

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

The Role of Multiple Learning Systems in Sensorimotor Adaptation of Human Reaching

Permalink

<https://escholarship.org/uc/item/49m484dg>

Author

Morehead, John Ryan

Publication Date

2017

Peer reviewed|Thesis/dissertation

The Role of Multiple Learning Systems in Sensorimotor Adaptation of Human Reaching

by

John Ryan Morehead

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Psychology

in the

Graduate Division

of the

University of California, Berkeley

Committee in Charge:

Professor Richard Ivry, Chair

Professor Jose Carmena

Professor Joni Wallace

Summer 2017

Abstract

The Role of Multiple Learning Systems in Sensorimotor Adaptation of Human Reaching

by

John Ryan Morehead

Doctor of Philosophy in Psychology

University of California, Berkeley

Professor Richard Ivry, Chair

Humans are very good at learning to make new movements, whether this is to practice a skill that many other people can perform or to overcome a new situation that they have never encountered. For instance, astronauts learn to maneuver in zero gravity and skydivers learn to precisely control falling with the poise of an acrobat. The same learning is evident in everyday life, as people regularly adjust for the small changes to their movements caused by articles of clothing, such as the additional weight of a watch on the forearm or the differences in gait necessary for many types of shoes. In motor learning research, it has been an open question whether learning a new skill, such as the controlled fall from skydiving, is the same as altering an existing motor skill, such as reaching, to compensate for the weight of a watch. In my dissertation work I have focused on the question of how and why people alter their existing motor skills, a type of learning called adaptation.

Adaptation is a specific subset of motor learning that occurs when the sensory outcome of motor commands is systematically altered. In order to induce this adaptation in the laboratory, we manipulate the visual feedback that human participants see when they are performing reaching tasks. It is thought that this type of learning, visuomotor adaptation, is driven by the difference between the feedback that was predicted to occur and the actual feedback. This discrepancy in feedback is known as a sensory prediction error. If present, these errors indicate that the sensorimotor system is not properly calibrated, and future motor commands (and their predicted sensory outcomes) are adjusted to bring the system back into alignment. Adjustments made to the motor commands by this process are historically believed to be independent of other factors that commonly affect learning, such as reward and punishment. It is becoming increasingly accepted, however, that the behavior observed in sensorimotor adaptation tasks may not only be the output of error-based adaptation. In the work that forms my dissertation, we attempted to characterize the effect of three different systems on behavior in visuomotor adaptation tasks.

In Chapter 1, we examined savings-upon-relearning in visuomotor adaptation tasks. Savings is the phenomenon of faster relearning after something has been forgotten. Visuomotor adaptation seems to be an ideal form of learning to study savings-upon-relearning, as participants can rapidly learn to compensate for altered visual feedback while also decaying fully to baseline behavior within a single experimental session. Following this “forgetting” of the motor memory, participants can then be re-exposed to the same visual perturbation; savings would be evident if

they compensate for the perturbation faster during re-exposure compared to the first learning event. This has been a conundrum for models of sensorimotor adaptation that function solely on sensory prediction errors, as the error size is the same for both learning episodes. If learning was only driven by these errors, it should proceed at the same rate both times. Here we examine the idea that this faster relearning comes from outside of the motor system and is not driven by sensory prediction errors, but rather an impetus to restore good task performance. Specifically, the results indicate that savings comes about because participants learn to implement a cognitive aiming strategy that helps them hit the target again. The difference in the rate of behavioral change arises because participants require time to develop the strategy when first encountering the altered visual feedback, but can then immediately implement it upon re-encountering the altered feedback.

In Chapter 2, we attempted to isolate the effects of error-based adaptation with a novel experimental manipulation. Participants were exposed to altered visual feedback and, unlike traditional adaptation studies, were fully informed of the nature of this alteration and explicitly told to ignore it. The specific visual feedback manipulation employed is known as a “visual error clamp,” where the visual cursor is set to a fixed heading angle. This means that no matter where the participant moves in the workspace, the feedback will always move in this direction instead of the direction of movement. We carefully manipulated the offset of the heading angle for this feedback relative to the direction participants were reaching in order to induce task-irrelevant sensory prediction errors. The only reason participants should adjust for these error clamps is if error-based learning is taking place given that they were told to ignore the feedback. We observed very robust adaptation in response to this manipulation. Surprisingly, the adaptation was consistent with that observed in typical adaptation studies in every way but one: the size of the change in behavior was not related to the size of the error. This is potentially a substantial challenge for theories of error-based adaptation, as they predict that there is either a linear or curvilinear relationship between error size and the magnitude of the adaptive response.

In Chapter 3, we explore the consequences that intrinsic biases have on visuomotor adaptation studies. When participants move without visual feedback, they often exhibit individual biases in the direction of their reaches. Here we show that there is a systematic bias for all participants, varying with the reach direction, and that it cannot be fully eliminated through visuomotor adaptation. This is because learning at any given reach direction is not fully independent of learning in other directions given that learning generalizes locally in the workspace. Furthermore if feedback is removed (a common manipulation in adaptation tasks), participants will drift back to this bias over time. If unaccounted for, this systematic bias (or its re-emergence) can be misinterpreted as a learning effect in adaptation tasks. We outline a few experimental and analytical techniques that can help account for this bias in these tasks so that future researchers can study adaptation without this contaminant.

Taken together, these studies show that many different processes contribute to the behavior of participants in sensorimotor adaptation tasks. These processes function with considerable independence and affect behavior in response to distinct stimuli. We have made an attempt to dissociate these processes primarily at a psychological level, a critical step for the investigation of the neural underpinnings of such processes.

To Mom, Dad, Robin, Drew,
Mallory and Gizmo

Table of Contents

List of Figures	iii
List of Tables	iv
Acknowledgements	v
Chapter 1. Savings Upon Re-Aiming in Visuomotor Adaptation	01
Chapter 2. A Limited Role for Sensory Prediction Errors in Visuomotor Learning	22
Chapter 3. Practical and Theoretical Consequences of Intrinsic Reaching Biases in Visuomotor Adaptation Experiments	44
References	64

List of Figures

1.1 Experimental task description	02
1.2 Savings is only observed with large perturbations	07
1.3 Gains from savings are abolished by instructions to aim directly for the target	09
1.4 Savings assessed with aiming report task in Experiment 3	12
1.5 Control of aiming strategy by an arbitrary color cue	14
1.6 Inappropriate recall of aiming strategies in response to a novel perturbation	17
2.1 A new method to induce implicit visuomotor adaptation	25
2.2 Aftereffects and generalization are similar following exposure to a standard visuomotor rotation and clamped visual feedback	31
2.3 Adaptation to task-irrelevant clamped visual feedback is attenuated in patients with cerebellar degeneration	33
2.4 Adaptation to clamped visual feedback does not scale with error size	34
2.5 Visual prediction error, contingent or not, drives adaptation	36
3.1 Human behavior and model performance	46
3.2 Initial bias and model performance after a large number of feedback trials	51
3.3 Simulation of target-specific washout effects with and without feedback	54
3.4 Simulated effects of target location on generalization functions	56
3.5 Simulated effects of subtracting the feedback and no-feedback baselines	58
3.6 Interaction of bias with generalization task trial designs, and the effects of visual feedback during baseline	60

List of Tables

2.1 Diagnostic information for patients with cerebellar ataxia	30
--	----

Acknowledgements

Science is a collaborative effort, and a great many friends and mentors helped me with this work. I would like to especially thank a few of my research assistants: Salman Qasim, Elizabeth Marrone, Dashel Thompson, and Shabnam Bonyadi.

My co-authors are listed below.

Chapter 1

John Ryan Morehead, Salman Qasim, Matthew Crossley, and Richard Ivry

Chapter 2

John Ryan Morehead, Jordan Taylor, Darius Parvin, and Richard Ivry

Chapter 3

John Ryan Morehead and Richard Ivry

Chapter 1

Savings Upon Re-Aiming in Visuomotor Adaptation¹

Abstract

Sensorimotor adaptation has traditionally been viewed as a purely error-based process. There is, however, growing appreciation for the idea that performance changes in these tasks can arise from the interplay of error-based adaptation with other learning processes. The challenge is to specify constraints on these different processes, elucidating their respective contributions to performance, as well as the manner in which they interact. We address this question by exploring constraints on savings, the phenomenon in which people show faster performance gains when the same learning task is repeated. In a series of five experiments, we demonstrate that error-based learning associated with sensorimotor adaptation does not contribute to savings. Instead, savings reflects improvements in action selection, rather than motor execution.

Introduction

Humans do a remarkably good job in compensating for natural and artificial perturbations of the sensorimotor system. This type of motor learning, adaptation, has served as a cornerstone in understanding motor control and motor memory. Studies of adaptation have highlighted an error-based form of learning in which change is driven by the difference between predicted and actual sensory feedback (Wolpert et al., 1995). The core feature of these algorithms is that they minimize sensory prediction error, dynamically moving toward a steady state in which the adapted motor commands will result in the intended sensory outcome.

On the whole, standard motor learning models that respond only to error do a good job of describing behavior in a range of adaptation tasks. However, these models cannot account for a well-described phenomenon in learning, savings (Krakauer et al., 2005; Zarahn et al., 2008). Savings is defined as faster learning when information is presented a second time after being forgotten, compared with the rate of initial learning (Ebbinghaus, 1913; Krakauer, 2009). This faster relearning has canonically been thought to reflect the existence of a “hidden” trace, or residual representation of the initial learning experience, even if the effects of this experience are no longer manifest. Variants of this idea have been advanced to describe changes in behavior (Ebbinghaus, 1913) or to characterize the functional role of physiological mechanisms of learning, such as the latent facilitation of synaptic potentiation (Medina et al., 2001).

Several classes of error-based models have been proposed to account for savings in sensorimotor adaptation tasks. One class centers on the idea that multiple error-based adaptation mechanisms operate in parallel, but at different rates (Smith et al., 2006; Zarahn et al., 2008). A second posits

¹ Previously published in *The Journal of Neuroscience*, Volume 35, Issue 42; October, 2015. doi.org/10.1523/JNEUROSCI.1046-15.2015. All figures and text reproduced with permission from The Society for Neuroscience

that people learn to switch between different internal models, one associated with the baseline context and at least one other associated with the perturbation context (Haruno et al., 2001; Lee and Schweighofer, 2009), or that performance reflects the combined output of multiple internal models (Pearson et al., 2010). A third class builds on the context idea but assumes that the learning process is modified by the context. For example, the large error encountered when a perturbation is reintroduced may entail the recruitment of additional error-based processes (Berniker and Körding, 2011) or lead to a change in the gain on the learning rate based on the direction and magnitude of past errors (Herzfeld et al., 2014). These models share the idea that savings arises from faster modification, or latent storage and subsequent recall, of internal models that specify how intended sensory goal states are translated into joint torques. In contrast, recent work has shown that the pattern of behavior observed in visuomotor savings experiments is often inconsistent with an error-based process, and may be better explained as reflecting changes in action selection that facilitate performance (Huang et al., 2011; Haith et al., 2015; Huberdeau et al., 2015; McDougle et al., 2015; Orban de Xivry and Lefèvre, 2015).

Behaviorally, it is often difficult to distinguish between performance gains that are due to faster adaptation of an internal model, recall of a strategy, or the combined effects of strategic and adaptation processes. Here, in a series of experiments using a converging set of measures, we test the hypothesis that savings arises from an action selection process rather than error-based adaptation. Specifically, we propose that, when a previously experienced perturbation is re-encountered, participants may achieve savings by recalling an aiming strategy that facilitated

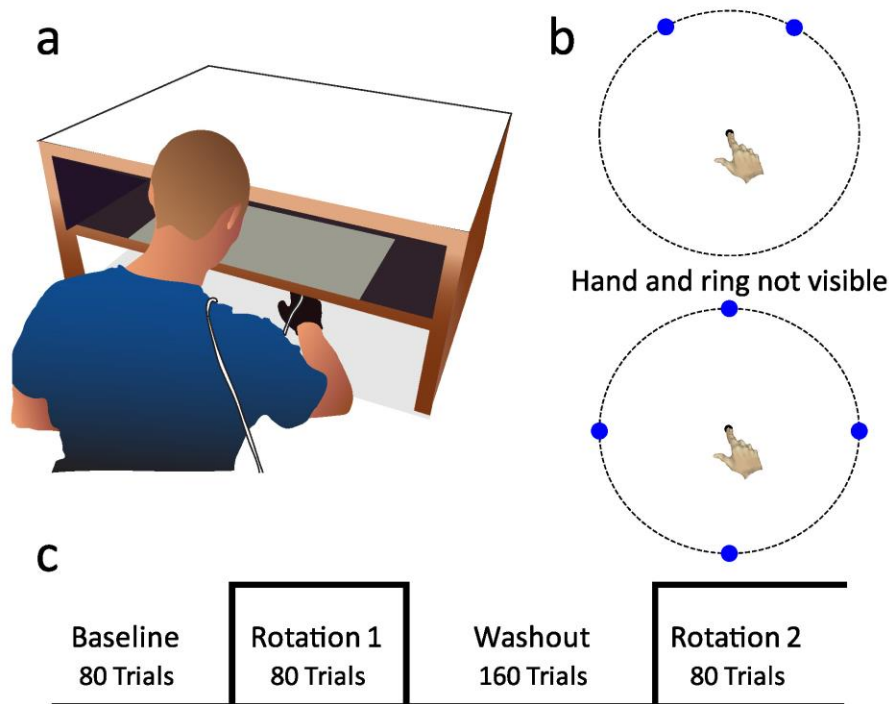


Figure 1.1 - Experimental task. a, Each participant slid his hand across a tabletop to hit visual targets. Vision of the hand was occluded by a mirror, which reflected a back-projected visual image to create the illusion that the hand and visual targets were coplanar. b, Two- (top) and four- (bottom) target sets used in experiments. On each trial, one target was pseudorandomly selected. The participant reached from a central start location, attempting to “slice” through the target with the cursor. A trial was complete when the radial distance of the movement exceeded 10 cm, indicated here by a dotted ring. c, Task design in Experiment 1 to assess savings. Veridical feedback was provided in the baseline and washout blocks. The visual feedback was perturbed during the two rotation blocks, with the size of the perturbation the same in each of these blocks.

performance when participants first encountered the perturbation. Recalling this strategy could allow for rapid improvements in performance, even in the absence of “true” adaptation (recalibration of a sensorimotor mapping).

Materials and Methods

Participants

A total of 130 participants (43 males, age 21 ± 2 years) were recruited from the undergraduate population at the University of California, Berkeley through the Department of Psychology research participation pool. All participants were right handed, as measured by the Edinburgh Handedness Inventory (Oldfield, 1971). Participants received either class credit or financial compensation for their participation. The experimental protocol was approved by the institutional review board of the University of California, Berkeley.

Experimental Apparatus

Participants reached to visual targets by sliding their hands across a table surface. Kinematic data were obtained via a six-axis magnetic position sensor (miniBird, Ascension Technology) attached to the index finger. A ceiling-mounted LCD projector projected the visual stimuli onto a screen that, through a reflecting mirror, created the illusion that the targets appeared in the same plane as the table surface (Fig. 1.1a). The experimental task was implemented using custom software written in Python 2.7.

Reaching task

On each trial, the participant reached from a central starting position (1.2 cm diameter circle) to a blue target (1.6 cm diameter circle). Targets were always centered at a distance of 10 cm from the start position but could appear at either one of two positions, 30° to the left or 30° to the right of midline, or at one of four target positions positioned 90° increments around the workspace (Fig. 1.1b). The target position for a given trial was pseudorandomly selected so that each target was presented 4 times out of every 8 trials (two-target conditions) or twice out of every 8 trials (four-target conditions). The projection screen and mirror precluded vision of the arm. Feedback of the position of the index finger was indicated by a small white cursor (0.8 cm diameter). The participant was instructed to “hit the blue target with your white dot, and make sure you move fast.” The feedback cursor was visible until movement amplitude exceeded 10 cm. At this point, the cursor froze for 2 s (Experiments 1, 2, 4, 5) or 0.5 s (Experiment 3), providing spatial feedback of the accuracy of the reach at the target amplitude. Participants were required to traverse the 10cm reach in 300 ms. If this criterion was exceeded, the computer played an auditory message, “too slow.” At the end of the feedback period, the screen blanked and the participant was required to return to the start position without positional feedback from the cursor. To guide the participant to the start position, a white ring with a radius equal to the distance of the hand from the start position was presented. This ring turned into the feedback cursor when the hand made contact with the central start position. The index finger had to remain within the start position for 200 ms before the presentation of the next target.

Experiment 1

Experiment 1 was designed to assess the effect of perturbation size on savings. Participants ($n = 10/\text{group}$) were randomly assigned to one of four perturbation groups: 15° , 30° , 45° , or 60° . Within each group, 10 participants completed 400 trials, divided into four phases (Fig. 1.1c). The initial block of 80 trials served as a baseline to assess reaching performance in the absence of a perturbation. This was followed by an 80-trial initial adaptation block in which the cursor feedback was rotated in the counterclockwise direction (Rotation 1), with the size of the perturbation varied between groups. In the washout block, the rotation was turned off for 160 trials, allowing performance to return to baseline. The final block served as our probe of savings; here, the perturbation was reintroduced for another 80 trials (Rotation 2).

Experiment 2

Experiment 2 examined whether the savings observed with large perturbations results from faster adaptation of an internal model or from changes in action selection. For this experiment, we tested three groups ($n = 10/\text{group}$): one exposed to a 45° rotation at the two target locations as Experiment 1, whereas the other two were exposed to either a 30° or 45° rotation at four targets spaced 90° apart around 360° . There were five blocks in this experiment. For the two-target group, the first four blocks were the same length and order as in Experiment 1, except that the Washout 1 block was only 100 trials long, and there was a 60 trial washout block after the Rotation 2 block. The perturbation size was fixed at 45° for all participants in the Rotation 1 and Rotation 2 blocks. The fifth block was a second washout block (Washout 2), composed of 60 trials without a perturbation. In the four-target condition, the structure of the trial blocks was identical, except that the number of trials was adjusted in each block to accommodate the greater number of targets while keeping the total number of trials similar. The four-target tasks began with a Baseline of 40 trials, followed by Rotation 1 (100 trials), Washout 1 (120 trials), Rotation 2 (100 trials), and finally Wash out 2 (40 trials). During the rotation blocks, participants were exposed to either a 30° or 45° visuomotor rotation. All other features of the four-target condition were identical to the two-target condition.

For all groups, we assessed savings over the initial trials of Rotation 2. To evaluate whether the participants were using an aiming strategy, the savings assessment was followed by a verbal intervention in the Rotation 2 block. After the seventh trial of this block in the two-trial condition, the program was paused. The experimenter instructed the participant that, “The perturbation that just turned on will now be off for the next two trials (see Fig. 1.3a). I want you to aim directly for the blue target and reach directly for the blue target on these trials.” Following these two reaches, the participant was told that “the perturbation is back on” and the participant was allowed to finish the task without further interruption. For the four-trial conditions, the verbal intervention was introduced after trial 12 in the Rotation 2 block and participants performed a series of four trials, one to each target. The rotation was then reimposed for the remainder of the block.

Experiment 3

This experiment was designed to directly assay the contribution of an aiming strategy to savings in visuomotor adaptation tasks. To this end, we used the aim report technique of Taylor et al. (2014). Briefly, numbers were displayed along a ring (10 cm circumference) in 5.625° increments, with the value increasing in the clockwise direction and decreasing in the counterclockwise direction, relative to the target position (see Fig. 1.4a). Beginning with the 21st trial in the baseline block, participants were required, at the beginning of each trial, to report their intended aim, verbally indicating the number closest to where they intended to reach. Aiming directly to the target was reported as “target” or “0.” Participants were randomly assigned ($n = 10/\text{group}$) to either a 15° or 45° rotation condition.

Each group completed five blocks of trials while reaching to four targets spaced in 90° increments around 360° . The number of trials in each block was the same as the four-target conditions of Experiment 2: Baseline of 40 trials, followed by Rotation 1 (100 trials), Washout 1 (120 trials), and Rotation 2 (100 trials). At the end of Rotation 2, the program was paused and the numbers were removed from the screen. Participants were instructed to aim and reach directly to the visual location of the target for the remainder of the experiment (Washout 2, 40 trials).

Experiment 4

This experiment was designed to assess whether compensatory aiming could be controlled via an arbitrary cue, and whether the formation of such an action selection policy is dependent on the size of the perturbation. The Rotation blocks included both rotation and no-rotation trials. On trials in which the visual feedback was rotated, the cursor and return ring were presented in red. Interspersed within these rotation trials were no-rotation trials; here the cursor and return ring were white. Participants were told at the beginning of the session, “Later on, your cursor will occasionally turn red. This means something weird will happen on that trial but you are still trying to do the same thing, hit the target with your cursor.”

Participants ($n = 10/\text{group}$) were randomly assigned to either a 15° or 45° rotation group. Each group completed five blocks: Baseline, Rotation 1, Washout 1, Rotation 2, and Washout 2. The baseline was 80 trials long. The rotation blocks consisted of 70 rotation (“cued”) trials in which the cursor was red and rotated and 10 (“uncued”) trials in which the cursor was white and the rotation was turned off (see Fig. 1.5a). The 10 uncued trials were distributed unpredictably throughout the first rotation block. The same order of cued and uncued trials was repeated in the Rotation 2 block. The two rotation blocks were separated by 100 trials of washout. The final washout block was composed of 60 trials.

Experiment 5

In the final experiment, we examined how participants responded when presented with a novel perturbation, comparing conditions in which we expect an aiming strategy to be operative or not operative. Participants ($n = 10/\text{group}$) were randomly assigned to either a 15° or 45° rotation group, with the expectation that the 45° group would use an aiming strategy. The trial structure was 80 baseline trials, 80 counterclockwise rotation trials (Rotation CCW), 160 trials of washout, and 80 trials of a 30° clockwise rotation (Rotation CW) (see Fig. 1.6a). The second rotation is

different in both direction and magnitude from the first rotation, providing a probe on “savings” in a novel context.

Data analysis

Movement duration was defined as the interval between when the finger exited the starting region and when the radial amplitude passed 10 cm. We determined the reach trajectory between these two time samples. For each trial, position data was interpolated to 150 data points using a cubic spline. The Cartesian position data were converted to polar coordinates, and the radial component was numerically differentiated to determine movement speed independent of hand angle. Peak speed was calculated by smoothing the radial speed with a 75-window loess function and taking the maximum value of this function. The hand angle (relative to target angle) at the position of peak speed was the primary dependent variable in all of the experiments. In Experiment 3, we also performed analyses on the verbal aiming report data and the estimate of implicit adaptation.

