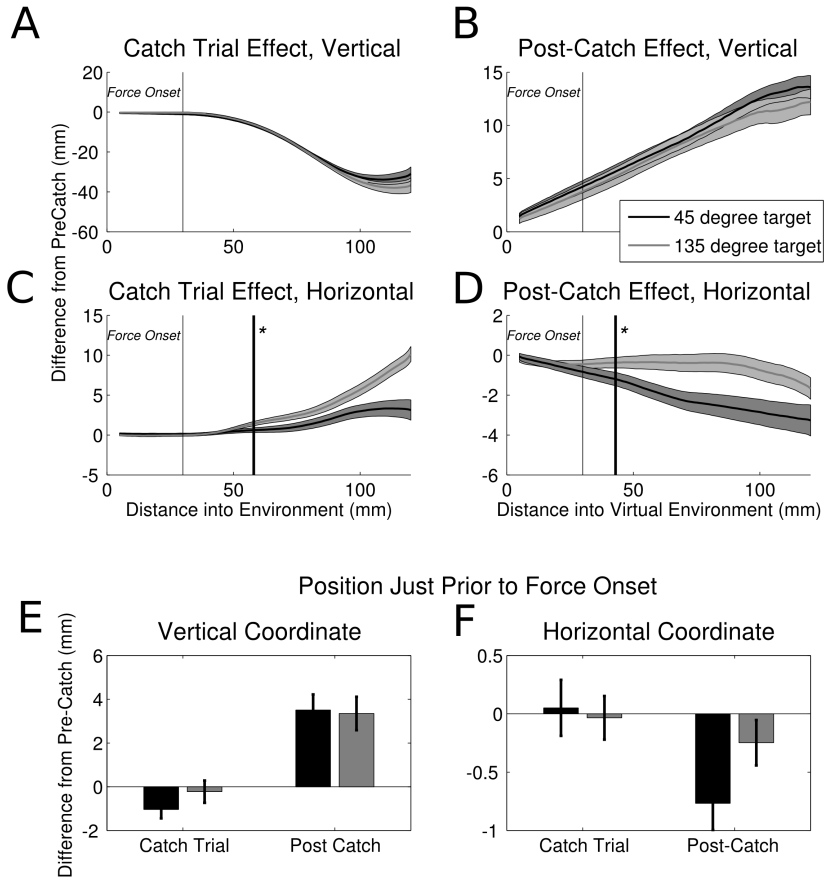

Figure 3.4: Average trajectories are shown for catch trials and post-catch trials. Prior to averaging, trajectories were standardized by subtracting the trajectories produced on the trials immediately preceding catch trials. Shown are the Y (A, B) and X (C, D; negative numbers indicated a leftward change from the pre-catch trial) coordinates as a function of Z (depth into the workspace). Arm positions 30 mm into the reach demonstrated the strategic changes following a catch trial for the vertical coordinate (E) and showed a target dependence for the X coordinate (F).

Type (catch or post-catch). For the vertical coordinate, a highly significant main effect of Trial Type was observed ($F(1,20)=76.8$, $p<0.001$). The main effect of Target Angle was not significant, ($F(1,20)=0.55$, $p=0.44$), nor was the interaction ($F(1,20) = 2.27$, $p = 0.11$). For the X coordinate, there was again a main effect of Trial Type ($F(1,20)=18$, $p<0.001$), and not Target Angle ($F(1,20)=1.4$, $p=0.24$). However, the interaction was reliable ($F(1,20)=6.46$, $p<0.02$). The interaction here suggests that participants were adjusting their strategy to account for the previous error in task coordinates rather than simply addressing the mismatch in predicted force.

As noted above, statistical analysis of the lateral component consistently reveals an effect of target orientation, consistent with the hypothesis that responses to expected and unexpected perturbations incorporate information about behavioral redundancies. A final question of interest concerns the latency of these effects. To answer this question, we looked for the point at which the trajectories (specifically, the difference from the pre-catch trials; Figure 3.3a-d) diverged for the two targets, considering both catch and post-catch trials separately. For each point along the Z axis (the reach direction), we computed the difference between the average position (across subjects) when reaching toward the 45 degree target compared to the 135 degree target. A null distribution was then synthesized by taking the observed trajectories, and assigning them at random to either the 45 degree target or the 135 degree target (independent of the actual target) and computing the difference. This procedure was repeated 10,000 times. Assuming a null hypothesis of no real difference in trajectory, the actual differences should fall somewhere in the middle of the synthesized distribution. We defined significant divergence by requiring the actual difference to fall within a tail of the null distribution (taking an alpha of 0.05) for more than 10 mm.

Examining the vertical coordinates, the average catch trial trajectories toward the two targets did not diverge. In contrast, the lateral components of the average catch trial trajectories diverged noticeably, and this divergence reached significance 58 mm into the reach (Figure 3.4c). On average, participants reached this point 179 ms (sd 31 ms) after initiating the movement. When considered from the point of expected force onset, the latency is 75 ms (sd 10 ms). If only the first three test blocks for each subject are considered, however, the point of significant divergence occurs at 103 mm, which subjects arrive at 200 ms (sd 25 ms) after the time of the expected force onset. The average post-catch trajectories did not differ when the vertical coordinates were considered. The lateral components of these trajectories diverged quite soon, reaching

significance 43 mm into the reach (152 ms from start, 35 ms from force onset).

3.4 Discussion

Studies of sensorimotor adaptation have generally involved reaching movements that are restricted to a two-dimensional plane with the goal defined as a single point. This paradigm has been adopted in a variety of tasks including visuomotor adaptation (Fishbach & Mussa-Ivaldi, 2008; Mazzoni & Krakauer, 2006; Sober & Sabes, 2003; Tseng *et al.*, 2007) and force field learning (Shadmehr & Mussa-Ivaldi, 1994; Taylor & Thoroughman, 2007; Thoroughman & Shadmehr, 2000). Successful learning requires the formation of an accurate internal model of the disturbance. In this context, the current study contains a replication of these results. Participants in our task were quite sensitive to the vertical perturbation, and rapidly adjusted their motor plan to anticipate this force over the initial exposure trials. During catch trials, in which the perturbing vertical force was absent, the endpoints of the movements were much lower than on the standard trials, reflecting the fact that the participants were generating a downward correction in anticipation of the upward perturbing force. Moreover, the resulting error on the catch trials led to an adjustment on the next trial such that participants now appeared to be aiming for a higher location than on standard trials. This pattern of results is consistent with what has been observed in studies of planar force field adaptation (Fine & Thoroughman, 2007; Taylor & Thoroughman, 2007; Thoroughman & Shadmehr, 2000).

Within such adaptation tasks, a second strategy is also possible. Participants could adjust their movements from trial to trial with the goal of minimizing the behavioral error. Importantly, both learning an internal model and minimizing task error would lead to very similar behavioral changes when the movements are performed in a two-dimensional space to a point-like target. To investigate whether participants are able to correct for task error, we required participants to move in a three-dimensional space and expanded the goal region. The three-dimensional workspace allowed us to create a dimension that was completely redundant with respect to the force perturbation task. The expanded goal region made this redundant dimension task-relevant. Assuming that participants' goal is to terminate the movement within the target region (as opposed to an implicit singular point at the middle of the target region), behavioral error can be defined as the distance from the major axis defining the tar-

get. In this way, we can ask if corrective movements are sensitive to variation in the orientation of the target, even when a perturbing force is held constant. By using only an upward force perturbation, we would expect to see adjustments restricted to the vertical axis if learning is based on a force compensation strategy. In contrast, if corrective processes are sensitive to task redundancies, then we would expect to see adjustments along the lateral axis, and most importantly, these corrections will be in opposite directions for the two targets. The results, both when assessed in terms of trial-by-trial adjustments as well as on-line corrections within a trial, indicate that participants incorporate goal-based information when moving within a novel environment.

3.4.1 Between-trial changes

Over the course of the experiment, participants adjusted their movements such that the endpoints would result in better performance. This was evident during the initial exposure block. The unexpected upward perturbing force resulted in an endpoint much higher than expected. Since the target was an oblique line, the endpoint was thus to the left for the 45 degree target and to the right for the 135 degree target. Thus, the changes observed over the course of the initial exposure block moving lower and to the right for the 45 degree target, and lower and to the left for the 135 degree target were in a direction that attenuated both the force field and the behavioral error.

During the test trials, when the force field was unexpectedly removed, participants adjusted their performance on the subsequent trial. During typical force field adaptation, this effect is believed to reflect rapid changes in an internal model (Fine & Thoroughman, 2007; Thoroughman & Shadmehr, 2000). In the current study, a compatible effect is observed when the vertical component of the trajectory is considered. For the lateral component, however, the effect of a task-based strategy can be evaluated. The results confirm that participants utilize the target information in order to improve performance, as the orientation of the target significantly affected the magnitude of the adjustment on trials following catch trials.

When learning a new task, it has been demonstrated that participants are able to exploit tolerance afforded by the task space (Müller & Sternad, 2004). The current investigation extends this idea by showing adjustments that occur when well-learned movements are produced in a novel context, indicating that utilizing goal-based re-

dundancies may be a general feature of motor control.

3.4.2 Within-trial changes

Within-trial changes are important for evaluating whether the exploitation of redundant task dimensions occurs via rapid feedback mechanisms or results from feedforward, strategic processes. These effects can be assessed with the catch trials, comparing performance on these trials to that observed when an expected perturbation is presented (precatch trials). We observed that the lateral component of the trajectory was influenced by the target in a manner consistent with the hypothesis that the participants were using task goal information to minimize endpoint error. Given that the absence of the upward perturbing force would result in an endpoint lower than expected, an online goal-based correction would entail a leftward lateral component when reaching to the 45 degree target and a rightward component when reaching to the 135 degree target. This behavior was observed for the 135 degree target: the movement endpoints were consistently displaced to the right on catch trials. However, we also observed a shift to the right on the 45 degree target trials, albeit to a significantly reduced degree than for the 135 degree target. Inspection of the trajectories on catch trials (Figure 3.4) indicates an incomplete, late-onset leftward correction when reaching to the 45 degree target.

While the manipulandum applied a purely vertical force field, the consistent rightward shift of the endpoints on catch trials suggests that the force field induces a lateral deviation. Such a displacement is also apparent during the initial trials of the exposure block. This likely reflects a biomechanical effect arising from the interaction of the manipulandum and arm. Consistent with this hypothesis, a separate group of participants were tested with their left arm. These participants showed an overall leftward shift in their endpoints on catch trials (see Figure A.2).

The catch trials provide a window into the timing of these corrective movements in the absence of any external, perturbing force. When all six test blocks are considered, the trajectories to the two targets diverge approximately 80 milliseconds after the expected onset time of the missing perturbation. This delay is faster than what would be expected given the time required for supraspinal feedback corrections (Allum, 1975), particularly if one considers the necessary delay between muscular activation and resultant limb movement. Interestingly, if we consider only the first three test blocks during which the movements were limited to a single target, the divergence

occurs around 200 ms from the expected onset time, within the expected time for supraspinal feedback loops. At present, a conservative and parsimonious account of the on-line corrections favors the hypothesis that the adjustments here reflect feedback processing, possibly caused by changes within the brainstem (Jacobs & Horak, 2007). However, these results may also suggest that, with training, goal-based information may lead to additional feedforward contributions. Indeed, the cost of recomputing trajectories during the movement in order to reduce an imminent error may be greater than the costs associated with adjusting feedback gains in a feedforward manner.

3.4.3 Summary

The current results confirm that an essential factor in motor behavior is the maximization of expected reward. This aspect of optimal control has been highlighted in recent studies examining aiming strategies when people reach for a target surrounded by asymmetric penalty zones. Under such conditions, participants aim for an optimal point that accommodates uncertainty related to their own movement control (Trommershäuser *et al.*, 2005) or uncertainty inherent in the environment (Trommershäuser *et al.*, 2003a). The current study provides an important extension of this concept, showing that feedback processes also incorporate goal-relevant information. Indeed, it appears that feedback processes are adapted to maximize the expected reward, doing so in a task-specific way.

Acknowledgments

We wish to acknowledge Kimberly Koike, Marisa Whitchurch, and Christina Merrick for their assistance in data collection and discussion of preliminary results.

Chapter 4

Unmasking error signals in the cerebellum by correcting for heart rate

Abstract

A central tenet of motor neuroscience is that the cerebellum exploits sensory prediction errors to support sensorimotor learning. However, functional magnetic resonance imaging (fMRI) studies have generally failed to reveal definitive signatures of error processing in the cerebellum. Furthermore, neurophysiological results have argued that the cerebellum codes only for the unexpected presence, but not for the unexpected absence of sensory stimuli, rending this signal unsuitable for learning. Here, we used fMRI to compare reach errors that resulted in the unexpected presence or absence of sensory stimulation. Surprisingly, we observed an error-related decrease in heart rate that artificially lowered the BOLD signal. Only after removing this influence, significant error-related activity was found in the arm area of anterior cerebellum; the signal was similar for the unexpected presence and absence of sensory stimulation. These results show that prediction errors are encoded symmetrically in the cerebellar BOLD signal and stress the importance of monitoring physiological processes during fMRI.

4.1 Introduction

The human cerebellum plays an important role in motor learning. Following cerebellar damage, learning deficits have been observed during saccade (Golla *et al.*, 2008), visuomotor (Martin *et al.*, 1996; Tseng *et al.*, 2007), force field (Maschke *et al.*, 2004; Smith & Shadmehr, 2005), and gait adaptation tasks (Morton & Bastian, 2006). Performance improvements on such tasks depend on the appropriate utilization of motor errors. Specifically, cerebellar-dependent adaptation appears to be driven by sensory prediction errors (e.g., observing a movement outcome that is different from the expected outcome), rather than by the act of making a corrective movement (Tseng *et al.*, 2007).

How sensory prediction errors are coded within the cerebellum remains poorly understood. The Marr-Albus-Ito model of cerebellar function (Albus, 1971; Marr, 1969; Ito, 1972) postulates that climbing fibers from the inferior olive convey an error signal to the cerebellar cortex, attenuating the Purkinje cell response to the parallel fiber input that is active at the time of the error. This model provides a clear account of cerebellar function during eyeblink conditioning (for review, see Ito, 2001; Kim & Thompson, 1997) and gain modulation of the vestibuloocular reflex (for review, see Ito, 2001, 1998). However, for volutional movements, such as saccade adaptation or reaching, the role of climbing fibers is less clear (Catz *et al.*, 2005). Furthermore, it has been argued that climbing fiber activity would not provide a signal appropriate for learning, because it codes sensory prediction errors in an asymmetric way: Horn *et al.* (2004) recorded from the olivary nuclei in cats who had been trained to perform a reaching task. They induced two types of sensory prediction errors: in one case the animal encountered an unexpected stimulation of the forepaw, in the other situation the expected forepaw stimulation was absent. The inferior olive only responded when an unexpected stimulus was presented, remaining silent when an expected stimulus was absent. Since learning is required in both conditions, this asymmetry challenges the idea that the climbing fiber response serves as the only teaching signal for the cerebellar cortex. Indeed, some studies have failed to find any relationship between complex spike activity and error during volitional movements (Catz *et al.*, 2005; Ebner *et al.*, 2002), whereas others have shown only a weak relationship (Kitazawa *et al.*, 1998).

Human neuroimaging studies have also not been able to conclusively shown correlates of sensory prediction errors in the cerebellum. Consistent with a reduction

in sensory prediction errors over the course of learning, the hemodynamic responses in the cerebellum generally decrease as performance improves (Luauté *et al.*, 2009; Imamizu *et al.*, 2000). However, these studies often have the confound that movement errors are correlated with corrective movements, which have been show to drive the cerebellar BOLD signal substantially (Diedrichsen *et al.*, 2005b). Importantly, two studies that explicitly accounted for online corrections failed to observe error-related activation within the cerebellum (Diedrichsen *et al.*, 2005b; Krakauer *et al.*, 2004).

In the current study, we re-examined the representation of sensory-prediction errors within the cerebellar cortex. Specifically we asked, whether prediction errors code the presence or absence of sensory stimulation symmetrically or asymmetrically. Using an MR-compatible robotic manipulandum, human participants produced rapid out-and-back reaching movements, attempting to intercept a visual target. Using such fast movements eliminated any online corrections. The visual display consisted for a circular arc of dots, one of which the target differed in color from the distractors (Figure 4.1). If the movement passed through a red circle, the robot delivered a short, resistive force pulse, while no external somatosensory stimulation was provided if the movement passed through a white circle. By manipulating the color of the target, we thus controlled whether the participant expected a sensory stimulus. In this way, sensory prediction errors could be caused by the presence of an unexpected stimulus or the lack of an expected stimulus. In this manner, we assessed whether the cerebellar BOLD response was related to sensory prediction errors in general (diagonal) or prediction errors that are signaled by sensory stimulation only (upper right corner).

4.2 Methods

4.2.1 Participants

Ten adults (mean age 23 years, SD 4 years, 3 females) from the student population at Bangor University with no history of neurological injury or disorders served as participants. An 11th participant was excluded from the study after he indicated in a post-session debriefing difficulty in resolving the target locations without corrective lenses. Participants were compensated for their time with cash payment. The study was approved by the Ethics committee of the School of Psychology, Bangor University.

4.2.2 Apparatus

Participants held the handle of a nonmagnetic two-joint robotic manipulandum (http://fmrirobot.org). This device allows for low-friction, two-dimensional movements in any direction within the horizontal plane. Linear optical encoders on the elbow and shoulder joint provided position readings with endpoint accuracy better than 0.01 mm to a control computer outside the room. The robotic device is capable of applying forces to the hand. Forces were applied via air pistons supplied with 100 psi pressure of from a compressor outside the MR room. Time constant of the pistons response was 60 ms to a step input. A filter panel in the wall of the scanning room prevented leakage of radio frequency noise. Position and velocity of the hand and the generated forces were recorded at 200 Hz. Visual stimuli, including targets and any feedback of hand position, were projected from outside the scanner room onto a back screen, which was viewed by the participants through a mirror.

4.2.3 Behavioral Task

Participants lay in a supine position, used a custom-fit bite bar for head stabilisation, and grasped the robot handle with their right hand while the elbow was supported by a cushion. The center of movement was adjusted for each participant prior to scanning such that they could perform the required movement comfortably and without contacting the bore of the scanner. Participants were presented with an array of 21 colored dots (1 cm diameter), arranged on a semicircle 8 cm away from a starting location. One dot, the target, differed in color from the remaining distractors. When the cursor passed through a region of the semicircle containing a red dot, the hand received a brief resistive force pulse consisting of a half-sinusoid with an amplitude of 1.5 N and a duration (half-period) of 60 ms. No pulse was delivered when moving through a white dot nor when the hand moved through any of the dots on the return path. Participants were instructed to produce a rapid out-and-back movement that carried the cursor beyond the distance of the target. If the cursor passed through the target dot, a counter incremented by one point after the hand returned to the center region. A feedback dot was presented on the display at the point of the farthest extent of the movement. This feedback provided both direction- and amplitude-based error information that was expected to influence performance on the subsequent trial.

