

**From the holdings of the Department of Special Collections , Princeton University Library,
Princeton, NJ**

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyright law. We do not charge any permission or usage fees for the publication of images of material in our collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff. For more information regarding copyright, credit, and citation please review the information on the Copyright, Credit and Citation Guidelines page of our website.

<https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines>. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. <https://library.princeton.edu/special-collections/ask-us>

Doctoral Dissertations, 2011-2026

Psychology

2021

Implicit recalibration and strategic compensation in learned and de novo motor skills

Wilterson, Sarah A.

Princeton, NJ : Princeton University

<http://arks.princeton.edu/ark:/88435/dsp01hq37vr66c>

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. <https://www.copyright.gov/title17/title17.pdf>

IMPLICIT RECALIBRATION AND STRATEGIC COMPENSATION
IN LEARNED AND *DE NOVO* MOTOR SKILLS

Sarah A. Wilterson

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE BY
THE DEPARTMENT OF PSYCHOLOGY

Advisers: Jordan Taylor, Nathaniel Daw

May 2021

© Copyright by Sarah A. Wilterson, 2021. All rights reserved.

Abstract

The mechanisms underlying motor skill learning are at once well established and fiercely debated. Implicit recalibration, the process by which a motor skill is adjusted to account for unexpected feedback, has received the most attention. In fact, the processes underlying implicit recalibration was considered to account for most of skill learning. However, recent advances have shown a second process which plays a strong role in adapting an established skill, strategic compensation. It seems obvious that a strategy would be developed for dealing with unexpected perturbations to skill. However, when strategies are used, what influences their shape, and how they interact implicit recalibration is relatively unknown. We explore these adaptive processes in Chapters II – IV.

While adaptation of a learned skill has preoccupied much of the motor learning literature, and indeed much of this dissertation, how we learn new is skills equally important. Termed *de novo* learning, the acquisition of new action-outcome mappings is essential for common skills, such as typing on a keyboard, or playing a musical instrument. In Chapter V, we describe a new paradigm for investigating *de novo* learning, and provide evidence that reinforcement learning is a useful model for understanding skill learning.

Acknowledgements

I would like to acknowledge and thank the following people, who made this work possible. Jordan Taylor, my advisor and mentor, for providing invigorating conversations, scientific advice, and ceaseless support. Michael Graziano and Casey Lew-Williams, who stuck with me as members of my general exam, dissertation proposal, and reading committees. Their genuine interest in my work gave me confidence that I was doing something right. Stacey Sinclair and Nathaniel Daw, my oral exam committee, who modeled academic excellence throughout graduate school.

My family, for their ceaseless support, and for letting me ramble on about science.

Andrew Wilterson, my husband and colleague.

Thank you for being my rock, for your understanding and support, for being there throughout.

Table of Contents

Abstract	iii
Acknowledgements.....	iv
Chapter 1: Introduction	1
Chapter 2: Sensitivity and Limitations of Implicit Recalibration	6
Chapter 3: Longitudinal Tests of Implicit and Explicit Learning in Adaptation Tasks.....	30
Chapter 4: Cognitive Strategies in Mirror Reversal Tasks	64
Chapter 5: Model-Free and Model-Based Learning in <i>de Novo</i> Motor Skills.....	87
Chapter 6: Conclusion.....	130
References	136
Prior Publication	148

Chapter I

Introduction

The very first time one learns to use a keyboard, there is no necessary *a priori* association between the keys pressed and the letter produced on the screen. This skill must be learned from the ground up. At first, one types the word “tiger” by finding and pressing “t”, then finding and pressing “i”, and so on. Eventually, mastery of the skill allows one to think “I would like to type tiger”, and the fingers do the rest. Learning this novel (i.e. *de novo*) skill requires a fundamental shift from laboriously considering and performing individual actions, to representing only the eventual goal and automatically performing the actions necessary to achieve that goal.

It takes considerable time to be able to quickly and reliably produce the action sequences necessary for keyboarding - on the order of months to years. However, once learned, adapting one's finger movements to account for small changes between different keyboards is remarkably fast - on the order of minutes to hours. Adaptation to movement errors, as would occur when a finger misses the key on a new keyboard, is critical to the motor learning process (Cheng and Sabes, 2006; Jordan and Rumelhart, 1992; Miall and Wolpert, 1996; Pouget and Snyder, 2000; Thoroughman and Shadmehr, 2000). Characterization of the driving forces behind adaptation reveal fundamental principles and constraints of motor learning (Fine and Thoroughman, 2006, 2007; Marko et al., 2012; Scheidt et al., 2001; Semrau, Daitch, and Thoroughman, 2012; Wei and Körding, 2009).

Adaptation processes have been characterized using system identification techniques borrowed from engineering (Cheng and Sabes, 2006; Thoroughman and Shadmehr, 2000). These techniques generally consist of introducing a small perturbation to the system and observing the behavioral response (Chapter II; Fine and Thoroughman, 2006, 2007; Scheidt et al., 2001; Wei

and Körding, 2009). These perturbations simulate a small error in the motor system, such as when the typing finger misses the e-key by a moving too far to the right. The system recalibrates by adjusting future movements, to the e-key, toward the left.

This fine-grained recalibration is achieved through an implicit process. In fact, a key trait of adaptive learning is that it does not require the learner to know about the error which catalyzes recalibration. The learner being aware of the error also has little effect on the process. Adaptation is highly stereotyped and proceeds in the same way even when the recalibration routine increases the size of the error and the learner knows that their performance is degrading.

The stereotyped implicit recalibration process is also sensitive to only a small range of perturbation sizes (Chapter II). When errors are small, recalibration is proportional ($< \sim 8^\circ$ of rotational perturbation). However, once the error exceeds this threshold, the response saturates and all large perturbations are treated in the same way (Kim et al., 2018; Marko et al., 2012; Morehead et al., 2017; Wei and Körding, 2009). The implicit recalibration process is very slow and imprecise when responding to large errors (Fine and Thoroughman, 2006; Butcher and Taylor, 2018). To entirely overstretch our example, missing the e-key by 10 mm would result in the same amount of implicit recalibration as missing it by 1mm. Even with prolonged training, implicit recalibration cannot produce enough learning to compensate for a large or complicated perturbation (Chapter III; Morehead and Smith, 2017).

Larger errors tend to induce the learner to employ a cognitive strategy. Unlike implicit adaptation, strategies are flexible, responding to the demands of unusual environments with a variety of tactics (Chapters III and IV). Cognitive strategies are often defined as the portion of learning that can be reported by the learner. They may take the form of a rule: “Press each key a bit harder than usual because this keyboard has sticky keys.” One may also have very specific

strategies for particular sub-goals: “Use the middle finger of the left hand to press the e-key.” The knowledge of the learner is generally ignored. That is, it is not necessary for the learner to know why they are employing a cognitive strategy for it to be labeled as such. It does not matter if the learner presses each key harder because they think the keys are sticky from spilled juice, or if they think the keys require greater actuation force, or if they think their muscles have gotten weaker.

Although cognitive strategies are often seen under experimental conditions, our experience with everyday challenges argues that we eventually account for large perturbations without continuous strategic compensation. In the late 19th century, George Stratton (1896, 1897) noted that radically altered visual input could be overcome and experienced as normal. These experiments were conducted with prism goggles, which flipped the visual world upside down and transposed it left to right. Stratton (1896, 1897) and Kohler (1941, 1951a, 1951b; review in Sachse et al., 2017), documented their adjustment to the world of prism goggles, including an early stage of physical mishaps and visual oddities (e.g. dodging the wrong direction to avoid an object, holding cups upside-down to be filled, stepping over a street sign that appears to be on the walkway). Kohler (1951b) reported that after five days of continuous prism goggle use, the visual world suddenly became upright once more. He noted that his initial strategies were no longer necessary – the world was once again traversed simply by deciding what he wanted to do and doing it.