Some of the experiments reported here used two targets spaced 60° apart, while others used four targets, spaced in 90° increments. Because of these differences, savings cannot be directly compared across these experiments. However, to standardize our savings metric, we adopted a common convention by creating a savings difference score using a measure that averaged over three reaches to each target: In two-target conditions, this was six trials; in the four-target conditions, this was 12 trials. The difference score for each condition was calculated by subtracting the mean heading angle for these three trials in the Rotation 1 block from the corresponding three trials in the Rotation 2 block. Before computing each difference score, the values were adjusted by subtracting the mean hand angle on the 10 null trials that preceded the associated rotation block. All *t* tests are two-tailed unless otherwise noted. Statistical calculations were made with MATLAB 2009b (The MathWorks) (*t* tests) and SPSS 17 (ANOVAs).

In Experiments 1 and 5, we also determined the learning rate for each participant in the two rotation blocks by fitting an exponential function of the following form: $a - b * e^{x*c}$. Parameter fitting was implemented with the built-in MATLAB function `fminbnd` to minimize squared error. The asymptotic performance parameter, *a*, was set to the mean of the last 5 trials. The parameter corresponding to the total amount of learning, *b*, was fixed as the difference between the mean of the last 5 trials and the first trial. The only free parameter was the learning rate constant, *c* (Huang et al., 2011).

Results

Experiment 1

This experiment was designed to test whether savings varied as a function of the size of the perturbation, with different groups exposed to visuomotor rotations of 15°, 30°, 45°, or 60°. Participants in all four groups learned to compensate for the rotation by moving in a direction opposite of the visual perturbation (Fig. 1.2a). The visual error was substantially reduced by the end of the first rotation block (blue lines), although it fell short of full compensation.

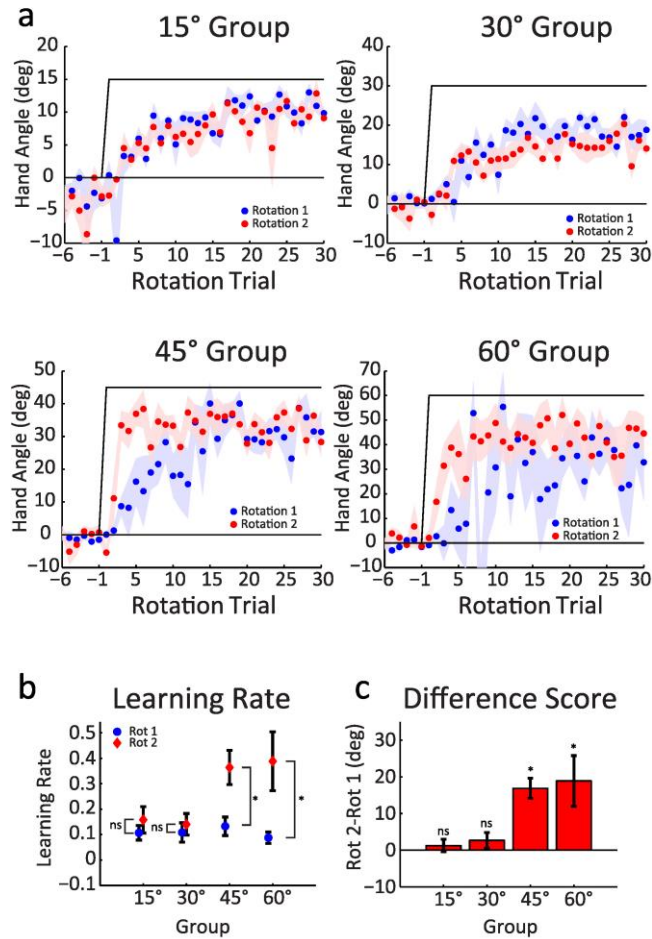


Figure 1.2 - Savings is only observed with large perturbations. (A) Heading angle of the hand relative to the target during last five baseline and first 20 trials of Rotation 1 (blue) and Rotation 2 (red) in Experiment 1. The data for the two rotation blocks are displayed in a superimposed manner for the four perturbation sizes. Note the different scaling of the vertical axis for each group. The 45° and 60° groups show faster relearning for the second rotation; this savings effect is absent for the 15° and 30° groups. (B) Mean learning rate from an exponential function independently fit to all trials in Rotation 1 and 2. (C) Difference score measure of savings (Rotation 2 – Rotation 1). Data are mean \pm SEM. Asterisk or "ns" indicate significant or insignificant differences, respectively.

As one measure of savings, we compared the rate of adaptation for the two rotation blocks (Fig. 1.2b). A mixed-factor ANOVA with the factors Rotation Size and Block revealed a within-subject effect of Block $F(1,36) = 19.38$, $p < .001$, and a significant Block by Rotation Size interaction $F(3,36) = 3.62$, $p = .022$. Bonferroni-corrected post hoc analyses revealed a faster learning rate in the second rotation for the 45° and 60° groups (one-tailed t-test $t(9) = 4.3$, $p = .001$; $t(9) = 3.1$, $p = .007$), but not the 15° and 30° groups ($t(9) = 0.94$, $p = .187$; $t(9) = 0.50$, $p = .313$). Thus, savings by this measure was limited to the two groups who experienced the largest perturbations. The same dissociation is observed if we use the difference score to estimate savings, limiting the analysis to six initial trials of the rotation blocks (Rotation 2 – Rotation 1 over these trials). Here there was a significant effect of Rotation Size, $F(3,36) = 6.76$, $p = .001$, Block, $F(1,36) = 25.2$, $p < .001$, and interaction of these factors, $F(3,36) = 5.45$, $p = .003$). As can be seen in Figure 1.2c, savings was limited to the 45° and 60° groups.

The fact that savings was limited to the two groups presented with the largest perturbations is problematic for error-based models of visuomotor adaptation (with the exception of Berniker & Körding, 2011). These models predict that savings should be observed across different perturbation sizes; our results suggest a qualitative difference. Instead, we hypothesize that the differences in savings may arise from a process that is not adaptation, namely the redeployment of an aiming strategy. It is increasingly recognized that participants may initially employ strategies to compensate for salient perturbations in sensorimotor adaptation tasks (Taylor et al., 2014; Redding & Wallace, 1993; Osu et al., 2002). Building on this idea, it has recently been proposed that savings in these tasks can be attributed to an action selection process (Huang et al., 2011; Orban de Xivry & Lefèvre, 2015; Haith et al., 2015; Huberdeau et al., 2015; McDougle et al., 2015). The current results are consistent with this view and indicate, more specifically, that a change in the selection of an aiming strategy causes savings. Indeed, our results raise the possibility that savings in sensorimotor adaptation tasks may be restricted to conditions that feature the prominent use of an aiming strategy.

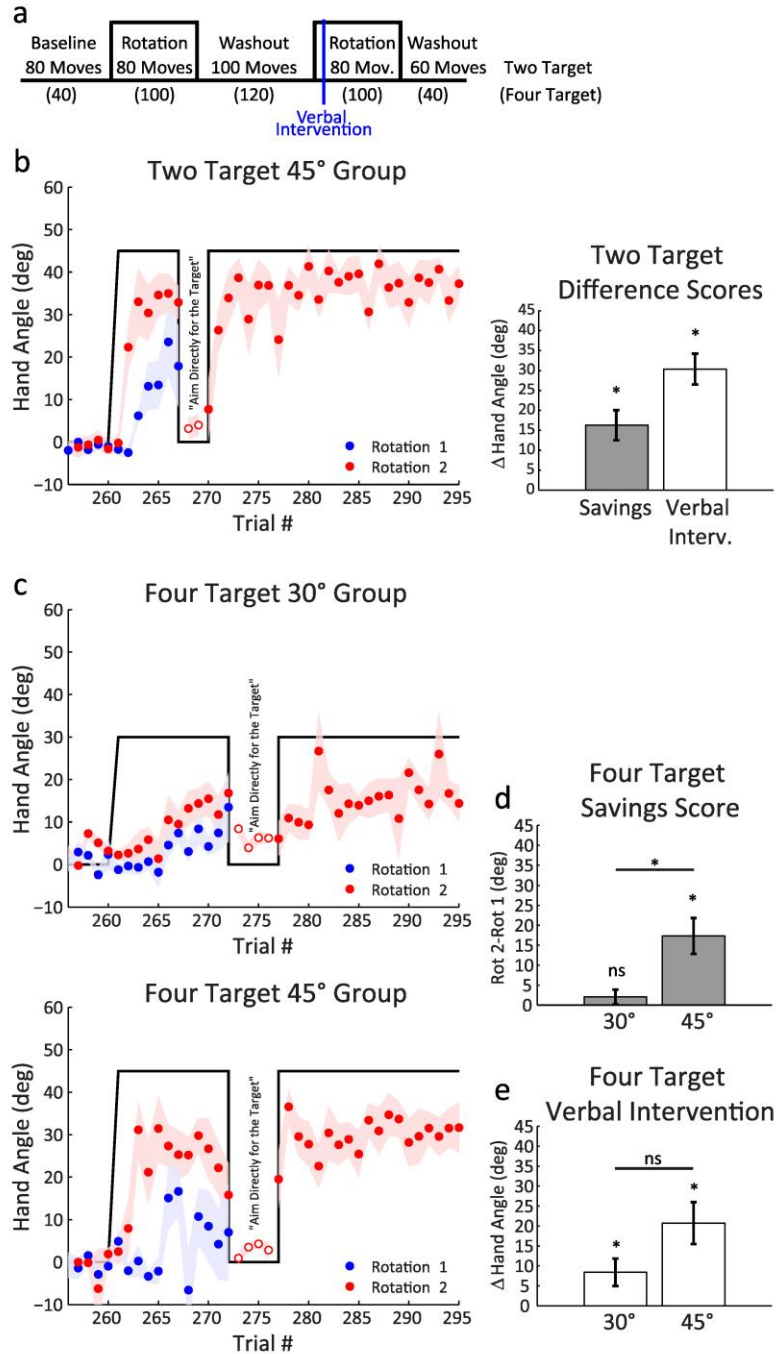


Figure 1.3 - Gains from savings are abolished by instructions to aim directly for the target. (A) Trial schematic for Experiment 2, denoting the verbal intervention following the 7th trial of the Rotation 2 block. For the following 2 trials (blue bar), participants were instructed to aim and reach directly for the target. (B) Left: Behavior during the Rotation 2 block for the two-target 45° group. There is an initial period of rapid relearning, followed by a sharp change in hand angle in response to the verbal intervention. Right: Difference score between the first 6 trials of Rotation 1 and Rotation 2 blocks (red) and between the average heading angle for the two intervention trials and the two trials preceding the intervention (blue). (C) Time course of hand angles for the 45° (top) and 30° (bottom) groups in the four-target groups. Savings score (D), defined by the difference between the first 12 trials of the Rotation 1 and Rotation 2 blocks and (E), verbal intervention scores (using the 4 intervention trials and 4 pre-intervention trials). While savings was only reliable for the 45° group, there was a drop in hand angle following the verbal intervention for both groups. Data are mean \pm SEM.

Experiment 2

Experiment 2 was designed to indirectly assess the contribution of aiming to savings, building on the idea that aiming, an action selection process, can be under volitional control. To this end, the experiment was paused shortly into the second rotation block, and the participants were explicitly instructed to aim directly at the target. Our intention was to direct the participants to terminate the use of an aiming strategy if one was being used. In this manner, the initial trials of Rotation 2 allow us to assess savings, comparing performance here with that observed in the initial trials of Rotation 1. The probe trials allowed us to assess the degree of true adaptation. If savings is due to adaptation, then the hand angle on these trials should remain similar to that seen in the initial trials of Rotation 2 (with a slight decay due to the pause required to give the instructions). In contrast, it would be parsimonious to attribute savings to re-aiming if participants exhibit a precipitous change in hand angle following the instructions.

In the two-target condition, participants showed faster relearning of the 45° rotation (Fig. 1.3b). Over the first six trials of Rotation 2, the mean hand angle was 26.8°, a value that is 16.8° larger than the value for the first six trials of Rotation 1 (paired t-test $t(9) = 4.38$, $p = .002$). When the participants were told that the perturbation was turned off and instructed to aim directly at the target, there was a 30.4° drop in heading angle between the two trials preceding the verbal intervention and the two intervention trials (Fig. 1.4b-c, paired t-test $t(9) = 7.86$, $p < .001$). The magnitude of this decline is considerably larger than what would be expected from temporal decay, which is only 20% of the adapted state (Hadjiosif & Smith, 2013). Rather, the flexible behavior observed here is consistent with the hypothesis that savings is related to a process under volitional control; participants recalled a previously successful aiming strategy to rapidly reduce target error without a large change in their sensorimotor mapping. The small, residual 4° shift in hand angle observed during the probe trials, relative to the null trials preceding Rotation 2 (95% CIs [0.4°, 7.5°]), likely reflects the degree of true adaptation that took place during the initial trials of Rotation 2.

To ensure that strategy use is not limited to tasks in which reaches are restricted to a limited part of the workspace, we repeated this experiment with four target locations spaced 90° apart. Two groups were tested, one exposed to a 30° rotation and the other exposed to a 45° rotation (Fig. 1.3c). The 45° group showed reliable evidence of savings in the initial trials of Rotation 2, with a 17.4° increase in the change in hand angle compared to the same trials in Rotation 1 (Fig. 1.4d; paired t-test $t(9) = 3.92$, $p = .004$). In contrast, the 30° group did not show reliable savings, with a modest 2.7° change in hand angle over the first rotation (paired t-test $t(9) = 1.41$, $p = .20$). Furthermore, the two groups were different from one another (two-sample t-test $t(18) = 3.06$, $p = .007$). Thus, consistent with Experiment 1, savings was present for the 45° group and absent or minimal for the 30° group.

Both groups showed a significant decline in hand angle following the verbal intervention (Fig. 1.4e; 45° group mean = 20.7°, paired t-test $t(9) = 3.94$, $p = .003$; 30° group mean = 8.3°, paired t-test $t(9) = 2.43$, $p = .037$). The decline was larger for the 45° group compared to the 30° group, although this difference was only marginally significant (two-sample t-test $t(18) = 1.96$, $p = .066$). Similar to the two-target condition, both groups showed small residual aftereffects (Fig.

1.4c, blue trials; 45° group mean = 2.7°, paired t-test $t(9) = 2.59$, $p = .029$, 30° group mean = 3.4°, $t(9) = 4.02$, $p = .003$). The small aftereffects indicate that, not only can participants flexibly deploy an aiming strategy, but the performance gains from aiming dwarf those of adaptation.

The results from the verbal intervention in the 30° group are puzzling: While these participants showed a drop in hand angle, indicative of aiming, they did not show reliable savings. Previous studies have reported savings with this size perturbation (Huang et al., 2011; Hadjiosif & Smith, 2013; Herzfeld et al., 2014; Orban de Xivry & Lefèvre, 2015; Haith et al., 2015; Huberdeau et al., 2015; McDougle et al., 2015). In the current results, the participants did exhibit savings, apparently because they aimed the same amount in the early phase of both rotations. In contrast, the 45° groups showed a larger change in aim from Rotation 1 to Rotation 2. We return to this issue in the Discussion. Taken together, the results of Experiment 2 indicate that the savings observed early in a second rotation can be attributed to aiming behavior that is under volitional control. The small residual aftereffects observed during the intervention trials suggest that aiming and adaptation occur in parallel.

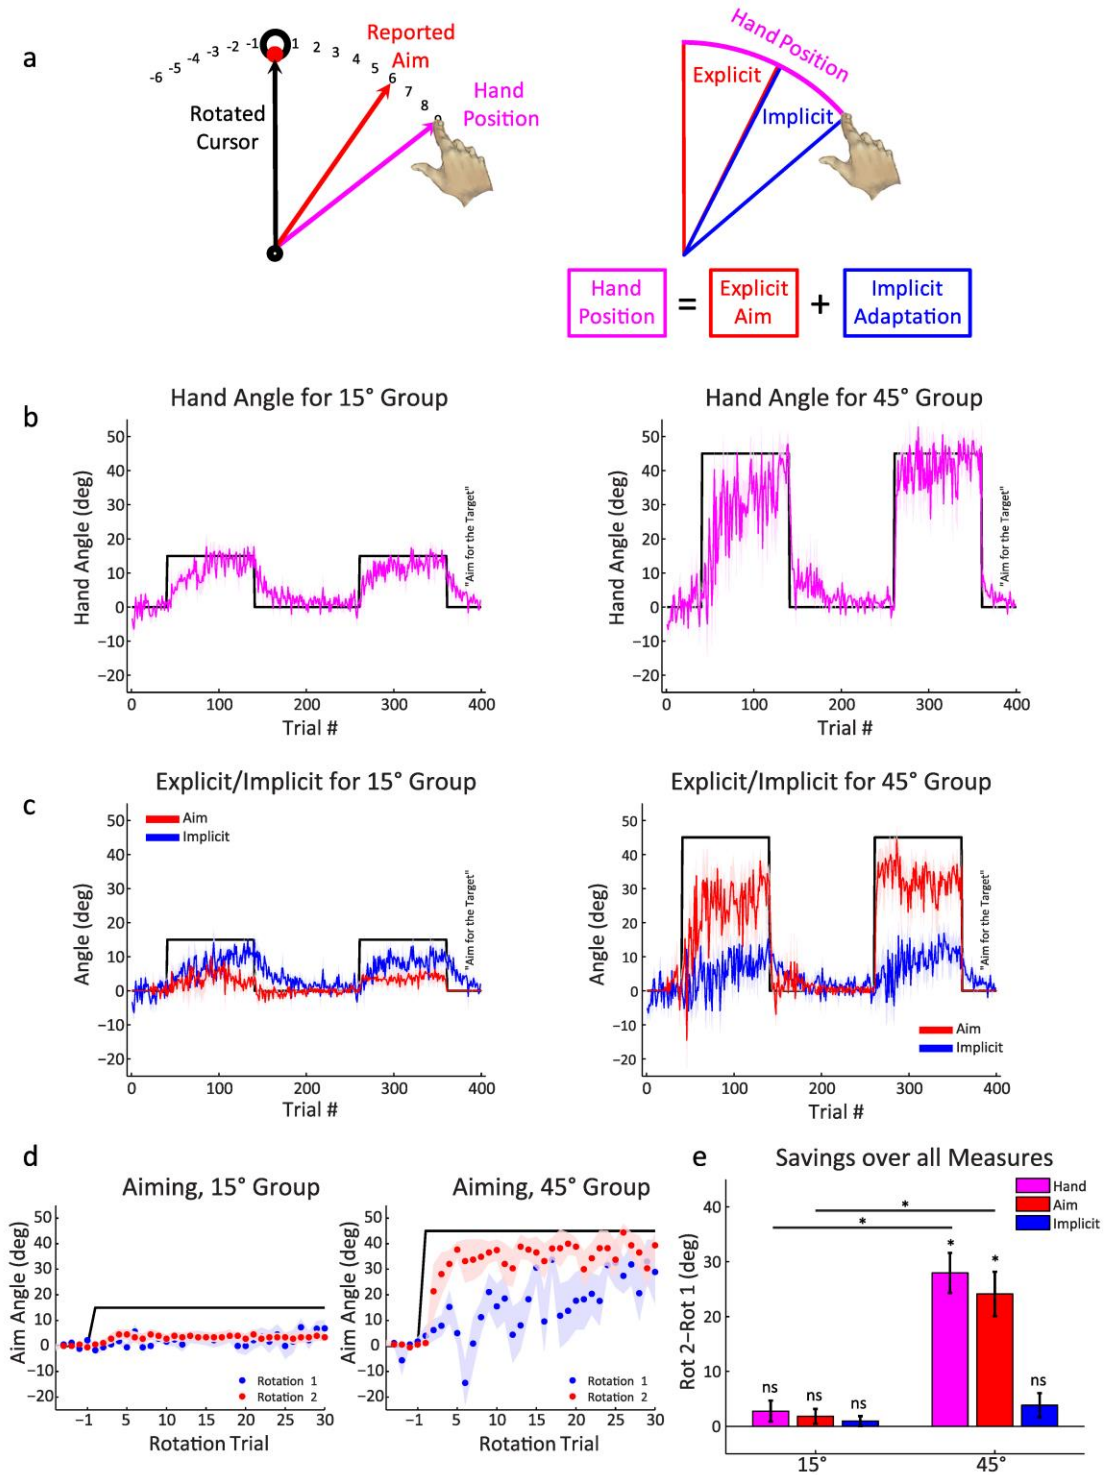


Figure 1.4 - Savings assessed with aiming report task in Experiment 3. (A) Left: Task schematic. Numbers were visible in ascending and descending manner around the target. Participant verbally reported the number corresponding to his/her aiming location and then performed the reach. Right: Assuming performance (heading angle) is the composite of an explicit aiming strategy and implicit adaptation, the magnitude of the latter can be estimated via subtraction. (B) Overall behavior for 15° (left) and 45° (right) groups. (C) Mean aiming reports (red) and implicit adaptation (blue, estimated via subtraction). Participants were instructed to aim directly for the target in the final washout block. (D) Aiming reports over the first 30 trials, superimposed for Rotation 1 and Rotation 2. (E) Savings difference scores for each measure. Data are mean \pm SEM.

Experiment 3

This experiment directly measured aiming behavior in the context of savings, asking participants to verbally report their aiming location before each reach (Taylor et al., 2014). Participants were randomly assigned to either a 15° or 45° group in order to assess changes in aim across different magnitudes of visuomotor rotation. The trial design was similar to the four-target conditions of Experiment 2, with two 100 trial rotation blocks separated by a 120 trial washout. Aiming reports were initiated halfway through the baseline block, and were assessed on every trial until the end of the second rotation. In a final washout block, the numbers were removed from the screen and participants were instructed to move their hands directly to the target (Taylor et al., 2014).

We excluded trials where the participant reported an aiming location that was more than two times the perturbation size, as well as trials in which the discrepancy between the reported aiming location and actual hand position was greater than 45°. On average, this excluded 3.0% of trials in the 45° group and 2.2% of trials in the 15° group. Most of these trials came from three individuals (1 in the 45° group, 2 in the 15° group), but even for these individuals, the percentage did not exceed 10% of their total aiming trials. Statistical tests of savings were not meaningfully affected by the exclusion of these trials.

As in Experiment 1, the hand angle difference score indicated reliable savings for the 45° group (Fig. 1.4b,e, mean = 28°, paired t-test $t(9) = 7.67$, $p < .001$), but not for the 15° group (mean = 2.7°, paired t-test $t(9) = 1.43$, $p = .19$). This savings pattern was mirrored in the aiming report data: Participants in the 45° group showed a reliable increase in the reported aiming location for Rotation 2 compared to Rotation 1 (mean = 24°, paired t-test $t(9) = 5.94$, $p < .001$). Participants in the 15° group did not show a consistent change in the reported aiming location between the two rotation blocks (mean = 1.8°, paired t-test $t(9) = 1.33$, $p = .21$).