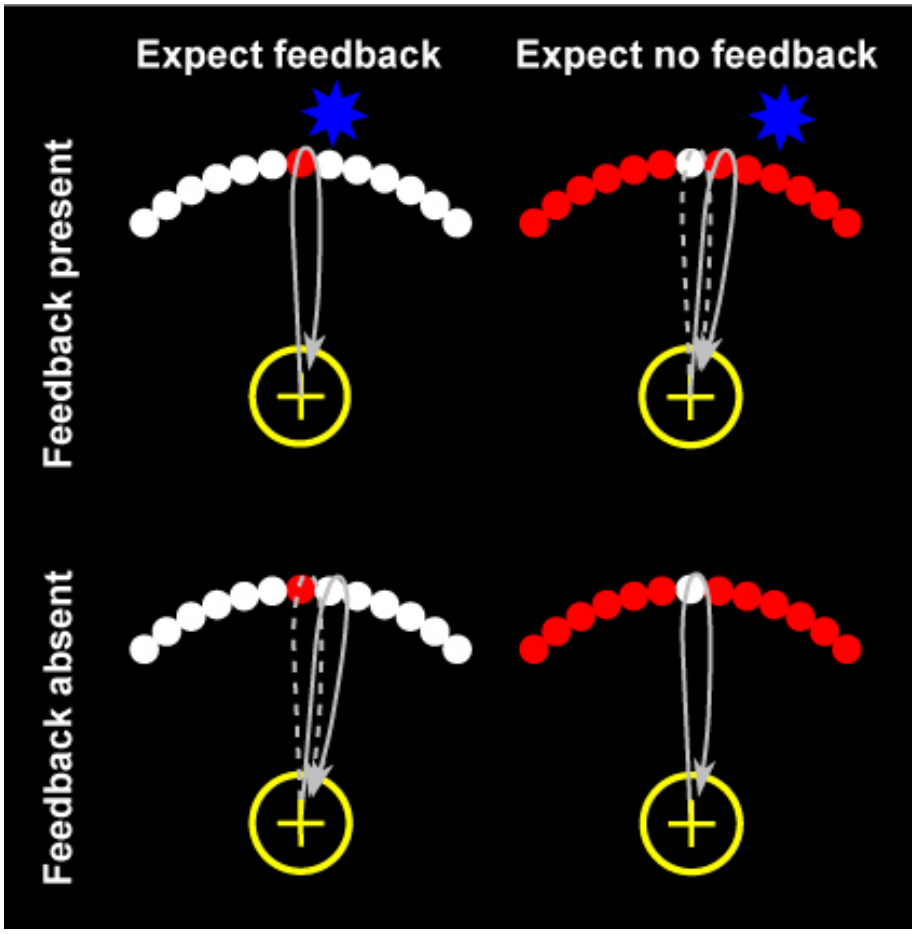

Figure 4.1: Participants controlled a cursor with a nonmagnetic robot. The goal is to move the cursor from the yellow starting circle through the target: the dot with a different color. If a red dot is touched, a small force pulse is delivered. No force pulse is given when a white dot is touched. Errors are induced by rotating the cursor on the screen by one dot to the left or right. Errors can be signaled by the presence of an unexpected stimulus (white target) or the lack of an expected stimulus (red target).

Each trial began with the presentation of the target and distractors. These stimuli only appeared when the cursor was within a 2 cm diameter region at the center of the workspace. Starting position was indicated by the color of a crosshair pattern: the crosshair was blue if the cursor was within its boundary and yellow if the cursor was outside the starting region. The expected sensory outcome for a trial was varied by controlling the color of the target and the distractor dots. When the target was a red dot with white distractors, participants expected to receive a somatosensory stimulation, and an error was signaled by the absence of this perturbation. When the target was a white dot, the unexpected presence of the perturbation signaled the error.

To increase the occurrence of errors, the cursor was randomly rotated by seven degrees, exactly the angular width of a single target, to the left or the right around the starting position. All participants were informed that a rotation would be applied randomly. During the fMRI acquisition, these rotations were present on 33% of trials, selected at random. To determine the relative contributions of the visual and the proprioceptive error signal, we withdrew the visual feedback (cursor and reversal point) on half the trials (rotated or unrotated). These trials were randomly intersperse with trials in which visual feedback was provided.

Participants completed a training session 1-2 days before scan acquisition. The training session was conducted in a mock scanner with a setup identical to the real scanning environment. The training session was used to familiarize the participants with the task and train them to complete the movements with a consistent movement speed and amplitude across conditions. During training, participants completed four blocks of 80 trials each.

In the fMRI session, a target was presented every 4 seconds. The target color, and thus the expected outcome, was held constant for blocks of 15 trials, followed by ten seconds of rest. During each run, participants performed two blocks with each expectation condition, with the order counterbalanced across participants. Each participant completed eight runs, with the exception of one subject who was unable to complete the eighth run. One subject had to complete a ninth run to make up for a technical error with the apparatus.

Participants were instructed to maintain fixation at the center location and not saccade to the target. Eye movements were monitored during the fMRI session, and participants were informed at the end of the run if they made any eye movements.

Very few saccades were detected during the task. Scan acquisition

Data were acquired on a 3 T Philips Intera system (Philips Medical Systems, Best, The Netherlands). For functional runs, we used an echo planar imaging (EPI) sequence with sensitivity-encoded MRI (Pruessmann *et al.*, 1999) and a sensitivity encoding factor of 2. MRI slices were prescribed to cover the cerebellum and brainstem, including the inferior olive. 24 oblique slices oriented approximately 45 degrees from horizontal (1.8 mm thickness; 0.3 mm gap; repetition time, 2 s) were collected. Each slice was acquired as a 96-96 matrix (field of view was 24.0 x 24.0 cm) with a voxel size of 1.88 x 1.88 x 2.1 mm. Each run contained 150 volumes that were analyzed offline. For anatomical localization, high-resolution T1-weighted structural images were acquired with 0.67 x 0.67 x 0.7 mm resolution using a magnetization-prepared rapid-acquisition gradient echo sequence.

Because physiological variables are known to influence the BOLD response (Glover *et al.*, 2000; Chang *et al.*, 2009; Birn *et al.*, 2008; Shmueli *et al.*, 2007), we recorded cardiac and respiration rate during the functional runs. Heartrate was recorded at 125 hz using a 4-lead ECG system integrated with the Phillips scanner. Respiration was recorded at 60 hz using a pneumatic compression belt.

4.2.4 Imaging Analysis

Functional images were converted to four-dimensional Nifti files (using the free program MRIcron: http://www.cabiatl.com/mricro/mricron/dcm2nii.html) and analyzed using SPM5 (http://www.fil.ion.ucl.ac.uk/spm/software/spm5). Images were corrected for slice timing, realigned to correct for residual head movement (using rigid-body realignment), and coregistered to the anatomical image. Following this preprocessing, images were analyzed using a general linear model. For group analysis, the anatomical images were normalized to a high-resolution cerebellar template (SUIT, see Diedrichsen, 2006). The resulting transformations were used to reslice individual contrast images into a common template space. Normalized contrast images were smoothed with a three-dimensional Gaussian kernel (5mm FWHM) prior to group analysis.

Event related regressors were created for eight trial types (two expectation conditions x two outcomes x two visual feedback conditions) as stick function at the time of target onset. These functions were convolved with the canonical hemodynamic response function.

To correct for the influence of heart rate changes, we needed to know how heart rate affects the BOLD signal. For the neocortex, the impulse response function of heart rate on the BOLD signal peaks after a delay of 4 s and lasts for up to 24 s (Chang *et al.*, 2009). However, as the blood supply of the cerebellum differs from that in the neocortex, we might expect that a different response function for the cerebellum. We elected to adopt an empirical approach to address this discrepancy. We used the heart rate time series directly, shifted in time by 0-11TR. The resulting 12 regressors were used directly as covariates in the general linear model.

The heart rate timeseries were computed using the Physiological Log Extraction for Modeling toolbox (http://sites.google.com/site/phlemtoolbox/), following the methods outlined in Chang et al. (4). After marking the time of the peak of the QRS waveform, all heartbeats within a six-second window centered at the time of a single volume in the fMRI time series were selected. The average inter-beat interval was computed, inverted, and multiplied by 60 to compute beats per minute.

4.3 Results

4.3.1 MRI results before rate correction

Areas that are sensitive to prediction errors involving both the unexpected absence and presence of sensory stimulation should show a difference in a contrast that compares error trials and correct trials, irrespective of error type. An initial analysis (Figure 4.2a) failed to reveal any areas that showed an increase in the BOLD signal related to prediction errors. In fact, the results revealed a rather puzzling picture in that errors were associated with a strong decrease in the BOLD signal. A number of clusters with large signal decreases were found in lobules Crus I and II i.e., areas that are typically not associated with sensorimotor control.

4.3.2 Cardiac Rate

Figure 4.3a shows the normalized heart rate across a number of representative movements. A drop in heart rate is visible during the inter-trial interval and likely reflects cardiac deceleration associated with the anticipation of the forthcoming target (Damen & Brunia, 1987). Importantly, when participants committed an error, there was a notable decrease in heart rate compared to that observed on correct trials. The

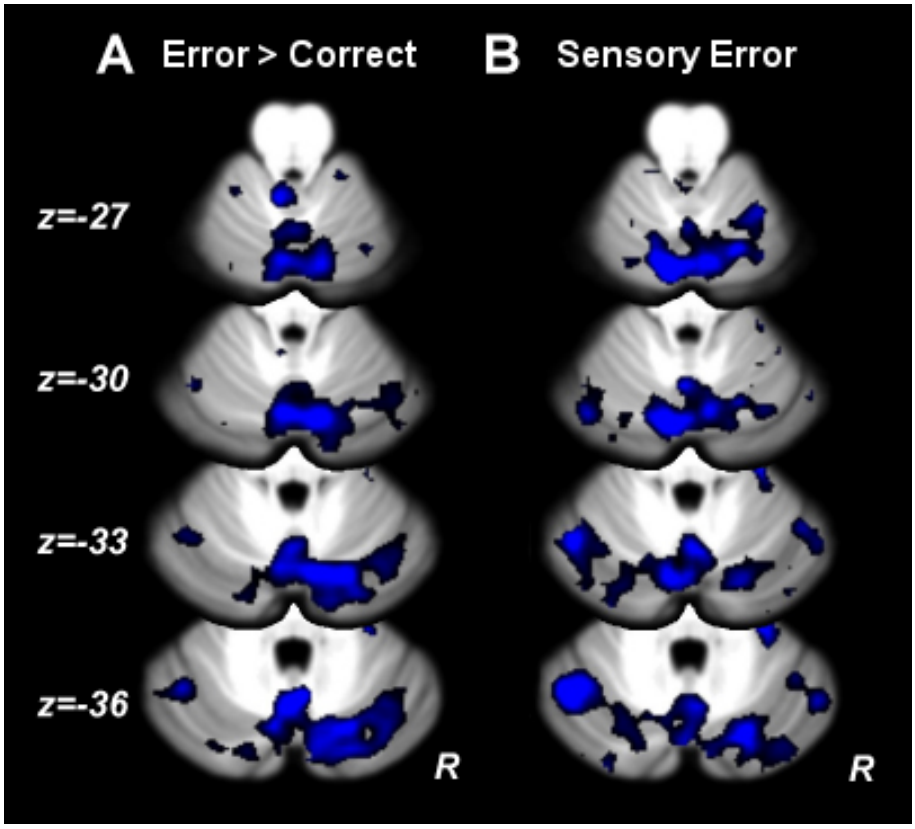

Figure 4.2: MRI without heartrate correction. A) The effect of errors. In blue are regions more active for correct trials than error trials. No regions are more active during errors. B) Errors conveyed by a sensory signal. In blue are regions more active during trials in which no sensory stimulus was present, or an expected sensory stimulus was present, compared to trials in which the sensory stimulus conveyed an error signal. No regions are more active during trials in which the sensory stimulus conveys the error.

magnitude of this difference (Figure 4.3b) reached about 1% of the mean heart rate and lasted for approximately 8 seconds after the error.

To statistically evaluate these effects, we averaged the normalized heart rate over the 4 s following a response (though a similar analysis over an 8 second window revealed qualitatively identical results), and analyzed these data with a 2x2 repeated measures ANOVA involving the factors Feedback Expectation and Movement Outcome. Neither main effect was reliable (Expectation: $F(1,9)=0.27$, $p=0.62$; Outcome: $F(1,9)=4.02$, $p=0.076$), but a highly significant interaction was obtained ($F(1,9)=19.83$, $p=0.0016$). The interaction reflects the fact that error-related heart rate deceleration was highly consistent across participants, and was associated with an unexpected stimulus as well as with the absence of an expected stimulus. No comparable effect on the rate of breathing was found.

4.3.3 MRI results, corrected for cardiac rate

After removing the influence of heart rate, we then repeated our primary analysis, comparing the cerebellar BOLD signal for trials with and without prediction errors (Figure 4.5). The broad decrease in BOLD activation on error trials in the uncorrected analysis was now strongly attenuated. Importantly, a single region within the cerebellar cortex was now found to show an increase in activation on error trials (1000 mm3 with a t-value above 2.82, $p<0.01$, cluster-wise p-value corrected for multiple test: $p<0.002$). The cluster lies in lobule V of the right anterior lobe, a region associated with movement as well as somatosensation of the right hand (Rijntjes *et al.*, 1999; Grodd *et al.*, 2001; Fox *et al.*, 1985). Critically, a map-wise test of the activation following unexpected sensory stimulation compared to the other three conditions failed to reveal any reliable areas (Fig 4.5b), suggesting that the signal may be similar for both types of error signal.

4.3.4 Influence of Heart Rate on BOLD

Heart rate had a significant effect on the BOLD signal for 57% of the cerebellar voxels ($F(12,108)>2.354$, corresponding to $p<0.05$ corrected for multiple comparisons using false discovery rate). Figure 4A shows the reduction in overall variance achieved by including heartrate regressors in the fMRI analysis. Overall, the addition of cardiac regressors systematically explained approximately 3% of the variance of the raw

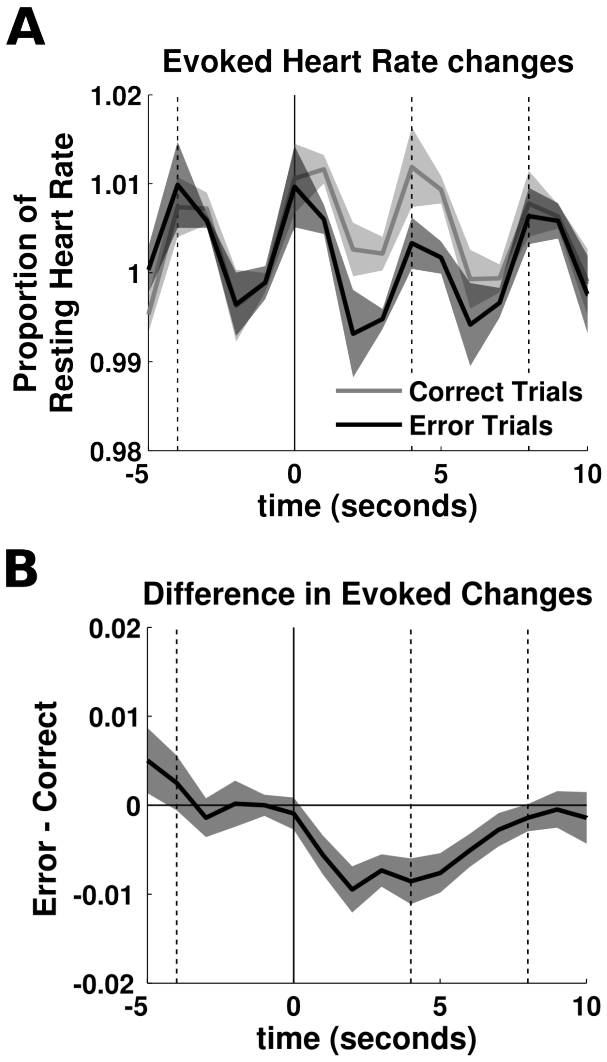

Figure 4.3: Event-related changes in heart rate. Each movement is indicated as a vertical line. The instantaneous heart rate (expressed as percent change from individual baseline) is aligned to the movement at 0s (solid line), and shown depending on whether this movement was a correct or error trial. A) The heart rate decelerates before every movement. B) The difference between error and correct trials reveals a depression of the cardiac rate that lasts for 8 seconds.

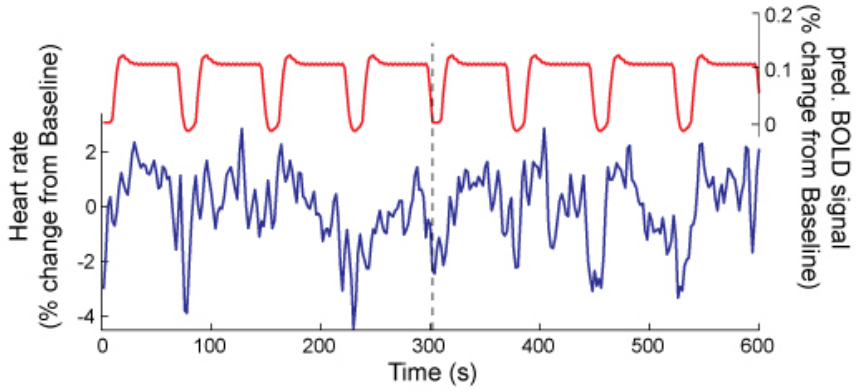

Figure 4.4: Exemplar heart rate changes of one subject across two imaging runs. The blue curve indicates deviations of the heart rate from the average for that run, in beats per minute. The red curve is the sum of all behavioral regressors (i.e. task vs. rest contrast). The curve indicates the 4 blocks of 15 trials per run with the short breaks between blocks.

BOLD signal across the cerebellar cortex.

Adopting an empirical approach, we included 12 regressors, each consisting of the instantaneous heart rate, shifted by 0-11TRs in time. The coefficients for these regressors give an estimate of the cardiac response function (Figure 4B). The response was relatively homogenous across three subdivisions of the cerebellum that were chosen to correspond to the distribution regions of the three main cerebellar arteries. The cerebellar cardiac-response function was similar to that observed for the neocortex (Chang *et al.*, 2009), with the difference that the neocortical response function demonstrated a more pronounced dip at 12 s.