So then, implicit adaptation works to maintain established skills by adjusting for small errors but does not have the capacity to learn a new skill. New skills *can* be learned and do not forevermore require the constant application of a strategy. Perhaps then, learning a *de novo* skill is fundamentally different from adapting an old one. Few studies have sought to determine how a skill is learned *de novo*. So, we have limited understanding of how *de novo* skills, and the ability

to adapt those skills, initially develop. To address this question, we've taken inspiration from the reinforcement learning literature (Chapter V).

There is some experimental evidence that model-free and model-based learning underlie *de novo* skill learning (Fermin et al, 2010). Indeed, the model-based stage of learning is often characterized as slow and effortful, requiring a computationally intensive tree search of action choices (Daw, Niv, and Dayan, 2005). This parallels the deployment of cognitive strategies. Both model-based learning and use of cognitive strategies are effortful and attention-demanding processes. (Adams, 1982; Dickinson, 1985; Killcross and Coutureau, 2003). On the other hand, model-free learning results in automatic retrieval of an output for each input. This results in characteristically rigid performance that is quickly and automatically accessible without need for effortful processing or attention (Daw, Niv, and Dayan, 2005). This as a close parallel to a learned skill, which can be completed rapidly without effortful control.

In this dissertation, I begin with experiments of the implicit recalibration and cognitive strategies associated with adaptive learning. I then turn to the processes underlying *de novo* learning of a new skill. In Chapter II, we investigate the sensitivity of implicit recalibration and cognitive strategies to varying perturbation sizes in an adaptation task. We provide evidence that the combination of these processes in trial-by-trial learning is dependent, in part, on the perturbation size. Implicit adaptation scaled relative to the visual error, but only for small perturbations – replicating prior work. In contrast, explicit re-aiming scaled linearly for all tested perturbation sizes. Additionally, people stopped using strategies when the errors were so inconsistent that learning was impossible, but implicit recalibration continued on a trial-by-trial basis. This comports with the basic idea that implicit adaptation is obligatory and inflexible, while explicit re-aiming is volitional and flexible

In Chapter III, we used a multi-day experiment to expose how implicit recalibration contributes to adaptation over the long term. The recalibration process plateaued at a value far short of complete learning. Along with the standard adaptation task, we used a task which creates an environment where the canonical recalibration response is counterproductive to learning. In this task, implicit recalibration drives performance away from the goal. We saw that this inappropriate recalibration could be suppressed over many days of training, providing evidence that implicit recalibration can be available.

In Chapter IV, we return to the strategies that are employed in adaptation tasks when implicit recalibration is insufficient for optimal performance. Given an angular perturbation greater than 8° or so, the majority of learning is credited to cognitive strategies. However, the nature and mechanism of these strategies has received little attention. We present evidence of two fundamentally different types of strategy, one parametric in nature and the other resembling a cached S-R mapping.

In Chapter V, we introduce a new paradigm for studying *de novo* skill learning. We propose that *de novo* skill learning is the product of an interaction between two systems, analogous to model-based and model-free learning. A process like model-based planning, which may be in the form of a cognitive strategy, provides rapid improvement in performance. Meanwhile a model-free-like learning process works under the surface to capture the action-outcome mapping needed for skilled performance. We also show that we can measure adaptation of the skill learned in this paradigm, providing a link between the adaptation and *de novo* learning literatures.

Chapter II

Sensitivity and Limitations of Implicit Recalibration

Introduction

Compensating for movement errors is critical to the motor learning process (Cheng and Sabes, 2006; Jordan and Rumelhart, 1992; Miall and Wolpert, 1996; Pouget and Snyder, 2000; Thoroughman and Shadmehr, 2000). Thus, characterizing the sensitivity of the behavioral response to these errors should reveal fundamental principles and constraints of the motor system (Fine and Thoroughman, 2006, 2007; Marko et al., 2012; Scheidt et al., 2001; Semrau, Daitch, and Thoroughman, 2012; Wei and Körding, 2009). For over twenty years, the motor control field has sought to characterize this sensitivity function using system identification techniques borrowed from engineering (Cheng and Sabes, 2006; Thoroughman and Shadmehr, 2000). These techniques generally consist of imposing a transient and often varying perturbation on the system to observe the behavioral response (Fine and Thoroughman, 2006, 2007; Scheidt et al., 2001; Wei and Körding, 2009).

Despite employing this theory-driven and elegant approach, the observed sensitivity functions have been highly variable and appear to depend on a number of experimental factors. In a seminal study, Scheidt and colleagues (2001) found that the motor system adapted to transient and random force perturbations on a trial-by-trial basis in a force-field-adaptation task. This adaptive response appears to be sensitive to the direction of the perturbation, but insensitive to both the timing and magnitude (i.e., strength) of the perturbation (Fine and Thoroughman, 2006). However, if the perturbations are drawn from a non-zero mean distribution (Fine and Thoroughman, 2007), occur frequently (Fine and Thoroughman, 2007), or are applied in a consistent fashion (Castro et al., 2014), then the adaptive response becomes more sensitive.

Additionally, experience with an error in a more stable environment has been shown to increase response sensitivity to reoccurrences of that error (Herzfeld et al., 2014).

Similar results have been found in studies of visuomotor rotations: Implicit adaptation appears to be highly sensitive to the direction of the rotation, but less sensitive to its magnitude (Butcher and Taylor, 2018). In fact, the time course of implicit adaptation begins to saturate in response to rotations greater than $\sim 6^\circ$ (Kim et al., 2018; Marko et al., 2012; Morehead et al., 2017; Wei and Körding, 2009), despite prolonged periods of training with both non-zero mean and consistent rotational perturbations (Morehead and Smith, 2017). This was made clear by employing a task-irrelevant error clamp task, in which cursor feedback of the supposed hand's position is consistently offset from the target path by a fixed angular value regardless of the hand's true location (Morehead et al., 2017). Despite the task-irrelevance of this feedback, robust adaptation is observed even when movement angles are uncorrelated with the angle of the rotational perturbation.

At first glance, the lack of sensitivity of this adaptive response is puzzling given the variety of motor behaviors and skills we can employ. However, when subjects have control of the angular position of the cursor – thus, making it task-relevant – sensitivity is restored (Morehead et al. 2017). This suggests that additional learning processes, such as explicit re-aiming, may play a role in making the appropriate behavioral response given the demands of the task (Abeeel and Bock, 2001; Bock et al 2003; Bond and Taylor, 2015; Heuer and Hegele, 2008; Taylor et al., 2014; Werner and Bock, 2010). Indeed, the varying sensitivity functions measured in the aforementioned studies vacillated around conditions when the subjects could potentially affect the outcome in the task. For example, Wei and Körding (2009) used a task in which visual error changed randomly on every trial. In this situation, aiming anywhere except straight toward the target would actually

be counter-productive. In contrast, the adaptive response become more proportional when force perturbations have a non-zero mean (Fine and Thoroughman, 2007). Similarly, when force perturbations were made to be more consistent, by increasing the number of trials in a row for which a particular perturbation is present, the learning rate increased (Castro et al., 2014).