In order to estimate the amount of implicit adaptation, we subtracted each participant's reported aim from the hand angle. We then calculated savings using the difference score measure with these data. Interestingly, neither group showed savings in the estimate of implicit adaptation (45° group, paired t-test $t(9) = 1.74$, $p = .115$; 15° group, paired t-test $t(9) = 1.06$, $p = .31$). Moreover, for both the overall performance and aiming report measures, savings in the 45° group was different from the 15° group (reaching: two-sample t-test $t(18) = 6.1$, $p < .001$; aiming: two-sample t-test $t(18) = 5.2$, $p < .001$), but not in the measure of implicit adaptation (two-sample t-test $t(18) = 1.2$, $p = .242$). Note that all t-tests here were Bonferroni corrected.

The results of Experiment 3 provide direct evidence that savings in visuomotor adaptation is associated with how participants change their aiming strategy when re-encountering a perturbation, and not associated with faster modification or recall of an internal model. Additionally, the results indicate that the larger performance changes observed with large perturbations is primarily due to changes in the explicit aiming process. Indeed, the magnitude of implicit adaptation appears to be relatively constant for the small and large perturbations (Bond & Taylor, 2015).

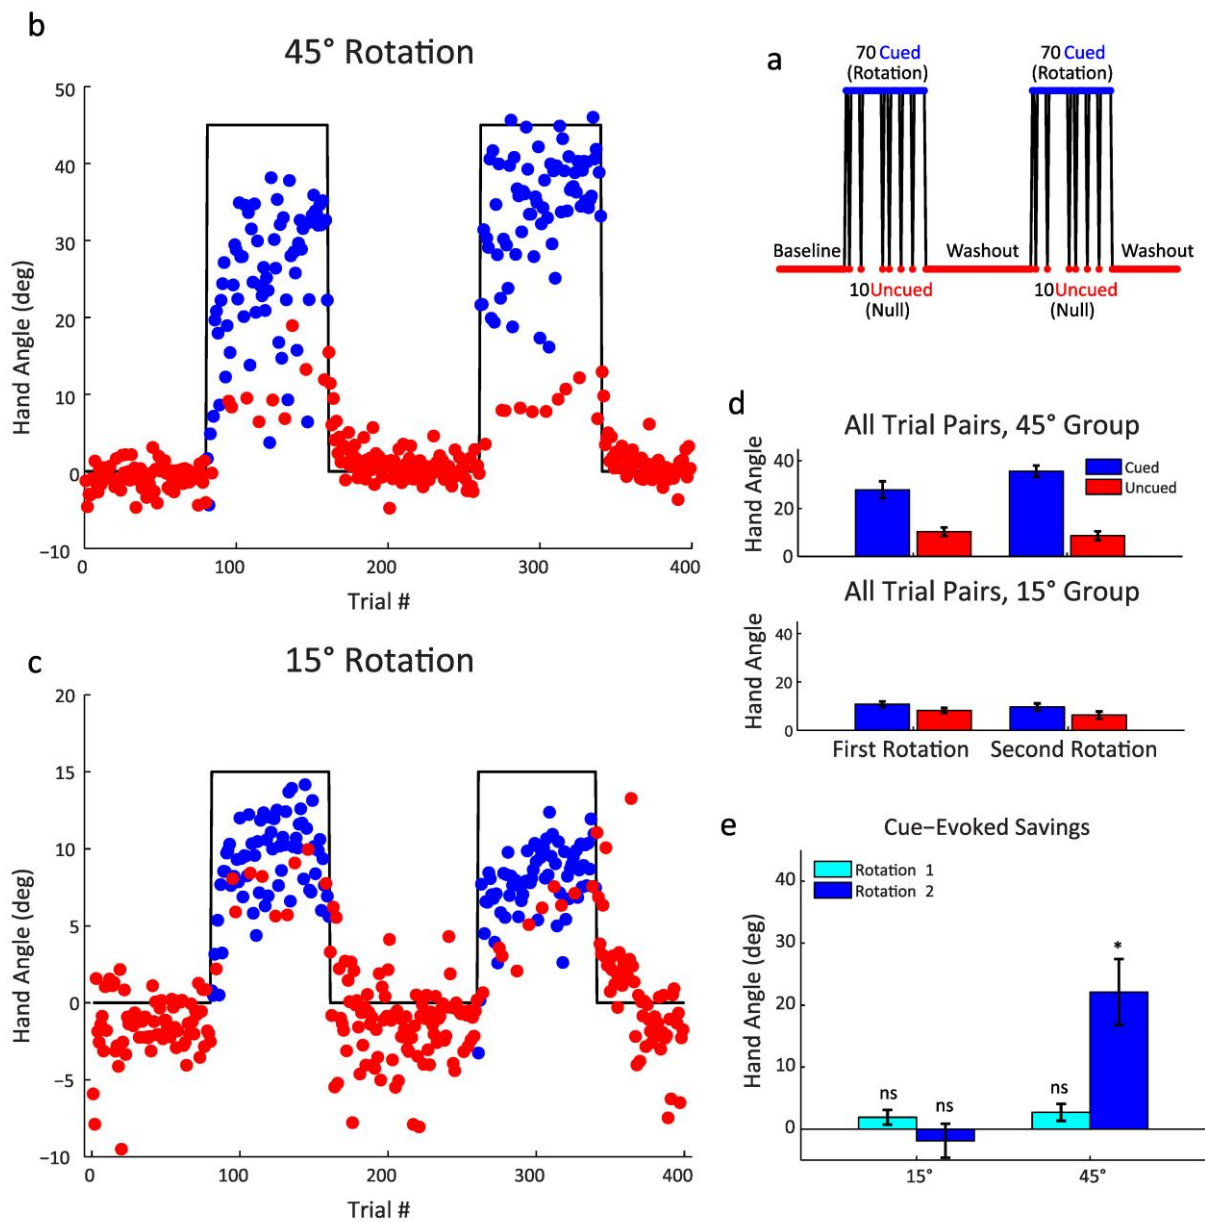


Figure 1.5 - Control of aiming strategy by an arbitrary color cue. (A) Interleaved design of Experiment 4 in which the cursor changed color on rotation trials (depicted in blue). The cursor remained white on no-rotation trials (depicted in red). (B, C) Heading angle for groups exposed to 45° and 15° rotations, respectively. Note the difference in the scale of the y-axis (here and in 5D). (D) Blue bars indicate mean of the cued trials, limited to those that immediately preceded an uncued trial. Red bars are mean of uncued trials. (E) Baseline subtracted hand angle on the first trial for Rotation 1 and Rotation 2. The color cue prompted the use of an aiming strategy for the 45° group, even before the participants had experienced errors caused by the perturbation in Rotation 2. Data are mean \pm SEM.

Experiment 4

Prior studies have shown that, without extensive training, processes involved in sensorimotor adaptation are insensitive to arbitrary static cues (Cunningham & Welch, 1994; Miall et al.,

2004; Howard et al., 2013). For example, people have great difficulty simultaneously learning internal models for clockwise and counterclockwise rotations when the context is cued by color; under such conditions there is substantial interference between the two contexts. However, with minimal training, people are extremely good at using novel cues to select arbitrary actions.

We exploited this difference between action selection and adaptation by designing an experiment where the presence of a visuomotor rotation was indicated by an arbitrary color cue. These perturbation trials were interspersed with trials in which the rotation was absent (Fig. 1.5a). We reasoned that if participants were employing an aiming strategy, they should be able to use the color cue to determine whether or not to use the strategy. In contrast, the color cue should have minimal effect on behavioral changes caused by adaptation.

We again compared two groups, one exposed to a 15° rotation and the other exposed to a 45° rotation. Based on the results of Experiment 1, we assume that the 45° group would use a combination of error-based adaptation and aiming. As such, their performance should switch rapidly between the cued and uncued trials during the rotations and, more important, that the color cue will result in savings on the second rotation block. Indeed, we would expect this savings to be evident on the first trial of the second rotation block, even before the perturbation has been re-experienced. In contrast, we expected to observe little flexibility in the performance of the 15° group.

Figure 1.4b shows the time series for the two groups. There is a clear separation between hand angle on the cued and uncued trials for the 45° group, whereas there is considerable overlap for the 15° group. We quantified the effect of the cue by creating trial pairs, with each pair composed an uncued trial from the rotation blocks and its immediately preceding cued trial. These data were evaluated in a Rotation Size x Block x Trial Type (Cued/Uncued) mixed ANOVA (Fig. 1.5d). The main effects of Rotation Size, $F(1,18) = 46.64$, $p < .001$, and Block, $F(1,18) = 48.88$, $p < .001$, were reliable, whereas the effect of Trial Type was not, $F(1,18) = .73$, $p = .405$. A three-way Block x Rotation Size x Trial Type interaction was reliable, $F(1,18) = 6.54$, $p = .02$, driven by simple effects of Block and Trial Type in the 45° group ($F(1,9) = 76.07$; $p < .001$, $F(1,9) = 5.78$, $p = .03$) but not the 15° group ($F(1,9) = 1.36$, $p = .259$; $F(1,9) = 1.43$, $p = .247$). These results indicate a difference between the cued and uncued trials for the 45° group, but not the 15° group, as well as an improvement in performance on cued trials in the 2nd rotation for the 45° group.

Performance on the uncued trials for both groups shows a gradual increase in hand angle within a rotation block. Moreover, the magnitude of this effect is similar across both rotation blocks and across the two groups. Coupled with the results of Experiment 3, this pattern is consistent with the on-going operation of adaptation, operating with relative indifference to the size of the perturbation. We assume that this adaptation occurs even on cued trials because there remains a substantial error between the strategic aiming location and the position of cursor feedback (Mazzoni & Krakauer, 2006). The current design provides a demonstration of the parallel operation of strategic and adaptation processes (Taylor & Ivry, 2011) in a task setting that does not overtly instruct participants to aim somewhere other than the target.

We assume that the cued trials are composed of both re-aiming and adaptation, while the uncued trials reflect adaptation alone. If aiming strategies are being flexibly invoked by this arbitrary cue, we should be able to trigger savings with the cue alone. To test this prediction, we focused on the first trial of the second rotation. This trial offers a powerful test of the re-aiming hypothesis since the participants have yet to experience the error associated with the reintroduction of the visuomotor rotation. Therefore if savings is observed on this trial, it can only represent a change in the selected action and not error-based adaptation.

To assess this cue-evoked savings, we computed a difference score using only the first trial of Rotation 1 and Rotation 2 (Fig. 1.5e). The 15° group did not show evidence of cue-evoked savings (mean = -3.8°, paired t-test, $t(9) = 1.41$, $p = .19$). In contrast, the 45° group displayed significant cue-evoked savings (mean = 19.4°, paired t-test, $t(9) = 3.51$, $p = .007$). Furthermore, the two groups were different from each other (two-sample t-test using Satterthwaite's approximation for unequal variances, $t(18) = 3.77$, $p = .002$). This cue-evoked savings in the 45° group was not the result of online feedback corrections: Individual reach trajectories were straight (not shown), hand position was measured at peak radial speed, and the average movement time was 208 ms, a value unchanged from the mean movement time of the 10 null trials previous to Rotation 2 onset (paired t-test, $t(9) = 0.26$, $p = .80$). In summary, this “first trial” analysis provides a powerful demonstration that savings in visuomotor adaptation reflects the use of an aiming strategy and that this process is not contingent on experiencing previously encountered errors.

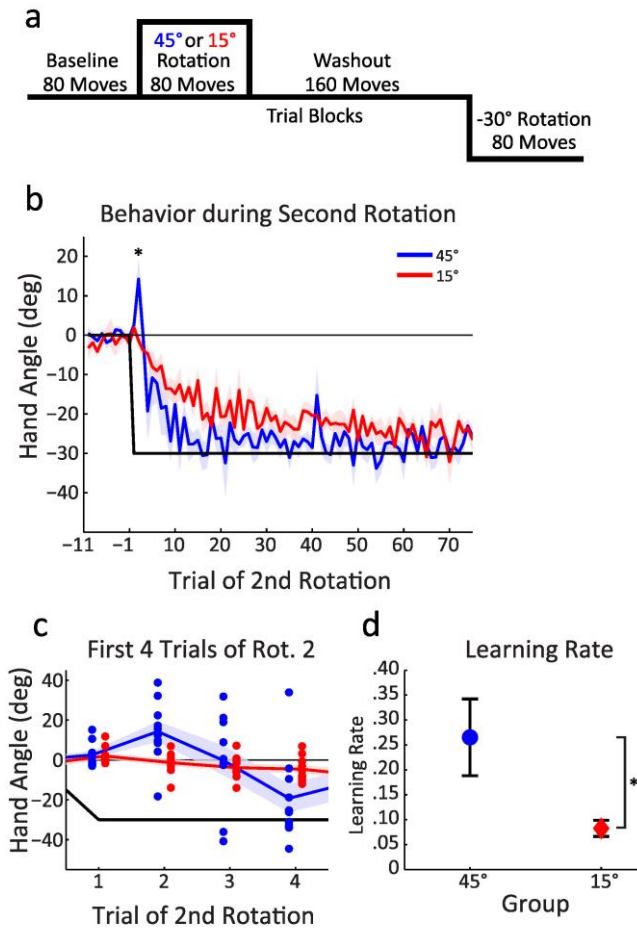


Figure 1.6 - Inappropriate recall of aiming strategies in response to a novel perturbation. (A) In Experiment 5, all participants were exposed to a novel -30° rotation after completing a training block with either a 45° or 15° rotation, followed by a washout block. (B) Heading angle for the last 10 trials of the washout block and all trials of the Rotation 2 block. The 45° group initially responded to the perturbation by moving in the wrong direction and, after a few trials, reversed direction, reducing error faster than the 15° group. (C) Enhanced picture of behavior over the first four trials of the Rotation 2 block, showing performance of each individual (dots). Note the expanded vertical axis relative to (B). (D) Mean learning rate from an exponential fit of all trials in Rotation 2 block, showing faster learning of the rotation for the 45° group. Data are mean \pm SEM.

Experiment 5

In the final experiment, we examined how an aiming plan generalizes to a novel perturbation. We exposed participants to either a 15° or 45° counterclockwise rotation in the first block, and then tested both groups with a clockwise 30° rotation in the second rotation block (Fig. 1.6a). Given the absence of savings for a 15° rotation observed in Experiments 1, 3 and 4, we expect this group will provide a baseline measure of learning a reversed 30° rotation. The key question is what happens when the 45° group is tested with the reversed rotation. It may be that, given the novel context, there will be little benefit from the initial experience with a rotation; as such, performance will be similar to that exhibited by the participants who had been exposed to an initial 15° rotation. Alternatively, having developed a strategy to compensate for the initial 45°

rotation, these participants may rapidly invoke a strategy when encountering the new perturbation.

Both groups showed good learning in the first rotation block and returned fully to baseline before the second rotation (Fig. 1.6b, trials -11 to -1). Their performance in the Rotation 2 block was quite different. The group which had initially been exposed to the 15° rotation showed a standard adaptation function, with visual error decreasing in a monotonic fashion. In contrast, there were two striking features in the performance of the group who had been exposed to the 45° rotation in the Rotation 1 block. First, on the second trial of Rotation 2, nine out of ten participants moved in a direction that amplified, rather than reduced the visual error (Fig. 1.6c-d, one-sample t-test $t(9) = 2.87$, $p = .02$). This behavior is the opposite of what would be expected from an error-based adaptation process. Rather, when faced with the new perturbation, these participants appeared to redeploy an aiming strategy that had been useful in the previous context. Second, the participants quickly abandoned this strategy and moved in a direction that rapidly compensated for the 30° rotation. Indeed, relative to the 15° group, these participants subsequently showed faster learning of the 30° rotation.

We quantified learning rate by fitting a single exponential function over all trials of the Rotation 2 block. The learning rate parameter for the 45° group was larger than the learning rate of the 15° group (two-sample t-test $t(18) = 2.97$, $p = .04$, Fig. 1.5e), indicating a faster change in behavior over the same number of trials. We propose that the participants who were initially trained with a 45° rotation were sensitized to invoke and adjust an aiming strategy upon encountering the novel rotation.

Discussion

Recent work shows that the behavior in sensorimotor adaptation tasks arises from both error-based adaptation and other learning processes and that these other processes may be responsible for savings (Huang et al., 2011; Herzfeld et al., 2014; Taylor et al., 2014; Haith et al., 2015; Huberdeau et al., 2015; McDougale et al., 2015). Here we examined the contribution of aiming strategies to savings upon relearning in visuomotor adaptation tasks. Our results converge on two key ideas. First, savings reflects recall in the action selection domain, with participants invoking an aiming strategy when a perturbation was reintroduced. Second, error-based adaptation does not appear to make any contribution to savings in visuomotor adaptation. This dissociation was observed whether we used indirect or direct methods to partition behavior into strategic and adaptation components.

Savings as action selection

Our results are best described by the hypothesis that savings in visuomotor tasks arises from the recall of an aiming strategy. In a sense, the term “savings” here is a misnomer, given the classic definition of faster relearning after forgetting (Ebbinghaus, 1913). We assume that the participants never forgot the strategy but determined when or when not to use the strategy. Experiments 4 and 5 shed light on factors that can evoke aiming upon re-encountering a rotation.

In both experiments, the strategy was recalled without re-exposure to the errors that drove its development in the first rotation. In Experiment 4, the aiming strategy was evoked by a color cue; in Experiment 5, it was an error of opposite sign to the previously experienced error. This suggests that, in typical savings designs, such as Experiment 1 (A-Null-A), the large error can become a cue for the recall of aiming strategies (Huberdeau et al., 2015). The Experiment 3 results are also consistent with this idea, as re-aiming in the 45° group occurred immediately after a single perturbation trial (Fig. 1.4*d*).

Huang et al. (2011) have also proposed an action selection account of savings, arguing that savings arises from the memory of hand positions associated with successful performance during earlier presentations of a perturbation. This conceptual model would offer a different interpretation of some of our results. For example, in Experiment 5, the 45° group produced an initial shift in movement direction that had been successful in the first rotation block. We attribute this to inappropriate recall of an aiming strategy rather than a memory for hand position, but we cannot discriminate between these hypotheses with our results. Other results from the current study, however, do not align with predictions derived from the Huang model. First, their model predicts savings for all perturbation magnitudes and fails to predict the changes in aim observed in Experiment 3. Moreover, the 45° group in Experiment 5 showed faster learning of a novel rotation with a hand space solution that had never been experienced by the participants. More recently, Haith et al. (2015) have shown that savings is only observed when participants are given adequate time to prepare a movement and that savings can be observed with a small number of initial learning trials (Huberdeau et al., 2015). These findings are consistent with our conclusion that aiming strategies are responsible for savings.

The role of perturbation size

We only observed savings with large perturbations. It is tempting to conclude that this is because aiming does not occur when the perturbation is small. However, the aiming report data for the 15° group in Experiment 3 revealed small, yet consistent, deviations from the target location. Bond and Taylor (2015) have also shown that participants use aiming strategies across a large range of rotations. Given that some form of aiming occurs with small rotations, we surmise that use of an aiming strategy alone is not responsible for savings. Rather, savings may depend on a substantial change in aiming strategy between the beginning and end of the first rotation, the absolute angular magnitude of the aiming strategy, or the relative proportion of aiming and implicit adaptation.

Unlike the results from the current experiments, other studies have reported savings with 30° rotations (Caithness et al., 2004; Krakauer et al., 2005; Huang et al., 2011; Hadjiosif and Smith, 2013; Herzfeld et al., 2014; Haith et al., 2015; Huberdeau et al., 2015; Orban de Xivry and Lefeuvre, 2015). We are unable to offer a straightforward explanation for this discrepancy. There can be differences between the experimental setups of different laboratories as well as differences in the instructions. These factors can affect how participants understand the task, which may in turn affect how they respond to a perturbation. Our data suggest that participants in our setup actively use aiming strategies in response to 30° rotations (see also Bond and Taylor, 2015), but their strategy was of similar size for both Rotations 1 and 2. Importantly, our central

claim, that savings is associated with action selection rather than error-based adaptation, is not dependent on the idea that there is an absolute point at which aiming or savings becomes operative.

Relationship to error-based models of savings

Our account of savings in visuomotor adaptation tasks stands in contrast to many error-based models of savings. Smith et al. (2006) presented a model in which error-based adaptation operates at two time scales. Savings comes about as a result of the interaction of these two processes. This model makes clear predictions that the degree of savings will be independent of rotation size and only observed when the slower process has not returned to a baseline state. Counter to these predictions, we observed savings in groups exposed to large rotations after a long washout where the slow process should have returned to baseline (based on parameters from Zarahn et al., 2008). Berniker and Körding (2011) also proposed a two-process model, one in which small and large errors are treated differently. Critically, the Berniker model cannot account for savings observed when an aiming strategy is evoked by a cue rather than re-exposure to an error. As such, although we also argue that savings reveals the operation of multiple processes, we favor a model entailing the operation of two qualitatively distinct processes: an error-based adaptation process to ensure accurate response execution and an aiming process to optimize response selection.

Herzfeld et al. (2014) recently published an alternative error based model of savings. The core idea is that errors are used to adapt an internal model of the correct movement, but also use error information to modulate future learning. Savings comes about because the system is sensitized to familiar errors. This model fails to predict a number of the results from the current study. First, the model cannot account for our finding that savings depends on rotation magnitude. Second, the model does not account for the volitional nature of aiming and its impact on savings. Third, unlike the behavior in Experiment 4, the model requires re-experiencing errors to exhibit savings. Finally, although the model does offer a novel account of generalization in savings, it would not produce an inappropriate response to a novel rotation as observed in Experiment 5.

Awareness and aiming

We have intentionally skirted the role of awareness with respect to the contrast of aiming and adaptation. In our experiments, perhaps most salient in Experiment 2 and 3, aiming is under volitional control. Indeed, we assume that the processes involved in the use of aiming strategies are generally accessible to awareness even when not directly probed, whereas adaptation always operates in an implicit manner. Although aiming strategies and adaptation differ in terms of their accessibility to awareness, we do not believe this is the defining distinction between them, choosing to focus instead on their mapping to action selection and movement execution, respectively. We expect there are features of action selection that are not, or need not be, explicit. For example, action selection may be modified through operant mechanisms in the absence of awareness (Pessiglione et al., 2008). Similarly, it is likely that most participants were aware that the environment was perturbed following the introduction of 15° or 30° rotations in our experiments, but this did not lead to the rapid reinstatement of an aiming strategy when the

perturbation was reintroduced. Therefore, we do not consider awareness to be the defining distinction between adaptation and action selection in the context of savings.