There was substantial variation in heart rate independent of behavior (see Figure 4.6). As can be seen in Figure 4.4, the heart rate drops every time there is a break between blocks of trials. To determine how much the heart rate correlated with the behavioral regressors (red curve in Figure 4.4), we used a multiple regression analyses, with the heart rate at specific lag as a dependent variable and the 8 behavioral regressors as independent variables. Maximally, the heart rate with a lag of 6s could be explained with R^2=0.22 (between person SD=0.075) by the behavioral regressors combined (and respectively lower for other lags). In turn, a multiple regression analysis of each behavioral regressor onto the 12 heart rate regressors showed an R^2 of 0.10 (SD=0.09). Of critical interest, however, is the contrast between error trials and

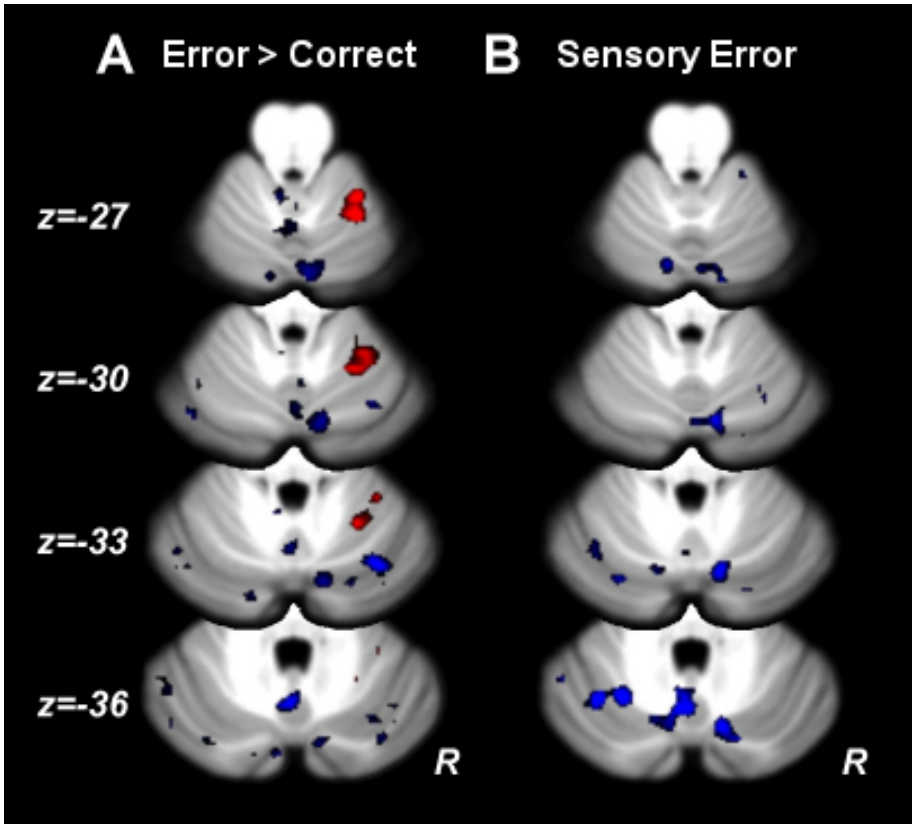

Figure 4.5: MRI after heartrate correction. A) The effect of errors. Shown in red is the region in Right Lobule V that is more active during error trials than during correct trials. In blue are regions more active for correct trials than error trials. B) Errors conveyed by a sensory signal. In blue are regions more active during trials in which no sensory stimulus was present, or an expected sensory stimulus was present, compared to trials in which the sensory stimulus conveyed an error signal. Once again, no regions are selectively active during trials in which the sensory stimulus conveys the error.

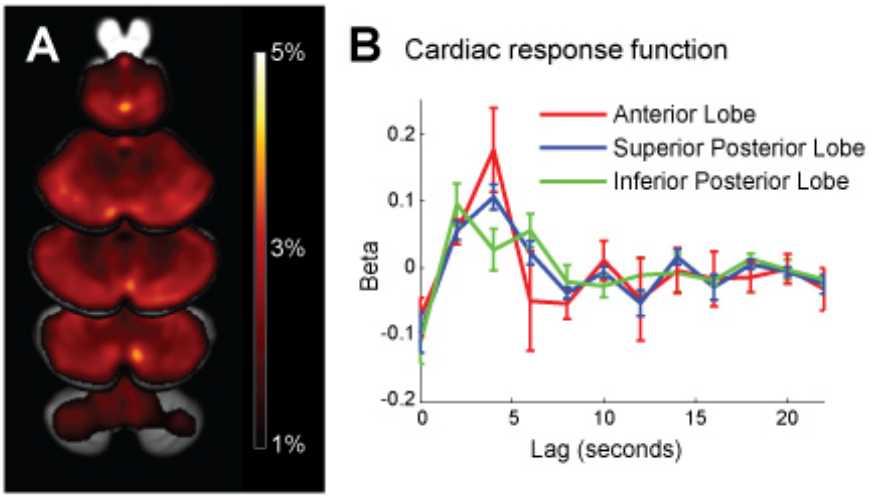

Figure 4.6: Effects of heart rate on cerebellar BOLD signal. A) Percentage of the high-pass filtered BOLD signal that is explained by the inclusion of heart rate regressors. The statistical map is shown on horizontal cerebellar slices, as in Figure 4.2. B) Average cardiac response function computed for three areas cerebellum, which roughly correspond to the areas of vascular supply.

correct trials. This contrast alone explains only 2% of the variance of the heart rate regressors with 6 second lag (R^2=0.02).

Does this dependence allow for a good estimation of the regression coefficients? The expected variance-covariance matrix of the regression coefficients is proportional to $(X^tX)^{-1}$, where X represents the design matrix. When one introduces nuisance regressors that are co-dependent with regressors of interest, the expected variance of the regressors of interest will increase. The expected variance of the error contrast increased on average by 16% (SD=13%), when comparing a linear model with and without the inclusion of the 12 heart rate regressors. Thus, while it is clear that parameter estimates will become slightly more variable, the heart rate showed enough variation independent of whether the preceeding trial contained an error, enabling us to estimate both accurately. Importantly, the second-level analysis accounts for any increases in parameter uncertainty. The denominator of the t-test (the between participants SE of the regression coefficients) is driven both by parameter uncertainty and true inter-subject variability, ensuring that our inferences are valid.

4.3.5 Behavioral Results

Because subtle differences in kinematic variables can induce substantial changes in the cerebellar BOLD signal, here we attempt to rule out any such behavioral confounds. The kinematic variables movement time, peak velocity, and movement amplitude were measured during fMRI acquisition and analyzed using a 2x2 repeated measures ANOVA with factors of Feedback Expectation (expecting a stimulus or not) and Movement Outcome (receiving a stimulus or not). Movement errors will manifest as an interaction in such an analysis. No significant main effects or interactions were observed for movement time (expectation: $F(1,9) = 1.7$, $p=0.23$; outcome: $F(1,9) = 0.04$, $p=0.85$; interaction: $F(1,1)=1 * 10^{-}9$) or peak velocity (expectation: $F(1,9)=1.22$, $p=0.30$; outcome: $F(1,9)=0.07$, $p=0.8$; interaction: $F(1,1)=4 * 10^{-}8$). For movement amplitude, a significant main effect of movement outcome was observed ($F(1,9)=66$, $p<0.0001$), as the force pulse provided a small decelerating impulse to the hand, lowering the mean amplitude by 1.6 cm on average (sd=0.64). The main effect of expectation was marginally reliable ($F(1,9)=4.32$, $p=0.06$), but, critically, the two factors did not interact ($F(1,1)=7 * 10^{-}8$). Thus, not only were the interactions terms non-significant, but they were very close to zero, indicating that the comparison of error vs. no-error trials was well-matched with respect to all basic kinematic variables.

4.4 Discussion

The current investigation set out to examine error-related activity within the cerebellar cortex during a simple reaching task. Despite evidence that the cerebellum is critical for learning from motor errors (Martin *et al.*, 1996; Tseng *et al.*, 2007; Maschke *et al.*, 2004; Smith & Shadmehr, 2005; Morton & Bastian, 2006), previous neuroimaging studies have failed to reveal activation patterns consistent with error processing when online corrections are carefully controlled for (Diedrichsen *et al.*, 2005b; Krakauer *et al.*, 2004).

An initial linear analysis, similar to that performed in most fMRI studies to date, revealed a dramatic and widespread decrease in BOLD signal following error trials. The diffuse spatial distribution of the decrease in BOLD response on error trials raised the possibility of a physiological artifact. The BOLD signal measured with fMRI is highly dependent on blood flow and oxygenation (Hoge *et al.*, 1999a,b), making it

vulnerable to variations in breathing and heart rate. Recent methodological advances have characterized these effects and demonstrated that these extraneous physiological influences can be reliably corrected (Glover *et al.*, 2000; Birn *et al.*, 2008; Chang *et al.*, 2009; Shmueli *et al.*, 2007). Indeed, behavioral events have been observed to reliably impact physiological processes. Heart rate deceleration is observed when people anticipate making a response (Damen & Brunia, 1987) or terminate a prepared response (Jennings *et al.*, 1991, 1992). In the current study, movement errors induced a reliable decrease in heart rate. Previous investigations have also observed pronounced cardiac deceleration following movement errors (Jennings & van der Molen, 2002) or negative feedback about cognitive errors (Crone *et al.*, 2005). We thus recorded heartrate using ECG electrodes and respiration using a compression belt during imaging, and then examined these variables for changes that may be related to our behavioral conditions. Notably, the consequences of this deceleration have the opposite effect on the BOLD signal than the hypothesized changes in neural activity. Indeed, had we failed to record heart rate, we may have falsely concluded that movement errors lead to widespread depression of activity in the cerebellar cortex.

After correcting for the effects of changing heart rate, a very different picture emerged. We now observed an increase in the cerebellar BOLD response on error trials, limited to the hand area of lobule V. Moreover, the pattern of activity suggests the symmetric encoding of sensory prediction errors, with increased activity for both the presence of unexpected sensory stimulation or the absence of an expected sensory input. The identification of a neural correlate of prediction errors within the cerebellar cortex agrees well with prominent role of the cerebellum in error-based learning, a notion that is consistently supported by behavioral studies. Future investigations will be required to determine whether the cerebellum simple processes these errors, or whether the activity reflects adaptive changes in a predictive forward model of motor processes.

The findings by Horn *et al.* (2004) would have predicted that the error signal in the cerebellar cortex should have been asymmetric. Several possible explanations for the discrepancy between the two studies exist. First, in their study, the cat was moving the forepaw on trials in which a sensory stimulus was unexpectedly missing (the condition that failed to elicit climbing fiber activity), while the unexpected sensory stimulation occurred in trials where the movement was artificially shortened by inserting an obstruction and when no movement occurred. This difference in kinematics

may underlie the apparent asymmetry. In contrast, our study carefully matched the kinematics between these two types of errors.

Second, it is important to note that fMRI does not provide a direct assay on the underlying neural events. BOLD signal changes in the cerebellar cortex may reflect a combination of metabolic processes related to climbing fiber inputs, mossy fiber inputs, and intracortical processing. Studies in the anaesthetized rat have shown that stimulation of parallel fibers leads to an increase in the BOLD signal. Less is known about whether and how climbing fiber activity is reflected in the BOLD response. Lesion studies suggest that activity-dependent blood flow increases in the cerebellar cortex are reduced by 50% when the inferior olive is inactivated (Zhang *et al.*, 2003). In contrast, energy calculations indicate that complex spikes may only account for a small proportion of the total energy utilization within the cerebellum (Howarth *et al.*, 2009) and may not produce a measurable change in blood flow (Thomsen *et al.*, 2009). Based on these studies it is not unlikely that our results mainly reflect changes in mossy fiber rather than climbing fiber input.

A more direct test of the predictions of Horn *et al.* (2004) would be to examine activation within the inferior olive directly. Two recent studies have used fMRI to observe activity in an area close to the inferior olive (Xu *et al.*, 2006; Liu *et al.*, 2008). While we selected our MR slice angle to maximize coverage of the olive, we failed to detect error-related activity in this region (see B.4). This null result is tempered by the fact that nearly 25% of the variance of the imaging signal in this area was accounted for by the phase of the cardiac and respiration signals. Particularly in light of possible behaviorally evoked physiological effects, this makes it quite challenging to detect true neural-based variation in the BOLD response within the olivary complex.

There are numerous reports of the modulation of cardiac output during cognitive and movement tasks. A phasic deceleration of heart rate is the most common effect, presumably as a direct result of vagal innervation. Deceleration has been reported during various phases of movement (Damen & Brunia, 1987; Jennings *et al.*, 1991, 1992), including in response to movement errors (Jennings & van der Molen, 2002). In more cognitive tasks, negative feedback has been shown to produce a deceleration in heart rate, provided that feedback provides reliable information relevant to learning (Crone *et al.*, 2005). The current results raise a cautionary flag for interpreting fMRI results when heart rate is likely to vary with the experimental condition. Careful monitoring of breathing and heart rate should become a standard in the field of

functional neuroimaging.

Chapter 5

When size doesn't matter: The cerebellum is equally important for large and small visuomotor errors

Abstract

The cerebellum has long been recognized to play an important role in motor adaptation. Following cerebellar damage, impaired learning is observed in visuomotor adaptation tasks such as prism adaptation and force field learning. Both types of tasks involve the adjustment of an internal model in order to compensate for an applied perturbation. Recently, it was demonstrated that the way in which a force perturbation is applied can affect learning, with cerebellar partients able to learn more when the perturbation is applied gradually (Criscimagna-Hemminger *et al.*, 2010). In the current investigation, we asked whether participants with cerebellar ataxia might respond similarly to visuomotor perturbations. We introduced a rotation either all at once or in small steps. We observed a comparable deficit for the ataxics regardless of how the rotation was introduced, indicating that visuomotor learning may occur using different mechanisms than force field learning.

5.1 Introduction

The production of smooth, accurate movements in an unstable world is a challenging problem for the brain. The maintenance and flexible adjustment of an internal model is critical for maintaining such performance. Insight into the process of such adaptation comes from two major experimental tasks: force field learning (Fine & Thoroughman, 2006, 2007; Lackner & Dizio, 1994; Shadmehr & Mussa-Ivaldi, 1994; Taylor & Thoroughman, 2007; Thoroughman & Shadmehr, 2000) and visuomotor adaptation (Cheng & Sabes, 2007; Fishbach & Mussa-Ivaldi, 2008; Grafton *et al.*, 2008; Krakauer *et al.*, 2000, 2004; Mazzoni & Krakauer, 2006). In such tasks, after learning the production of accurate movements is possible despite the presence of an external perturbation applied by the experimenter.

Performance in such learning tasks is often approximated by the linear dynamical systems class of models (Cheng & Sabes, 2006). Being linear, these models suggest that the amount of learning that occurs from one trial to the next is directly proportional to the size of the error experienced (Thoroughman & Shadmehr, 2000). An efficient and neurally plausible solution to problems of this form given noisy input is the Kalman Filter. Indeed, evidence suggests that during adaptation tasks, human participants perform very similarly to the expected performance of the Kalman Filter (Burge *et al.*, 2008).

In typical adaptation experiments, following a series of movements in a baseline condition the external perturbation is suddenly applied. However, there is evidence that more stable performance in the perturbed environment can be obtained if the perturbation is applied gradually (Kagerer *et al.*, 1997). Notably, participants can learn to move more accurately under a 90 degree rotation if the perturbation is applied in steps of 10 degrees compared to all at once, and such learning occurs without explicit awareness of the presence of this transformation (Kagerer *et al.*, 1997).

In adaptation tasks, optimal performance seems to require the cerebellum. Participants with cerebellar damage show impaired learning in force field learning tasks (Maschke *et al.*, 2004; Smith & Shadmehr, 2005) as well as visuomotor adaptation (Martin *et al.*, 1996; Tseng *et al.*, 2007; Werner *et al.*, 2009), as well as other motor learning tasks (Diedrichsen *et al.*, 2005a; Morton & Bastian, 2006). In these studies, participants with cerebellar damage typically show some learning, though considerably less than controls.

Recently, Criscimagna-Hemminger *et al.* (2010) showed that when a force pertur-

bation was administered gradually, participants with cerebellar damage were able to learn to make straight trajectories, and even demonstrated an aftereffect. In contrast, when the force perturbation was administered suddenly, the same participants failed to produce straight movements. Such a finding may shed light on the specific cerebellar contribution to motor learning, suggesting that the cerebrum alone is capable of learning an internal model given the proper conditions. As such, the role of the cerebellum may be to assist cerebral structures in maintaining a novel mapping between desired outcome and muscle commands. When subtle changes alone are required, as is the case in the gradual adaptation, the cerebellum may not be necessary. Should this notion hold, then the cerebellar contribution to visuomotor adaptation may be quite similar: gradually introducing a visual perturbation should result in improved performance among participants with cerebellar damage. In the current study, we examined the cerebellar contribution to visuomotor adaptation applied either abruptly or gradually.

5.2 Experiment 1

5.2.1 Methods

Participants

8 individuals with cerebellar ataxia (see Table 5.1) and 8 age and education matched controls volunteered to participate in this study. Data from one control participant were excluded from analysis due to technical malfunction during data collection. All participants provided informed consent, and were compensated for their time in accordance with the University of California at Berkeley's Committee for the Protection of Human Subjects.

Apparatus

Participants made reaching movements on a table surface while wearing a cotton glove, which reduced friction between the hand and surface. A magnetic sensor (Ascension Flock of Birds, http://www.ascension-tech.com) was fixed to the glove near the nail of the index finger. Position (XYZ) of the sensor was continuously recorded to the nearest 0.5 mm at 138 Hz. The participants' view of their hand was occluded by a horizontal mirror 23.5 cm above the table's surface. A rear projection

Table 5.1: Information on ataxic participants

Subject	Age	ICARS	Diagnosis	Participated in Exp. 1	Participated in Exp. 2
1	45	36.5	Possible SCA2	yes	yes
2	58	29	Unknown	yes	yes
3	72	17.75	OPCA	yes	yes
4	55	47	SCA 3	yes	yes
5	62	8	SCA 6	yes	yes
6	49	23	Possible SCA 17	yes	—
7	61	23	SCA 6	yes	yes
8	48	n/a	Possible SCA 6	yes	—
9	50	51	SCA 6	—	yes
10	76	65.5	Unknown	—	yes
11	55	50	SCA 3	—	yes

screen (100 cm wide x 75 cm deep) was suspended 23.5 cm above the mirror, upon which stimuli were displayed from a cieling-mounted LCD projector. The visual display was calibrated such that the cursor displayed on the screen matched the position of the sensor on the table. Stimulus presentation and data collection were performed on a PC workstation running Windows 98 in DOS mode.

Task On every trial, participants began by locating the cursor within a 15 mm diameter region 21.5 cm from the table's edge. After maintaining the cursor within this start position for 1000 ms, a 15 mm white target appeared 12 cm away. The target appeared on average 10 degrees to the left of straight ahead, and was jittered on every trial by a random angle between -10 and 10 degrees. Participants were instructed to make a single reach attempting to end within the target. Online feedback of the cursor was present in a set of training trials (not recorded), and then withheld throughout the remainder of the experiment. Upon terminating the reach (tangential velocity below 4 mm/sec), a red dot appeared to provide feedback of the final position.

Participants performed 20 blocks of 20 trials each, with a short break provided between blocks. In the Single Step condition, participants first performed two blocks of the baseline rotation, followed by six blocks with the cursor rotated by a counterclockwise angle of 25 degrees (Figure 5.1). Two blocks with no rotation were then provided as a washout. In the Multi Step condition, the two baseline blocks were followed by a series of five blocks during which the same 25 degree rotation was gradually applied in steps of 5 degrees. A sixth adaptation block with a 25 degree rotation

was presented for comparison with the single step condition. Finally, two washout blocks with no rotation were performed, during which the aftereffect was assessed. All participants performed both conditions, with the order counterbalanced.

Data Analysis

Data analysis was performed in Matlab. Three kinematic measurements were computed for each block separately to assess basic features of the movement. Movement Time was defined as the duration between the moment the participant first exceeded a tangential velocity of 4mm/sec and the moment the participant fell below this velocity. Reach Amplitude was defined as the Euclidean distance between the position of the feedback dot and the initial hand position. Angular Variability was defined as the variance in the angular error (defined as the angle between the target, the start location, and the visual endpoint feedback) across a block. To assess the amount of learning, two measures were used. Residual Error was computed for each participant as the average angular error in the final 20 trials of the adaptation phase. Aftereffect was defined as the difference between the average angular error on the first 5 trials of the washout phase and the average angular error during the baseline phase.