While we note that there are other substantial differences between these aforementioned studies, such as rotational versus force perturbations, we hypothesize that controllability over performance is a critical factor, which engages explicit re-aiming to restore performance. Because we cannot directly or, at least reliably, manipulate the subjective experience of control, we operationalize controllability as environmental consistency. We reason that a more consistent environment gives greater opportunity for the deployment of effective strategy for controlling errors, whereas it is not possible to control error in an environment that has constantly changing movement-feedback contingencies. Specifically, we propose that explicit re-aiming may only play a role when the perturbations are consistent and the visual errors are quite large (and/or implicit adaptation is saturated). Further, given the number of studies showing that implicit adaptation appears to be highly stereotyped (Bond and Taylor, 2015; Morehead et al., 2017; Wei and Körding, 2010), we hypothesize that implicit adaptation will not respond to environmental consistency. Here, in two experiments, we set out to test these ideas by examining the sensitivity functions of both explicit re-aiming and implicit adaptation to a range of perturbations that varied in their degree of consistency.

Materials and Methods

Participants

Eighty human subjects [57 women, mean age 20.4 (SD 3.24) yrs.] were recruited for Experiment 2.1, and 26 subjects [18 women, mean age 19.69 (SD 1.35) yrs.] were recruited for Experiment 2.2. The sample size for Experiment 2.1 was guided by the typical convention of 10-20 subjects per condition (4 conditions) for visuomotor adaptation tasks. The sample size for Experiment 2.2 was determined by an *a priori* power analysis using the slope of explicit re-aiming sensitivity function from Experiment 2.1 between the Consistent-2 (mean slope = -0.374) and Consistent-7 conditions (mean slope = -0.152). Robust power (80%) was obtained using the calculated effect size of 0.88 for a two tailed t-test between two independent means. This analysis determined that we needed 13 subjects per condition. All subjects were drawn from the research participation pool maintained by the Department of Psychology at Princeton University and received either course credit or monetary compensation for participating. Subjects were right handed, as verified by the Edinburgh handedness Inventory (Oldfield, 1971), and self-reported having normal or corrected-to-normal vision. The experimental protocol was approved by the Princeton University Institutional Review Board and all subjects provided written, informed consent.

Apparatus

Subjects performed horizontal movements in a center-out reaching task. These movements were recorded with a digitizing pen and Wacom tablet, with the tablet sampling movement trajectories at 60 Hz. All stimuli were displayed by a 17-in., Planar touch sensitive monitor with a refresh rate of 60 Hz and computed by a Dell OptiPlex 7040 machine running Windows 7. The

touch sensitive monitor allowed subjects to report their intended movement by simply tapping on the screen (Figure 2.1; Bond and Taylor, 2017). Visual feedback of the hand was obstructed by the monitor, which was mounted 25 cm above the tablet. A small circular cursor (0.15 cm radius) provided feedback information to subjects. The game was controlled by custom software coded in Matlab, using Psychtoolbox extensions (Brainard, 1997, Kleiner, Brainard, and Pelli, 2007).



Figure 2.1. Experimental setup for both experiments. The subject holds a digitizing pen in her right hand, which is covered by the touch screen. The left hand is used to tap the touch screen on the blue ring to report aiming location before each trial. After the subject taps on the screen, the blue aiming ring would disappear and the orange target would turn green.

Procedure

Subjects began each trial with their right hand at the center of the visual workspace. After holding this position for 300ms, a circular orange target (0.25 cm radius) appeared 7 cm from the starting position. Appearing along with the target, an “aiming” ring consisting of a blue circle that was centered on the starting location and had a radius of 7 cm (Figure 2.1). Subjects were instructed to report their intended aiming location by tapping the aiming ring on the surface of the touch screen with their left hand. Once a touch was recorded, the target turned from orange to green, the aiming ring disappeared, and subjects were able to reach with their right hand. If a subject

attempted to reach with their right hand prior to tapping the touchscreen with their left hand, the message “Remember to report aim” was displayed and the trial was restarted.

Subjects were instructed to make a fast, straight “shooting” movement through the target with their right hand. They were informed that it was not necessary to stop on the target, but that they should be careful to move far enough to pass completely through the target. Subjects were provided with continuous, online feedback of the cursor (0.15 cm radius) throughout the movement. Once the subject’s hand passed 7 cm, endpoint feedback was displayed for 1 s. If the final position of the cursor overlapped with the target (<1 degree of angular deviation), subjects heard a pleasant “ding” sound; otherwise, they heard an unpleasant “buzz.” If the time from leaving the start position to reaching out 7 cm exceeded 800 ms, the feedback “too slow” was given (this occurred on approximately 1% of trials, and these trials were excluded from further analysis. The average movement time was well below 250 ms for all conditions). Following feedback presentation, subjects were guided back to the start position by a white ring that was centered on the starting location and whose radius represented the distance between subjects’ hand position and the starting location. Veridical feedback of the cursor was restored once the hand was within 1 cm of the starting position.

We pseudorandomized the target locations across the workspace and across subjects so that any potential visual or biomechanical biases would average out (Ghilardi and Ghez, 1995; Jiang, Yuan, and Wei, 2018). On each trial, the target could appear in one of eight locations, spaced 45° apart to cover the full circular space defined by the aiming ring. The exact angular configuration of the target locations differed across subjects. There were five sets of target configurations, where the target set could be rotationally shifted 0° , 9° , 18° , 24° , or 36° . Each subject was exposed to only one configuration of target locations.

To assay the sensitivity function of implicit adaptation and explicit re-aiming, an angular rotation of $\pm 0^\circ$, 2° , 4° , 8° , and 16° was imposed on the cursor in Experiment 2.1. These rotation sizes were chosen to correspond to the set of lateral displacements used in Wei and Körding (2009). Furthermore, the rotations were counterbalanced such that the mean rotation size over the experiment was 0° . For experiment 2.2, a 32° rotation was exchanged for the 2° rotation, such that the rotational perturbations imposed during the task were $\pm 0^\circ$, 4° , 8° , 16° , and 32° .

A secondary goal of this experiment was to determine if the sensitivity function changed based on the consistency of the rotation (i.e., how frequently the rotations changed during training), which has been reported by prior studies (Castro et al., 2014; Fine and Thoroughman, 2007). To this end, in Experiment 2.1 subjects were equally divided into four groups: Consistent-1, Consistent-2, Consistent-3, and Consistent-7. In the Consistent-1 condition, the rotation changed on every trial, effectively making this an inconsistent condition; though the target location remained the same for 7 trials. In the Consistent-2, Consistent-3 and Consistent-7 condition, each “mini-block” consisted of 2, 3, or 7 trials, respectively, where the rotation changed after each mini-block. For these conditions, the target location also changed at the onset of each mini-block. Although changing the target location following every mini-block causes those conditions with greater error consistency to also have greater target consistency, we felt that it was important to ensure that the change in perturbation size occurred in the same way both within a condition (between mini-blocks) and between groups. The exception is the Consistent-1 condition, where we did not sync the target and perturbation changes because a pilot experiment showed that there is no learning when both the target and perturbation change on every trial – likely due to the generalization function of implicit adaptation. In all cases, visual perturbations were pseudo-randomly generated such that no rotation size was immediately repeated and each rotation size

occurred at each target location at least once. Experiment 2.2 consisted of only the Consistent-2 and Consistent-7 conditions, which had mini-block lengths of 2 and 7 trials, respectively.