In conclusion, it is becoming clear that motor learning results from a panoply of learning processes that allow flexible behavior in varied environments (Wolpert et al., 2011; Taylor et al., 2014; Huberdeau et al., 2015). The current results highlight that within-session savings for a visuomotor adaptation task is accounted for by action selection (Huang et al., 2011; Haith et al., 2015; Huberdeau et al., 2015; McDougle et al., 2015; Orban de Xivry and Lefevre, 2015), arising from the recall of an aiming strategy. As striking, we failed to observe any savings in conditions in which learning appeared to be restricted to error-based adaptation.

Chapter 2

A Limited Role for Sensory Prediction Errors in Visuomotor Learning²

Abstract

Sensorimotor adaptation occurs when there is a discrepancy between the expected and actual sensory consequences of a movement. This learning can be precisely measured, but its source has been hard to pin down because standard adaptation tasks introduce two potential learning signals: task performance errors and sensory prediction errors. Here we employed a new method that induces sensory prediction errors without task performance errors. This method combines the use of clamped visual feedback that is angularly offset from the target and independent of the direction of motion, along with instructions to ignore this feedback while reaching to targets. Despite these instructions, participants unknowingly showed robust adaptation of their movements. This adaptation was similar to that observed with standard methods, showing sign dependence, local generalization, and cerebellar dependency. Surprisingly, adaptation rate and magnitude were invariant across a large range of offsets. Collectively, our results challenge current models of adaptation and demonstrate that behavior observed in many studies of adaptation reflect the composite effects of task performance and sensory prediction errors.

Introduction

Sensorimotor learning is often described as a process of error reduction. By this view, the motor system generates a movement that it judges most likely to achieve the desired sensory outcome (Jordan & Rumelhart, 1992). Subsequently, the difference between the desired feedback and actual feedback arising from the motor command indicates the form and magnitude of the error. Learning occurs when this information is used to modify future performance. The preeminent theoretical framework for this form of learning involves a process where internal models are retuned to minimize errors via gradient descent (Cheng & Sabes, 2006; Pouget & Snyder, 2000; Thoroughman & Shadmehr, 2000; Miall & Wolpert, 1996; Jordan & Rumelhart, 1992).

The empirical base for this theoretical framework primarily comes from studies of short-term motor learning where a perturbation is introduced to disrupt the sensory feedback and/or alter the limb dynamics during reaching movements (Shadmehr, Smith, & Krakauer, 2010; Krakauer, Pine, Ghilardi, & Ghez, 2000; Scheidt et al., 2000; Thoroughman & Shadmehr, 2000; Held & Hein, 1958). Perturbations such as these introduce a sensory prediction error, the discrepancy between the expected and observed sensory consequences of a given motor command (Wolpert, Ghahramani, & Jordan, 1995a). For example, in visuomotor adaptation tasks, participants are trained to reach to visually defined targets, with feedback limited to a cursor that is, at least initially, veridical with hand position. A perturbation here might involve rotating the feedback cursor by 45°, thus creating a difference between the predicted location of the cursor and its actual position, the sensory prediction error. Models in which learning is driven solely by the

² Previously published in the *Journal of Cognitive Neuroscience*, Volume 29, Issue 6; June, 2017. doi:10.1162/jocn_a_01108. All figures and text reproduced with permission from *MIT Press*.

magnitude and direction of these errors have provided an excellent account of short-term sensorimotor adaptation.

Sensory prediction error is distinct from task performance error. The latter refers to any source of information that signals task success or failure. Task performance error can be presented symbolically; for example, the pitch of a tone might indicate if the movement resulted in a hit or miss, or the participant might see a point value, indicating the distance from the target. In most studies of sensorimotor adaptation, task performance feedback is provided. However, the most common form of task performance feedback in adaptation tasks is the distance between the cursor and target given the standard experimental instructions to “hit the target with the visual cursor.” As such, task performance errors are confounded with sensory prediction errors. When feedback is perturbed, the participant no longer achieves their intended goal. As adaptation proceeds and sensory prediction errors become smaller, task performance also improves. This situation makes it difficult to determine if sensory prediction errors, task performance errors, or a combination of these signals drives adaptation.

Behavioral changes have been observed in response to either error type, and historically, it has been argued that greater changes occur when both types of error are present (Schaefer, Shelly, & Thoroughman, 2012; Wolpert, Ghahramani, & Jordan, 1995b; Howard, 1968; Held & Freedman, 1963). However, these early studies did not use methods to clearly discriminate between different kinds of learning processes. Recent work has shown that task performance errors drive both explicit and implicit changes in behavior and that these changes are distinct from adaptation driven by sensory prediction error (Miyamoto, Wang, Brennan, & Smith, 2014; Taylor & Ivry, 2011).

The use of feedback that is contingent upon the participant’s behavior has made it difficult to dissociate the impact of task performance errors and sensory prediction errors. For instance, participants can be given a specific aiming strategy to counter a feedback perturbation, a manipulation that immediately reduces the large task performance error that typically accompanies the abrupt introduction of a perturbation (Mazzoni & Krakauer, 2006; Welch, 1969). However, this strategy does not eliminate the imposed sensory prediction error, which will drive implicit adaptation. Moreover, the implicit response to this error will reintroduce a task performance error after only a few reaches, forcing the sign of the two error signals into opposition. This results in interesting interactions between explicit aiming and implicit adaptation, but the dual operation of both processes makes it difficult to characterize their individual features (Taylor, Krakauer, & Ivry, 2014; Taylor & Ivry, 2011). Without characterizing each process in isolation, we cannot know whether their combined operation is simply the sum of the two independent processes, or whether there is modulation of either process by the other.

Here we set out to better characterize the sensory prediction error learning process by measuring its operation when isolated from task performance errors. To bridge the gap between traditional techniques and our new method, we first demonstrate how task instructions can dramatically affect the behavioral response to a sensorimotor perturbation without changing a primary hallmark of implicit adaptation: the aftereffect observed when the perturbation is removed.

We then employ a new approach that combines two existing techniques to selectively engage the implicit learning process that produces these aftereffects. Specifically, we used time-locked visual feedback that followed an invariant trajectory, independent of the participant's direction of movement (Shmuelof et al., 2012). This eliminated the typical contingency between angular hand position and visual feedback and provided a constant sensorimotor discrepancy over the course of the entire perturbation block. Critically, participants were fully informed of the nature of the manipulation and were instructed to ignore the feedback (Welch, 1969). By emphasizing that the task goal was to move the hand to the target and informing the participants of the nature of the visual manipulation, we eliminated signals that conveyed task performance error. Over a series of experiments, we exploit this method to reveal behavioral changes that result purely from sensory prediction errors.

Figure 1.

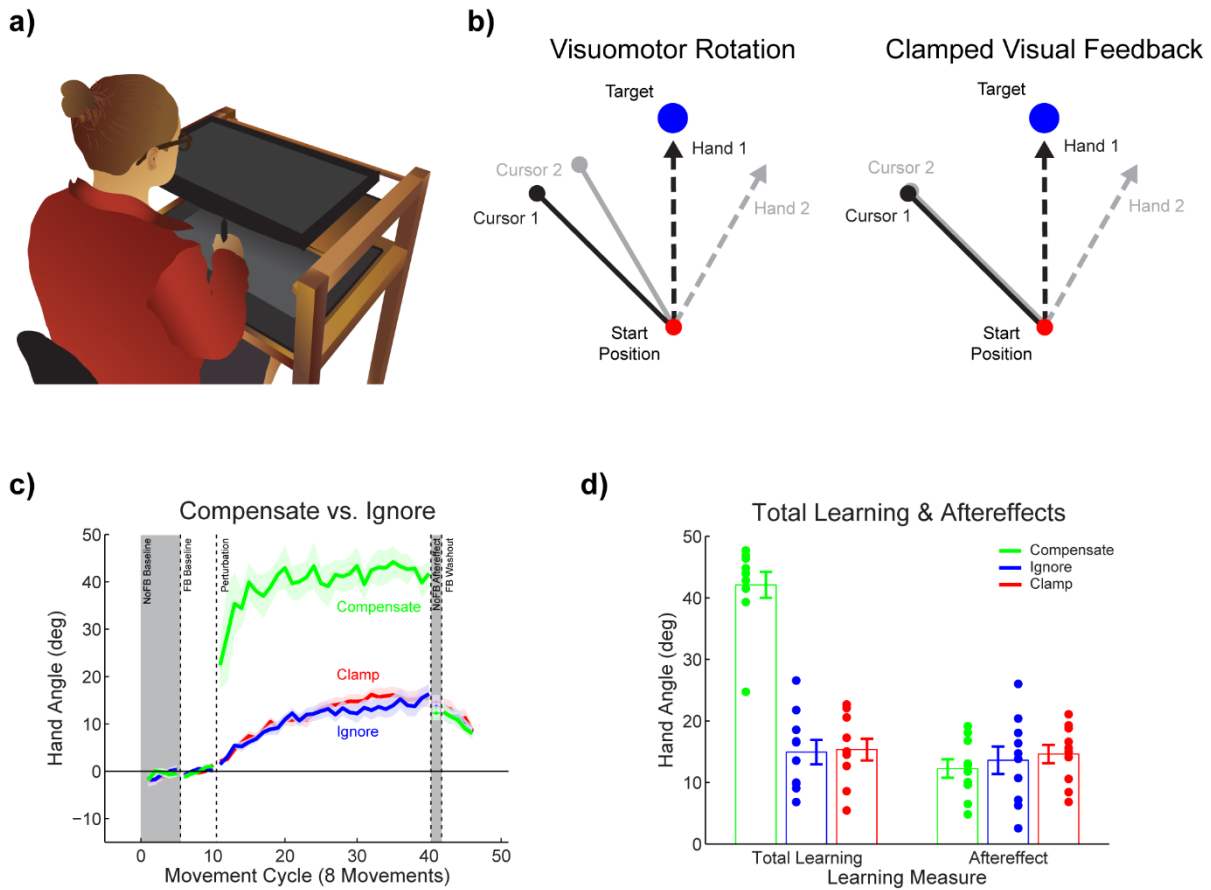


Figure 2.1 - A new method to induce implicit visuomotor adaptation. (A) Participants made center-out reaches on a digitizing tablet while grasping a stylus. (B) Illustration of the two visual feedback manipulations. For visuomotor rotations, the visual feedback is contingent on the hand movement. Clamped visual feedback breaks this contingency, with the feedback the same for all hand movements. Both types of feedback can be made irrelevant to the task with explicit instructions, but the error signal remains invariant with clamped visual feedback. (C) Behavior over all movement cycles for the Compensate, Ignore, and Clamp groups. (D) Baseline-adjusted mean hand angle over the last cycle of the perturbation block, and the no-feedback aftereffect cycle. Dots are individual means, lines indicate group mean \pm SEM. Gray shading denotes cycles without visual feedback.

METHODS

Participants

Undergraduate students ($n = 160$, 115 women, age = 22 ± 2 years) were recruited from the University of California, Berkeley, Department of Psychology research participant pool. Cerebellar degeneration patients and age-matched controls ($n = 10$ /group) were recruited from the Ivy Lab neurological patient database and communities around the San Francisco Bay Area (see Table 1). Patients with ataxia were scored on the International Cooperative Ataxia Rating Scale (ICARS), and a formal medical history was taken to confirm the etiology of each patient's

disease. When available, the outcome of genetic tests for spino-cerebellar ataxia (SCA) subtypes is provided. Only right-handed participants (Oldfield, 1971) were recruited. The institutional review board at the University of California, Berkeley, approved the research protocol. We did not perform a power analysis to predetermine our group sizes for each condition, as our experiments used a novel technique to induce the adaptation we measured. We did not know the magnitude of the group level effect it would have nor how it would affect the trial-by-trial variability of reaching. We did not, however, expect a wildly different outcome relative to standard techniques, leading us to choose a standard group size for the field, with 10 participants per condition. The only exception was Experiment 2, where we increased the group sizes to 20 participants because of high variance at generalization probe targets.

Experimental Apparatus

All participants performed center-out reaches on a horizontal surface while seated at a custom-made tabletop. Participants held a modified air hockey “paddle” that contained a stylus, and moved this device across a digitizing tablet (49.3 cm × 32.7 cm; Intuos 4XL; Wacom, Vancouver, WA). An LCD screen, 53.2 cm × 30 cm (ASUS, Taipei, Taiwan), was suspended 27 cm above the tablet. The experimental software was custom written in Python 2.7 with Pygame 1.9 modules (Beaverton, OR).

Reaching Task

Participants made reaches from the center of the workspace to targets positioned at a radial distance of 8 cm. The center location and target location were indicated by a white circle (1.2 cm diameter) and filled blue circle (1.6 cm), respectively. Direct vision of the hand was occluded by the monitor, and the lights were extinguished in the room to minimize peripheral vision of the arm. Position of the hand was indicated by a filled white cursor (4-mm circle). Participants were instructed to rapidly “slice” through the blue target to hit it with the white dot. If movements were not fast enough (300 msec), the words “too slow” were generated by the sound system of the computer. Visual feedback of the cursor was provided during the reach until the movement amplitude exceeded 8 cm. The cursor then froze in position for 1 sec to provide feedback concerning the spatial accuracy of the movement. The participant was free to begin moving back to the center during this time. After the spatial feedback period, the cursor disappeared and a white ring appeared, indicating the radial distance between the hand and center start position. The ring was displayed to aid the participant in returning to the center, without providing angular information about hand position. Once the participant maintained the digitizing stylus within the central start position for 200 msec, the target for the next trial was displayed.

Experimental Feedback Conditions and Procedure

Depending on the specific experimental condition and block, there were four types of visual feedback: no feedback, veridical feedback, rotated feedback, and clamped visual feedback. The cursor was not visible on no-feedback trials, whereas on veridical feedback trials, the cursor accurately showed the position of the hand. Rotated feedback involved an angular rotation of the cursor from the true hand position about the central start position. For clamped visual feedback trials, the feedback followed a trajectory that was fixed along a specific heading angle (Shmuelof

et al., 2012). The radial position of the cursor was based on the radial extent of the participant's hand (up to 8 cm amplitude) but was independent of the angular position of the hand.

Before the introduction of clamped visual feedback trials, participants were fully briefed on what they would see and were informed that they could not control the direction of the cursor movement. They were instructed to ignore the cursor and reach directly to the visual location of the target. These same instructions ("ignore the cursor and move your hand directly to the target location") were employed in the rotated feedback conditions of Experiment 5. Note that in the Ignore conditions, the visuomotor rotation was explicitly described to participants.

Experiment 1

This experiment had two goals. First, we sought to assess how task instructions affected implicit adaptation. Second, we introduce the clamped visual feedback method to examine changes in performance in the absence of task performance errors. There were three groups of participants, Compensate, Ignore, and Clamp. All groups were briefed on the nature of the perturbation, whether it was a 45° rotation or clamped feedback offset by 45°. Both the Ignore and Clamp groups were told to ignore the feedback while moving their hands directly to the target. The Compensate group was told to compensate for the imposed rotation to make the cursor hit the target. The Compensate participants were not given instructions on how, specifically, they should compensate for the rotation. Our primary measure of implicit adaptation was the aftereffect participants exhibited following removal of the perturbation. In this block, the participants were instructed to move their hands directly to the target in the absence of visual feedback. There were eight targets, spaced in 45° increments around 360° of the workspace. One target was presented on each trial, and the order was constrained such that all eight locations were tested in cycles of eight trials. The session started with two baseline blocks of 40 trials without visual feedback, and then 40 more trials with veridical feedback. The experimenter then informed the participant that the visual feedback would no longer be veridical, either rotated by 45° (Compensate, Ignore) or clamped 45° from the target location (Clamp). The perturbation block was composed of 240 trials, 30 to each of the eight target locations. This block was followed by two additional blocks: first, an eight-trial aftereffect block (1 reach/location) in which there was no visual feedback and then a 40-trial block with veridical visual feedback.

Experiment 2

Experiment 2 was designed to examine performance changes from extended exposure to task-irrelevant clamped visual feedback, comparing performance in this condition to that observed in a standard visuomotor rotation condition. "Learning" and generalization functions were obtained during and following exposure to clamped visual feedback or a visuomotor rotation. The visual target appeared at 1 of 24 locations, spaced in 15° increments around 360°. The test session began with a no-feedback block (96 trials) to measure intrinsic angular biases when reaching to each of the 24 target locations, followed by a block of trials with veridical feedback (96 trials). These two baseline blocks were followed by the perturbation block in which all of the reaches were to the 90° location (50 trials). For participants in the Rotation group (n = 20), visual feedback of the cursor was rotated in the counterclockwise direction by 22.5° during this block.

For participants in the clamped feedback group ($n = 20$), the visual feedback was invariant across trials, restricted to a trajectory that was offset from the reach target by 22.5° . The Rotation group was not informed of the perturbation; they were instructed to “do your best to hit the target.” The clamped feedback group was fully informed that the direction of the feedback cursor would be invariant and independent of their movement direction. These participants were instructed to move their hand to the visual location of the target and ignore the cursor. Following the perturbation block, the participants completed a no-feedback block (24 trials, 1 reach to each of 24 locations), with the instructions emphasizing for both groups that they should move directly to the visual target. This was followed by a veridical feedback block (96 trials) and, finally, a no-feedback block (96 trials).

Experiment 3

This experiment examined how the adaptation to task-irrelevant clamped visual feedback is affected by cerebellar degeneration, a neurological disorder that is consistently associated with sensorimotor adaptation deficits. A group of patients with varying severity of cerebellar ataxia (see Table 2.1 for details) and age-matched controls ($n = 10/\text{group}$) were exposed to a 45° -offset clamped visual feedback while reaching to four targets spaced in 90° increments around 360° . The trial structure began with a no-feedback and veridical feedback baseline blocks (40 trials each). The program was then paused, and the experimenter described the task-irrelevant clamped visual feedback, instructing the participants to always reach directly to the target while ignoring the feedback. The participants then completed a clamped visual feedback block (160 trials), following by a washout block (40 trials) with veridical feedback.

Experiment 4

This set of experiments examined how adaptation to task-irrelevant clamped visual feedback varies as a function of the magnitude of the angular offset of feedback from the target. Participants ($n = 90$, $10/\text{group}$) were randomly assigned to one of nine groups that differed in terms of the size of the counterclockwise clamped visual feedback: 0° , 7.5° , 15° , 30° , 45° , 60° , 95° , 135° , and 175° . To ensure that adaptation was sign dependent, half of the participants in the 30° and 60° groups were exposed to clockwise clamped visual feedback offsets.

The procedure was the same for all nine groups. On each trial, a visual target appeared at one of eight locations, spaced in 45° increments around 360° . The session started with two baseline blocks: first, 40 trials without visual feedback and then 40 more trials with veridical feedback. The experimenter then informed the participant that the visual feedback would no longer be veridical but would be clamped at a fixed angle from the target location and that this feedback should be ignored. The clamped visual feedback block was composed of 240 trials, 30 to each of the eight target locations. This block was followed by two washout blocks: first, an eight-trial block (1 reach/location) in which there was no visual feedback and then a 40-trial block with veridical visual feedback.

Experiment 5

This task was designed to assess whether the invariant response to clamped visual feedback of varying size (Experiment 4) was specific to conditions in which the feedback was not contingent on hand position or whether it was a general property of adaptation. To evaluate these hypotheses, we employed standard visuomotor rotation feedback (feedback direction based on reach direction) but informed the participants about the feedback perturbation. The participants were instructed to ignore the cursor and move their hand directly to the target location, similar to the instructions provided in task-irrelevant clamped visual feedback conditions.

Participants ($n = 30$, 10/group) were randomly assigned to one of three rotation sizes: 7.5° , 45° , or 95° . These rotations were selected to test the range over which we observed performance changes from clamped visual feedback in Experiment 4. The statistical comparisons for this experiment included three groups from Experiment 4 that were exposed to clamped visual feedback of the same magnitude of angular offset. The test session was organized in an identical manner, consisting of 40 no-feedback baseline trials, 40 feedback baseline trials, a 240-trial perturbation block, and two washout blocks, one with no visual feedback (8 trials) and one with veridical feedback (40 trials). Target number and spacing were the same as Experiment 4.

Data Analysis

The primary dependent variable in these experiments was the heading direction, defined by the position of the hand with respect to the target when the radial amplitude of the movement reached 8 cm. We describe this position in terms of “hand angle,” taking the difference between the heading direction and target position. Reaches with endpoint hand angles $\pm 90^\circ$ from the target location were excluded from analysis (less than 1% of trials on average, maximum 7%). Separate analyses were performed in which heading direction was identified at peak radial speed. There were no meaningful differences in the two behavioral measures, and thus, we only present the data from the 8-cm endpoint measure. All t tests were two-tailed unless otherwise noted. Statistical calculations were made with MATLAB 2009b (t tests) (The MathWorks, Natick, MA) and SPSS 17 (ANOVAs) (IBM, Armonk, NY).

Table 2.1

Patient #	Age	Gender	ICARS	Diagnosis
1	34	Male	17	SCA3
2	74	Male	20	Family History
3	53	Male	39	Family History
4	68	Female	31	Unknown
5	51	Male	22	Family History
6	44	Male	10	SCA5
7	44	Male	31	Unknown
8	77	Male	N/A	Unknown
9	63	Male	N/A	SCA10
10	55	Male	48	SCA6

RESULTS

Experiment 1

In many visuomotor adaptation studies, participants are aware, at least initially, of the perturbation (Werner et al., 2015; Hwang, Smith, & Shadmehr, 2006; Kagerer, Contreras-Vidal, & Stelmach, 1997). For example, with a 45° visuomotor rotation, participants recognize that the feedback has been manipulated and may adjust their movement goals by implementing an aiming strategy to compensate for the error (Taylor et al., 2014). To regulate the use of aiming strategies, we borrowed a technique that has been used in the prism adaptation literature to experimentally control for both awareness and changes in movement intent (Welch, 1969). We informed two groups of participants that there would be a 45° visuomotor rotation of their feedback. One group was instructed to compensate for the perturbation, with a goal to make the cursor hit the target (Compensate). The other group was told to ignore the cursor and continue moving their hand directly to the target (Ignore). Immediately following the perturbation block, both groups were instructed to aim directly to the target and no visual feedback was provided. These two conditions represent extremes in how naive participants might employ an aiming

strategy in a visuomotor rotation task, allowing us to explore how this affects behavior when the perturbation is present and, more importantly, after it has been removed (Figure 2.1B).

Both groups showed a significant change in hand angle between the feedback baseline and the last five cycles of the perturbation block (Figure 2.1C; Compensate = 42.3°, $t(9)=16.7$, $p<.0001$; Ignore=14.1°, $t(9)=7.4$, $p<.0001$). The change was much greater for the Compensate group compared with the Ignore group (95% CIs [21°, 34°], $t(9) = 8.9$, $p < .0001$). Following the perturbation block, both groups showed significant aftereffects when instructed to move their

Figure 2.