5.2.2 Results

We first consider basic movement features during task performance (Figure 5.2), and then turn to learning. On average, control participants completed each reach in 296 ms (sd=76 ms), while ataxics completed each reach in 396 ms (sd=104 ms). We analyzed performance using a 2x10x2 repeated measures ANOVA, using within subject factors of Condition and Block and between subjects factor of Group. The only reliable effect was the group effect ($F(1,13)=4.76$, $p<0.05$). The Group x Condition interaction approached significance ($F(1,247)=3.566$, $p=0.06$), with ataxics moving slightly faster during the Multi step condition (390 ms for Multi Step compared to 403 ms for Single), and controls moving slightly slower (296 ms for Multi Step compared to 291 ms for Single). All remaining main effects and interactions were not significant (all $F<1.5$).

Reach amplitude did not differ grossly as a function of group, with ataxics reaching 12.7 cm (sd=1.0 cm) and controls reaching 12.4 cm (sd=1.0 cm). The same

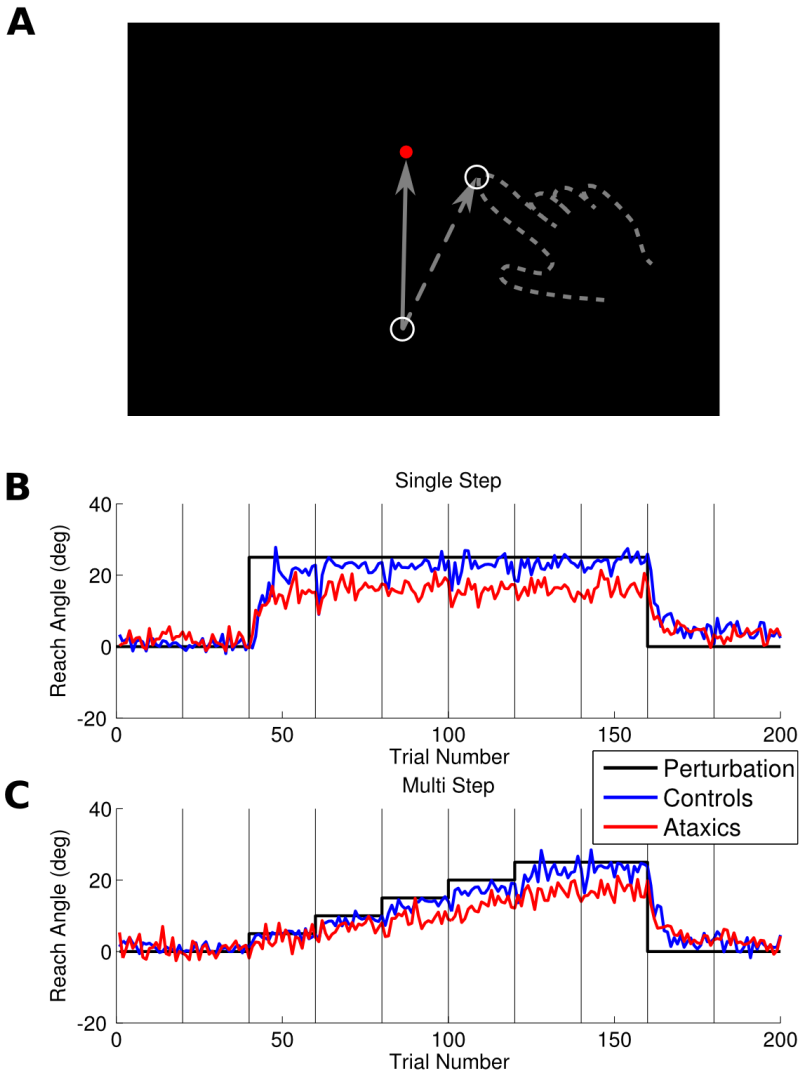

Figure 5.1: Experiment 1. A) Reaching task. Participants attempted to reach toward a circular target without vision of their hand or online cursor feedback. Upon reach termination, participants observed a red cursor. When rotation was turned on, this cursor was displaced counterclockwise from the hand. B) Trial order during the Single Step condition, where the rotation was suddenly applied. C) Trial order in the Multi Step condition, where the rotation was applied in steps of 5 degrees. In both B and C, Vertical black lines indicate the time of short breaks between blocks, which last for 20 trials. Shown in blue is the average performance of controls on this task, and shown in red is the average performance of ataxics.

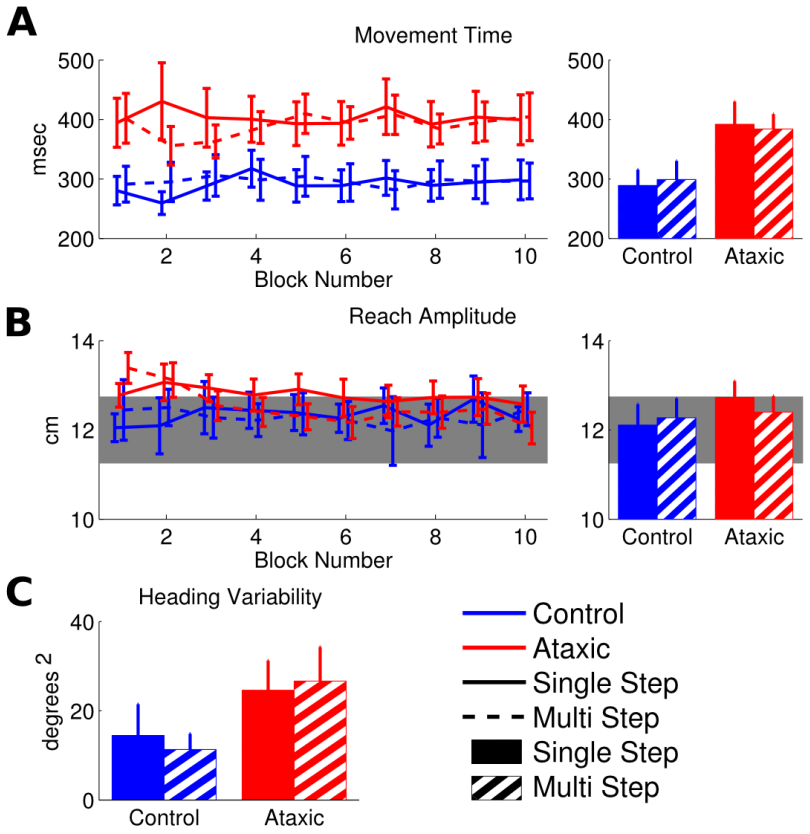

Figure 5.2: Kinematic measures from Experiment 1. A) Movement time over blocks. No change was observed over blocks, though on average ataxics take longer to complete the reach (bar graph). B) Reach amplitude over blocks. Again, no consistent changes were observed over blocks. The gray region represents the boundaries of the target. All participants were on average moving the proper distance. Ataxics tended to reach slightly farther, though this difference was not significant. C) Heading variability during the baseline phases. As no rotation was applied, no differences were observed across conditions. Controls show considerably less variable movement direction than ataxics, a feature which may lead to slower learning.

2x10x2 repeated measures ANOVA revealed a significant main effect of condition ($F(1,247)=5.158$, $p<0.024$). The difference, however, was quite slight, with participants reaching on average 0.17 cm further on the Multi Step block (sd=0.55 cm), and a post-hoc test (Tukey's HSD) failed to confirm that this difference was significant ($p>0.4$).

As Angular Variability is expected to be quite sensitive to learning, we only statistically analyzed the variability measurements during the baseline blocks. As there is no difference between Single and Multi Step conditions during these blocks, we collapsed across condition. Controls performed with Angular Variability of 9.2 degrees2, while ataxics performed with Angular Variability of 25.6 degrees2. This difference was significant ($t(13)=2.166$, $p<0.05$).

We next turn to the learning performance (Figure 5.3). During the Single Step condition, the ataxics performed worse than controls. Controls showed a residual error of 1.1 degrees (SD=0.8 deg), while Ataxics showed a residual error of 8.5 degrees (sd=3.4 deg), indicating that ataxics adapted less. The aftereffect was quite noisy, with controls showing an aftereffect of 12.0 degrees (sd=6.9 deg) and ataxics showing an aftereffect of 5.7 degrees (sd=3.7 deg). We next consider the Multi Step condition. Surprisingly, given the findings of Criscimagna-Hemminger *et al.* (2010), ataxics failed to adapt as much as controls in this condition. Controls showed a residual error of 1.9 degrees (sd=1.6 deg), while Ataxics showed a residual error of 8 degrees (sd=4.6 deg). Overall residual error was thus very close to that observed during the Single Step condition, with ataxics demonstrating less adaptation. Controls show an aftereffect of 10.5 degrees (sd=4.6 deg), while ataxics show an aftereffect of 7.6 degrees (sd=6.5 deg).

To statistically assess learning, we used a 2x2 repeated measures ANOVA, using Condition (Single or Multi) as a within subject factor and Group (Ataxic or Control) as a between subject factor. We compared the Residual Error and Aftereffect separately. For Residual Error, we observed a significant main effect of Group ($F(1,13)=25.1$, $p<0.001$), but no effect of condition ($F(1,13)=0.09$) nor interaction ($F(1,13)=1.1$, $p>0.3$). The same repeated measures ANOVA performed on the aftereffect revealed no significant effects of group ($F(1,13)=0.49$, $p>0.4$) nor condition ($F(1,13)=0.31$, $p>0.5$), and the interaction similarly failed to reach significance ($F(1,13)=0.01$).

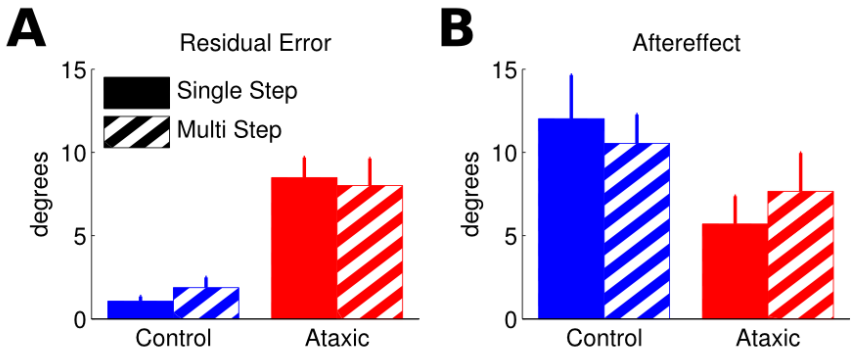

Figure 5.3: Learning measures in Experiment 1. A) During the final 20 trials at the maximum 25 degree rotation, ataxics show a markedly higher residual error than controls (shown as a positive value for convenience). No differences were observed as a function of the task, with participants performing similarly for both single and multi step conditions. B) During the first five trials after the rotation is removed, controls show a stronger aftereffect than the ataxics (nonsignificant difference). This measure is in general inversely correlated with residual error, as bars with higher residual error have a lower aftereffect, and vice-versa.

5.2.3 Discussion

Gradual introduction of a force perturbation has been shown to lead to better learning in patients with cerebellar ataxia (Criscimagna-Hemminger *et al.*, 2010). Surprisingly, this was not the case for the visuomotor rotation applied in the current study, as ataxics remained equally impaired. This may indicate a different role of the cerebellum in these two types of adaptation.

Previous work with visuomotor perturbations suggests that the gradual introduction of a visuomotor rotation results in greater aftereffects (Kagerer *et al.*, 1997), a prediction which was similarly not confirmed in the current investigation. Indeed, aftereffects were in general quite comparable for our participants across both conditions. Importantly, however, the former work used rotations of much greater magnitude (90 degrees), thus the failure to observe a difference may be a consequence of using a smaller rotation. A second important difference between the current experiment and both previous investigations is the number of targets used. Both previous investigations used more than one target. Using multiple targets should slow overall learning, as learning does not generalize perfectly across different targets (Donchin *et al.*, 2003). Importantly, the use of a single target in the current experiment may have facilitated

unlearning of the rotation, resulting in the aftereffects being too transient to measure with certainty. Though not statistically significant, there was a trend toward a relationship between the magnitude of the aftereffect and the residual error during the final 20 trials, with smaller residual errors leading to larger eftereffects. The negative correlation between these two learning measurements is often observed (see e.g. Tseng *et al.*, 2007). As we are numerically in good agreement with previous literature, then, we expect that the residual error provides a sufficient assessment of learning to explore group differences.

5.3 Experiment 2

In Experiment 1, the variability within the ataxics reaching performance was quite large compared to controls. This raises an important question about error attribution: when an error is made, is the error due to a change in the environment or is it random, due to uncontrollable factors such as implementation noise? If the motor system assumes that an error is caused by a change in the environment, then it would be advantageous to correct for it. If the error is caused by a random event within the arm itself, on the other hand, correcting it will not be as helpful. For the motor system to perform optimally, then, the motor system has to determine the proper correction to both compensate for a change in the environment and reduce the influence of random, uncontrollable events. If an error is more likely to result from uncontrollable factors, learning should be slower. Thus, the observed group difference in learning may be a side effect of the noisier motor apparatus. This prediction is made by the Kalman Filter: given an identical input signal, an optimal system should learn more slowly in the presence of greater output variability. As the Kalman Filter has been shown to provide a good model of human performance (Burge *et al.*, 2008), we performed a follow-up experiment to investigate whether the ataxics' increase in implementation variability could account for the observed learning deficits.

Experiment 2 was a variation on Experiment 1, with several changes implemented to facilitate detailed statistical modeling of behavior. To reduce possible effects of transient forgetting, we removed breaks between blocks of trials. We also altered the target and feedback to require a slicing movement, and provided feedback solely on the angle of the slice. Reach amplitude feedback was withheld to remove potentially distracting information. Previous work has found that this type of manipulation is

effective in reducing kinematic differences between ataxics and controls (Tseng *et al.*, 2007).

5.3.1 Methods

Participants

Nine individuals with cerebellar ataxia (see Table 5.1) and ten age and education matched controls volunteered to participate in this study. 6 of the ataxics and 1 control also participated in experiment 1, with more than three months separating the two experimental sessions.

Apparatus

Similar to Experiment 1, participants made reaching movements on a table surface while wearing a magnetic sensor. In contrast to the projector used in Experiment 1, the participants' view of their hand in Experiment 2 was occluded by a 19" LCD monitor placed horizontally approximately 17 cm above the table's surface, and calibrated such that the cursor displayed on the monitor matched the two-dimensional position of the sensor on the table. The monitor provided minimal interference with the operation of the magnetic transmitter, which was placed under the table.

Task

On every trial, participants began by locating the cursor within a 1 cm region approximately 15 cm from the table's edge. After maintaining a fixed position for 500-1000 ms, a white ring centered on the hand with a radius of 10 cm appeared. A white 1 cm radial line (the target) was displayed on this ring. Participants were instructed to make a single reach trying to slice through the ring at the position of the target. Online feedback of the cursor was present in a set of training trials (not recorded), and then withheld throughout the remainder of the experiment. When the subject's hand crossed the ring, a red line appeared to provide angular feedback of the final position.

During this experiment, participants performed trials in four phases, with a short rest between each phase. Following initial practice, participants performed in a baseline phase during which they produced 100 reaches without any applied rotation.

Next participants performed a No Feedback phase, during which they produced 100 reaches without any feedback about their final position. Participants then performed the Multi Step phase, in which they performed 32 reaches without rotation followed by 16 reaches at 5 degrees, 16 reaches at 10 degrees, 16 reaches at 15 degrees, 32 reaches at 20 degrees, 16 reachs at 15 degrees, 16 reaches at 10 degrees, 16 reaches at 5 degrees, and 32 reaches with no applied rotation (Figure 5.4). This was done to minimize awareness of the applied perturbation, though it does interfere with the ability to test for an aftereffect. Finally participants performed the Single Step phase, which required 32 movements with no rotation, 48 movements with a 20 degree rotation, 32 movements with no rotation, 48 movements with the same 20 degree rotation, and a final 32 movements with no rotation. In order to control awareness as much as possible, all participants performed Experiment 2 in the same order, completing the Multistep condition prior to the Singlestep condition.

Statistical Modeling

All modeling was performed in Matlab 2007b on PC workstations running Linux. Estimation of the external perturbation (the rotation) was modeled using a Kalman Filter. On each trial, the estimate of the rotation ($\hat{d}$) was updated based on the input (u) as follows:

$$\begin{aligned} \hat{d}_{n|n} &= \hat{d}_{n|n+1} + Ku_n \\ \hat{d}_{n+1|n} &= A\hat{d}_{n|n} \end{aligned} \tag{5.1}$$

Parameter A was treated as a free parameter, representing the state transition or memory of the system. The Kalman Gain (K) is based on the uncertainty at each step, defined as follows:

$$K = \frac{\sigma^2_{n+1|n}}{\sigma^2_{n+1|n} + \sigma^2_x} \tag{5.2}$$

The uncertainty parameter depends on the process variance (or equivalently, environmental instability) parameter σ_d, and is updated as follows:

$$\begin{aligned} \sigma^2_{n+1|n} &= A^2\sigma^2_{n|n} + \sigma^2_d \\ \sigma^2_{n|n} &= \sigma^2_{n|n-1} - K\sigma_{n|n-1} \end{aligned} \tag{5.3}$$

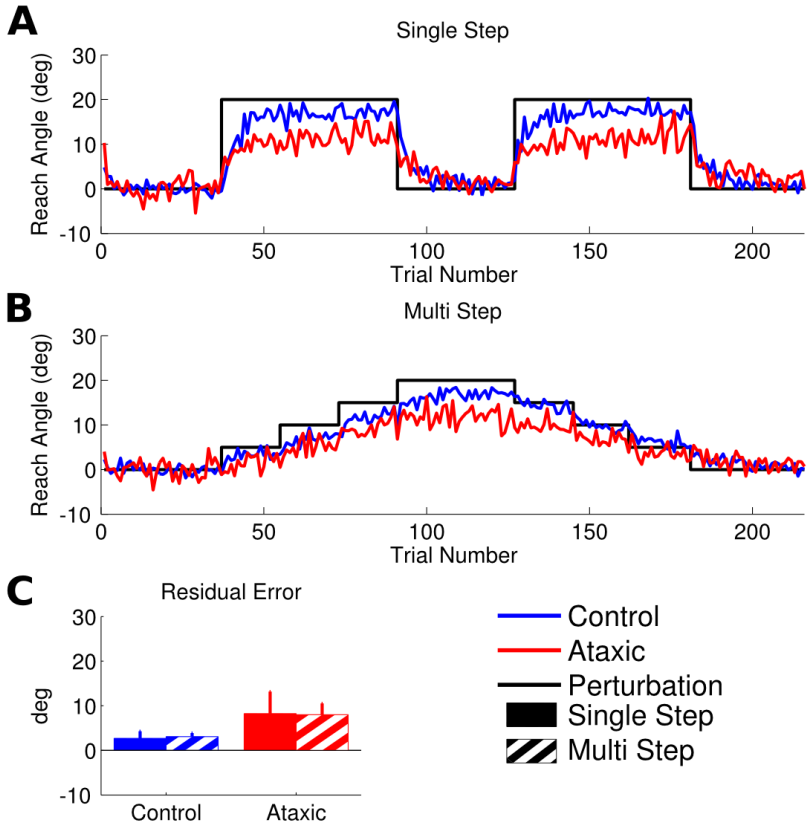

Figure 5.4: Experiment 2 and learning measures. A) The single step condition is shown. No breaks were provided between trials, and participants experienced a 20 degree rotation twice. Ataxics (red line) show attenuated learning compared to controls (blue line). B) The multi step condition is shown. The rotation is gradually introduced in steps of 5 degrees, and then gradually removed. Ataxics continue to show attenuated learning. C) The residual error during the final 16 trials at 20 degrees are compared. Ataxics show a larger residual error than controls, and no differences are observed across single and multi step conditions.