The experiments proceeded by first providing veridical feedback for eight familiarization trials and then pseudo-randomly generated visual rotations of the cursor for 504 trials in the Consistent-1 and Consistent-7 conditions and 432 trials in the Consistent-2 and Consistent-3 condition. The discrepancy in trial length is the result of complete counter-balancing. The entire experiment took approximately one hour.

Note, subjects were informed that the mapping between their hand and the cursor may change during the experiment and that they should tap on the aiming ring where they intend to aim in order to hit the target. Subjects were not told the nature of the visual perturbation nor when the visual disturbance would be in effect during the experiment.

Data and statistical analyses

All data and statistical analyses were performed in Matlab. The digitizing tablet recorded the trajectory of the right hand and the touchscreen monitor recorded the terminal position of the location tapped with the left hand. These data were transformed to define heading hand angles and aiming angles during training as follows: The hand angle trajectories and aiming locations were transformed from Cartesian to polar coordinates and rotated to a common axis with the convention that the target was positioned at 0° (directly to the right). As our primary interest concerns only the feedforward portion of the reach, we focused on the initial heading of the hand angle by examining the average angle of the hand between 1 and 3 cm into movement. Aiming angle was defined as the angle between the target and the tapped location on the touchscreen. Implicit adaptation was calculated by subtracting the subjects' aiming angle from their hand angle (Bond

and Taylor, 2017; Taylor et al., 2014). For all measures, positive angles indicate a counterclockwise deviation from the target. As we have only two measured values, with implicit adaptation being computed as the subtraction of aiming angle from hand angle, we performed statistical analyses only on aiming angles and implicit adaptation angles – our primary variables of interest – and we report hand angles only for completeness.

In the following statistical analyses, unless otherwise specified, the average of the second trial of every mini-block was used. This allowed us to control for the confounding additive effects inherent in having different length mini-blocks for each condition. For predictive purposes, the rotation size is considered to be the rotation size of the mini-block. Thus, subjects experienced the rotation on the first trial of the mini-block and we evaluated their response on the next (second) trial. For the Consistent-1 condition, the rotation size is considered to be the rotation experienced on the previous trial ($n-1$, where n is the trial being evaluated). While counterbalancing the perturbations resulted in a different number of trials between groups (432 trials in Consistent-2 and 3, 504 trials in Consistent-1 and 7), the total number of trials should not affect trial-by-trial adaptation since the study is designed to prevent the accumulation of learning. Additionally, equating the number of trials analyzed across the conditions did not significantly change the result.

To quantify the sensitivity function for each consistency condition, we fit separate linear functions to each subject's aiming angles and implicit adaptation angles with respect to the rotation size. A significant slope indicates that the subject changed their behavior in response to the error. Differences in slopes between consistency conditions were evaluated by submitting the slopes to a one-way ANOVA; post-hoc t-tests were conducted when appropriate and corrected using the Bonferroni method. We also sought to determine if the overall slope of the sensitivity function was similar across rotation sizes. Previous studies have reported that the sensitivity of the response

scales with the rotation size before reaching a saturation point at higher rotation magnitudes (Morehead et al., 2017; Wei and Körding 2009). To assess this possibility, we adopted the method of Wei and Körding (2009) and fit a second linear function to the range from -4 to 4° , which corresponds to lateral displacement between -2 and 2 cm in their study. The slopes of this second function and the overall function were compared with a pairwise t-test. In Experiment 2.2, an additional analysis is conducted on the range from -8 to 8° . This range was chosen *post hoc* to preserve the ratio of 1:4 between the smaller test range and the full range experienced in the experiment.

The intercept (or offset) of these functions is of less interest, but significant intercepts could be viewed as an accumulation of learning throughout training or the development of a more general bias during training (Ghilardi et al., 1995).

Results

Experiment 2.1

In this experiment, we sought to assess the sensitivity of implicit adaptation and explicit re-aiming as a function of the magnitude and consistency of rotational perturbations, which ranged from 0 - 16° within a subject and changed every 1, 2, 3, or 7 trials across subjects. Subjects attempted to counteract these perturbations for all consistency conditions, as can be seen by the change in the angle of the hand in response to the imposed error (Figure 2.2a). For the majority of conditions, these changes in hand angle are the result of the combined output of implicit adaptation (Figure 2.2b) and explicit re-aiming processes (Figure 2.2c). To quantify the sensitivity of these processes, we fit a linear function to each process for each subject over the imposed rotations. For all consistency conditions, we found that the slope of the sensitivity function was significant for

both implicit adaptation ($p < 0.01$) and explicit re-aiming ($p < 0.001$; see Table 2.1). The intercept of these linear fits was not different from zero, except in the case of Consistent-1 which had a small yet significant positive shift for both implicit adaptation (mean = 0.209° ; $t(19) = 2.573$, $p = 0.02$) and explicit re-aiming (mean = 0.29° , $t(19) = 3.218$, $p = 0.005$), suggesting that learning largely did not accumulate throughout the experiment.

Despite finding significant slopes for all the implicit adaptation sensitivity functions, they do not appear to be perfectly linear as can be seen in Figure 2.2b. The sensitivity appears to saturate for large rotations, which has been observed in previous studies (Morehead and Smith 2017; Wei and Körding 2009). To determine if implicit adaptation sensitivity is better described as a piecewise function, we followed the method of Wei and Körding (2009) to compare the slope for small rotations (-4 to 4°) versus the overall function. We find that these slopes are different for the Consistent-1 ($t(19) = 2.971$, $p = 0.008$), Consistent-2 ($t(19) = 3.227$, $p = 0.004$), and Consistent-3 conditions ($t(19) = 2.941$, $p = 0.008$). Note, the Consistent-1 condition is nearly identical to the study by Wei and Körding (2009), replicating their findings. However, the slopes were not different for the Consistent-7 condition. Given the visual similarity in the functions between consistency conditions, we suspect that this is likely attributable to noise – an issue we will address in Experiment 2.2. We performed the same comparisons between slopes within each consistency condition for explicit re-aiming and found no significant differences (all p values > 0.05), although it appears that there may be differences in the slopes between consistency conditions.




Figure 2.2. Average response of second trial in each mini-block in degrees to imposed visual perturbations of -16 to 16° for each consistency condition: Consistent-1 (red), Consistent-2 (green), Consistent-3 (blue), and Consistent-7 (purple). The analyses include twenty subjects per condition. The black diagonal line represents unity between input and output, a participant who learned fully would fall on this line. Shaded regions represent standard error. A) Angle of the hand heading direction. B) Angle of calculated implicit adaptation (hand angle – explicit re-aiming). C) Angle of explicit re-aiming, which was reported by touching screen with left hand.

Implicit Adaptation - Slope of Linear Fit				
	Mean	Confidence interval	Mean r-value	Percent of p-values < 0.05
Consistent 1	-0.192	-0.214 to -0.170	0.894	95.0%
Consistent 2	-0.298	-0.365 to -0.231	0.824	90.0%
Consistent 3	-0.335	-0.423 to -0.248	0.869	95.0%
Consistent 7	-0.300	-0.363 to -0.237	0.789	80.0%

Explicit Re-Aiming - Slope of Linear Fit				
	Mean	Confidence interval	Mean r-value	Percent of p-values < 0.05
Consistent 1	-0.028	-0.045 to -0.012	0.488	25.0%
Consistent 2	-0.374	-0.508 to -0.240	0.794	75.0%
Consistent 3	-0.230	-0.352 to -0.108	0.702	60.0%
Consistent 7	-0.152	-0.232 to -0.071	0.651	55.0%

Table 2.1. Statistics [mean, confidence interval, mean r-value] for the slope of the linear fit between imposed rotation and response: implicit adaptation or explicit re-aiming. The linear equation was fit to the data of each subject for the whole range (-16:16°) and only the small rotations (-4:4°).