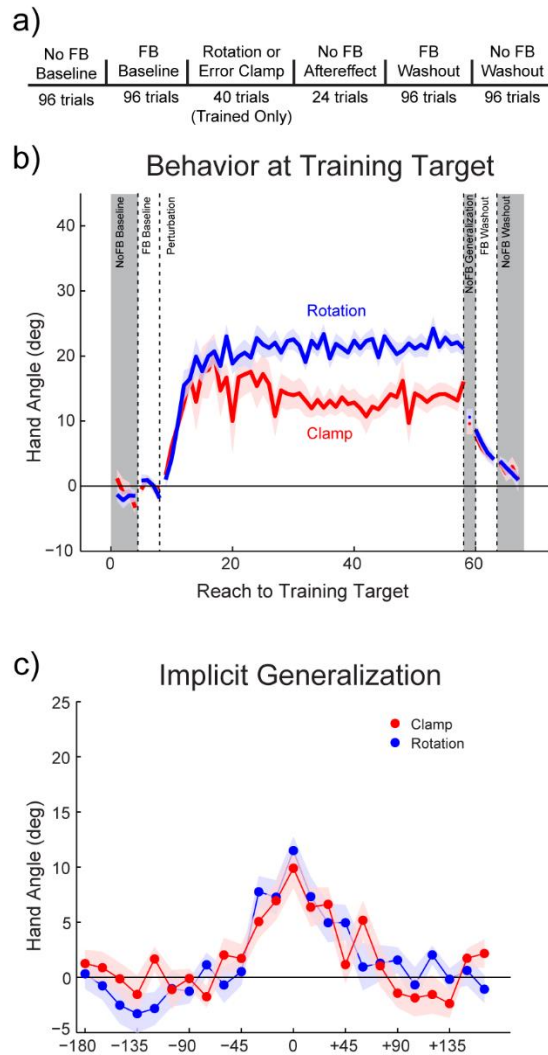


Figure 2.2 - Aftereffects and generalization are similar following exposure to a standard visuomotor rotation and clamped visual feedback. (A) Task trial design. Participants reached to all targets for every phase of the experiment, except for the perturbation block in which all reaches were to a single target. (B) Behavior at the training target during each phase of the experiment. Both groups show a change in heading angle in response to the visual feedback perturbation, with the rotation group changing hand angle by an additional 8°. However, there was no difference once the perturbation was removed. (C) Baseline subtracted aftereffect at all targets. The generalization functions are remarkably similar for the two groups around the workspace. Values are group means, shading and error bars denote SEM. Gray shading denotes cycles without visual feedback.

hands directly to the target (Figure 2.1D; Compensate = 12.3°, $t(9) = 8.2$, $p < .0001$; Ignore =

14.1°, $t(9) = 6.1$, $p < .0002$). Strikingly, the magnitude of the aftereffect was not different between the two groups despite the large difference in behavior during the perturbation block (95% CIs $[-7^\circ, 4.3^\circ]$, $t(9) = -0.4$, $p = .62$). The large decrease in hand angle for the Compensate group could be expected, given that the participants likely used aiming strategies to compensate for the rotation, and were instructed to move their hands directly to the target when the aftereffect was measured. Assuming that the aftereffect provides a probe of implicit adaptation, these results suggest that an implicit aftereffect is minimally affected by aiming strategies.

The similarity of the aftereffect for the Compensate and Ignore groups is particularly interesting given the difference in performance errors between the two groups. Both groups saw their feedback change over the course of the perturbation block as implicit adaptation progressed. For the Compensate group, continuous changes arising from implicit adaptation would actually lead to poorer performance if the participant's initial aiming strategy sufficiently countered the perturbation. The implicit learning would therefore necessitate adjustments of the aiming strategy to maintain good performance (Taylor et al., 2014; Taylor & Ivry, 2011). Things are not as clear for the Ignore group, but the contingency between the movement and cursor position may have induced participants, either explicitly or implicitly, to modify their behavior because of perceived performance errors (they were instructed that the feedback would always be rotated by 45°). This is especially problematic as mechanisms associated with reward and punishment have been hypothesized to modulate the rate of learning and retention in visuomotor adaptation tasks (Galea, Mallia, Rothwell, & Diedrichsen, 2015; Nikooyan & Ahmed, 2015). Such a modulation of learning or retention could have affected our measures of implicit adaptation in either group. We therefore sought to experimentally control task performance errors while simultaneously measuring implicit adaptation to putative sensory prediction errors.

To isolate implicit adaptation from task performance errors, we tested a third group in a condition in which we clamped the direction of the cursor's motion to an invariant angular offset (45°) from the target, while locking the motion of the cursor to the radial extent of the hand (Figure 2.1B). Participants were fully informed about the nature of this manipulation and instructed to ignore the cursor on all trials, always moving their hand directly to the target location. We refer to this method as task-irrelevant clamped visual feedback (Clamp). Note that, although both the Clamp and Ignore groups were instructed to ignore the feedback, the contingency between behavior and the direction of motion was only broken for the Clamp group, making it the only group that is fully isolated from task performance errors.

The Clamp group showed an adaptation profile that was remarkably similar to the Ignore group (Figure 2.1C), with a 15.5° change in hand angle between baseline and the last five cycles of the perturbation ($t(9) = 9.5$, $p < .0001$). These participants also showed a significant aftereffect when asked to reach directly to the target without any visual feedback (14.6°, $t(9) = 9.8$, $p < .0001$). Interestingly, all of the Clamp participants changed their behavior implicitly, reporting surprise that their hands were not traveling directly to the target when veridical cursor feedback was reinstated in a final washout phase.

Figure 3 .

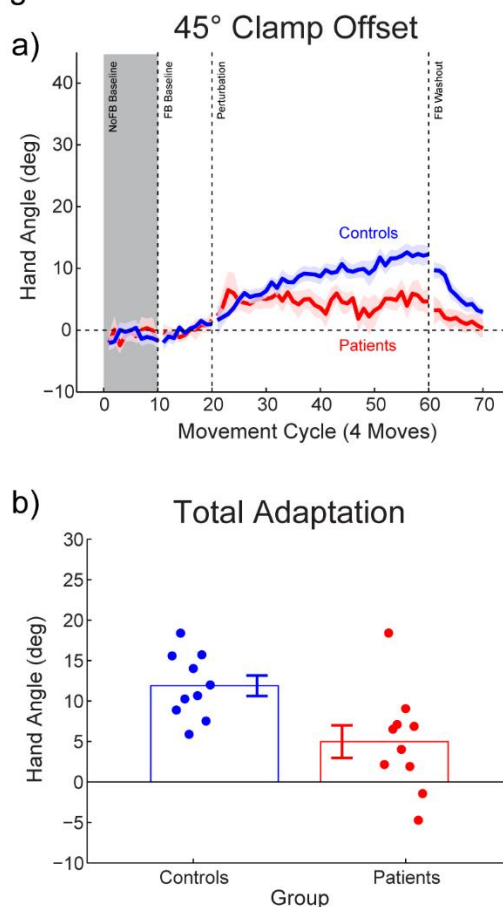


Figure 2.3 - Adaptation to a task-irrelevant clamped visual feedback is attenuated in patients with cerebellar degeneration. (A) Behavior for patients and controls over all movement cycles (4 reaches/cycle). (B) Adaptation over the last 10 cycles of the perturbation block showing both group lines indicate group mean \pm SEM and individual hand angle values (dots, horizontally displaced for visualization). Gray shading denotes cycles without visual feedback.

When the Clamp, Ignore, and Compensate groups were directly compared in a Condition \times Block mixed ANOVA, there was an effect of Condition (Figure 2.1D; $F(2, 27) = 29.49$, $p < .001$), block ($F(2, 54) = 262.5$, $p < .001$) and a significant interaction ($F(4, 54) = 52.0$, $p < .001$). The interaction was driven by the difference between the groups during the perturbation block ($F(2, 27) = 82.3$, $p < .001$). There was no difference between the groups in the aftereffect block ($F(2, 27) = 0.16$, $p = .852$).

These results suggest that the differences in task performance between the three groups had a negligible effect on implicit adaptation. The aftereffect was similar regardless of whether or not the participants compensated for a perturbation in which the feedback was contingent on task performance (Compensate vs. Ignore) and was also similar when the feedback was noncontingent (Clamp). Indeed, the rate of adaptation during the perturbation block was similar for the Ignore and Clamp groups, providing novel, trial-by-trial evidence that the sensorimotor system automatically adapts in response to sensory prediction errors, independent of task goals.

Figure 4.

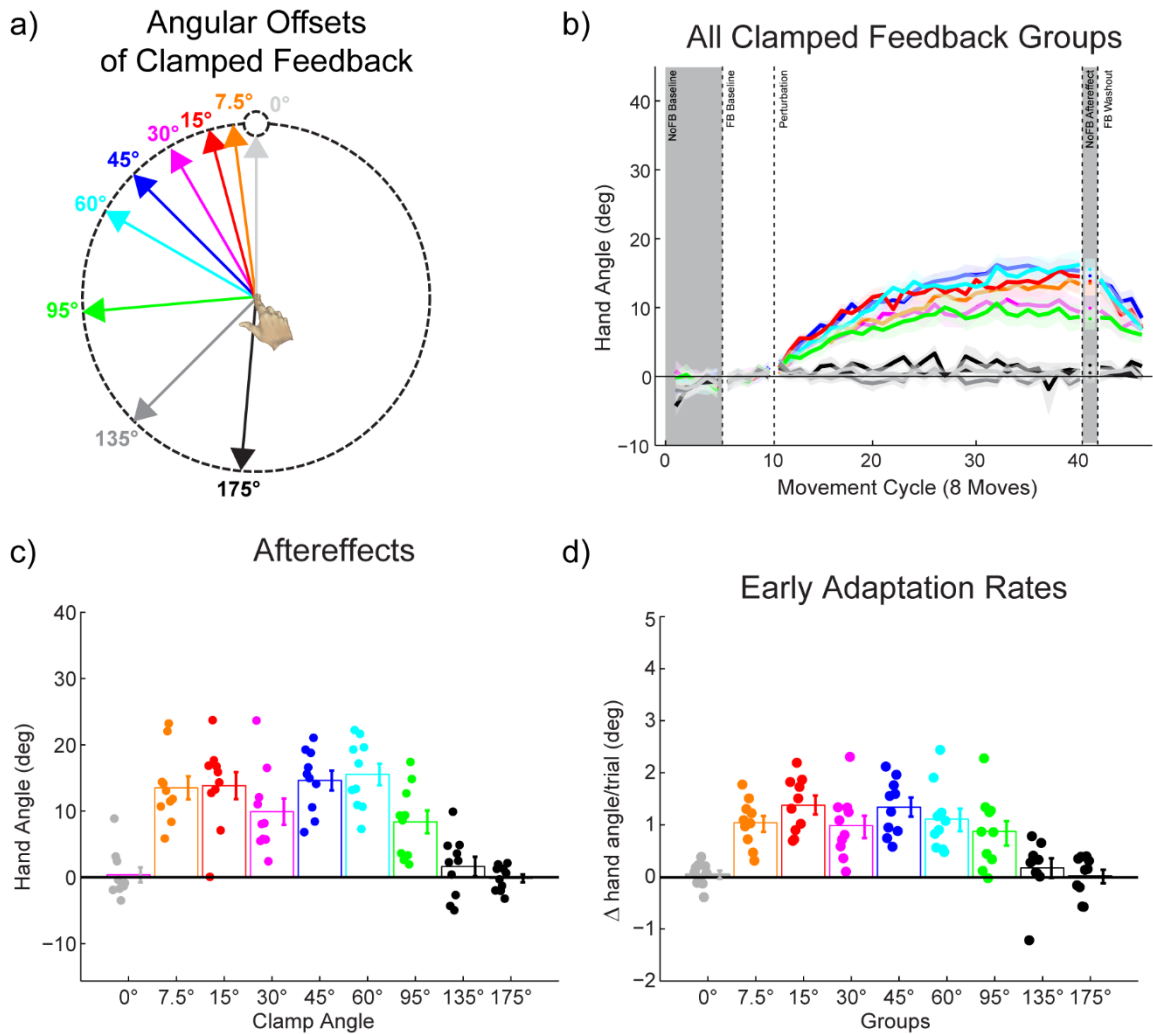


Figure 2.4 - Adaptation to clamped visual feedback does not scale with error size. (A) Color coding for clamped visual feedback of different offsets. Note that targets appeared at one of 8 positions on a given trial. (B) Behavior for all groups, with data averaged over movements to all 8 targets for each cycle. (C) Sensorimotor aftereffects, measured immediately following the perturbation block. (D) Per-trial adaptation rate over the first 10 cycles (80 trials) of the clamped visual feedback block. Dots are individual means, lines indicate group mean \pm SEM. Gray shading denotes cycles without visual feedback.

Characterizing the Adaptive Response from Sensory Prediction Errors

In the next two experiments, we further examined the changes in performance induced by task-irrelevant clamped visual feedback. In particular, we sought evidence that the learning process engaged by this type of feedback exhibits features associated with the adaptation of an internal model from sensory prediction errors.

Experiment 2

One hallmark of sensorimotor adaptation is that generalization is primarily local (Krakauer et al., 2000; Ghahramani, Wolpert, & Jordan, 1996; Imamizu, Uno, & Kawato, 1995). After adapting to targets at one location, the change in hand angle when reaching to other locations (without feedback) decreases as a function of the distance between the training location and probe locations, reaching a near-zero asymptote for targets that are more than 60° from the training location. We compared the generalization function following training with a 22.5° visual error clamp to that observed following training with a 22.5° visuomotor rotation. During training, all reaches were directed at a single location, with the feedback either non-contingent (Clamp) or contingent (Standard) on the participant's hand position (Figure 2.2A). The Standard group was naive to the perturbation, whereas participants in the Clamp group were briefed on the perturbation and instructed to ignore the feedback. In the subsequent generalization block, participants performed one reach to each of 24 locations, spaced at 15° around the workspace (Figure 2.2A). No feedback was provided in this block, and the participants in the Standard group were explicitly instructed to reach directly to the target, thus providing them with the same instructions as that given to the Clamp group for this phase of the task.

Both groups showed a shift in hand angle across the training block (Figure 2.2B; $F(1, 38) = 400.8, p < .001$). There was a reliable interaction between feedback condition and training block ($F(1, 38) = 19.4, p < .001$), with the Standard group exhibiting a larger shift in hand angle over the block relative to the Clamp group ($t(38) = 4.41, p < .001$). Despite these differences in performance during the perturbation block, the generalization functions for the two groups were remarkably similar: There was no difference between the two groups across the data from all probe locations (Figure 2.2C, $F(1, 38) = .09, p = .76$). Interestingly, even at the training location, the change in hand angle during the generalization block was not different for the two groups ($t(38) = .73, p = .47$). Similar to Experiment 1, the drop in hand angle at the training target for the Standard group, relative to the perturbation block, suggests that participants were aiming to compensate for errors during the perturbation block.

Experiment 3

Adaptation to sensory prediction errors is dependent on the integrity of the cerebellum (Schlerf, Xu, Kleffuss, Griffiths, & Ivry, 2013; Izawa, Criscimagna-Hemminger, & Shadmehr, 2012; Criscimagna-Hemminger, Bastian, & Shadmehr, 2010; Taylor, Kleffuss, & Ivry, 2010; Tseng, Diedrichsen, Krakauer, Shadmehr, & Bastian, 2007; Smith & Shadmehr, 2005; Martin, Keating, Goodkin, Bastian, & Thach, 1996; Weiner, Hallett, & Funkenstein, 1983). As such, if the behavioral change induced by clamped feedback is driven by sensory prediction errors, these changes should be attenuated in individuals with cerebellar pathology. To test this prediction, we compared the behavior of individuals with cerebellar degeneration and matched controls in response to a 45° visual clamp. As in the other experiments, all participants were fully briefed on the nature of the perturbation and told to ignore the visual feedback while reaching directly to the target.

On average, the patients took approximately 100 msec longer than the controls to complete the reaching movements, although this effect was not reliable at the group level ($t(18) = -1.44, p = .19$). Both groups responded to the visual clamp (Figure 2.3A), exhibiting a significant change in

hand angle relative to baseline over the last 10 movement cycles or 40 reaches (Figure 2.3B; controls: $t(9)=9.43$, $p<.001$; patients: $t(8)=2.4$, $p=.04$). However, the cerebellar group showed a smaller change (approximately 5° shift) than the controls (approximately 12° shift) ($t(18) = 2.58$, $p = .02$). Movement time was not correlated with the magnitude of the change in hand angle in either the controls ($r_8 = -.11$, $p = .77$) or patients ($r_7 = -.55$, $p = .12$).

Figure 5.

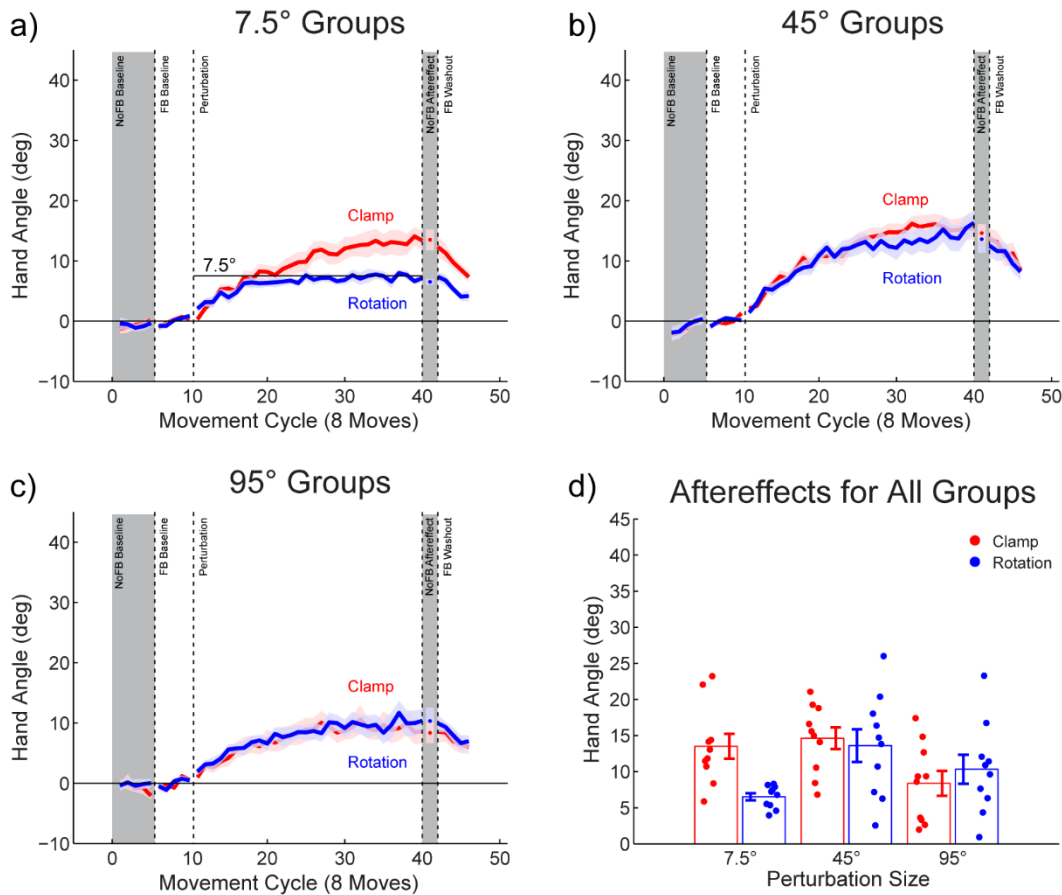


Figure 2.5 - Visual prediction error, contingent or not, drives adaptation. (A-C) Behavior over all movement cycles for each group with rotated feedback, presented with the corresponding clamped data from Experiment 3. (D) Sensorimotor aftereffects for all groups. Only the 7.5° group shows a difference between perturbation types. Dots are individual means, lines indicate group mean \pm SEM. Gray shading denotes cycles without visual feedback.

Experiments 2 and 3 provide converging evidence that the behavioral changes induced by noncontingent, clamped feedback are similar to those which occur from adaptation to sensory prediction errors in standard tasks. Generalization following noncontingent feedback was local and essentially identical to that observed with the standard method that involved contingent feedback and performance errors. The similarity in generalization between the two groups suggests that the clamped feedback method induces a similar change in an internal model (Krakauer et al., 2000; Ghahramani et al., 1996), rather than inducing a more idiosyncratic form of learning at the training location. Moreover, the behavioral response to the clamped feedback

was attenuated in patients with cerebellar degeneration. Taken together, the results of the first three experiments suggest that clamped feedback provides a powerful tool to examine implicit adaptation of an internal model that is driven by sensory prediction errors, without contamination from task performance error learning.

Sensitivity of the Adaptive Response to Sensory Prediction Error Size

In the final two experiments, we employ clamped visual feedback to assess constraints on sensorimotor adaptation. In standard visuomotor rotation studies, there is usually a positive correlation or a dose-dependent response between the size of the perturbation and the magnitude of compensation, at least when the perturbation is consistent during training (Abeele & Bock, 2001). However, in these tasks the amount of behavioral change required to reinstate good task performance is confounded with the size of the perturbation. As such, it is difficult to assess if dose dependency is a feature of adaptation from sensory prediction errors, a feature of learning from task performance errors, or both. Our clamped feedback method, by eliminating task performance errors, allows us to reexamine the impact of error size on adaptation from sensory prediction errors.

Experiment 4

We used a between-subject design, testing different groups of participants with clamped visual feedback in which the angular offset between the target and feedback ranged from 0° to 175° (Figure 2.4A). As expected, no change in hand angle was observed for the group exposed to a 0° clamp (i.e., where the cursor always went directly to the target). Surprisingly, we failed to observe dose-dependent adaptation for clamped feedback offsets between 7.5° and 95° (Figure 2.4B). For these groups, the change in hand angle during the perturbation block was similar, and the size of the aftereffect was largely invariant. For the two larger clamps, 135° and 175° , the presence of the clamp had no effect on hand angle.

These observations were confirmed by a single factor ANOVA of the aftereffect data. There was a significant effect of Offset size (Figure 2.3C; $F(8, 81) = 16.8, p < .001$), with the change in hand angle reliably different from baseline for clamped feedback offsets between 7.5° and 95° (paired t tests, $p_s \leq .0001$). In an analysis restricted to these groups, there was no correlation between the magnitude of the visual clamp and the size of the aftereffect ($r_{58} = -.15, p = .26$). Moreover, the per-trial rate of change in hand angle over the initial 10 movement cycles (80 trials) was not different between the groups that responded to the clamp (Figure 2.4D; $F(5, 54) = 1.17, p = .336$). Thus, we found no evidence of error size dependence in adaptation when task performance error was absent, suggesting that this is not a property of adaptation from sensory prediction errors.