Parameter σ_x represents the participants' own implementation noise, which we estimated on a subject-by-subject basis using performance during the Baseline block. To facilitate data fitting, we made the simplifying assumption that both A and σ_d are constant parameters and that the system has settled to a steady state. We evaluated the fits using likelihoods computed using the innovations form of Shumway & Stoffer (2006), comparing nested models using the likelihood ratio test. See Appendix C for the solution to the steady state variance as well as the likelihood estimates.

5.3.2 Results

The kinematic variables Movement Time, Reach Amplitude, and Angular Variability during the Baseline blocks were again considered. Movement time was assessed by a 2x3 repeated measures ANOVA with between subject factor Group (Ataxic or Control) and within subject factor Condition (Baseline, Single Step, or Multi Step). A significant main effect of group was observe ($F(1,16)=4.6241$, $p<0.05$). The main effect of block type and the interaction approached significance ($F(2,32)=2.97$, $p=0.07$ and $F(2,32)=2.69$, $p=0.08$, respectively). Controls completed each reach in 156 ms (sd=50 ms), while the ataxics completed each reach in 309 ms (sd=136 ms). Reach Amplitude was assessed by the same ANOVA. A significant main effect of group was observed ($F(1,16)=7.90$, $p<0.02$), but the type of block had no effect ($F(2,32)=0.33$) and the interaction was not significant ($F(2,32)=0.72$).

Angular Variability, being sensitive to learning, was assessed only for the baseline blocks. As visual feedback about the endpoint was withheld on certain trials as a way to obtain a pure estimate of variability in the absence of any possible trial-by-trial adjustments, we analyzed the variability data using a 2x2 repeated measures ANOVA with within subject factor Feedback (Present or Not) and between subject factor Group. A significant main effect of group was observed ($F(1,16)=6.13$, $p<0.05$), with controls showing an average variability of 10.5 degrees^2 (sd=7.5 deg^2), and ataxics showing an average variability of 29.2 degrees^2 (sd=19.3 deg^2). Feedback did not significantly affect endpoint variability ($F(1,16)=0.26$), and the interaction between feedback and group did not reach significance ($F(1,16)=2.56$, $p=0.13$).

Examining the residual error during the final 16 trials at the full rotation (trials 185 through 200 for the Single Step condition, and trials 111 through 126 for the Multi Step condition; see Figure 4), we replicated the findings of Experiment 1. Controls show a residual error of 2.7 degrees (sd=1.6 deg) during the Single Step condition, and

3.1 degrees (sd=0.8 deg) during the Multi Step condition. Ataxics show a residual error of 8.2 degrees (sd=4.9 deg) during the Single Step condition, and 8.0 degrees (sd=2.5 deg) during the Multi Step condition. A repeated measures 2 x 2 ANOVA with within subject factor condition (Single or Multi) and between subject factor group (Ataxic or Control) revealed a significant main effect of group ($F(1,16)=27.7$, $p<0.001$), and found no condition effect ($F(1,16)=0.01$) or interaction ($F(1,16)=0.09$). We turn now to the more detailed statistical regression using the Kalman Filter.

We used the estimate of Angular Variability as the implementation variance (σ_x) during model fitting. We tested four models. The first model, the null model, treated all participants as part of the same group, and did not distinguish between single- and multi-step conditions, using a single value of A and σ_d to model performance. We then tested two models against this null model. Each used a pair of values for A and σ_d. The first model fit separate parameter estimates to each group. The second model fit separate parameter estimates for the single and multi step conditions. A final model examined whether four values of A and σ_d performed significantly better than either pair of values, effectively splitting the data four ways. See Figure 5.5 for an overview of the model heirarchy.

Considering the first level of additional parameterization, the likelihood ratio test justified separating the data by group ($p<0.001$) as well as by condition ($p<0.001$). Moving to the next level, the likelihood ratio test further justified separating the data into the full group x condition separation (both $p<0.001$). The parameter estimates of the full model and their likely effect on behavior are considered below.

Parameter A represents the state memory across trials. The lower the value, the lower the asymptotic value of performance. The Ataxics, for both conditions, show a lower A than control participants, indicating that they have greater difficulty maintaining a new estimate of the displacement. During the Single Step condition, values of A were lower than during the Multi Step condition for both groups.

Parameter σ_d represents the added variance to the learning process. Increasing this value should lead to a higher Kalman gain, and thus faster learning (for fixed implementation variance). Interestingly, Ataxics show on average a higher σ_d than controls, and as Ataxics also show a higher Kalman gain this suggests that the σ_d increase is sufficient to outweigh the increased σ_x. The value of σ_d which fits the single step data best is greater than that for the multi step data, suggesting more rapid learning in this condition.

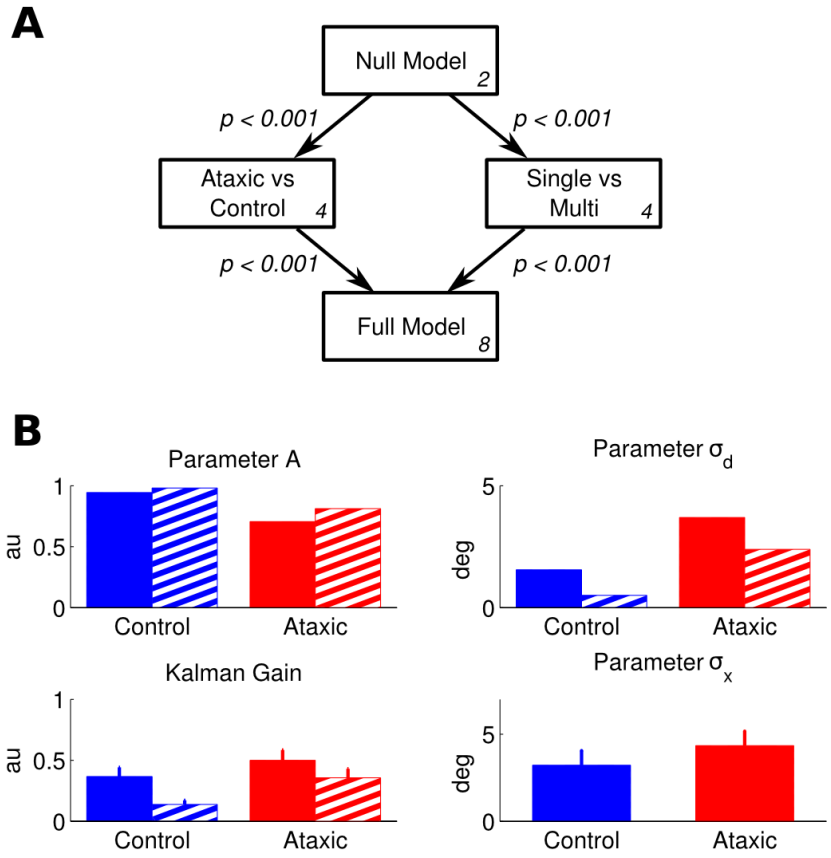

Figure 5.5: Statistical modeling of Experiment 2. A) The nested model heirarchy is shown. The null model (no group or condition differences), with two parameters, is at the top. At the next level, two four-parameter models are considered, splitting the data by group or by condition. Both divisions justified the additional parameters. At the final level, the full model with eight parameters, dividing by both group and condition, was significantly improved by the additional parameters. B) Most likely parameters from the full model. Shown on top are A and σ_d, the two free parameters in the model. On the lower row are the computed values of the Kalman Gain (using (5.2) and (C.5)) as well as the measured values of σ_x from the baseline block. Ataxics show a lower value of A, but a higher Kalman Gain than controls.

The values for both σ_d and A may trade off in some way, and have nonindependent effects on the likelihood. To investigate this further, we evaluated the log likelihoods for a set of values for σ_d and A using a grid search. Figure 5.6 shows heat maps for the log likelihood surfaces obtained (scaled to the maximal log likelihood for each condition for easy comparison), with contours representing proportions of the maximal log likelihood. Some correlation persists, as evident from the noncircular contours. If we examine the most likely σ_d as a function of A (Figure 5.6b), we see that this parameter suggest slower learning for the ataxics in both the single and multi step conditions for all levels of A, although the group difference in the multi step condition is smaller. If we similarly examine the most likely value of A as a function of σ_d, we observe that Ataxics show a lower value of A for all values of σ_d, suggesting that they always demonstrate incomplete learning.

5.3.3 Discussion

This experiment was conducted to examine the hypothesis that the ataxics' greater movement variability observed during the previous experiment may account for their lower performance in this visuomotor adaptation task. Within the assumptions of the Kalman Filter model, ataxics failed to perform optimally, as different values of the free parameters of interest were obtained between the groups for both single and multi step adaptation.

The detailed statistical modeling of participants' performance on this task can provide further insight into participants' performance. The relative magnitude of σ_d and σ_x is suggestive as to how participants will attribute their error. A relatively higher σ_x suggests that participants will attribute more of their error to their own system, while a higher σ_d suggests that participants will attribute more of their error to the environment and attempt to correct for it on the following trial. Optimal performance would require ataxics to attribute more of the error to their own system, resulting in a lower Kalman gain for ataxics. However, we actually observe a higher Kalman gain for ataxics than controls. There are two ways to interpret this finding.

First, a higher Kalman Gain suggests that the ataxics may be effectively underestimating their own motor noise during both single and multistep conditions. The value of σ_x which appears in the denominator of K (5.2) is more accurately defined as the system's *estimate* of implementation noise. We attempted to directly measure σ_x to facilitate model fitting. For our measured value of this parameter to be

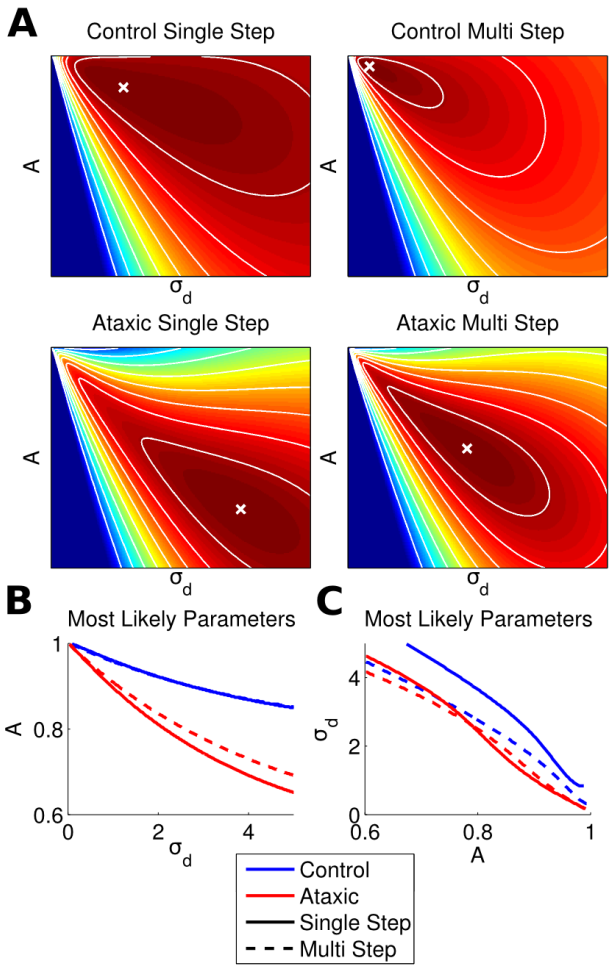

Figure 5.6: Exploration of the parameter space of Experiment 2. A) Likelihood is considered as a function of A and σ_d. Warmer colors indicate higher likelihood, cooler colors indicate less likely parameter combinations. The white x indicates the most likely parameter values, while the white contours represent equally likely combinations. As can be seen, the data are more likely explained by parameter combinations which fall along a diagonal. B) Searching over all values of $sigma_d$, the most likely values of the memory term A are plotted. Ataxics show a consistently lower value, with a slightly higher value during the multi step condition. C) Searching over all values of A, the most likely values of the stability parameter σ_d are plotted. For the same value of A, ataxics generally show a lower value of σ_d. During multi-step blocks, ataxics and controls are similar, but the fits are generally poor, as can be seen from the likelihood countours.

accurate, then, we necessarily assume that all participants have a consistent relationship between their *estimated* motor noise and their actual (measured) implementation variance. Should this assumption not be satisfied among the ataxics, then this model may not be valid to assess their performance. Nonetheless, the hypothesis tested in the current investigation was whether ataxics perform optimally given their increased motor noise. A misestimate of motor noise would be by definition suboptimal, and thus not in disagreement with our conclusions.

Alternatively, the higher Kalman Gain for ataxics may be a statistical artifact. A limitation of this analysis is the tradeoff between the parameters A and σ_d. This is quite evident when we consider the likelihood space as a whole (Figure 5.6). Interestingly, when we force the value of A to be the same across conditions and examine the most likely values of σ_d (Figure 5.6c), we see the predicted effect on learning, with ataxics showing a lower value of σ_d than controls, and thus necessarily a lower Kalman gain. This effect likely arises because of the nature of the linear systems model. In this model, a lower asymptotic value can arise through both lowering the value of A or a lowering the kalman gain. Indeed, fixing the value of σ_d across conditions reveals sharp differences in the most likely values of A (Figure 6 B), with ataxics showing a markedly lower value of A in order to obtain lower asymptotic values. Nevertheless, given the data collected, the ataxics remained impaired on both abrupt and gradual learning.

5.4 General Discussion

Participants with cerebellar ataxia generally show attenuated motor adaptation. This is observed in multiple domains, and taken as evidence for the functional role of this structure in the formation or adjustment of internal models of the environment. Recently, using a force perturbation, Criscimagna-Hemminger *et al.* (2010) suggested that cerebellar patients were better able to learn when the perturbation was introduced gradually. We performed two experiments to see whether these patients were similarly better able to learn a visuomotor perturbation when gradually introduced. In two experiments, we observed incomplete learning among participants with cerebellar ataxia, with no clear difference in overall performance related to whether the perturbation is introduced gradually or abruptly.

An important difference between the current investigation and the work of Criscimagna-

Hemminger *et al.* (2010) is the absence of online feedback in the current investigation. Online feedback is unavoidable in force field adaptation, while visuomotor transformation tasks can be run with only knowledge of results. Importantly, when online feedback is present, some amount of online error correction is possible. A previous investigation failed to demonstrate a substantial difference between the performance of cerebellar patients with and without online error correction (Tseng *et al.*, 2007), as these participants showed impaired learning in both conditions.

However, the opportunity to correct for errors online during visuomotor perturbations may affect learning in an important way. Recent work has suggested that online feedback can facilitate visuomotor learning by providing a secondary proprioceptive signal (Shabbott & Sainburg, 2010). This proprioceptive signal may perhaps enhance the ability of cerebral mechanisms to compensate for perturbations across trials. By this view, the correction of purely visuomotor errors across trials would be greatly facilitated by the cerebellum.

Visuomotor planning requires integration between cortical regions, and depends heavily on the parietal cortex (Clower *et al.*, 1996; Diedrichsen *et al.*, 2005b; Krakauer *et al.*, 2004). The correction of a visuomotor error, particularly across a delay, may require some higher level top-down control, regardless of whether the error is large or small. Small proprioceptive errors, on the other hand, may be more easily corrected without higher level interactions. Indeed, the proprioceptive system is used when producing normal movements (Cheng & Sabes, 2007; Sober & Sabes, 2003), particularly for feedback corrections of ongoing movements. Learning a gradually applied force perturbation may involve these automatic mechanisms, and thus require only very local cerebral connections.

The cerebellum is believed to be interconnected with both parietal and prefrontal cortex (Kelly & Strick, 2003; Krienen & Buckner, 2009; Middleton & Strick, 2001; Prevosto *et al.*, 2010; Ramnani *et al.*, 2006). Such connections may underlie deficits in alternative motor learning tasks following cerebellar damage, such as sequence learning (Carbon *et al.*, 2008; Gómez-Beldarrain *et al.*, 1998; Molinari *et al.*, 1997; Nixon & Passingham, 2000). Importantly, a cerebellar deficit in sequence learning is observed only when stimuli are indirectly related to the required action. If movements are cued directly, then cerebellar ataxics do show evidence of sequence learning (Spencer & Ivry, 2008). Thus, the cerebellar contribution to such tasks is likely assistive in nature rather than central to the learning, and when demands are low enough

learning can proceed without an intact cerebellum. If the cerebellum contributes similarly to reach adaptation, it would imply that visuomotor adaptation may be a more demanding task than force adaptation.

In summary, the cerebellar contribution to visuomotor adaptation is not sensitive to the speed at which the perturbation is introduced when performing without online feedback. Further experimental evidence will be required to determine if this is a fundamental difference between visuomotor rotation and force field learning more generally, or if the presence of online feedback is sufficient to distinguish these tasks.

References

Albus, J. 1971. A theory of cerebellar function. *Mathematical biosciences*, **10**(1-2), 25–61.

Allum, John H. J. 1975. Responses to load disturbances in human shoulder muscles: The hypothesis that one component is a pulse test information signal. *Experimental brain research*, **22**(3), 307–326.

Birn, Rasmus M., Smith, Monica A., Jones, Tyler B., & Bandettini, Peter A. 2008. The respiration response function: The temporal dynamics of fmri signal fluctuations related to changes in respiration. *Neuroimage*, **40**(2), 644–654.

Blakemore, S. J., Frith, C. D., & Wolpert, D. M. 2001. The cerebellum is involved in predicting the sensory consequences of action. *Neuroreport*, **12**(9), 1879–1884.

Bo, J., Block, H., Clark, J. E., & Bastian, A. J. 2005 (November). Can cerebellar motor performance be improved by explicit timing information? *In: 2005 abstract viewer/itinerary planner*. Society for Neuroscience.

Brody, Carlos D., Hernández, Adrián, Zainos, Antonio, & Romo, Ranulfo. 2003. Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cerebral cortex (new york, n.y. : 1991)*, **13**(11), 1196–1207.

Burge, Johannes, Ernst, Marc O., & Banks, Martin S. 2008. The statistical determinants of adaptation rate in human reaching. *Journal of vision*, **8**(4).

Carbon, M., Ghilardi, M. F., Argyelan, M., Dhawan, V., Bressman, S. B., & Eidelberg, D. 2008. Increased cerebellar activation during sequence learning in dyt1 carriers: an equiperformance study. *Brain : a journal of neurology*, **131**(Pt 1), 146–154.

Catz, N., Dicke, P. W., & Thier, P. 2005. Cerebellar complex spike firing is suitable to induce as well as to stabilize motor learning. *Curr biol*, **15**(24), 2179–2189.