To determine if there were significant differences in the slope of the sensitivity function between consistency conditions, we submitted the slopes for both implicit adaptation and explicit re-aiming to separate one-way ANOVAs. For implicit adaptation, while we find a significant difference between conditions ($F(3) = 3.55$, $p = 0.02$), this effect is driven by the difference between the Consistent-1 condition and all other conditions: Consistent-2 ($t(38) = 2.958$, $p = 0.005$), Consistent-3 ($t(38) = 3.099$, $p = 0.003$), Consistent-7 ($t(38) = 3.166$, $p = 0.003$). This suggests an overall reduction in the sensitivity of implicit adaptation when there is no consistency in error (Figure 2.2b). Explicit re-aiming also differs with consistency ($F(3) = 8.15$, $p < 0.001$). Interestingly, by comparing the Consistent-1 and Consistent-2 conditions, a post-hoc t-test revealed that increasing the consistency by a single trial radically increases the sensitivity ($t(38) = 5.02$, $p < 0.001$, Figure 2.2c). However, as the consistency is further increased the sensitivity function tends to decrease, comparing Consistent-2 with Consistent-3 ($t(38) = 1.561$, $p = 0.13$), Consistent-2 with Consistent-7 ($t(38) = 2.789$, $p = 0.008$), and Consistent-3 with Consistent-7 ($t(38) = 1.05$, $p = 0.30$, Figure 2.2c). While all of these comparisons show a decrease in sensitivity with corresponding increase in consistency, only the test between Consistent-2 and Consistent-7 is statistically significant.

We compared the slope of implicit adaptation and explicit re-aiming within each condition to determine if one was more sensitive to the imposed perturbation. Implicit adaptation was significantly more sensitive to perturbation size than explicit re-aiming in the Consistent-1 condition ($t(19) = 11.620$, $p < 0.001$) and in the Consistent-7 condition ($t(19) = 2.433$, $p = 0.02$). Greater sensitivity of implicit adaptation over explicit re-aiming in the Consistent-1 condition suggests that explicit re-aiming can be “turned off” when it is not useful, while implicit adaptation proceeds regardless. There was no difference in the slopes of implicit adaptation and explicit re-

aiming in the Consistent-2 ($t(19) = 0.810$, $p = 0.43$) and Consistent-3 conditions ($t(19) = 1.078$, $p = 0.30$). The relative magnitude of explicit re-aiming and implicit adaptation are generally equivalent when learning is useful to task performance.

It is worth noting that our analysis of sensitivity focused only on the changes in behavior following the first experience with a new rotation in a mini-block. For the Consistent-2, -3, and -7 conditions, the same rotational perturbation continued for additional trials. Thus, somewhat trivially, subjects could continue to implicitly adapt and explicitly re-aim. This is apparent in Figure 2.3, although the response appears to decelerate with continued training in the mini-block, which is likely attributable to progressively decreasing visual errors. Consequently, these trials become increasingly contaminated by prior performance and thus provide an impure measure of the error sensitivity function. Therefore, we limited our error sensitivity function estimations to only the second trial of the mini-block.




Figure 2.3. Time course of A) implicit adaptation and B) explicit re-aiming throughout each mini-block. Each mini-block is the average of the trials in a row for which each rotation size was consistent (2 trials for Consistent-2, 3 trials for Consistent-3, 7 trials for Consistent-7). Shaded regions represent standard error with twenty subjects in each condition.

Although not considered in our statistical analyses, Figure 2.3a does highlight the degree of sensitivity that implicit adaptation shows for rotation magnitude. This can be seen in the separation of the time course for each rotation magnitude. There is even clear differentiation between rotations as close as 2° and 4° (Kim et al., 2018). Re-aiming, however, does not show as clean a separation (Figure 2.3b).

Next, we were interested in determining if the change in explicit re-aiming sensitivity as a function of consistency was due to a fundamental feature of the learning process or as a result of a statistical property of the training environment. One possible explanation is that explicit re-aiming is actually sensitive to the changes in visual error between trials. For example, in the Consistent-7 condition, visual error was progressively smaller within a mini-block, but quite large in-between mini-blocks. In contrast, in the Consistent-1 condition, every trial was effectively in-between mini-blocks and, thus, larger visual errors were experienced more frequently. Indeed, the cumulative distribution of visual errors (CDF) varied significantly by condition for both means ($F(3) = 87.98$, $p < 0.001$) and standard deviations ($F(3) = 8.11$, $p < 0.001$, Figure 2.4b). Most notably, the mean of the CDF for Consistent-2 is larger than the means of the Consistent-3 ($t(38) = 3.032$, $p = 0.004$) and Consistent-7 groups ($t(38) = 10.2$, $p < 0.001$). The standard deviation of the CDF for Consistent-2 condition was also larger than that of Consistent-3 ($t(38) = 2.251$, $p = 0.03$), and Consistent-7 ($t(38) = 3.348$, $p = 0.002$). These cumulative distributions suggest that decreasing the consistency of the perturbation increased the average change in visual error. This may account for the increase in the magnitude of the aiming response, with the caveat that this relationship breaks-down when the visual error is completely unpredictable. It should be noted that we cannot produce a causal claim with this experimental setup as changes in visual error are by definition influenced by aiming behavior.




Figure 2.4. Empirical cumulative distribution function for A) imposed rotation magnitude and B) visual error magnitude. The imposed rotations are jittered for clarity.

In sum, we found that error sensitivity of implicit adaptation was largely invariant across consistency conditions. These findings are consistent with previous studies employing procedures to isolate implicit adaptation, although using prolonged block designs (Bond and Taylor, 2015; Morehead et al., 2017). Furthermore, we found that sensitivity functions for three out of the four consistency conditions tended to saturate for implicit adaptation, which is largely consistent with results from previous studies (Kim et al., 2018; Morehead et al., 2017; Wei and Körding 2009). However, the sensitivity function for implicit adaptation at the most consistent condition did not significantly saturate. In contrast, the sensitivity of explicit re-aiming changed as a function of consistency but showed linearity as a function of rotation size. To clarify these issues, we conducted a follow-up study (Experiment 2.2) to both replicate our central findings and extend them by including larger rotation sizes. Here we limited the study to two consistency conditions: Consistent-2 and Consistent-7.

Experiment 2.2

Similar to Experiment 2.1, implicit adaptation and explicit re-aiming were sensitive to the rotational perturbations for both consistency conditions. The average slope of implicit adaptation was significant for the Consistent 2 ($t(12) = 7.389$, $p < 0.001$) and Consistent-7 ($t(12) = 3.827$, $p = 0.002$) conditions (Figure 2.5 and Table 2.2). Likewise, the slopes were significant for explicit re-aiming in the Consistent-2 ($t(12) = 6.907$, $p < 0.001$) and Consistent-7 conditions ($t(12) = 3.685$, $p = 0.003$). This replicates the sensitivity of implicit adaptation and explicit re-aiming to rotation size from the first experiment.

Unlike in Experiment 2.1, the intercepts of the linear fits to implicit adaptation were significantly larger than zero for both Consistent-2 and Consistent-7 conditions, although quite small (mean = 1.436° and 1.552° respectively). The intercept of the linear fits to explicit re-aiming were significantly below zero for the Consistent-2 condition (mean = -1.342° , $t(12) = 2.881$, $p = 0.01$) but not the Consistent-7 condition. These small deviations in the intercept may be the result of a previously held bias or, less likely, learning that accumulated during training.