Experiment 5

In the final experiment, we asked whether the absence of a dose-dependent response in Experiment 4 was idiosyncratic to clamped visual feedback or indicative of a more general property of implicit adaptation from sensory prediction errors. As emphasized above, clamped visual feedback breaks the spatial contingency between the direction of the hand movement and

the angular movement of the feedback cursor. It is possible that the motor system learns that the feedback is not contingent upon behavior, resulting in a habituation of learning from the clamped visual feedback and perhaps a lack of dose dependence.

To examine whether the lack of contingency caused the absence of a dose-dependent response, we again employed the Ignore Rotation (Rotation) condition from Experiment 1, where the direction of the feedback cursor was rotated and contingent on the angular position of the hand. Participants were informed of the size and direction of the feedback perturbation and instructed to actively ignore the feedback, attempting to reach directly to the visual location of the target. Three groups of participants were tested, with the size of the perturbation set to rotations of 7.5°, 45°, or 95°. We compared their performance to the groups instructed to ignore visual clamps of the same magnitude in Experiment 4, thus creating a data set of six groups (Rotation vs. Clamp, with three magnitudes) in the statistical analyses. (The data for both 45° groups were previously shown in Experiment 1.)

All groups showed a shift in hand angle during the training block (Figure 2.5A–C) and a positive aftereffect, relative to baseline (Figure 2.5D, $F(1, 54) = 271.9$, $p < .001$). The magnitude of the aftereffect was modulated by perturbation size ($F(2, 54) = 3.44$, $p = .04$), but not by feedback contingency ($F(1, 54) = 2.05$, $p = .16$). There was, however, a Size \times Feedback contingency interaction ($F(2, 54) = 5.04$, $p = .01$). This interaction was driven by a simple effect of Feedback contingency in the 7.5° groups ($F(1, 56) = 10.7$, $p = .002$), with the Clamp group showing a larger change in hand angle than the Rotation group. This is to be expected given that, for the 7.5° Rotation group, full cancellation of the perturbation has been achieved when the hand angle has shifted by 7.5°. In contrast, the visual discrepancy persists throughout the training block for the Clamp group. Interestingly, the magnitude of the change in hand angle was the same for the Rotation and Clamp groups who were exposed to either a 45° or 95° perturbation ($F(1, 56) = .13$, $p = .71$; $F(1, 56) = .29$, $p = .59$).

Taken together, the results of Experiments 4 and 5 indicate that the absence of a dose-dependent response to clamped visual feedback is not due to the lack of a contingency between the participants' behavior and the visual feedback. We surmise that the invariant magnitude of change in response to clamped feedback of varying size reflects a fundamental property of implicit adaptation to sensory prediction errors, with the dose dependency effects in standard tasks arising from the contribution of other learning processes that are sensitive to task performance errors (see Discussion).

The data from the 7.5° groups in Experiment 5 provide further evidence that our clamped visual feedback manipulation induces a sensory prediction error: The adaptation functions in the Clamp and Rotation conditions were nearly identical until the behavioral change for the Rotation group canceled the visual perturbation. Note that, per our instructions, in both conditions the feedback was not relevant to the participant's task performance. Nonetheless, the error was sufficient to drive a change in hand angle. Furthermore, the 7.5° group shows the limitations of the Rotation condition. In this case, the full adaptive response to a given error size may not be observed because the error is canceled by the participant's behavior. The performance of the 7.5° Clamp

group indicates that adaptation would have continued unabated if the error had not been canceled.

DISCUSSION

We examined the behavioral response to task-irrelevant clamped visual feedback, introducing a method to measure adaptation in the absence of contamination from processes that are sensitive to task performance error. Previous studies have used clamped visual feedback to induce adaptation, but their participants were not informed that feedback would be clamped; as such, the participants likely believed that changes in their behavior could affect this feedback to reinstate good task performance (Vaswani et al., 2015; Scheidt, Conditt, Secco, & Mussa-Ivaldi, 2005). A critical difference between previous clamp studies and our method is that we fully informed participants of the nature of the manipulation and asked them to ignore it. Because our participants knew what to expect, the feedback was task-irrelevant, in addition to being clamped.

In our clamped feedback conditions spanning 7.5° – 95° of angular offset, participants showed a robust change in reach direction. These changes in behavior appear to engage the same learning process that underlies implicit adaptation in standard visuomotor rotation paradigms: The change in reach direction was opposite in sign to the visual feedback, generalized locally, and was dependent on the cerebellum. Our method also provides the unique ability to directly measure implicit adaptation over an unlimited number of trials without changing the size of the imposed perturbation or relying on estimation techniques (e.g., subtracting out changes due to aiming strategies or using a model to estimate different components of learning). Taken together, our results provide a compelling demonstration of the automatic and obligatory nature of adaptation to sensory prediction errors and illuminate constraints on this process that are often obscured with traditional methods.

Invariant Adaptation across a Large Range of Error Sizes

Adaptation to the clamped visual feedback was relatively stereotyped, proceeding at a similar rate and reaching a similar final magnitude across a wide range of angular offsets (7.5° – 95°). The absence of a dose-dependent response here appears, at least superficially, at odds with the behavioral changes observed in standard sensorimotor adaptation studies. In those studies, the rate of change in behavior, as well as the final magnitude of that change, increases with the size of the perturbation (Abeele & Bock, 2001).

However, several studies have reported a breakdown in the linear relationship between the size of the error and the response to that error (Kasuga, Hirashima, & Nozaki, 2013; Marko, Haith, Harran, & Shadmehr, 2012; Wei & Körding, 2009; Fine & Thoroughman, 2006). Rather than use a constant perturbation, these studies have employed a method in which the size and sign of the perturbation varied from trial to trial, with an overall mean of zero (but see Robinson, Noto, & Bevans, 2003). One study (Wei & Körding, 2009) reported that the rate of adaptation was linearly proportional for small perturbation values and saturated somewhere between 5° and 15° . The current results do not bear on the question of proportionality for small errors, a question relevant for future work. Nevertheless, the results are consistent with the finding that learning rates saturate for larger errors. The learning rate, as revealed by our clamped feedback method,

indicates that implicit adaptation was invariant for errors between 7.5° and 95° and only dropped off for larger errors.

An important similarity between our clamped feedback method and zero-mean random studies is that both perturb feedback in a way that makes an adaptive response irrelevant to task performance, albeit for different reasons. With a zero-mean random perturbation, a change in behavior from one trial to the next is as likely to increase or decrease performance, forcing only “obligatory” adaptation to take place (Srimal, Diedrichsen, Ryklin, & Curtis, 2008). Tellingly, if the mean of the random distribution is shifted away from zero, a dose-dependent response is restored (Fine & Thoroughman, 2007; Scheidt, Dingwell, & Mussa-Ivaldi, 2001). We propose this is because ideal performance in this context requires compensation equal to the mean of the distribution and that dose-dependent adaptation emerges because corrections in response to the perturbation are relevant for task performance. That is, zero-mean perturbations engage only adaptation from sensory prediction errors, whereas nonzero mean perturbations also involve adaptation to sensory prediction and task performance errors.

Both our method and zero-mean random task designs provide converging evidence of a lack of error size dependence for implicit adaptation. We note that random perturbations have limited utility because the size and sign of the perturbation are constantly changing. This means that the method cannot be used to measure the full adaptive response to a given error over many continuous trials but rather is limited to the response to the error over just a few trials. Moreover, the measurement of adaptation to a given perturbation may be influenced by the context created by randomly varying perturbations (Castro, Hadjiosif, Hemphill, & Smith, 2014; Herzfeld, Vaswani, Marko, & Shadmehr, 2014; Wei & Körding, 2009; Scheidt et al., 2001). A strength of our method is that it allows the adaptive response to a fixed perturbation to be measured directly over a great number of trials, until learning from this specific error reaches an asymptote. In our experiments, implicit adaptation continued to increase over 40 reaches to each target, and the magnitude of this change was invariant to error sizes between 7.5° and 95° . We believe this invariance would continue until learning becomes asymptotic, although our experiments were not long enough to assess whether learning had fully achieved this state. It is important to emphasize that previous studies reporting insensitivity to error size focused solely on the rate of learning; the methods used in those studies could not address whether a similar insensitivity to error size also holds for learning that accumulates over many trials.

The Limits of Implicit Adaptation

Most theoretical frameworks of motor learning assume that implicit adaptation will compensate for the majority of a sensorimotor discrepancy if given a sufficient amount of training, even for very large errors (Cheng & Sabes, 2006; Pouget & Snyder, 2000; Thoroughman & Shadmehr, 2000; Jordan & Rumelhart, 1992). Our results suggest that this assumption is incorrect, at least with respect to implicit adaptation that arises from sensory prediction errors. Regardless of the size of the perturbation, the average change in hand angle was only around 12° across our various conditions. It is possible that performance had not reached asymptote, but this angular shift is substantially less than the raw change in hand angle observed in standard visuomotor rotation tasks of the same length. Interestingly, similar values have been observed for implicit

adaptation in studies that have employed methods designed to dissociate implicit and explicit components of adaptation (Bond & Taylor, 2015; Miyamoto et al., 2014). In light of these convergent results from very different task contexts, we believe it unlikely that the limit of implicit adaptation is related to our instructions to ignore the feedback in either the visual clamp or rotation conditions tested here.

At present, we can suggest possible hypotheses for this constraint on the extent of adaptation. It may be that this upper bound reflects an interaction between proprioception and vision: As adaptation from the clamped visual feedback produces a change in heading angle with respect to the target, the difference between expected and observed proprioception increases. Total learning could reflect the point that these two sensory prediction errors reach a state of dynamic tension. However, as traditionally described, such a process would be sensitive to the magnitude of visual and proprioceptive errors (Zaidel, Ma, & Angelaki, 2013; Henriques & Cressman, 2012; Ghahramani, Wolpert, & Jordan, 1997). This predicts differences in learning rate based on the magnitude of the visual error and, as such, is unlikely to account for our data given the invariant response to error magnitude. Furthermore, visuomotor adaptation is unaffected when visual feedback conflicts with proprioception (Marko et al., 2012). Nonetheless, it is possible that magnitude of total learning is limited by proprioception.

Another possibility for the invariant amount of total adaptation is that learning and forgetting have reached a state of dynamic tension (Cheng & Sabes, 2006; Thoroughman & Shadmehr, 2000). The present work suggests that the learning rate saturates for sensory prediction errors above 7.5° . This would also mandate that the forgetting term is independent of error size. This seems a reasonable assumption, although we cannot evaluate this quantitatively given that our experiments were not designed to make reliable measurements of forgetting. It will be important in future work to further assess the constraints on implicit adaptation; at present, our method has revealed important limitations with current theoretical models of this process.

Implications for Models of Sensorimotor Adaptation

Many models of adaptation have been developed to characterize learning from sensorimotor perturbations across extended blocks of trials. A core property of most of these models is a delta learning rule, whereby the amount of trial-by-trial change in behavior is the product of the error multiplied by a scalar learning rate, with the latter assumed to be constant across a range of perturbations (Cheng & Sabes, 2006; Thoroughman & Shadmehr, 2000). Such models fail to describe our data, as they predict an adaptive response that scales linearly with the size of the perturbation. Notably, models featuring fixed scalar learning rates also fail to predict behavior in tasks with random perturbations that flip sign and magnitude from trial to trial (Marko et al., 2012; Wei & Körding, 2009; Fine & Thoroughman, 2006). We propose that linear state-space models, and indeed any model that features error size multiplied by a scalar learning rate, will provide an inappropriate characterization of adaptation from sensory prediction errors.

Alternative models of adaptation have been developed to account for the behavior observed in response to random perturbations. One class of these models discounts errors based on their size, because of either the discrepancy between proprioception and visual feedback (Wei & Körding,

2009) or a subquadratic loss function for human sensorimotor errors (Marko et al., 2012; Körding & Wolpert, 2004). In their current form, both the proprioceptive and loss function models predict that the rate of learning will decline to zero, well below 95° of error. As such, these discounting models are unable to account for the absence of dose dependency observed in our results. Notably, the invariance was not an artifact of clamped visual feedback; it was also found when we used the ignore instructions to look at behavioral changes to task-irrelevant visuomotor rotations (Experiment 5).

Another type of model has been proposed in which the rate of adaptation is modulated over time, based on the history of experienced errors (Herzfeld et al., 2014). Of greatest relevance to the current study, this model predicts a sensitization to previously experienced errors. With clamped visual feedback, the angular offset remains invariant, ostensibly presenting the same sensory prediction error on every trial. We saw no sensitization of the adaptation rate in the clamped visual feedback experiments relative to rotation experiments where the error size decreased over trials. It may well be that this model better describes learning processes related to task performance errors, rather than sensory prediction errors. As noted above, we suspect that a major limitation with current models of learning from sensory prediction errors is that, though assumed to provide a characterization of adaptation, they actually model data reflecting behavioral changes that arise from a composite of sensory prediction errors and task performance errors (Huberdeau, Krakauer, & Haith, 2015; McDougle, Bond, & Taylor, 2015; Wolpert, Diedrichsen, & Flanagan, 2011; Redding, Rossetti, & Wallace, 2005; Weiner et al., 1983). In this context, processes other than error-based updating of internal models could dominate performance, especially in response to large perturbations. Indeed, recent work decomposing visuomotor learning into implicit and explicit components converges with our finding that implicit adaptation is invariant, as the dose-dependent learning was solely attributed to variation in the use of explicit aiming (Bond & Taylor, 2015). Even when performance errors are removed (as in aftereffect blocks without feedback), the effect of these other learning processes may persist, further obscuring the contribution of adaptation from sensory prediction errors (Diedrichsen, White, Newman, & Lally, 2010).

Conclusions

The current experiments introduce a method designed to measure implicit adaptation in the absence of task performance errors. The behavioral change observed in response to clamped visual feedback had many of the features that are historically associated with adaptation to sensory prediction errors: The response was monotonic, generalized in a limited manner, and was attenuated in individuals with cerebellar degeneration. By isolating learning from sensory prediction errors, we found that adaptation was invariant across a large range of perturbations, both in terms of the initial learning rate and total accumulated learning. Our results suggest that this lack of a dose dependency is a fundamental characteristic of implicit “obligatory” adaptation. Moreover, the results underscore the importance of carefully considering the source of behavioral changes in studies of sensorimotor adaptation, as different types of errors may lead to different types of learning. The task-irrelevant clamped feedback method can serve as a

powerful tool for future work on the role implicit adaptation in various motor learning phenomena.

Chapter 3

Practical and Theoretical Consequences of Intrinsic Reaching Biases in Visuomotor Adaptation Experiments

Abstract

When reaching without visual feedback, people show angular biases in their direction of reach. These biases are not idiosyncratic, rather, they are consistent across participants (Ghilardi et al., 1995; Vindras et al., 1997). Although most researchers are aware biases exist, it is not customary for these biases to be measured in sensorimotor learning experiments, and they are often not considered beyond baseline subtraction. This is problematic, as many adaptation metrics may be contaminated by these biases, especially when visual feedback is removed, as biases rapidly re-emerge in the absence of feedback (Smeets et al., 2006). We explored how intrinsic reaching biases can affect adaptation experiments in a series of experiments. We first measured the bias in a systematic manner with a large number of participants and, similar to previous studies, found that the sign and magnitude of the bias varied with reach direction. This bias function was similar for reaches with either the right or left hand, arguing against a biomechanical account and consistent with the hypothesis that the bias is primarily perceptual in origin (Ghilardi et al., 1995; Vindras et al., 1997). In terms of absolute size, the mean bias ranged from a low of 1° to a high of 13° , and showed two changes of sign around the workspace. We examined the impact of these biases on different protocols used to study visuomotor adaptation by modeling adaptation and generalization with a single process linear state-space model. Finally, we review different techniques that can be used to minimize the impact of these biases in studies of adaptation.

Introduction

When people make planar reaches without visual feedback, they exhibit persistent angular biases in the direction of reaches relative to visual landmarks (Ghilardi et al., 1995; Bock & Arnold, 1993; Desmurget et al., 1995; Rosetti et al., 1995; Vindras et al., 1997). These biases have several origins, but the primary contribution is thought to be from a perceptual bias which causes a participant to misestimate that his hand is offset from its true starting position (Vindras et al., 2005). This offset causes systematic errors in the planned trajectory from the start position to the target location, which lead to systematic errors in the reach endpoint. The claim that the primary bias is perceptual is bolstered by evidence that it can be greatly reduced if direct vision of the initial hand position is provided (Ghilardi et al., 1995). Regardless of the origin, all biases should eventually be reduced via adaptation if visual feedback of the reach is provided. However, direct vision of the hand and visuomotor adaptation merely reduce biases, instead of eliminating them. If not properly measured, it is possible for residual biases to contaminate measures of learning in visuomotor adaptation tasks without an experimenter's knowledge.

For instance, visuomotor adaptation will reduce the overall bias (mean over all directions), but it is not necessarily true that the bias will be reduced to the same asymptotic value for all target directions. Put differently, even after a long baseline with visual feedback there may be a

persistent residual error that differs as a function of target angle around the workspace. If feedback is never removed, this is unlikely to substantially affect the outcome or interpretation of most experimental designs in visuomotor adaptation, as these residual errors are likely to be small. Such residual biases could, however, affect designs where minor differences in learning magnitude are expected for different target locations because of generalization or differences in the learning stimulus.

Another concern is for experiments where feedback is removed for more than a few reaches, or if only non-positional feedback (such as reward) is provided. In such experiments, there is normally an initial baseline period where participants reach with visual feedback, which will cause intrinsic biases to be reduced. However, once this feedback is removed, participants will show a slow decay back to their intrinsic bias (Smeets et al., 2006). Such decay could be mischaracterized as learning in non-positional feedback experiments, or as a difference in forgetting for experiments where feedback is removed after a perturbation block. Given the seriousness of these concerns, it is remarkable that experimenters rarely measure intrinsic biases at the beginning of an experiment, before they are exposed to any visual feedback, so that the point they should decay to in the absence of feedback can be measured.

The purpose of this study is to examine how intrinsic biases can affect results in visuomotor adaptation experiments, as well as to consider methods that can be employed to measure and account for these biases. To examine this issue, we used a center-out reaching task to measure intrinsic biases in a relatively large sample of participants. In Experiment 1, we measured bias functions in the absence of visual feedback and then observed how this bias was reduced in a location-specific way when visual feedback was introduced. We then perform a series of model simulations using a single-process linear state space model with Gaussian generalization and parameters derived from fitting the data to individual participants. These simulations afford an understanding of how these biases interact with various task designs that are commonly employed to study visuomotor adaptation.

Methods

General Experimental Protocol

All participants performed center-out reaches on a digitizing tablet (49.3cm x 32.7cm, Intuous 5; Wacom, Vancouver, WA) while holding a modified air hockey "paddle" that contained a stylus. The table was oriented horizontally and mounted 27 cm below an LCD monitor (53.2cm x 30cm, ASUS), also oriented horizontally. With this arrangement, visual feedback of the hand was occluded by the horizontal monitor. The experimental software was custom written in Python 2.7 with Pygame 1.9 modules.

Reaching Task

Participants made reaches from the center of the workspace to targets positioned at a radial distance of 8 cm. The start position and target location were indicated by a white circle (1.2 cm diameter) and a filled blue circle (1.6 cm), respectively. Direct vision of the hand was occluded by the monitor, and the lights were extinguished in the room to minimize peripheral vision of the arm. When provided, visual feedback of hand position was indicated by a small white cursor (4 mm). Participants were instructed to make a rapid reaching movement, attempting to "slice" through the blue target. If the amplitude of the hand movement did not reach 8 cm within 300 ms, the message "too slow" was displayed, regardless of whether the trial was a feedback or no-feedback trial. On trials with visual feedback, the cursor was visible during the reach and then froze for 1 s when the amplitude reached 8 cm. This pausing of visual feedback emphasized the spatial accuracy of the movement. Once the cursor was extinguished, a white ring appeared, indicating the radial distance between the hand and center start position. The ring was displayed to aid the participant in returning to the center, without providing directional information about hand position. Note that participants were free to return to the start position at any time, but without feedback until the ring was presented. Once the participant maintained the digitizing stylus within the start position for 200 ms, the target for the next trial was displayed.

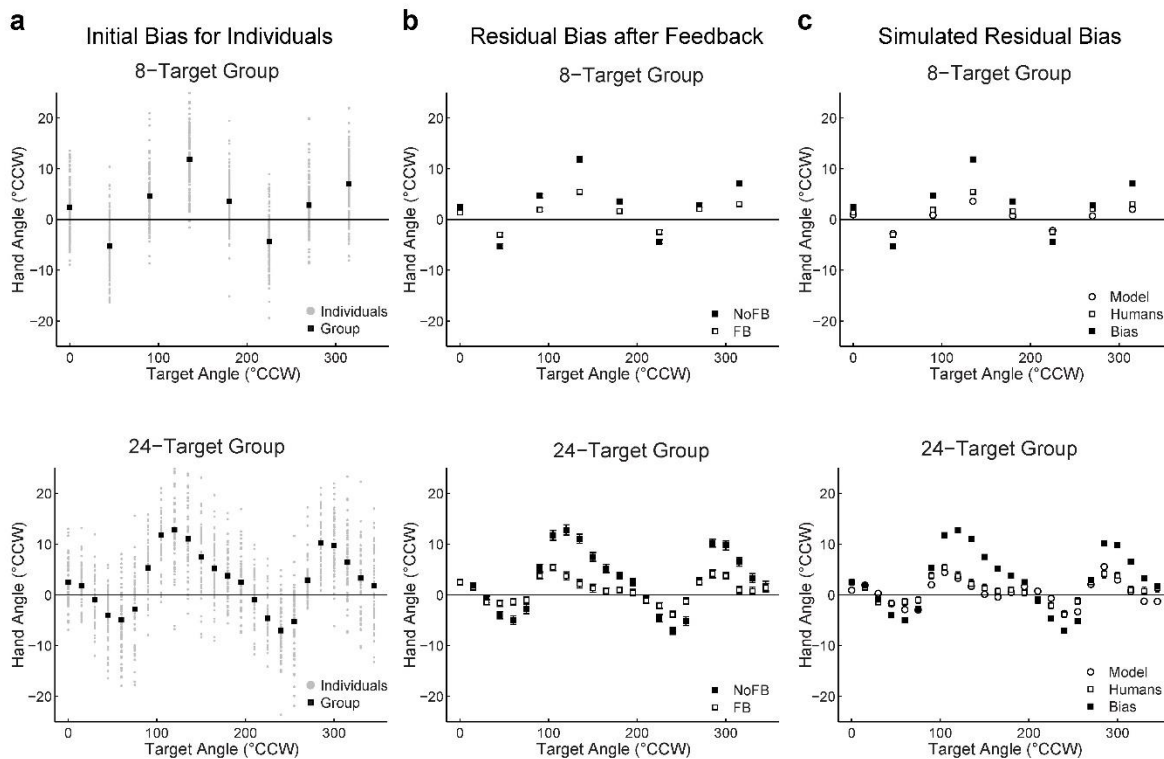


Figure 3.1 – Human behavior and model performance. a) Human reaching biases while reaching without visual feedback to either 8 (above) or 24 targets (below). b) Initial bias and performance after a short period with visual feedback for participants. Note that the bias is not fully eliminated. c) Simulated behavior of the biased state-space model after a short period of visual feedback.