Chang, Catie, Cunningham, John P., & Glover, Gary H. 2009. Influence of heart rate on the bold signal: The cardiac response function. *Neuroimage*, **44**(3), 857–869.

Cheng, Sen, & Sabes, Philip N. 2006. Modeling sensorimotor learning with linear dynamical systems. *Neural computation*, **18**(4), 760–793.

Cheng, Sen, & Sabes, Philip N. 2007. Calibration of visually guided reaching is driven by error-corrective learning and internal dynamics. *J neurophysiol*, **97**(4), 3057–3069.

Clower, D. M., Hoffman, J. M., Votaw, J. R., Faber, T. L., Woods, R. P., & Alexander, G. E. 1996. Role of posterior parietal cortex in the recalibration of visually guided reaching. *Nature*, **383**(6601), 618–621.

Collier, G. L., & Ogden, R. T. 2004. Adding drift to the decomposition of simple isochronous tapping: an extension of the wing-kristofferson model. *J exp psychol hum percept perform*, **30**(5), 853–872.

Criscimagna-Hemminger, Sarah E., Bastian, Amy J., & Shadmehr, Reza. 2010. Size of error affects cerebellar contributions to motor learning. *Journal of neurophysiology*, February.

Crone, Eveline A., Bunge, Silvia A., de Klerk, Phebe, & van der Molen, Maurits W. 2005. Cardiac concomitants of performance monitoring: context dependence and individual differences. *Brain research. cognitive brain research*, **23**(1), 93–106.

Cusumano, Joseph, & Cesari, Paola. 2006. Body-goal variability mapping in an aiming task. *Biological cybernetics*, **94**(5), 367–379.

Damen, E. J. P., & Brunia, C. H. M. 1987. Changes in heart rate and slow brain potentials related to motor preparation and stimulus anticipation in a time estimation task. *Psychophysiology*, **24**(6), 700–713.

Desmond, J. E., Gabrieli, J. D., Wagner, A. D., Ginier, B. L., & Glover, G. H. 1997. Lobular patterns of cerebellar activation in verbal working-memory and finger-tapping tasks as revealed by functional mri. *The journal of neuroscience : the official journal of the society for neuroscience*, **17**(24), 9675–9685.

Diedrichsen, J. 2006. A spatially unbiased atlas template of the human cerebellum. *Neuroimage*, **33**(1), 127–138.

Diedrichsen, Jörn, Verstynen, Timothy, Lehman, Steven L., & Ivry, Richard B. 2005a. Cerebellar involvement in anticipating the consequences of self-produced actions during bimanual movements. *Journal of neurophysiology*, **93**(2), 801–812.

Diedrichsen, Jörn, Hashambhoy, Yasmin, Rane, Tushar, & Shadmehr, Reza. 2005b. Neural correlates of reach errors. *The journal of neuroscience : the official journal of the society for neuroscience*, **25**(43), 9919–9931.

Diedrichsen, Jörn, Shadmehr, Reza, & Ivry, Richard B. 2010. The coordination of movement: optimal feedback control and beyond. *Trends in cognitive sciences*, **14**(1), 31–39.

Donchin, O., Francis, J. T., & Shadmehr, R. 2003. Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *J neurosci*, **23**(27), 9032–9045.

Ebner, Timothy J., Johnson, Michael T. V., Alexander Roitma, N., & Qinggong, F. U. 2002. What do complex spikes signal about limb movements? *Annals of the new york academy of sciences*, **978**(THE CEREBELLUM: RECENT DEVELOPMENTS IN CEREBELLAR RESEARCH), 205–218.

Fine, M. S., & Thoroughman, K. A. 2006. Motor adaptation to single force pulses: sensitive to direction but insensitive to within-movement pulse placement and magnitude. *J neurophysiol*, **96**(2), 710–720.

Fine, M. S., & Thoroughman, K. A. 2007. Trial-by-trial transformation of error into sensorimotor adaptation changes with environmental dynamics. *J neurophysiol*, **98**(3), 1392–1404.

Fishbach, Alon, & Mussa-Ivaldi, Ferdinando A. 2008. Seeing versus believing: conflicting immediate and predicted feedback lead to suboptimal motor performance. *The journal of neuroscience : the official journal of the society for neuroscience*, **28**(52), 14140–14146.

Fox, P. T., Raichle, M. E., & Thach, W. T. 1985. Functional mapping of the human cerebellum with positron emission tomography. *Proceedings of the national academy of sciences of the united states of america*, **82**(21), 7462–7466.

Glover, Gary H., Li, Tie-Qiang, & Ress, David. 2000. Image-based method for retrospective correction of physiological motion effects in fmri: Retroicor. *Magnetic resonance in medicine*, **44**(1), 162–167.

Golla, Heidrun, Tziridis, Konstantin, Haarmeier, Thomas, Catz, Nicolas, Barash, Shabtai, & Thier, Peter. 2008. Reduced saccadic resilience and impaired saccadic adaptation due to cerebellar disease. *The european journal of neuroscience*, **27**(1), 132–144.

Gómez-Beldarrain, M., García-Moncó, J. C., Rubio, B., & Pascual-Leone, A. 1998. Effect of focal cerebellar lesions on procedural learning in the serial reaction time task. *Experimental brain research*, **120**(1), 25–30.

Grafton, Scott T., Schmitt, Paul, Van Horn, John, & Diedrichsen, Jorn. 2008. Neural substrates of visuomotor learning based on improved feedback control and prediction. *Neuroimage*, **39**(3), 1383–1395.

Grodd, W., Hülsmann, E., Lotze, M., Wildgruber, D., & Erb, M. 2001. Sensorimotor mapping of the human cerebellum: fmri evidence of somatotopic organization. *Hum brain mapp*, **13**(2), 55–73.

Helmuth, L. L., & Ivry, R. B. 1996. When two hands are better than one: reduced timing variability during bimanual movements. *Journal of experimental psychology. human perception and performance*, **22**(2), 278–293.

Hoge, R. D., Atkinson, J., Gill, B., Crelier, G. R., Marrett, S., & Pike, G. B. 1999a. Linear coupling between cerebral blood flow and oxygen consumption in activated human cortex. *Proceedings of the national academy of sciences of the united states of america*, **96**(16), 9403–9408.

Hoge, Richard D., Atkinson, Jeff, Gill, Brad, Crelier, Gérard R., Marrett, Sean, & Pike, G. Bruce. 1999b. Investigation of bold signal dependence on cerebral blood flow and oxygen consumption: The deoxyhemoglobin dilution model. *Magnetic resonance in medicine*, **42**(5), 849–863.

Holmes, Gordon. 1939. The cerebellum of man. *Brain*, **62**(1), 1–30.

Horn, K. M., Pong, M., & Gibson, A. R. 2004. Discharge of inferior olive cells during reaching errors and perturbations. *Brain res*, **996**(2), 148–158.

Howarth, Clare, Peppiatt-Wildman, Claire M., & Attwell, David. 2009. The energy use associated with neural computation in the cerebellum. *Journal of cerebral blood flow & metabolism*, **30**(2), 403–414.

Imamizu, Hiroshi, Miyauchi, Satoru, Tamada, Tomoe, Sasaki, Yuka, Takino, Ryousuke, Putz, Benno, Yoshioka, Toshinori, & Kawato, Mitsuo. 2000. Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature*, **403**(6766), 192–195.

Ito, M. 1972. Neural design of the cerebellar motor control system. *Brain research*, **40**(1), 81–84.

Ito, M. 1998. Cerebellar learning in the vestibuloocular reflex. *Trends in cognitive sciences*, **2**(9), 313–321.

Ito, Masao. 2001. Cerebellar long-term depression: Characterization, signal transduction, and functional roles. *Physiol. rev.*, **81**(3), 1143–1195.

Ivry, R. B., Keele, S. W., & Diener, H. C. 1988. Dissociation of the lateral and medial cerebellum in movement timing and movement execution. *Experimental brain research*, **73**(1), 167–180.

Ivry, Richard B., & Keele, Steven W. 1989. Timing functions of the cerebellum. *Journal of cognitive neuroscience*, **1**(2), 136–152.

Ivry, Richard B., Spencer, Rebecca M., Zelaznik, Howard N., & Diedrichsen, Jörn. 2002. The cerebellum and event timing. *Annals of the new york academy of sciences*, **978**(December), 302–317.

Jacobs, J. V., & Horak, F. B. 2007. Cortical control of postural responses. *Journal of neural transmission*, **114**(10), 1339–1348.

Jennings, J. R., van der Molen, M. W., Brock, K., & Somsen, R. J. 1991. Response inhibition initiates cardiac deceleration: evidence from a sensory-motor compatibility paradigm. *Psychophysiology*, **28**(1), 72–85.

Jennings, J. R., van der Molen, M. W., Brock, K., & Somsen, R. J. 1992. On the synchrony of stopping motor responses and delaying heartbeats. *Journal of experimental psychology. human perception and performance*, **18**(2), 422–436.

Jennings, J. Richard, & van der Molen, Maurits W. 2002. Cardiac timing and the central regulation of action. *Psychological research*, **66**(4), 337–349.

Kagerer, Florian A., Contreras-Vidal, J. L., & Stelmach, George E. 1997. Adaptation to gradual as compared with sudden visuo-motor distortions. *Experimental brain research*, **115**(3), 557–561.

Kawato, Mitsuo, Kuroda, Tomoe, Imamizu, Hiroshi, Nakano, Eri, Miyauchi, Satoru, & Yoshioka, Toshinori. 2003. Internal forward models in the cerebellum: fmri study on grip force and load force coupling. *Progress in brain research*, **142**, 171–188.

Kelly, Roberta M., & Strick, Peter L. 2003. Cerebellar loops with motor cortex and prefrontal cortex of a nonhuman primate. *The journal of neuroscience : the official journal of the society for neuroscience*, **23**(23), 8432–8444.

Kim, Jeansok J., & Thompson, Richard F. 1997. Cerebellar circuits and synaptic mechanisms involved in classical eyeblink conditioning. *Trends in neurosciences*, **20**(4), 177–181.

Kitazawa, Shigeru, Kimura, Tatsuya, & Yin, Ping-Bo. 1998. Cerebellar complex spikes encode both destinations and errors in arm movements. *Nature*, **392**(6675), 494–497.

Klockgether, T., Skalej, M., Wedekind, D., Luft, A. R., Welte, D., Schulz, J. B., Abele, M., Bürk, K., Laccone, F., Brice, A., & Dichgans, J. 1998. Autosomal dominant cerebellar ataxia type i. mri-based volumetry of posterior fossa structures and basal ganglia in spinocerebellar ataxia types 1, 2 and 3. *Brain : a journal of neurology*, **121 (Pt 9)**(September), 1687–1693.

Krakauer, J. W., Ghilardi, M. F., Mentis, M., Barnes, A., Veytsman, M., Eidelberg, D., & Ghez, C. 2004. Differential cortical and subcortical activations in learning rotations and gains for reaching: a pet study. *J neurophysiol*, **91**(2), 924–933.

Krakauer, John W., Pine, Zachary M., Ghilardi, Maria-Felice, & Ghez, Claude. 2000. Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J. neurosci.*, **20**(23), 8916–8924.

Krienen, Fenna M., & Buckner, Randy L. 2009. Segregated fronto-cerebellar circuits revealed by intrinsic functional connectivity. *Cerebral cortex (new york, n.y. : 1991)*, **19**(10), 2485–2497.

Lackner, J. R., & Dizio, P. 1994. Rapid adaptation to coriolis force perturbations of arm trajectory. *J neurophysiol*, **72**(1), 299–313.

Leon, Matthew I., & Shadlen, Michael N. 2003. Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron*, **38**(2), 317–327.

Lincoln, J., McCormick, D., & Thompson, R. 1982. Ipsilateral cerebellar lesions prevent learning of the classically conditioned nictitating membrane/eyelid response. *Brain research*, **242**(1), 190–193.

Liu, T., Xu, D., Ashe, J., & Bushara, K. 2008. Specificity of inferior olive response to stimulus timing. *J neurophysiol*, **100**(3), 1557–1561.

Luauté, Jacques, Schwartz, Sophie, Rossetti, Yves, Spiridon, Mona, Rode, Gilles, Boisson, Dominique, & Vuilleumier, Patrik. 2009. Dynamic changes in brain activity during prism adaptation. *The journal of neuroscience : the official journal of the society for neuroscience*, **29**(1), 169–178.

Marr, David. 1969. A theory of cerebellar cortex. *The journal of physiology*, **202**(2), 437–470.

Martin, T. A., Keating, J. G., Goodkin, H. P., Bastian, A. J., & Thach, W. T. 1996. Throwing while looking through prisms. i. focal olivocerebellar lesions impair adaptation. *Brain*, **119 (Pt 4)**(August), 1183–1198.

Maschke, Matthias, Gomez, Christopher M., Ebner, Timothy J., & Konczak, Jurgen. 2004. Hereditary cerebellar ataxia progressively impairs force adaptation during goal-directed arm movements. *J neurophysiol*, **91**(1), 230–238.

Mazzoni, Pietro, & Krakauer, John W. 2006. An implicit plan overrides an explicit strategy during visuomotor adaptation. *The journal of neuroscience : the official journal of the society for neuroscience*, **26**(14), 3642–3645.

McCormick, D. A., & Thompson, R. F. 1984. Neuronal responses of the rabbit cerebellum during acquisition and performance of a classically conditioned nictitating membrane-eyelid response. *J. neurosci.*, **4**(11), 2811–2822.

Medina, J. F., Garcia, K. S., Nores, W. L., Taylor, N. M., & Mauk, M. D. 2000. Timing mechanisms in the cerebellum: testing predictions of a large-scale computer simulation. *J neurosci*, **20**(14), 5516–5525.

Middleton, F. A., & Strick, P. L. 2001. Cerebellar projections to the prefrontal cortex of the primate. *The journal of neuroscience : the official journal of the society for neuroscience*, **21**(2), 700–712.

Molinari, M., Leggio, M. G., Solida, A., Ciorra, R., Misciagna, S., Silveri, M. C., & Petrosini, L. 1997. Cerebellum and procedural learning: evidence from focal cerebellar lesions. *Brain*, **120**(10), 1753–1762.

Molinari, Marco, Leggio, Maria G., Filippini, Valeria, Gioia, Maria C., Cerasa, Antonio, & Thaut, Michael H. 2005. Sensorimotor transduction of time information is preserved in subjects with cerebellar damage. *Brain research bulletin*, **67**(6), 448–458.

Morton, Susanne M., & Bastian, Amy J. 2006. Cerebellar contributions to locomotor adaptations during splitbelt treadmill walking. *J. neurosci.*, **26**(36), 9107–9116.

Müller, H., & Sternad, D. 2004. Decomposition of variability in the execution of goal-oriented tasks: three components of skill improvement. *J exp psychol hum percept perform*, **30**(1), 212–233.

Nixon, P. D., & Passingham, R. E. 2000. The cerebellum and cognition: cerebellar lesions impair sequence learning but not conditional visuomotor learning in monkeys. *Neuropsychologia*, **38**(7), 1054–1072.

O'Boyle, D. J., Freeman, J. S., & Cody, F. W. 1996. The accuracy and precision of timing of self-paced, repetitive movements in subjects with parkinson's disease. *Brain : a journal of neurology*, **119 (Pt 1)**(February), 51–70.

Paulin, M. G. 2005. Evolution of the cerebellum as a neuronal machine for bayesian state estimation. *Journal of neural engineering*, **2**(3), S219+.

Perrett, S. P., Ruiz, B. P., & Mauk, M. D. 1993. Cerebellar cortex lesions disrupt learning-dependent timing of conditioned eyelid responses. *J. neurosci.*, **13**(4), 1708–1718.

Pressing, J. 1998. Error correction processes in temporal pattern production. *Journal of mathematical psychology*, **42**(1), 63–101.

Prevosto, Vincent, Graf, Werner, & Ugolini, Gabriella. 2010. Cerebellar inputs to intraparietal cortex areas lip and mip: Functional frameworks for adaptive control of eye movements, reaching, and arm/eye/head movement coordination. *Cereb. cortex*, **20**(1), 214–228.

Pruessmann, K. P., Weiger, M., Scheidegger, M. B., & Boesiger, P. 1999. Sense: sensitivity encoding for fast mri. *Magnetic resonance in medicine*, **42**(5), 952–962.

Ramnani, Narender, Behrens, Timothy E., Johansen-Berg, Heidi, Richter, Marlene C., Pinsk, Mark A., Andersson, Jesper L., Rudebeck, Peter, Ciccarelli, Olga, Richter, Wolfgang, Thompson, Alan J., Gross, Charles G., Robson, Matthew D., Kastner, Sabine, & Matthews, Paul M. 2006. The evolution of prefrontal inputs to the cortico-pontine system: diffusion imaging evidence from macaque monkeys and humans. *Cerebral cortex (new york, n.y. : 1991)*, **16**(6), 811–818.

Rao, S. M., Harrington, D. L., Haaland, K. Y., Bobholz, J. A., Cox, R. W., & Binder, J. R. 1997. Distributed neural systems underlying the timing of movements. *The journal of neuroscience : the official journal of the society for neuroscience*, **17**(14), 5528–5535.

Ravizza, Susan M., Mccormick, Cristin A., Schlerf, John E., Justus, Timothy, Ivry, Richard B., & Fiez, Julie A. 2006. Cerebellar damage produces selective deficits in verbal working memory. *Brain*, **129**(2), 306–320.

Rijntjes, M., Buechel, C., Kiebel, S., & Weiller, C. 1999. Multiple somatotopic representations in the human cerebellum. *Neuroreport*, **10**(17), 3653–3658.

Robertson, S. D., Zelaznik, H. N., Lantero, D. A., Bojczyk, K. G., Spencer, R. M., Doffin, J. G., & Schneidt, T. 1999. Correlations for timing consistency among tapping and drawing tasks: evidence against a single timing process for motor control. *Journal of experimental psychology. human perception and performance*, **25**(5), 1316–1330.

Roux, Sébastien, Coulmance, Michèle, & Riehle, Alexa. 2003. Context-related representation of timing processes in monkey motor cortex. *The european journal of neuroscience*, **18**(4), 1011–1016.

Rüb, U., Gierga, K., Brunt, E. R., de Vos, R. A., Bauer, M., Schöls, L., Bürk, K., Auburger, G., Bohl, J., Schultz, C., Vuksic, M., Burbach, G. J., Braak, H., & Deller, T. 2005. Spinocerebellar ataxias types 2 and 3: degeneration of the pre-cerebellar nuclei isolates the three phylogenetically defined regions of the cerebellum. *Journal of neural transmission (vienna, austria : 1996)*, **112**(11), 1523–1545.

Schlerf, J. E., Spencer, R. M. C., Ivry, R. B., & Zelaznik, H. N. 2005 (November). The role of the cerebellum in the temporal control of continuous movements. *In: 2005 abstract viewer/itinerary planner*. Society for Neuroscience.

Schmahmann, J. D., & Sherman, J. C. 1998. The cerebellar cognitive affective syndrome. *Brain*, **121**(4), 561–579.

Scholz, J. P., & Schöner, Gregor. 1999. The uncontrolled manifold concept: identifying control variables for a functional task. *Experimental brain research*, **126**(3), 289–306.