Figure 2.5. Responses to imposed visual perturbations of -32 to 32° . A) Angle of the hand heading direction. B) Angle of calculated implicit adaptation (hand angle – explicit re-aiming). C) Angle of explicit re-aiming. Shaded regions represent standard error with thirteen subjects per condition.

Implicit Adaptation - Slope of Linear Fit				
	Mean	Confidence interval	Mean r-value	Percent of p-values < 0.05
Consistent 2	-0.176	-0.223 to -0.129	0.782	84.6%
Consistent 7	-0.196	-0.296 to -0.096	0.750	61.5%

Explicit Re-Aiming - Slope of Linear Fit				
	Mean	Confidence interval	Mean r-value	Percent of p-values < 0.05
Consistent 2	-0.451	-0.579 to -0.323	0.910	92.3%
Consistent 7	-0.310	-0.475 to -0.145	0.692	69.2%

Table 2.2. Statistics [mean, confidence interval, mean r-value] for the slope of the linear fit between imposed rotation and response: implicit adaptation or explicit re-aiming. The linear equation was fit to the data of each subject for the whole range (-32:32°), the range from Experiment 2.1 (-16:16°), the range of small rotations which captures the 1:4 ratio of small to large rotations (-8:8°), and the range of small rotations used in Experiment 2.1 (-4:4°).

The non-linearity of implicit adaptation previously seen in Experiment 1 conditions Consistent-1, 2, and 3, replicated in the Consistent-2 condition of Experiment 2. Non-linearity was found when comparing the average slope between rotations of $\pm 4^\circ$ and the full range ($\pm 32^\circ$; $t(12) = 3.124$, $p = 0.008$), as well as when comparing between $\pm 8^\circ$ and the full range ($t(12) = 3.886$, $p = 0.002$). As in Experiment 1, the Consistent-7 condition did not show this effect for the comparison of the partial range $\pm 4^\circ$ to the full range ($t(12) = 1.815$, $p = 0.09$). However, the range from $\pm 8^\circ$ did have a significantly different average slope from the full range ($t(12) = 3.878$, $p = 0.002$). The same tests performed on explicit re-aiming produced no significant results, consistent with the findings from Experiment 1: Implicit adaptation shows a strong tendency to saturate at relatively larger perturbation sizes, while explicit re-aiming continues to contribute proportionately to learning throughout the whole range.

To compare sensitivity as a function of consistency, we submitted the slopes of implicit adaptation and explicit re-aiming to separate two-sample t-tests. As in Experiment 2.1, we found that the sensitivity of implicit adaptation did not change as a function of consistency ($t(24) = 0.350$,

$p = 0.73$). Unlike Experiment 2.1, however, explicit re-aiming behavior between Consistent-2 and Consistent-7 was not significantly different ($t(24) = 1.327$, $p = 0.20$). This result held when we restricted the comparisons of the two conditions to the range of rotational perturbations tested in Experiment 2.1 ($-16:16^\circ$, $t(24) = 1.661$, $p = 0.110$). This suggests that the subtle scaling of sensitivity of implicit adaptation and explicit re-aiming as a function of consistency, which we observed in Experiment 2.1, is not a robust effect.

A within-condition comparison between implicit adaptation and re-aiming showed that explicit re-aiming was more sensitive to these large perturbations than was implicit adaptation in the Consistent-2 condition ($t(12) = 3.844$, $p = 0.002$). However, this did not hold for the Consistent-7 condition ($t(12) = 0.932$, $p = 0.37$). As in the first experiment, we do not see a reliable pattern indicating that implicit adaptation is more or less sensitive than explicit re-aiming when adaptive learning is relevant to task performance.

Finally, implicit adaptation shows less separation of responses to the larger perturbation sizes used in this experiment (Figure 2.6a). This is in stark contrast to what was seen for the smaller perturbations in the previous experiment (Figure 2.3a). Explicit re-aiming, on the other hand, is sensitive to the range of perturbations in this experiment and shows clear differentiation throughout the time courses for large perturbation sizes ($>8^\circ$, Figure 2.6b).




Figure 2.6. Time course of A) implicit adaptation and B) explicit re-aiming throughout each mini-block. Shaded regions represent standard error with thirteen subjects per condition.

Discussion

To determine the sensitivity function of implicit adaptation and explicit re-aiming we probed the motor system with small visual perturbations. In addition, we investigated the impact of error consistency on the aforementioned sensitivity of implicit and explicit processes by manipulating the number of trials in a row for which the perturbation was consistent. In two experiments, we perturbed visual feedback during center-out reaching movements. Both explicit re-aiming and implicit adaptation are sensitive to and respond differentially depending on the size of the visual errors. By varying the consistency of the perturbation, we found that the sensitivity of implicit adaptation to small visual errors is impeded when the environment is completely unpredictable but stereotyped over all other levels of consistency. Likewise, the sensitivity of explicit re-aiming was practically null when the environment was inconsistent but stabilized with increased consistency in the perturbations. These results suggest that both implicit adaptation and explicit re-aiming are sensitive to very small perturbations, although implicit adaptation saturates

for larger perturbations, and only minimal consistency of the perturbations is needed to stabilize these sensitivity functions.

Implicit adaptation was most sensitive to changes in error at very small error magnitudes and saturated between eight and sixteen degrees, in keeping with previous results (Kim et al., 2018; Marko et al., 2012; Morehead et al., 2017; Wei and Körding, 2009). This non-linearity of implicit adaptation largely replicated across both experiments in most consistency conditions, with the only exception being the Consistent-7 condition. We suspect that the failure of this condition to replicate is likely the result of sampling error, as the general shape of its sensitivity curve is similar to those of the other conditions. Unlike implicit adaptation, explicit re-aiming showed no non-linearity and we infer that it is similarly sensitive to the full range of visual perturbations that we tested. These findings are consistent with previous studies suggesting that implicit adaptation is more stereotyped than explicit re-aiming and does not scale its response to a broad range of large perturbations (Bond and Taylor 2015, 2017).

We hypothesized that previous accounts of environmental consistency influencing motor learning could be fully attributable to explicit processes. However, we found that implicit adaptation was also affected by consistency. Changing error magnitude on every trial produces less implicit adaptation than when there is some consistency to allow for predictive responses. This suggests that there is a minimum amount of error consistency necessary to achieve the full possible amount of implicit adaptation. Once this minimum consistency requirement is met, however, additional consistency does not produce greater adaptation. This finding modifies previous work suggesting that implicit adaptation is extremely stereotyped and insensitive to environmental features or task demands (Morehead et al., 2017).

Consistency was important to the magnitude of the re-aiming response. Moving from no error consistency to two trials in a row with the same error magnitude produced a sharp increase in explicit re-aiming. Surprisingly, further increases in error consistency decreased this response, as found in Experiment 2.1 (a similar trend was seen in Experiment 2.2).