Experiment 1

The goal of this experiment was two-fold, to Measure the intrinsic angular reaching bias before exposure to visual feedback and to measure how bias was attenuated after visual feedback was provided. In order to collect a large number of participants, we added two short block of trials for this purpose to the beginning of unrelated experiments. There were two different conditions (between-subject). For 150 participants, there were 8 targets, spaced in 45° increments around the workspace; for 56 participants, there were 24 targets, spaced in 15° increments. Participants in the 8 target conditions made 5 reaches per target (40 trials/block) during a no feedback and feedback block, while the 24 target group had 4 reaches per target (96 trials/block).

Data Analysis

The primary dependent variable in these experiments is the hand angle, defined as the position of the hand with respect to the target when the radial extent of the movement exceeded 8 cm. As such, this measure defines the angular difference between heading direction and target position. We also performed these measurements based on the time at which the hand reached peak radial speed (approximately 4 cm into the movement). The results were essentially the same for the two measures, indicating that online corrections were minimal. As such we only report the hand angle at 8cm reach extent.

All t-tests were two-tailed unless otherwise noted. Statistical calculations were made with Matlab 2009b (t-tests) and SPSS 23 (ANOVAs).

Modeling

We used a single-process linear state-space model to assess learning and generalization in Experiment 1 (Thoroughman & Shadmehr 2000; Cheng & Sabes, 2006; Tanaka et al., 2009). The output of the model (y) on a given trial (t) was determined by the internal state (x) for a given aiming direction (a) and its associated bias (b) for that aiming location:

$$y_t = x_{a,t} + b_a \quad (1)$$

The internal model (x) was simulated as a matrix with values for all aiming directions on every trial, and the bias (b) as an array of intrinsic biases for every aiming location.

$$x = \begin{bmatrix} a_{1,1} & \cdots & a_{1,t} \\ \vdots & \ddots & \vdots \\ a_{n,1} & \cdots & a_{n,t} \end{bmatrix} \quad b = [a_1 \cdots a_n] \quad (2)$$

The error on a given trial (ε) was computed as the motor output minus the fixed angular perturbation, modulated by the visual feedback on that trial, f , which was either a 0 or 1.

$$\varepsilon = (y_n - r_n) \cdot f_n \quad (3)$$

We simulated generalization by computing a Gaussian function that spanned all reach directions, k , centered on the aiming direction with a peak value of 1. The Gaussian's width, σ , was set to 30° to approximate empirically derived Gaussian generalization widths (Pine et al., 1996; Ghez et al., 2000; Krakauer et al., 2000; Brayanov & Smith, 2012).

$$k_a = 1 \cdot e^{-\frac{D_a^2}{2\sigma^2}} \quad (4)$$

The height of the generalization at each aiming location is determined by D_a , an array defined as the absolute angular distance between all aiming locations and the aiming location (always the target position) on the current trial.

$$D_a = |a_t - \begin{bmatrix} a_1 \\ \vdots \\ a_n \end{bmatrix}| \quad (5)$$

This generalization function was then used to modulate both learning, B , and forgetting, A , of the internal state, x , for the next trial at each possible aiming direction.

$$x_{a,t+1} = k^T A x_{a,t} + k^T B \varepsilon_t \quad (6)$$

Note that the amount of forgetting is a function of the internal state for each location and the generalization value, k , while learning is a function of the error at the current aiming location combined with k . T is a transpose operator. The forgetting and learning rates, A and B , were free parameters bounded between 0 and 1 or -1 and 1, respectively.

We fit this model to individual participants using the built-in Matlab function `fmincon`, leaving only the learning and forgetting rates as free parameters, bounded between -1 and 1. Aiming directions were always fixed as the target direction for this study, and other aiming locations were limited to the targets in the experiment.

Results

Experiment 1

Two large groups of participants were tested on two short blocks of trials, the first to measure intrinsic bias in the absence of visual feedback and the second to observe how this bias was attenuated by visual feedback. For the group tested in the 8-target condition, there was a clear and unambiguous angular bias that varied with reach direction (Fig. 3.1a). It is common for experimenters to report only the mean (signed) bias over all target directions, so we calculated this value for all the participants and found was different from zero (mean 2.8° , 95% CIs [2.35° , 3.3°], $t_{149} = 11.85$, $p < .001$). Furthermore, turning feedback on for 40 trials (5/target) did not fully eliminate the mean bias (Fig. 3.1c, mean 1.2° , 95% CIs [$.91^\circ$, 1.57°], $t_{149} = 7.49$, $p < .001$).

An average over all targets will, however, obscure the differences between the targets. It is especially important in this case because the bias function is not a fixed angular offset, but an oscillating function that differs in sign and magnitude depending on the target direction. To assess the effect of both target direction and feedback, we performed a mixed-design ANOVA. There was a main effect of feedback type ($F_{1,101} = 18.9$, $p < .001$); a within-subject effect of target direction ($F_{23,79} = 60.0$, $p < .001$) and a feedback type by target direction interaction ($F_{7,292} = 14.8$, $p < .001$). Taken together, these statistical tests indicate that participants reduced the bias through visuomotor adaptation, but were not able to fully eliminate it with this small number of reaches.

We also performed the same experiment with a more fine-grained target spacing of 15° , for a total of 24 targets around 360° (Fig. 3.1b, d). Participants only had four reaches per target in the two blocks of this experiment, but a greater number of overall trials compared to the 8-target group (192 vs 80). As with the 8 target group, there was a main effect of feedback type ($F_{1,298} = 29.89$, $p < .001$); a within-subject effect of target direction ($F_{7,292} = 144.7$, $p < .001$) and a feedback type by target direction interaction ($F_{7,292} = 16.8$, $p < .001$). Furthermore, the mean bias was different from zero (mean 2.8° , 95% CIs [2.1° , 3.6°], $t_{55} = 7.99$, $p < .001$), and this was not reduced to zero by the last reach with feedback (mean 1.2° , 95% CIs [$.6^\circ$, 1.3°], $t_{55} = 5.15$, $p < .001$).

Importantly, we wanted to assess whether the measurement of the intrinsic bias was affected by the number of total reaching targets. To do this, we compared the average hand angle at the same eight target locations in both target number conditions. We found an effect of target direction ($F_{1,1428} = 193.7$, $p < .001$), but no effect of target number ($F_{1,204} = .06$, $p = .81$) and no interaction ($F_{7,1428} = .70$, $p = .67$), indicating that the bias at this subset of target locations was not different between the two conditions.

Because the bias was not eliminated by the end of the feedback block in both conditions, we fit a single process state-space model with Gaussian generalization to the individuals in both conditions of Experiment 1 (Eq. 1-6). The parameters derived from this process allow us, in subsequent simulations, to judge whether the bias can be reduced below motor noise thresholds. The model was provided with each individual's target-specific bias (mean of no-feedback baseline reaches), and was fit using the data from the feedback block.

The A value was $.089 \pm .03$ for the 8 target group, and $.035 \pm .01$ for the 24 target group, while the B value was $-.356 \pm .02$ for the 8 target group, and $-.241 \pm .015$ for the 24 target group. The parameter values for A were not different across the two groups ($t_{204} = 1.43$, $p = .15$), suggesting that the forgetting rate was independent of the number of targets. However, the B values were different ($t_{204} = 3.75$, $p < .001$). This difference in learning rate is probably due to the fact that the spacing between targets is different for the groups. The narrower spacing in the 24-target groups increases the amount of learning via generalization from reaches to neighboring targets. We believe it unlikely that the number of targets modulates learning rates, we interpret this discrepancy as a limitation of the modeling approach rather than a true difference in learning rate.

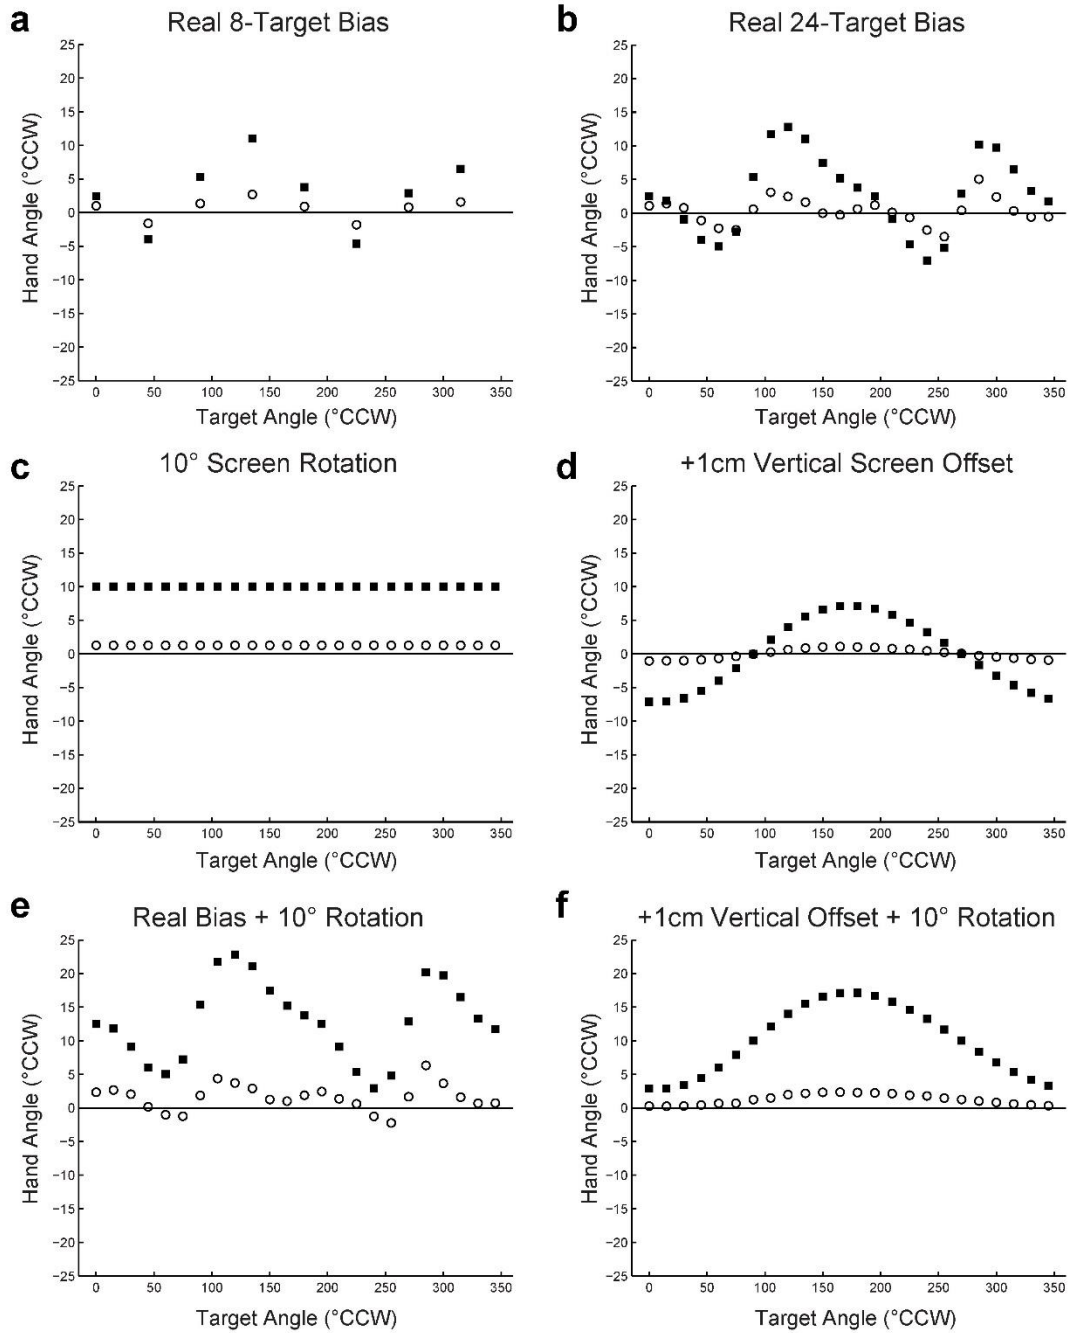


Figure 3.2 - Initial bias and model performance after a large number of feedback trials. a-b) Asymptotic performance of the model after a great number of trials for the 8- and 24-target conditions. c) Simulation of a perfect angular bias, like that introduced by a screen rotation. d) Simulation of a 1cm vertical offset of a the screen, relative to the reaching surface. e) The real bias, with the additional bias shown in c) added. f) Simulation of a 1cm vertical offset, with the additional bias of c) added.

Despite this caveat, the model does a fairly well at capturing the behavior of the two groups when provided with median parameters and the group mean bias (Figures 3.1e and 3.1f). The fit

accounts for 98.7% of the variance in the 8-target condition and 72% of the variance in the 24-target condition, and retains the idiosyncratic shape of the residual bias following adaptation in the feedback baseline. The greater R2 value in the 8-target condition comes about because it has fewer number of targets, or fewer degrees of freedom, making higher R2 values easier to achieve than in the 24-target condition. We believe that our model represents a reasonable approximation of human sensorimotor adaptation in this setting. This type of model is typically used to simulate a response to an imposed perturbation. Here we add the bias as a fixed, constant perturbation that the model will always try to eliminate in the presence of visual feedback. We will use the model later to explore questions that this experiment cannot address, such as whether the bias can be fully eliminated and how the bias interacts with a variety of visuomotor adaptation task designs.

Simulation Set 1

The purpose of this simulation was to assess the persistence of intrinsic biases once learning becomes asymptotic. Specifically, we examined whether our generalization model would eventually achieve a small, uniform, bias across the workspace after making 600 reaches to each target (14400 total trials). In addition, we considered how the asymptotic performance of our model would change if provided with different biases, such as those that may be introduced from improperly aligning a digitizing tablet and screen.

We used the median best-fitting parameters from the model fits in Experiment 1 to simulate learning. Each simulation used the parameters from the data set with the same number of targets. Additionally, we gave each model simulation the mean bias from each respective group. Targets were pseudo-randomly sampled so that each target was reached to once before the set repeated, similar to the procedure used in the actual experiment (see Methods). For all the simulations in this set, feedback was present on every trial.

As shown in Figure 3.2a and 3.2b, the model did not fully eliminate the bias for either the 8-target or 24-target simulations. Instead, the model's asymptotic performance retained much of the shape of the initial bias, especially near target locations where the initial bias flipped sign. We did not anticipate this result, but soon realized that it is a consequence of signed-error based learning and generalization. This is because the model's learning asymptote is not only determined by the equilibrium between learning and forgetting rates, but also the ratio between the width of Gaussian generalization and differences in magnitude of the bias between neighboring targets. Therefore, it is not the initial sign of the error at each location that determines whether the model can effectively reduce the bias to near zero, but rather the rate of change in the bias slope for neighboring targets. In areas of the workspace where the bias slope changes very rapidly, learning at neighboring targets can cause interference through generalization (assuming learning, forgetting, and generalization width are constant). These zones of interference appear as 'wrinkles' in the generalization function, where a substantial bias will persist.

To further investigate this effect, we assessed how our model would respond to other biases that do not have sharp changes of sign and magnitude around the workspace. We first simulated a misalignment of the screen and reaching surface, so that they were rotated relative to each other to induce a bias with a fixed angular offset of 10° (Fig 3.2c). With extended visual feedback training, our model reduced the error from this bias to create a near-uniform, non-zero asymptote at all target locations. We also simulated a 1 cm translation of the two surfaces (along the y-axis), which induces an asymmetric angular bias at every target except the 90° and 180° locations (Fig 3.2d). Here, the model reaches asymptotic values closer to zero than what is possible with the real bias. This suggests that it is the rate of change in the local slope of the bias that causes interference from generalization, not merely the changes in sign.

To confirm that it is the changes in local slope that determine residual errors from a bias, we added a fixed angular offset of 10° to the real bias so that all target directions would have an initial bias with the same sign (Fig. 3.2e). The model will reduce this overall bias until some of the target locations flip sign, reaching asymptotic performance that is nearly identical to that achieved with the true bias that does not have an additional angular offset (Fig. 3.2b). This dynamic is not observed when the model reduces the slowly changing bias that would occur if a 10° rotation of the screen was added to a 1cm screen translation (Fig. 3.2f). The difference between the model's asymptotic performance for these two simulated biases helps illustrate why the real intrinsic bias function cannot be fully eliminated.

Both biases begin at approximately the same distance from zero, and adaptation initially reduces errors in an equivalent way for both simulations, because both local learning and generalization work together when the overall function is far from zero. As adaptation proceeds, the error at some target directions will approach zero, while other directions will still have a substantial amount of error. At this point, generalization from target directions with substantial errors will begin to push the other target directions beyond zero degrees of error, causing negative errors at these locations. This learning and generalization will continue until the model reaches a steady state of dynamic tension, where the average error of all target angles is close to zero, and some target directions still show a substantial amount of error. This will not occur with the simulation shown in Fig. 3.2f, because there are not large differences in angular bias between neighboring targets. Taken together, these simulations demonstrate that the real bias is not eliminated because generalization interacts with the rapid cycling of the bias around the 90° and 270° target locations.

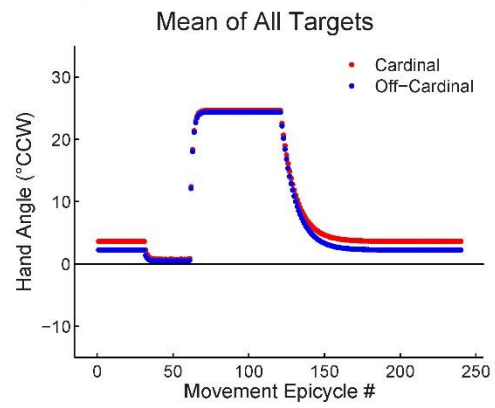
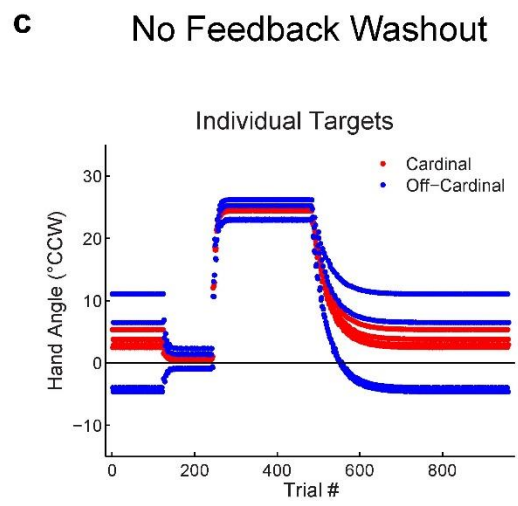
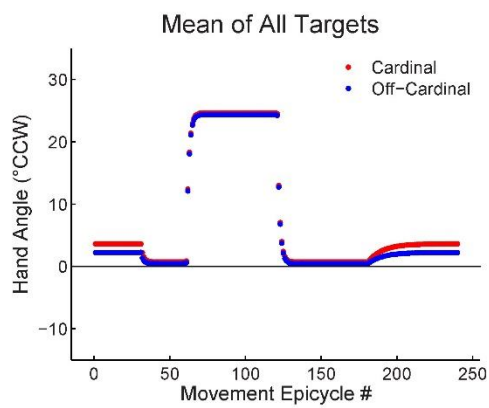
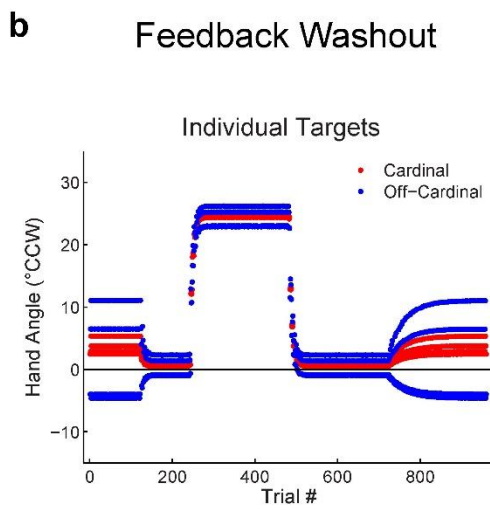
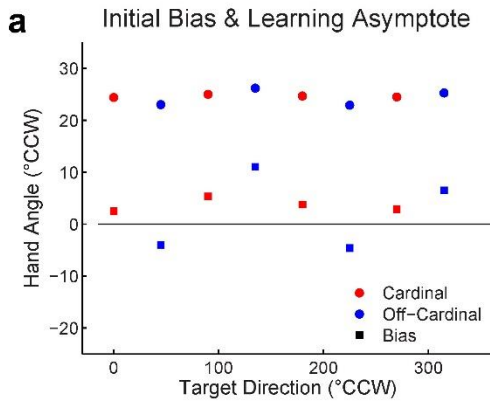


Figure 3.3 – Simulation of target-specific washout effects with and without feedback. a) Initial bias of the model, and asymptotic performance for each target location. b) Performance of the model with a feedback washout, for all targets (top) and the mean of all targets (bottom). c) Model performance during no-feedback washout.

Simulation Set 2

In visuomotor adaptation tasks, the direction of reach is often chosen for arbitrary reasons. In this set of simulations, we examined how target choice could affect the results of a learning experiment relative to the bias we measured. First we looked to see if there was a difference between experiments in which the targets were restricted to the cardinal directions (0° , 90° , 180° , 270°) or off-cardinal directions (45° , 135° , 225° , 315°). Note that, in our experimental data, the intrinsic bias has the same sign for the cardinal directions, while half the off-cardinal targets have an opposite sign from other two targets (see Fig. 3.3a).

We simulated learning at target sets in response to a 30° rotation, comparing two ways of probing learning: aftereffects with and without visual feedback. We examined these two types of feedback conditions during washout because they are often compared within a study (Galea et al., 2010; Kitago et al., 2013; Leukel et al., 2015; Galea et al., 2015), but may be differentially affected by biases. No feedback washouts are typically used to assess retention of the motor memory in the absence of error-based learning, while feedback washouts assess all aspects of error-based learning. For both simulations, we included the same three initial blocks: no-feedback baseline (120 trials), feedback baseline (120 trials), and a perturbation block with the rotation (240 trials). In one simulation, we assessed the aftereffect by first conducting a block of washout trials with feedback (240 trials) followed by a washout without feedback (240 trials). In the second simulation, feedback was removed for all trials after the perturbation block (480 trials).