Scott, Stephen H. 2004. Optimal feedback control and the neural basis of volitional motor control. *Nature reviews. neuroscience*, **5**(7), 532–546.

Shabbott, Britne A., & Sainburg, Robert L. 2010. Learning a visuomotor rotation: simultaneous visual and proprioceptive information is crucial for visuomotor remapping. *Experimental brain research. experimentelle hirnforschung. experimentation cerebrale*, March.

Shadmehr, R., & Mussa-Ivaldi, Fa. 1994. Adaptive representation of dynamics during learning of a motor task. *J. neurosci.*, **14**(5), 3208–3224.

Shmueli, Karin, van Gelderen, Peter, de Zwart, Jacco A., Horovitz, Silvina G., Fukunaga, Masaki, Jansma, J. Martijn, & Duyn, Jeff H. 2007. Low-frequency fluctuations in the cardiac rate as a source of variance in the resting-state fmri bold signal. *Neuroimage*, **38**(2), 306–320.

Shumway, Robert H., & Stoffer, David S. 2006. *Time series analysis and its applications: With r examples (springer texts in statistics)*. 2nd edn. Springer.

Smith, Maurice A., & Shadmehr, Reza. 2005. Intact ability to learn internal models of arm dynamics in huntington's disease but not cerebellar degeneration. *J neurophysiol*, **93**(5), 2809–2821.

Sober, S. J., & Sabes, P. N. 2003. Multisensory integration during motor planning. *J neurosci*, **23**(18), 6982–6992.

Spencer, Rebecca M., Ivry, Richard B., & Zelaznik, Howard N. 2005. Role of the cerebellum in movements: control of timing or movement transitions? *Experimental brain research. experimentelle hirnforschung. expérimentation cérébrale*, **161**(3), 383–396.

Spencer, Rebecca M. C., & Ivry, Richard B. 2005. Comparison of patients with parkinson's disease or cerebellar lesions in the production of periodic movements involving event-based or emergent timing. *Brain and cognition*, **58**(1), 84–93.

Spencer, Rebecca M. C., Zelaznik, Howard N., Diedrichsen, Jorn, & Ivry, Richard B. 2003. Disrupted timing of discontinuous but not continuous movements by cerebellar lesions. *Science*, **300**(5624), 1437–1439.

Spencer, Rebecca M C M., & Ivry, Richard B B. 2008. Sequence learning is preserved in individuals with cerebellar degeneration when the movements are directly cued. *Journal of cognitive neuroscience*, August.

Taylor, J. A., & Thoroughman, K. A. 2007. Divided attention impairs human motor adaptation but not feedback control. *J neurophysiol*, **98**(1), 317–326.

Thach, W. T. 1998. A role for the cerebellum in learning movement coordination. *Neurobiol learn mem*, **70**(1-2), 177–188.

Thomsen, Kirsten, Piilgaard, Henning, Gjedde, Albert, Bonvento, Gilles, & Lauritzen, Martin. 2009. Principal cell spiking, postsynaptic excitation, and oxygen consumption in the rat cerebellar cortex. *J neurophysiol*, **102**(3), 1503–1512.

Thoroughman, K. A., & Taylor, J. A. 2005. Rapid reshaping of human motor generalization. *J neurosci*, **25**(39), 8948–8953.

Thoroughman, Kurt A., & Shadmehr, Reza. 2000. Learning of action through adaptive combination of motor primitives. *Nature*, **407**(6805), 742–747.

Todorov, Emanuel, & Jordan, Michael I. 2002. Optimal feedback control as a theory of motor coordination. *Nature neuroscience*, **5**(11), 1226–1235.

Trommershäuser, Julia, Maloney, Laurence T., & Landy, Michael S. 2003a. Statistical decision theory and the selection of rapid, goal-directed movements. *J. opt. soc. am. a*, **20**(7), 1419–1433.

Trommershäuser, Julia, Maloney, Laurence T., & Landy, Michael S. 2003b. Statistical decision theory and trade-offs in the control of motor response. *Spatial vision*, **16**(3), 255–275.

Trommershäuser, Julia, Gepshtein, Sergei, Maloney, Laurence T., Landy, Michael S., & Banks, Martin S. 2005. Optimal compensation for changes in task-relevant movement variability. *The journal of neuroscience : the official journal of the society for neuroscience*, **25**(31), 7169–7178.

Trouillas, P., Takayanagi, T., Hallett, M., Currier, R. D., Subramony, S. H., Wessel, K., Bryer, A., Diener, H. C., Massaquoi, S., Gomez, C. M., Coutinho, P., Hamida, M. Ben, Campanella, G., Filla, A., Schut, L., Timann, D., Honnorat, J., Nighoghossian, N., & Manyam, B. 1997. International cooperative ataxia rating scale for pharmacological assessment of the cerebellar syndrome. the ataxia

neuropharmacology committee of the world federation of neurology. *J neurol sci*, **145**(2), 205–211.

Tseng, Ya-Weng W., Diedrichsen, Joern, Krakauer, John W W., Shadmehr, Reza, & Bastian, Amy J J. 2007. Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J neurophysiol*, May.

Vorberg, D., & Wing, A. M. 1996. *Modeling variability and dependence in timing.* San Diego: Academic. Pages 181–262.

Werner, Susen, Bock, Otmar, & Timmann, Dagmar. 2009. The effect of cerebellar cortical degeneration on adaptive plasticity and movement control. *Experimental brain research. experimentelle hirnforschung. expérimentation cérébrale*, **193**(2), 189–196.

Wing, A. M., & Kristofferson, A. B. 1973. Response delays and the timing of discrete motor responses. *Perception & psychophysics*, **14**(1), 5–12.

Wolpert, Daniel M., Miall, R. Chris, & Kawato, Mitsuo. 1998. Internal models in the cerebellum. *Trends in cognitive sciences*, **2**(9), 338–347.

Xu, Duo, Liu, Tao, Ashe, James, & Bushara, Khalafalla O. 2006. Role of the olivocerebellar system in timing. *J. neurosci.*, **26**(22), 5990–5995.

Zelaznik, H. N., Spencer, R. M., & Doffin, J. G. 2000. Temporal precision in tapping and circle drawing movements at preferred rates is not correlated: further evidence against timing as a general-purpose ability. *Journal of motor behavior*, **32**(2), 193–199.

Zelaznik, H. N., Spencer, R. M., Ivry, R. B., Baria, A., Bloom, M., Dolansky, L., Justice, S., Patterson, K., & Whetter, E. 2005. Timing variability in circle drawing and tapping: probing the relationship between event and emergent timing. *J mot behav*, **37**(5), 395–403.

Zelaznik, Howard N., Spencer, Rebecca M., & Ivry, Richard B. 2002. Dissociation of explicit and implicit timing in repetitive tapping and drawing movements. *Journal of experimental psychology. human perception and performance*, **28**(3), 575–588.

Zhang, Yi, Forster, Colleen, Milner, Teresa A., & Iadecola, Costantino. 2003. Attenuation of activity-induced increases in cerebellar blood flow by lesion of the inferior olive. *Am j physiol heart circ physiol*, **285**(3), H1177–1182.

Appendix A

Left arm Performance from Chapter 3

Participants in the experiment reported in Chapter 3 performed all movements with the right arm. In so doing, the introduction of the upward force perturbation resulted in a leftward shift of the trajectory. When this upward force was missing on catch trials in the test blocks, the trajectories deviated to the right. To investigate the nature of this horizontal displacement, we conducted a follow up experiment. A new group of 18 individuals participated in a nearly identical experiment to that reported in Chapter 3, except the movements were produced with the left arm. Half of the participants began with the 45 degree target, and half began with the 135 degree target. If the lateral deviation in response to the upward force perturbation was due to a biomechanical interaction between the arm and the manipulandum, then the direction of the deviation should be reversed when moving with the left arm.

Figure A.1 shows the results, following the format of Figure 3 in the main text. Comparing Figure 3.3c to Figure A.1c, the overall picture is quite similar, as the overall effect of the target angle and the trial type is preserved. However, the main effect of target type is reversed (see below, Figure A.3). Looking at the effect of the catch trial by comparing endpoints on catch trials to endpoints on pre-catch trials, when moving with the right arm an overall rightward shift is present. When moving with the left arm, we now see a leftward shift. Thus, the lateral displacement observed on catch trials is reversed.

Following conventions in Chapter 3, endpoint data when reaching with the left hand were analyzed statistically using a 2x2 repeated measures ANOVA, with within-

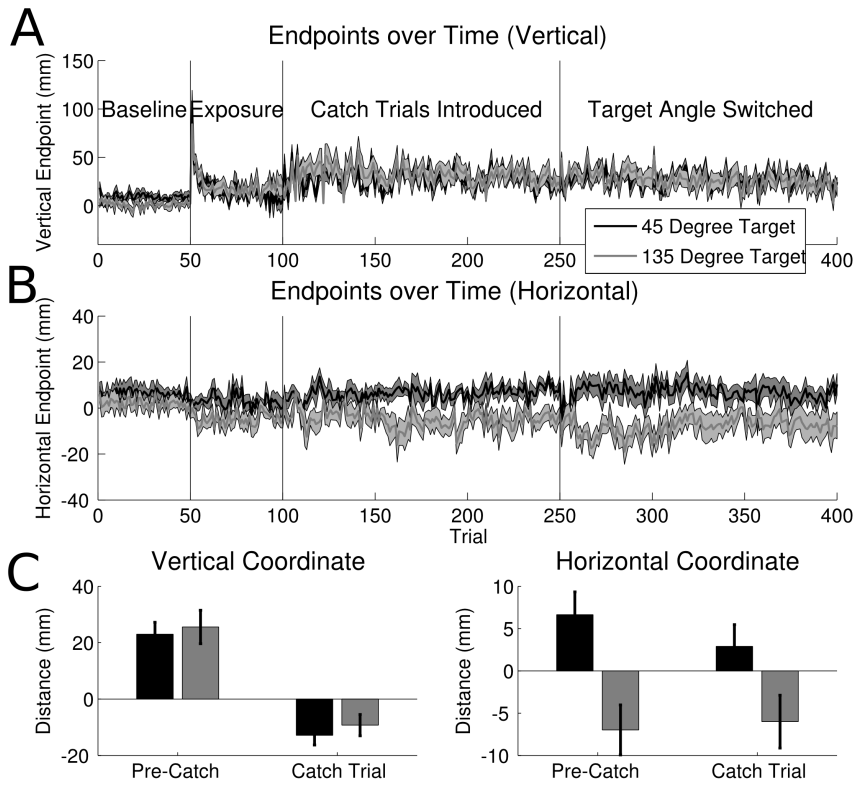

Figure A.1: When reaching with the left hand, participants demonstrate similar learning effects as was seen in the main experiment when participants reached with the right hand. An important difference occurs on catch trials (D), where participants end up further to the left, suggesting a biomechanical interaction with the manipulandum.

subject factors Target Angle (45deg/135 deg) and Trial Type (Precatch/Catch). For the vertical dimension, this analysis revealed a main effect of Trial Type ($F(1,17)=106.4$, $p<0.001$), confirming that the vertical endpoints were lower on catch trials compared to force field trials. There was no significant effect of target angle ($F(1,17)=1.96$, $p=0.18$) nor interaction ($F(1,17)=0.11$, $p=0.74$). For the horizontal dimension, a significant main effect of Target Angle was observed ($F(1,17)=8.64$, $p<0.01$). The main effect of Trial Type was not reliable ($F(1,17)=2.51$, $p=0.13$), but the interaction was highly significant ($F(1,17)=48.7$, $p<0.001$), providing a replication of a key finding from the main experiment.

Figure A.2 presents the trajectories during catch trials following the format of Figure 3.4. In Figure 3.4c and Figure A.2c, the overall effects of target orientation are similar. To facilitate the comparison between the two hands, we collapsed over the two targets in Figure A.3. An early rightward shift in the horizontal component is observed when reaching with either hand. The overall rightward bias observed in the main experiment when reaching with the right hand is now reversed when reaching with the left hand. The initial rightward bias may reflect interaction torques within the manipulandum. Nonetheless, the latter effect indicates a biomechanical interaction between the arm and manipulandum.

For completeness, we assessed the normalized position at 30 mm into the reach (the first point at which the force field, if present, would be nonzero) for both catch and post-catch trials when moving with the left hand. For this analysis, we used 2x2 repeated measures ANOVAs, with within-subjects factor of Target Angle and Trial Type (catch or post-catch). For the vertical coordinate, a highly significant main effect of Trial Type was observed ($F(1,17)=120.2$, $p<0.001$). The main effect of Target Angle was not significant, ($F(1,17)=1.11$, $p=0.31$), nor was the interaction ($F(1,17)=0.18$, $p=0.68$). For the horizontal coordinate, there was no significant main effect of Trial Type ($F(1,17)=0.59$, $p=0.45$) or Target Angle ($F(1,17)=1.4$, $p=0.26$). However, the interaction was significant ($F(1,17)=19.9$, $p<0.001$), replicating the main finding of Chapter 3.

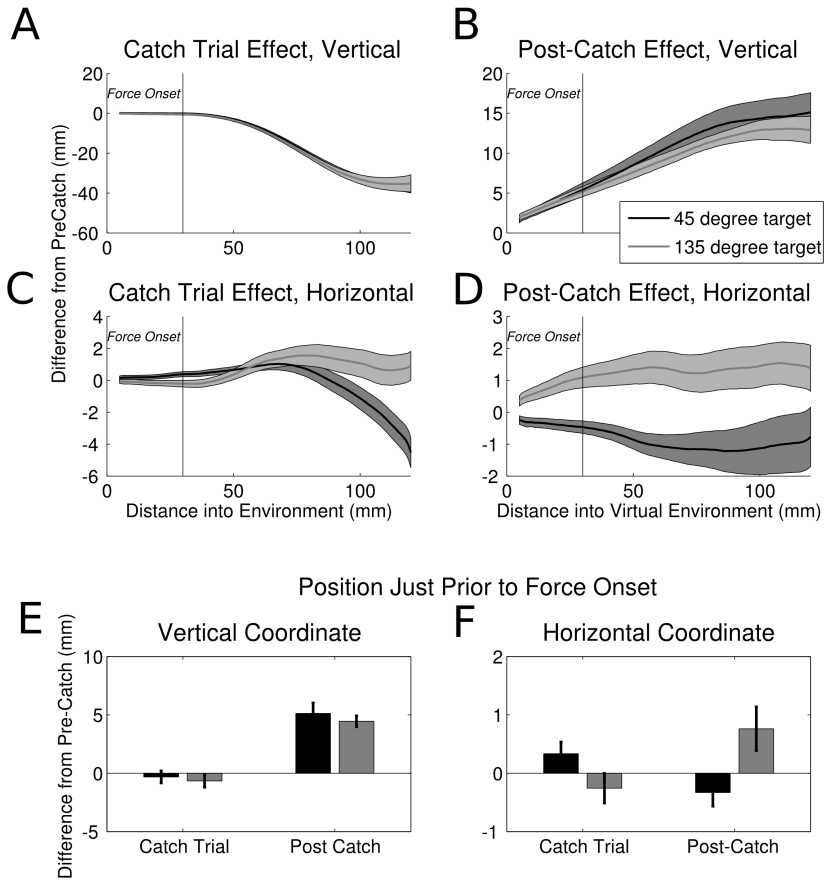

Figure A.2: When reaching with the left arm, the effects of the force field on the vertical coordinate (A, B) closely resemble those observed for the right arm. For the horizontal coordinate, stronger leftward corrections were observed on catch trials, despite a similar initial drift to the right (C).

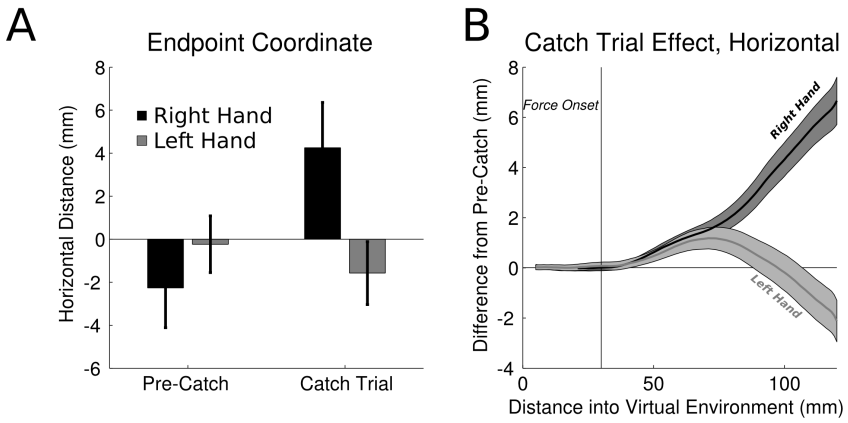

Figure A.3: Unexpected removal of the force field produced opposite effects on the horizontal component of the movement when considered independent of the target, shown here for both the endpoint (A) and the trajectory (B).

Appendix B

Additional MRI results from Chapter 4

Additional MRI results are presented here, to provide complete presentation of the data obtained. See Figures 4.2 and 4.5 in Chapter 4 for the main contrasts of interest, Errors vs. Correct Trials, as well as the contrast of Sensory Errors compared with all other trial types. Tables of the location of significant activation clusters (the three largest are presented when none are significant) from each of the fMRI contrasts are also provided. Table B.1 details activation clusters before heartrate is included in the model. Shown are the corrected statistical probability of observing a cluster of that size by chance, the volume of the cluster, the peak voxel, and the location (in mm from the anterior commissure) of the peak voxel. Table B.2 details activation clusters for the same contrasts after heartrate is included in the model.

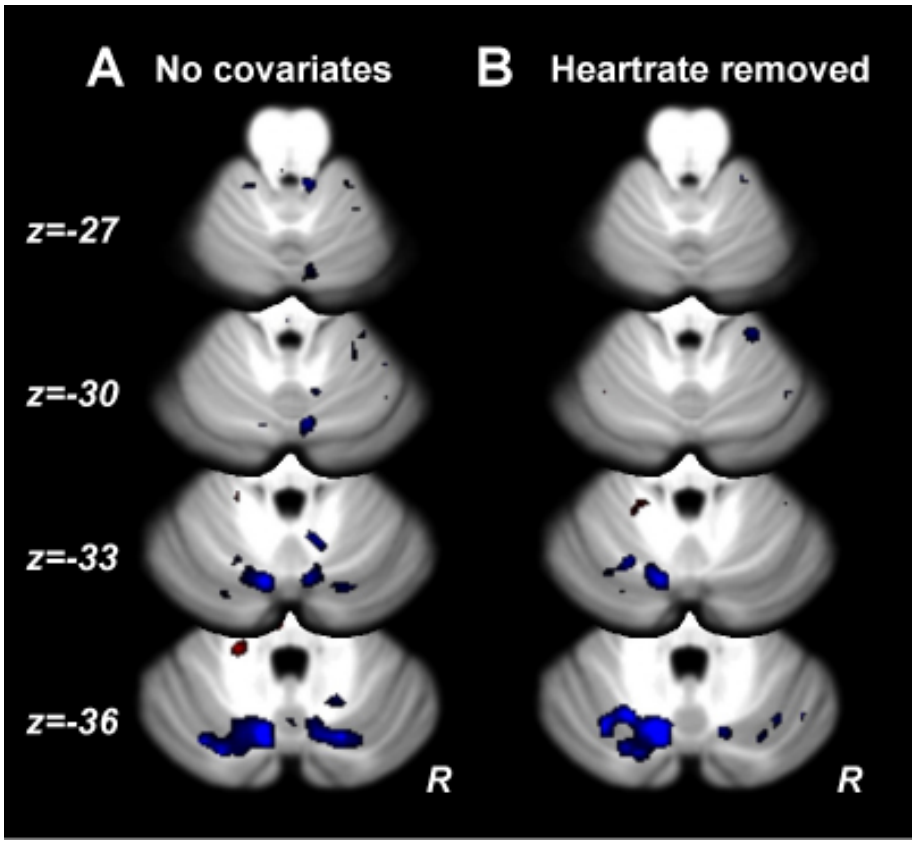

Figure B.1: Sensory stimulus vs. absence of stimulus. Regions in blue indicate greater activation when the somatosensory stimulus is absent, regardless of whether that absence indicates correct or incorrect performance. No regions were more active when the stimulus was present. However, this result is confounded by the increased movement amplitude when the sensory stimulus is absent.