As we analyzed only the second trial of each mini-block in order to control for accumulation of learning, we believe that this reduction in the sensitivity of explicit re-aiming with increasing error consistency is the result of the statistical properties of our task: While the mean of the errors and the mean change in error from one trial to the next was zero for all conditions, the relative size of the change in error was larger for conditions with lower environmental consistency (see Figure 2.4). This resulted in more circumstances where the error magnitude was very high in conditions with low error consistency as compared to those with high error consistency. Herzfeld et al. (2014) previously found that under certain conditions sensitivity to an error increases with repeated exposures. Thus, the increased exposure to large perturbations in the less-consistent conditions could have influenced sensitivity. However, Herzfeld and colleagues did not separate implicit and explicit learning. Additionally, both results could be explained if explicit re-aiming is more engaged when there is a sudden change to the error signal in a sufficiently consistent environment, as we previously argued when discussing the CDF results in Experiment 2.1, this would account for the modest increase in explicit re-aiming that we see in low-consistency conditions.

Although not tested in these experiments, the sensitivity function of explicit re-aiming may be related to possible reward in the environment. Recent work has shown that explicit re-aiming is sensitive to bivalent feedback (Holland et al., 2018). One reason that explicit re-aiming is seemingly not sensitive to very small perturbations could be that the system categorizes these

errors as “close enough” and treats them as correct (i.e. within the subject’s natural motor noise), allowing implicit adaptation to “clean up” the residual error. Under this framework, explicit re-aiming would be expected to begin adjusting when movement is perceived to be sufficiently off-target. This phenomenon would likely be different across individuals and contexts.

The trade-off between implicit adaptation, which is most sensitive at small error sizes, and explicit re-aiming, which is sensitive to a wide range of errors, suggest that the two may be working in tandem to produce the appropriate response for any particular visual error. Specifically, we speculate that implicit adaptation allows for small updates to the internal model, necessary to avoid accumulation of small errors and to avoid drift in the system. The explicit system is simultaneously available for fast learning of large perturbations. Whereas slow, gradual adaptation is sufficient for most learned tasks, sometimes a leap in learning is necessary to accommodate a radical environmental change. Whether implicit adaptation and explicit re-aiming are independent of one another, or if they necessarily function in tandem, is currently an open question. Implicit adaptation continues even when an explicit strategy has already fully compensated for any visual perturbation and when the adaptation conflicts with task goals, suggesting that the two processes are independent (Mazzoni and Krakauer, 2006; Taylor and Ivry, 2011). However, there is also evidence that implicit adaptation is affected by performance or reward (Kim et al 2018; Leow et al 2018; Reichenthal et al., 2016). However, we did not observe a reliable change in the sensitivity of implicit adaptation as a function of the degree of consistency or explicit re-aiming: The sensitivity function of implicit adaptation was identical between the Consistent-2 and -7 groups, despite these groups having different degrees of consistency and explicit re-aiming. Nonetheless, this remains an important question for future research.

Given mounting evidence, in this paper and others, that implicit adaptation saturates at quite small perturbation sizes and well before reaching target learning when the perturbation is large, we suggest that the main function of implicit adaptation is calibrating to small errors. How then can we accomplish the full remapping necessary for skill learning? This re-mapping has been experimentally shown to be possible: Long-term adaptation studies conducted with subjects practicing throwing movements while wearing prism goggles showed that after weeks of practice subjects are able to make similarly precise movements immediately after donning the disrupting prism goggles as when they are not wearing them (Martin et al., 1996). Similarly, researchers have anecdotally found that they could interact normally with their environment after weeks acclimating to wearing inversion goggles (Stratton, 1896, 1897; Review in Sachse et al., 2017). In a visuomotor experiment, Semrau, Daitch, and Thoroughman (2012) found that adaptation is influenced by environmental factors introduced days earlier. While these experiments do not account for explicit re-aiming and implicit adaptation separately, our experience with everyday visual perturbations argues that the mapping is eventually implicit. For example, a new pair of glasses essentially imposes a scaling on the world, but we do not feel as though we are using an explicit strategy to compensate for this.

One possibility is that implicit adaptation does eventually constitute the fully remapped motor behavior, but that this process proceeds on a far longer time-scale than usually studied in visuomotor adaptation experiments. However, this is contradicted by recent work examining the time series of implicit adaptation under error-clamp conditions. Adaptation in this paradigm does not show the continuous, if slow, rise expected if implicit adaptation was eventually fully responsible for learning. An extended, multi-day experiment that can examine explicit and implicit processes separately is needed to provide insight into this open question (see Chapter III).

Chapter III

Longitudinal Tests of Implicit and Explicit Learning in Adaptation Tasks

Introduction

A significant perturbation must be overcome when using a computer mouse to navigate around a monitor screen. If sliding the hand forward moves the cursor up the screen, a 90-degree rotation has been applied to the visual feedback of the hand. Despite this discrepancy, we do not feel as though we are explicitly compensating for the shift between hand and computer mouse. The movement is made automatically, with little conscious effort necessary to achieve the desired goal. Intuitively, this suggests an important role for the implicit process, compensation without strategic intervention, in motor skill. Recent research has focused on characterizing the mechanisms giving rise to the implicit processes used in motor skill learning, mostly through sensorimotor adaptation paradigms. A common approach is to ask participants to make point-to-point reaching movements to nearby targets while the hand is obscured from view and visual (or proprioceptive) feedback is artificially perturbed (Held and Schrank, 1959; Cunningham, 1989; Imamizu et al., 1995; Pine and Krakauer, 1996; Krakauer, 2009).

The results of these sensorimotor adaptation experiments have been interpreted under the framework that the implicit process is driving a significant proportion of learning (e.g. Scheidt, Dingwell, and Mussa-Ivaldi, 2001; Baddeley, Ingram, and Miall, 2003; Fine and Thoroughman 2006, 2007; Wei and Körding, 2009). However, a number of more recent studies have challenged this view. Initially, these studies disassociated explicit compensation strategies and the implicit process with post-experimental assays and questionnaires (Heuer and Hegele, 2008; Hegele and Heuer, 2010), instruction (Mazzoni and Krakauer, 2006; Benson et al, 2011), and modeling

(Taylor and Ivry, 2014). This important early work showed that the slow, gradual learning of the implicit process is not necessary to compensate for visuomotor rotations.

In parallel to the recognition that strategies play an important role in adaptation tasks, we see rapidly emerging evidence that the implicit process is highly stereotyped regardless of the particular task demands. The implicit process asymptotes to a value far less than complete learning and is largely insensitive to perturbation magnitude (Bond and Taylor, 2015; Morehead et al., 2017). Morehead and colleagues (2017), revealed a stereotyped implicit learning curve when the incentive to strategize was completely removed. Surprisingly, asymptotic implicit learning remained insufficient to account for the perturbation after many hours of practice. (Morehead and Smith, 2017, Morehead et al., 2017; Kim et al., 2018). Finally, recent evidence suggests that the contribution of the implicit process decreases with extended training (Avraham, et al., 2020).

The essential role of the implicit process may be in fine adjustments to very small errors. A few recent reports show that the implicit process is proportional for rotations under approximately 8 degrees and there may be an effect of reward (Kim, Morehead et al., 2018; Kim et al., 2019; Leow et al., 2018; Chapter II). Given all of this, it is difficult to see the role of asymptotically constrained the implicit process in long-term motor learning.

However, a rest period, following exposure to a perturbed training environment, allows for continued learning in the absence of the stimuli (Brashers-Krug, Shadmehr, and Bizzi, 1996). If a rest period is important for learning, a single session may be insufficient to determine how the implicit process contributes over the long-term. The full effect of adaptation seen by Stratton (1896, 1897) and Kohler (1941, 1951a, 1951b), complete adaptation to an inverted world without need for explicit compensation, occurred following five full days of continuous exposure. Implicit learning may require extensive consolidation periods in order to fully compensate for a large

visuomotor rotation (Krakauer and Shadmehr, 2006; Criscimagna-Hemminger and Shadmehr, 2008; but see Caithness et al. 2004).