Although there were pronounced differences in the initial bias across target sets, this had only a minor effect on the state of asymptotic learning during the perturbation block (Fig 3.3a). These minor differences persisted into the washout phase, but the differences across targets remained minor if the initial washout block included feedback (Fig. 3.3b, c). Furthermore, the mean hand angle over all the targets, gave a faithful representation of the behavior at individual targets. In contrast, in the no-feedback washout condition (Fig. 3.3d, e), there were more pronounced differences between the two target sets. This occurred because each target location decayed to its intrinsic bias rather than a state of zero visual error. In this condition, the mean behavior across the targets did not accurately convey the behavior at each target location. Furthermore, because of the mean difference in the bias that the model was decaying to, this effect could be misinterpreted as a difference in retention of the motor memory, when it was entirely due to target direction-dependent biases.

This simulation indicates that when the critical phases of an experiment are performed with feedback (and the same perturbation is delivered to all reaching locations), the bias we measured in Experiment 1 is unlikely to cause problems of interpretation because there are not major differences across target locations. However, caution is advised if feedback is removed during the washout block, because the re-emergence of intrinsic biases will impact where adaptation will decay to, and this effect can vary with the target angle.

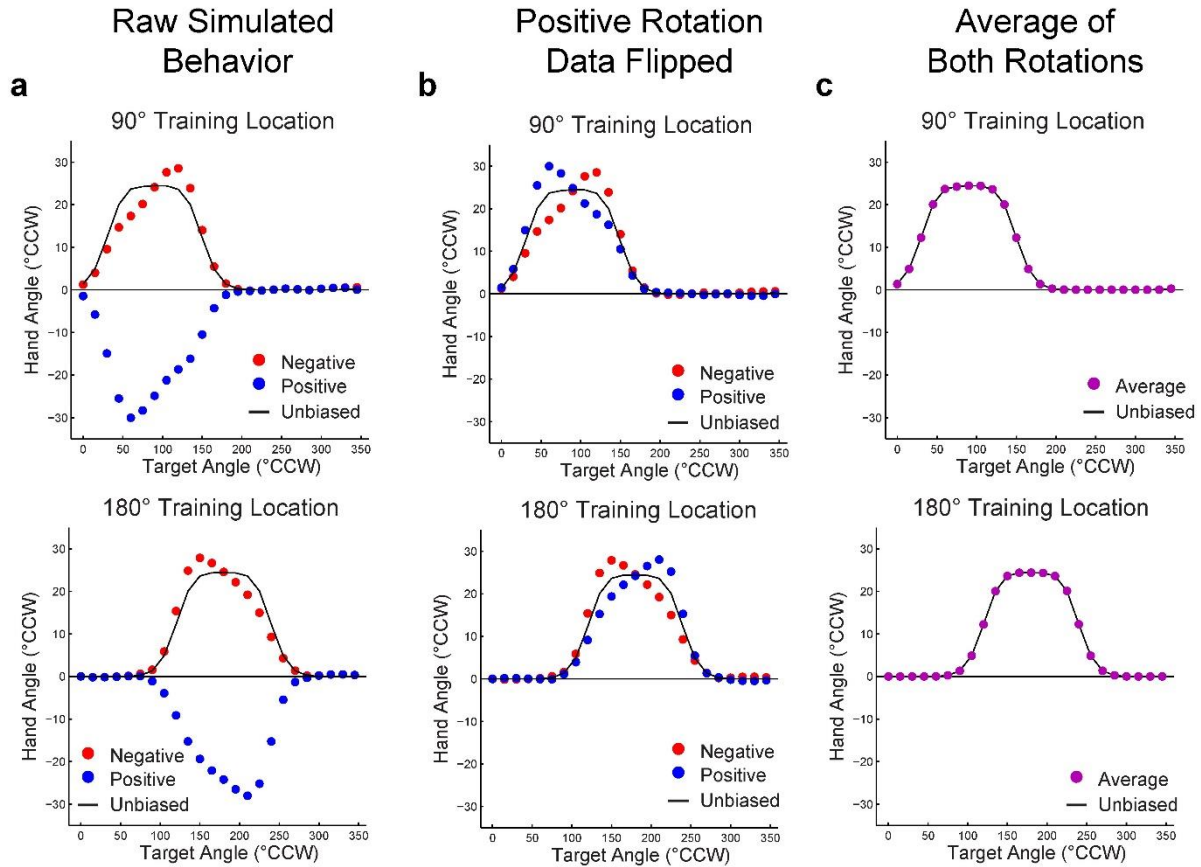


Figure 3.4 – Simulated effects of target location on generalization functions. a) Raw feedback-baseline subtracted generalization functions for positive and negative rotations, at two different targets. b) Generalization functions after flipping the sign of the hand angles for one group. c) Average function after the data from both groups are combined. Note that this achieves the same result as an unbiased model.

Simulation Set 3

We further explored target location effects by simulating generalization from training at one of the four cardinal directions, which are the most common training locations used to study generalization of visuomotor learning (Ghahramani et al., 1996; Krakauer et al., 1999; Brayanov & Smith, 2012).

This simulation had an initial no feedback baseline (240 trials), a feedback baseline (120 trials), a training/rotation block (120 trials), and finally a no feedback washout (480 trials). As shown above (Fig. 3.3a), the bias at each of the cardinal locations is of similar magnitude and sign, but the bias at neighboring locations is not the same. The 180° location has a negatively sloped bias for adjacent targets, while the 90° target has a positive slope. The simulation was conducted to ask how these differences could affect measures of generalization.

We simulated training at all the cardinal locations with both a positive and negative 30° visuomotor rotation, although for brevity we only show simulations that were trained at the 90° and 180° targets. Generalization was assessed by looking at the first reach to each target after the

training block, after subtracting the feedback baseline (as is commonly done in learning studies). At first glance (Fig. 3.4a), the learning appears similar for both negative and positive 30° rotations. However, when the hand angles are sign flipped for comparison (Fig. 3.4b), the generalization functions are clearly shifted in different directions. Furthermore, the direction of the shift is reversed when generalization from learning at the 90° target is compared with that from the 180° target.

This simulation demonstrates that the training target location and sign of the perturbation interact to form a skew in the generalization function, which could be misinterpreted if it is not controlled with counter-balancing. For instance, a skewed generalization function is consistent with use-dependent learning of a hand space attractor (Diedrichsen et al., 2010; Huang et al., 2011; Fernandez et al., 2012), and plan-based generalization (McDougle et al., 2017), but in our case the skew is purely a result of intrinsic biases that are not eliminated by baseline subtraction. Fortunately, this skew can be averaged out if the perturbation sign is counter-balanced and the data is combined by flipping the sign of one group's data and taking the mean of both (Fig. 3.4c). We strongly recommend counter-balanced perturbations as a standard practice in generalization experiments, and suggest using different training locations for each participant. However, counter-balancing will also cancel out true use-dependent learning. If this process is of interest, it would be better to counter-balance training target locations while holding the perturbation constant.

Simulation Set 4

The last set of simulations showed that measures of generalization are systematically skewed by intrinsic biases, despite subtraction of the feedback baseline. Here we further explore how baseline subtraction affects generalization, and assess whether subtraction of the initial, no feedback baseline might have less of an effect on the generalization function. In assessing these factors, we simulated two common generalization task designs, blocked and trial-by-trial (Ghahramani et al., 1996; Donchin et al., 2003). The blocked design follows the same trial structure as the generalization simulation in Simulation Set 2, a no feedback baseline (240 trials), feedback baseline (120 trials), a training/rotation block (120 trials), no feedback washout (480 trials). The trial-by-trial design has the same number of total trials, but is otherwise different: no feedback baseline (120 trials), feedback baseline (120 trials), trial-by-trial generalization block (240 trials, 7 training/23 generalization target reaches per 30 trials), and a no feedback washout (480 trials). For both simulations, generalization was assessed on the first reach to each target in the washout block. For all simulations, the training target was always at the 180° target location.

In a blocked design, baseline performance is measured at all target locations, followed by exposure to a perturbation while reaching to a single target location. After the perturbation training block is over, generalization is then assessed by again reaching to all the targets (Fig. 3.5a). This is the same generalization design as used in Simulation Set 2. Trial-by-trial designs, by contrast, interleave probe trials to non-training targets (without feedback) among perturbation trials (with feedback) at a training target. Unlike Simulation 2 where we subtracted the feedback

baseline in measuring generalization, here we subtracted either the feedback baseline or the no feedback baseline in order to show the effect both could have.

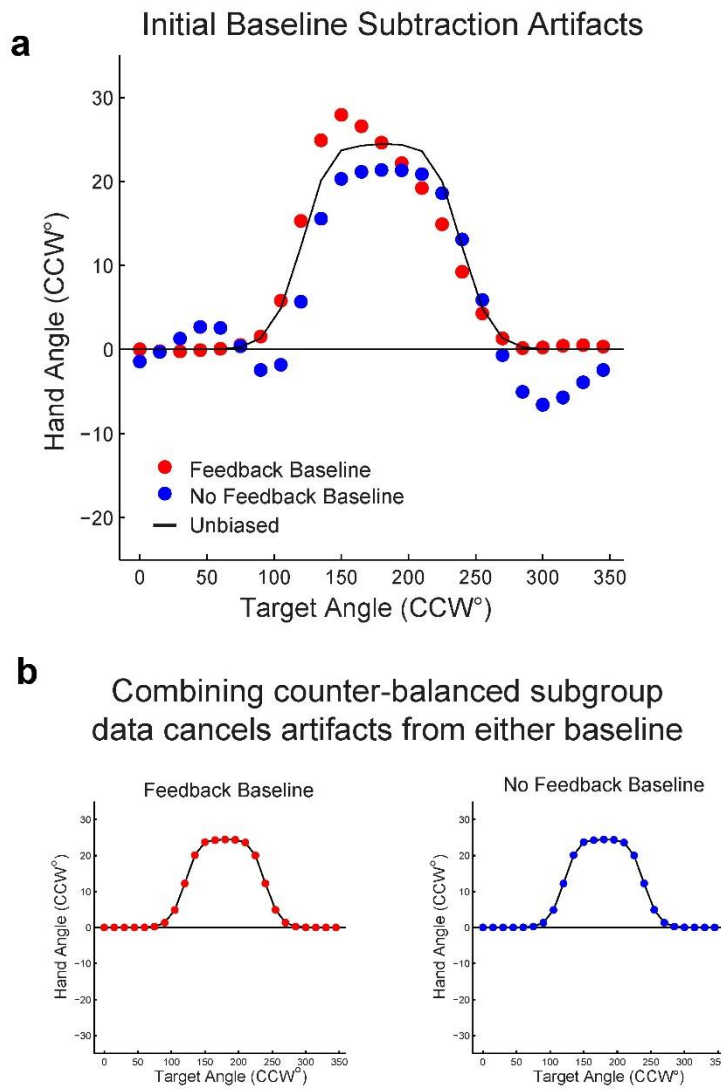


Figure 3.5 – Simulated effects of subtracting the feedback and no-feedback baselines. a) Artifacts are caused by the subtraction of either baseline. b) Combining data from counter-balanced groups will cancel out the artifacts from either baseline.

The aftereffect did not shift when the No Feedback baseline was used for the subtraction. Instead, it introduced artifacts for locations far away from the training location (Fig. 3.5b). These artifacts, like the skew introduced by feedback baseline subtraction, were effectively canceled if we combined data from simulations with clockwise and counterclockwise rotations (Fig. 3.5c, d). Subtraction of the No Feedback baseline, followed by flipping the sign of one group and averaging the both groups together accurately calculates changes from learning, as does the same operation calculated with subtraction of the feedback baseline. More importantly, this simulation reveals limitations with both types of baselines in terms of estimating generalization, as both induce artifacts that can only be eliminated via counter-balancing.

Given that neither baseline subtraction effectively estimates learning from the perturbation, we simulated learning and generalization when feedback was never provided at the generalization targets at any point in the experiment. We show a side-by-side comparison with an experiment where feedback was provided at these locations during the baseline. With this method, the trial-by-trial design (Fig. 3.6b) provided an accurate measure of generalization. However, because of the small counter-clockwise mean offset of the bias near the training location, there was a small difference in the magnitude of learning for both clockwise and counterclockwise perturbations, relative to an unbiased model (Fig. 3.6d). The magnitude of learning was slightly larger for a biased learner exposed to a counter-clockwise rotation, and slightly smaller for the clockwise rotation. The mean of these combined conditions, however, formed an accurate estimate of learning and generalization.

The same pattern was observed in a blocked generalization trial design, though this design also resulted in a larger aftereffect at all target locations (Fig. 3.7). This difference in the magnitude of learning between blocked and trial-by-trial simulations comes about because of forgetting. In our model, there is some forgetting at all target locations on every reach. The forgetting rate is maximal at the reach target, and generalizes in the same manner as learning (i.e., Gaussian function centered at target). This will result in much more retention of learning for blocked designs relative to trial-by-trial designs. This is because blocked designs do not include no feedback reaches, so they do not generalize any forgetting. Any model that does not forget equally at all locations on each reach should show similar behavior. The simulations presented here indicate that baseline performance in generalization tasks is best measured without visual feedback. We again recommend counter-balancing the perturbations, but this is only necessary if the absolute magnitude of adaptation is of interest.

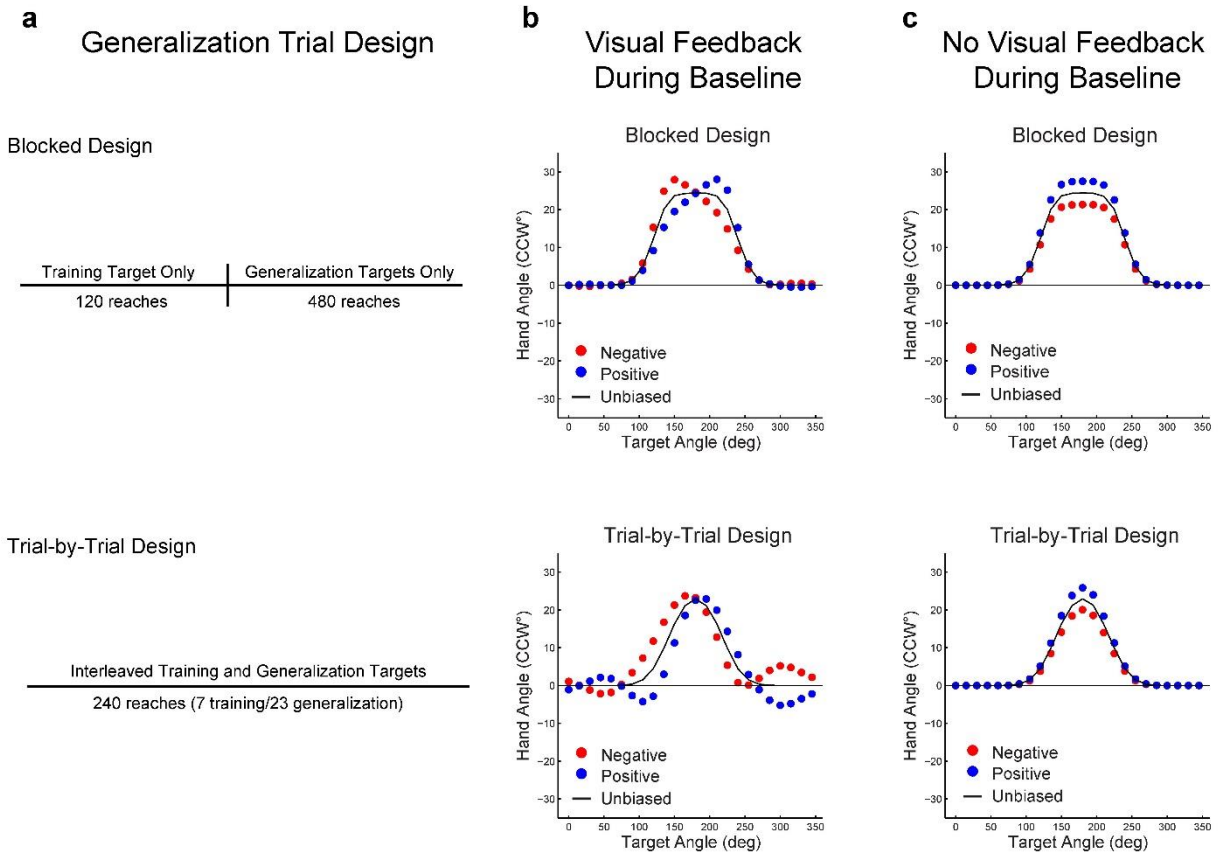


Figure 3.6 – Interaction of bias with generalization task trial designs, and the effect of visual feedback during baseline. a) Trial design schematics for blocked designs (top) and trial-by-trial (bottom). b) The subtraction of a feedback baseline causes irregular artifacts in both trial designs, although these can be canceled with counter-balancing. c) If visual feedback is never provided, it will not cause irregular artifacts, however the magnitude of adaptation will be misestimated at all target locations.

Discussion

Here we explored intrinsic biases in the context of reaching to visual targets with and without visual feedback. We first showed that these biases are similar across participants, systematically varying with the target angle, in line with previous work (Ghilardi et al., 1995; Vindras et al., 1997). Furthermore, following a feedback baseline period, participants still showed a substantial residual error at some target locations. Using model simulations, we showed that these local errors will persist with no matter the number of trials because they are a consequence of generalization from different targets interacting with rapid changes in the local slope of the bias. We explored how such a bias could affect visuomotor adaptation experiments, given its systematic nature across participants, its persistence in the face of veridical visual feedback, and the tendency for participants to decay to back to their intrinsic bias when feedback is removed (Smeets et al., 2006).

In a second set of simulations, we assessed how the bias we measured would affect the decay of learning from a perturbation in the absence of visual feedback. Here we saw that target-specific biases could easily be misinterpreted as differences in retention, if the researcher was unaware of the bias. In a third set of simulations, we showed that target-specific biases can differentially skew generalization functions, which could easily be misinterpreted as evidence for evidence of use-dependent learning (Diedrichsen et al., 2010; Huang et al., 2011; Fernandes et al., 2012), or plan-based generalization (McDougle et al., 2017). In our final set of simulations, we showed that this skew results from subtraction of the feedback baseline, and that an un-skewed generalization function can be calculated by subtracting the initial baseline that did not have feedback. However, this alternative baseline subtraction misestimates the magnitude of adaptation that occurred. Taken together, we believe our experiments and simulations demonstrate that motor control researchers should be wary of these biases because they are omnipresent and, in the certain contexts, can easily be misinterpreted as evidence for distinct learning mechanisms.

Intrinsic Biases Are Never Fully Eliminated

Visuomotor adaptation will quickly reduce intrinsic reaching biases, but as we showed in Experiment 1 and across our simulations, there is a systematic residual error that cannot be eliminated if the bias has differences that are sufficiently large between adjacent targets. This occurs because learning at nearby targets will be in opposition if the targets have errors of opposite sign. Each time one of the targets is reached to, it will generalize learning to adjacent targets, increasing their error. This results in a sort of tug-of-war, or dynamic equilibrium, where generalization from adjacent targets does not allow error to be reduced fully at any of the targets. The simulation shown in Figure 3.2e showed that this effect does not require errors to have initial errors of different sign, instead it depends on the relative change in the size of the bias over nearby targets.

Measuring Intrinsic Biases

Our simulations provided a concrete example of the interpretational complications that intrinsic biases create, but also allowed us to test strategies for ameliorating the impact of these biases on learning metrics. Our first recommendation is to simply measure intrinsic biases at the very beginning of the task by having participants reach without any visual feedback (as we did in Experiment 1). Importantly, the bias should be measured before any reaches are made with visual feedback. This is because any reach made with visual feedback will generate a visual error and induce visuomotor adaptation, which will drive behavior away from the true bias. This means that a faithful measurement of the intrinsic bias is unlikely to be acquired, because it takes many reaches without feedback for the visuomotor adaptation memory to decay (Smeets, 2006).

Unfortunately, it is not customary in computer-based reaching experiments to begin with no feedback reaches. Instead, the task begins with full online visual feedback that ‘familiarizes’ the participants with the experimental setup, and if researchers measure baseline movements made without feedback, these are performed after the familiarization phase (Krakauer et al., 1999; Brayanov & Smith, 2012; Fernandes et al., 2012; Taylor & Ivry, 2013). We surmise that studies using such a task design underestimate the magnitude of the intrinsic biases, and it is possible that this underestimate meaningfully affected how their results were interpreted. Accurate measurement of the bias is the first step in deciding how to best deal with it.

Accounting for the Bias

Whether intrinsic biases pose a problem for the interpretation of learning experiments depends on several things, such as the size and location of the learning effect, or if perturbations are counter-balanced across participants. It is difficult to formulate a one-size-fits-all analysis plan because of the wide spectrum of factors that a researcher may assess in their study. The most important thing is to properly measure the bias, and then reason as to how it might influence the results. We simulated the two scenarios that we thought would be the most commonly-encountered situations where the bias could affect things. These are the removal of feedback to assess decay of a motor memory, and assessing generalization by comparing learning across a range of target angles.

The removal of feedback presents a potential problem because adaptation will not decay to zero visual error in this scenario, but rather to the intrinsic bias location. It is not common practice to measure the intrinsic bias in these types of experiment (Galea et al., 2010; Kitago et al., 2013; Leukel et al., 2015; Galea et al., 2015), making it difficult to judge whether target-specific biases affected the results. However, the between-group differences of decay in these studies are less than 10°, making it possible that a substantial proportion of the measured effect could be attributed to re-emergence of an intrinsic bias. This highlights the need to institute an initial no-feedback baseline, so that effects from learning and effects from biases can be disambiguated.

The third and fourth sets of simulations we performed focused on generalization, and how the training target location or choice of baseline to subtract will affect results. We found that counter-balancing perturbation sign is critical to average out target specific biases. The downside of this technique is that it also will average out workspace-related effects that may be of interest, such as use-dependent learning (Diedrichsen et al., 2010; Huang et al., 2011; Fernandes et al., 2012), or plan-based generalization (McDougle et al., 2017). If these are of particular interest, we recommend the use of several training target locations instead of counter-balancing the sign of the perturbation. This will average out target-specific biases while leaving hand space effects intact. Both of these methods can be employed in the same experiment, with subgroups analyzed separately depending on the factor of interest.

Conclusion

Here we have shown, through experiment and simulation, that intrinsic biases systematically affect behavior in reaching tasks. Our intent was to demonstrate that these biases should not be ignored. Instead, they should be measured at the very beginning of the experiment, so that their possible effects on behavior can be anticipated relative to the experimental context. If researchers take these steps, it will help reduce errors of interpretation in visuomotor adaptation experiments.