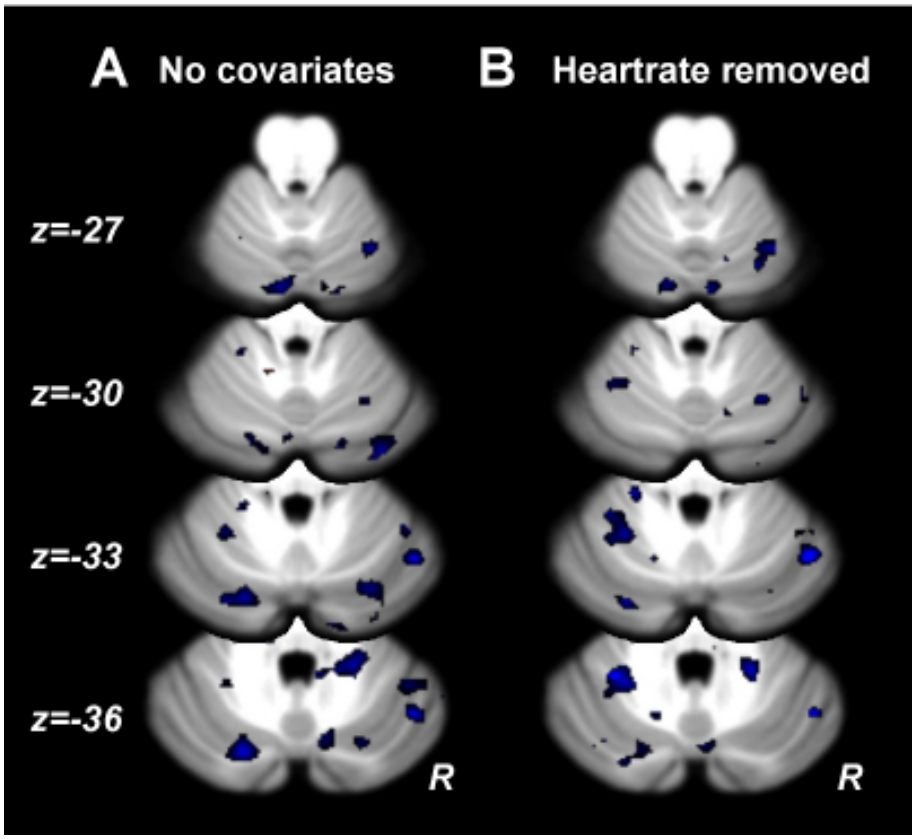

Figure B.2: Visible Cursor vs. Invisible Cursor. No regions were more active when the cursor and feedback were visible. Shown in blue are regions more active when the cursor and feedback were invisible. Very few regions were more active when the cursor was visible.

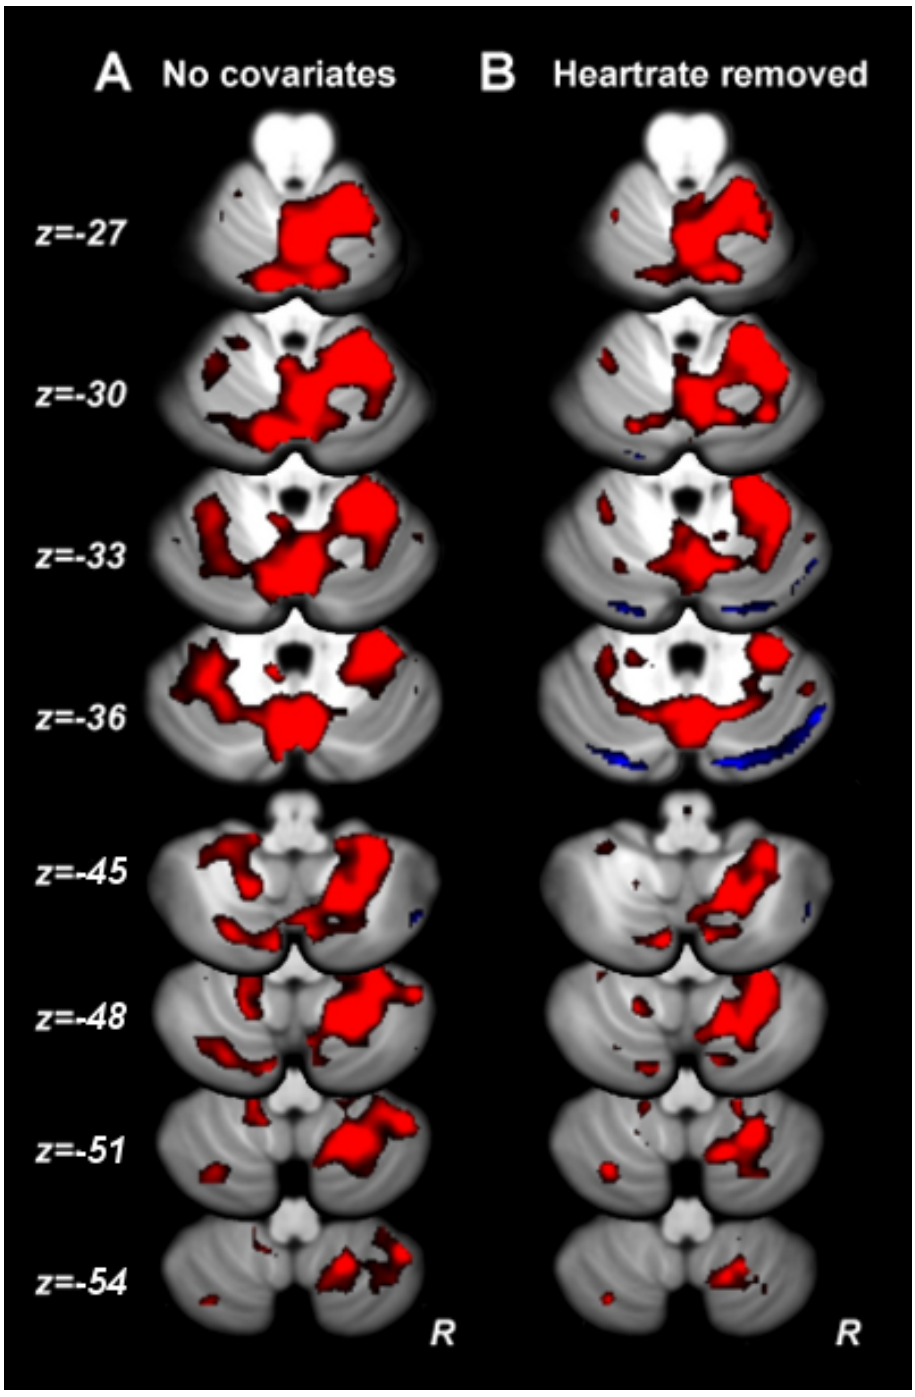

Figure B.3: Movement vs. Rest. Activation is observed in the anterior lobe as well as the inferior posterior lobe (lower slice). Activity is mainly observed in the right cerebellar hemisphere, ipsilateral to the moving limb.

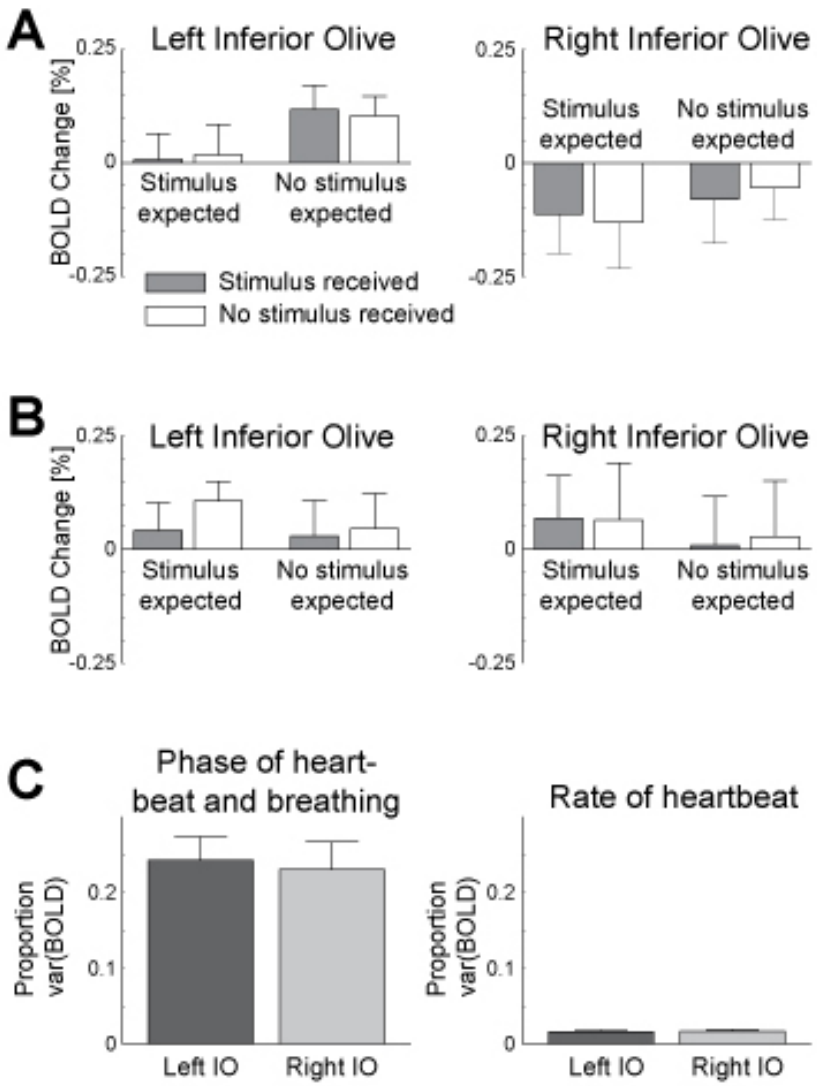

Figure B.4: ROI analysis of the inferior olive. No significant behavioral modulations on BOLD signal in this region were observed. A: Before heart rate correction, the right inferior olive showed an overall decrease in activation during movement, though this effect was not significant. B: Following heart rate correction, neither the right nor the left olive showed a significant increase or decrease in activation, though both showed a slight trend toward an increase. C: Including the phase of the heart and respiration reduced the BOLD variance by approximately 25%. Heart rate reduced the variance of the BOLD signal in this region by 2%.

Table B.1: Cluster-level statistics of group fMRI results without correcting for heart rate

Contrast	p(corr)	volume (mL)	Peak Voxel (Z)	MNI Coord (mm)
Error Trials > Correct Trials	1.000	16	2.47	(32, -40, -52)
	1.000	8	2.47	(14, -70, -50)
Correct Trials > Error Trials	0.000**	41960	4.88	(-22, -64, -42)
	0.844	264	3.71	(-20, -46, -18)
	0.925	216	3.50	(-36, -40, -48)
Unexpected Sensory Stim. > Other Trials	N/A	N/A	N/A	N/A
N/A				
Other Trials > Unexpected Sensory Stim.	0.000**	33040	5.38	(-6, -74, -22)
	0.256	600	3.85	(28, -36, -30)
	0.811	288	3.05	(-18, -42, -18)
Sensory Stimulus Trials > No Stimulus	0.961	176	3.13	(-18, -40, -30)
	1.000	64	2.75	(-16, -32, -28)
	1.000	40	2.83	(-4, -32, -30)
No Stimulus Trials > Sensory Stimulus	0.000**	3104	4.18	(-8, -72, -30)
	0.000**	2480	3.58	(4, -60, -40)
	0.624	336	3.22	(-38, -50, -42)
Visual Feedback > No Feedback	0.996	120	3.35	(22, -48, -44)
	1.000	32	2.61	(-14, -54, -62)
	1.000	24	2.67	(-10, -50, -24)
No Visual Feedback > Feedback	0.002**	1680	4.00	(18, -44, -32)
	0.004**	1488	3.73	(-18, -74, -34)
	0.095	768	3.86	(12, -70, -34)
Move > Rest	0.000**	49704	5.40	(22, -56, -52)
	0.997	120	2.53	(48, -56, -26)

Continued on Next Page...

Table B.1 – Continued

Contrast	**p(corr)**	**volume (mL)**	**Peak Voxel (Z)**	**MNI Coord (mm)**
Rest > Move	0.135	936	3.03	(46, -70, -44)
	0.999	88	2.78	(8, -22, -40)

Table B.2: Cluster-level statistics of group fMRI results after correcting for heart rate

Contrast	p(corr)	volume (mL)	Peak Voxel (Z)	MNI Coord (mm)
Error Trials > Correct Trials	0.001**	1000	3.66	(24, -48, -20)
Correct Trials > Error Trials	0.000**	2136	3.94	(-46, -62, -36)
	0.000**	1408	3.72	(48, -66, -36)
	0.000**	1192	3.62	(0, -62, -30)
	0.005**	808	3.10	(2, -50, -16)
	0.037*	560	3.47	(-2, -72, -40)
Unexpected Sensory Stim. > Other Trials	1.000	64	3.14	(-12, -26, -34)
	1.000	40	2.79	(20, -40, -34)
	1.000	8	2.41	(14, -78, -46)
Other Trials > Unexpected Sensory Stim.	0.000**	1664	3.70	(-6, -56, -36)
	0.000**	1880	4.14	(-24, -60, -30)
	0.001**	1088	4.01	(16, -76, -32)
	0.013*	744	3.31	(-28, -50, -34)
Sensory Stimulus Trials > No Stimulus	0.924	152	2.83	(-16, -44, -28)
	0.992	104	3.27	(24, -66, -40)
	1.000	48	3.18	(26, -50, -42)
No Stimulus Trials > Sensory Stimulus	0.000**	2840	4.01	(-12, -70, -32)
	0.308	336	3.15	(30, -72, -42)
	0.828	184	3.96	(-14, -84, -44)
Visual Feedback > No Feedback	1.000	16	2.89	(-10, -48, -50)
	1.000	16	2.82	(-22, -44, -42)
	1.000	8	2.72	(-10, -52, -20)
No Visual Feedback > Feedback	0.000**	2296	4.64	(-2, -32, -42)
	0.004**	896	3.61	(-28, -48, -30)

Continued on Next Page...

Table B.2 – Continued

Contrast	**p(corr)**	**volume (mL)**	**Peak Voxel (Z)**	**MNI Coord (mm)**
	0.016*	704	4.25	(14, -70, -36)
	0.024*	648	4.14	(-24, -78, -28)
Move > Rest	0.000	31032	5.31	(30, -42, -28)
	0.695	280	3.99	(-28, -68, -52)
	0.761	256	3.05	(-30, -44, -42)
Rest > Move	0.000**	3392	4.02	(16, -84, -30)
	0.004**	1288	4.04	(-22, -84, -30)

Appendix C

Kalman Filter Details from Chapter 5

C.1 Steady State Solution

Setting the model to steady state is equivalent to suggesting that the variance of both the time update and measurement update are constant, which is equivalent to (C.1):

$$\begin{aligned} \sigma^2_{n+1|n} &= \sigma^2_{n|n-1} \\ \sigma^2_{n|n} &= \sigma^2_{n-1|n-1} \end{aligned} \tag{C.1}$$

This removes the need to define a separate free parameter to represent the initial variance of the system. Ecologically, this assumption is reasonable since all participants are well-practiced in general at making reaching movements. Numerically, given reasonable values of σ_d, A, and σ_x any effect of any initial value of the system variance becomes miniscule after 10 trials. Thus, the system will have settled either during the unmodeled baseline phase or prior to the experimental initiation. Setting the system to steady state results in the following formula for the variance of the time update:

$$\sigma^2_{n+1|n} = A^2(\sigma^2_{n+1|n} - K\sigma^2_{n+1|n}) + \sigma^2_d \tag{C.2}$$

Substituting (5.2) for K gives:

$$\sigma^2_{n+1|n} = A^2(\sigma^2_{n+1|n}\frac{\sigma^2_{n+1|n}}{\sigma^2_{n+1|n} + \sigma^2_x}) + \sigma^2_d \tag{C.3}$$

Finally, solving (C.3) for $\sigma^2_{n+1|n}$ results in:

$$\sigma^2_{n+1|n} = \frac{A^2\sigma^2_x + \sigma^2_d - \sigma_x + \sqrt{(\sigma^2_x - A^2\sigma^2_x - \sigma^2_d)^2 + 4\sigma^2_x\sigma^2_d}}{2} \tag{C.4}$$

Similarly, solving for the measurement update results in:

$$\sigma^2_{n|n} = \frac{A^2\sigma^2_x + \sigma^2_d - \sigma^2_x + \sqrt{\sigma^4_d + 2\sigma^2_x\sigma^2_d(1 + A^2) + \sigma^4_d(1 - A^2)^2}}{2A^2} \tag{C.5}$$

C.2 Likelihood Estimation

The likelihood of receiving two particular observations on two trials is the product of the likelihood of receiving each observation individually (which is always less than 1). Since the current experiment involved hundreds of trials, the likelihood of any particular subject's performance is vanishingly small. The negative log likelihood was considered instead. This is computationally convenient, as the log likelihood of any two trials is the sum of the log likelihoods of the individual trials. Minimizing the negative log likelihood is equivalent to maximizing the likelihood.

Following the "Innovation Form" of Shumway & Stoffer (2006), the negative log likelihood of a single subject's performance ($-\log\Lambda_s$) is defined as follows:

$$-\log\Lambda_S = \frac{1}{2}\sum_{n=1}^{N}(\sigma^2_{n|n-1} + \sigma^2_x) + \frac{1}{2}\sum_{n=1}^{N}\frac{(y_n - \hat{d}_n)^2}{\sigma^2_{n|n-1} + \sigma^2_x} \tag{C.6}$$

y_n represents the subject's actual endpoint location on trial n, and $\hat{d}_n$ is the estimated rotation for the same trial as output from the Kalman Filter. Thus, $(y_n - \hat{d}_n)$ represents the error of the prediction. The negative log likelihood of the group performance is computed by summing over the individual subjects. To test a set of two nested models, one can use the likelihood ratio test. For sufficiently large samples of data, the negative of the logarithm of the likelihood ratio (the difference of the log likelihoods) follows a χ^2 distribution, with degrees of freedom equal to the number of additional parameters in the model.