In Experiment 3.1, we investigate adaptative learning to a visual perturbation over five days of training. If the implicit process, as measured in this task, will eventually contribute substantially to overall compensation, we expect to see a steady increase in the proportion of learning accounted for by the implicit process over time. We find that the implicit process, on average, does not change over the training period. However, we also find substantial between-subject differences that raise questions about the validity of the measure.

An additional complication, not considered in Experiment 3.1, is the directional nature of the implicit process. The implicit response to directional information appears to be automatic: the motor system adapts in a direction opposite to the perturbation even when such adaptation is task-irrelevant (Schaefer, Shelly, and Thoroughman, 2012, Morehead et al 2017; Butcher and Taylor, 2018). This response to directional information has led researchers to question the role that the implicit process might play in de-novo skill learning, such as the visuomotor mirror reversal task. This task requires a directional adaptation response *opposite* to that seen in rotation tasks (Figure 3.1; Gritsenko and Kalaska, 2010; Lillicrap et al., 2013; Telgen, Parvin, and Diedrichsen, 2014; Krakauer, 2009).

Indeed, there is evidence that, even after a few days of training, the implicit process works against performance in mirror reversal tasks (Telgen et al, 2014; Kasuga et al, 2015). However, complicating this picture, a mirror reversal task is exactly what bilateral hippocampal lesion patient H.M. was able to complete – presumably without the benefit of explicit processes (Corkin, 1968). How was this accomplished if the implicit process – as we know it from previous visuomotor rotation studies – is unable to solve a mirror reversal? In Experiment 3.2, we sought to determine

if sufficient practice would transform the mirror-reversal task from one of *de novo* learning, to an adaptation task. That is, if the implicit process could become useful in a mirror task given a sufficient training period. This would give deeper insight into the role of the implicit process in early and long-term learning of novel motor skills. We found that, on average, the implicit process gradually decreases over time – slowly becoming more appropriate for the mirror reversal task. However, we again find substantial between-subject differences.




Figure 3.1. Appropriate correction for both rotation and mirror reversal perturbations. If the goal is to reduce error between the cursor and the target, the adaptive correction, denoted by the blue arrow, is counterclockwise for the rotation-based correction but clockwise for the mirror-based correction. Prior evidence supports the implicit process being appropriate for the rotation, but not for the mirror. This has been used as evidence that visuomotor rotation tasks produce adaptations of an established skill (reaching), while mirror-reversal tasks are a *de-novo* skill that do not benefit from adaptation.

Materials and Methods

In two experiments, twenty-four subjects participated in a visuomotor rotation task for one hour each day for five consecutive days. Subjects attempted to compensate for a visual perturbation of 45° while reaching to targets that appeared on a circle centered on the start-position.

Participants

Seventeen first year graduate students, undergraduate students, and community members were recruited to participate in Experiment 3.1. Of the seventeen, 3 subjects were removed from the experiment due to equipment failure and 2 additional subjects dropped-out of the experiment without completing all five days. The remaining twelve subjects [5 female, age 24.04 ± 2.11 yrs.] successfully completed the full experiment and their data is reported here. To have an equivalent number of subjects in the second experiment, fourteen subjects were recruited in the same manner as in Experiment 3.1. Two of these subjects retired from the experiment before completion. Data is presented for the remaining twelve subjects [8 female, age 23.76 ± 3.04 yrs.].

Subjects were verified to be right handed by the Edinburgh Handedness Inventory (Oldfield, 1971) and self-reported having normal or corrected-to-normal vision. The experimental protocol was approved by the [Author University] Institutional Review Board and all subjects provided written, informed consent. All subjects received monetary compensation for their participation.

Apparatus

Subjects performed horizontal movements in a center-out reaching task similar to that first described in Bond and Taylor, 2017 (Figure 3.2). Stimuli were displayed on a 60 Hz, 17-in., Planar touch sensitive monitor (Planar Systems, Hillsboro, Oregon) and computed by a Dell OptiPlex 7040 machine (Dell, Round Rock, Texas) running Windows 7 (Microsoft Co., Redmond, Washington). Movements were recorded with a Wacom magnetic digitizing pen and tablet (Wacom Co., Kazo, Japan). Aiming locations were recorded by tapping the touch sensitive monitor, which was placed 25 cm above the Wacom tablet and obscured visual access to the right hand. The game was controlled by custom software coded in MATLAB (The MathWorks, Natick, Massachusetts), using Psychtoolbox extensions (Brainard, 1997, Kleiner, Brainard, and Pelli, 2007).

Procedure

In order to assay the operation of explicit strategies and the implicit process on a trial-by-trial basis, we asked participants to report their intended aiming location (i.e. the direction that they intend to move in order to compensate for the perturbation) by tapping the surface of a touch screen (following Bond and Taylor, 2017; Hutter and Taylor, 2018). Participants reported their intended aiming location and then make a reaching movement. The difference between the intended movement (explicit re-aiming) and the actual reach direction reveals any learning of which the subject was unaware (the implicit process).

Each trial began with subjects holding their right hand in the center of the workspace for 300 ms. Then, a circular orange target (0.25 cm radius) appeared 7 cm from the start position on a blue, 7 cm radius “aiming ring”. To begin a trial, the subject indicated their intended reach

direction by tapping on the aiming ring with their left hand. Once an aim was recorded, the target turned from orange to green, the aiming ring disappeared, and the subject was able to begin their reach with the right hand (see Figure 3.2). Subjects were instructed to proceed through the experiment as quickly as possible, however, there was no time limit for tapping the aim location or beginning the reaching movement.




Figure 3.2. Experimental paradigm for assaying the implicit process to an imposed visual perturbation. The target locations for the mirror task were chosen to maintain a 45° perturbation ($\pm 22.5^\circ$ from the mirror axis). Participants indicate their intended movement direction (re-aiming) by tapping on the “aiming ring”, which is shown here in blue. They then use the right hand to reach for the location that they think will bring the cursor to the target (ostensibly the re-aimed location). Any learning not accounted for by explicit re-aiming is presumed to be the implicit process.

If a subject attempted to begin moving their right hand before an aim location was registered, the message “Remember to report aim” was displayed and the trial restarted. While reaching, subjects were instructed to move as quickly as possible past their intended aim location and back to the center in one smooth movement. This “out-and-back” motion was encouraged in order to decrease the time necessary to return to the center of the workspace. After a successful movement beyond 7 cm, participants were guided back to the start position by a white ring. The

ring was centered on the start position and its radius represented the distance of the subjects' hand to the start position. However, subjects became highly accurate at stopping their out-and-back movement directly on the center of the workspace and used the guiding ring less overtime. This greatly increased the number of trials that could be completed in each one-hour session.

During each trial, once the right hand moved out past 7 cm, endpoint feedback was displayed for 500 ms in a position contiguous with the last point of on-line feedback. If the position of the cursor completely overlapped with the target (< 1 deg. of angular deviation), the subject heard a pleasant “ding”. Otherwise, an unpleasant “buzz” sounded. The feedback “Too Slow” was given if the reach time from start to 7 cm exceeded 600 ms ($< 1\%$ of trials, these trials were excluded from further analysis).

A familiarization period of 16 trials with online feedback began the experiment. Two baseline periods followed the familiarization. The first baseline period was composed of 80 trials without feedback to determine if any subject held strong biomechanical biases that would not be averaged out by the large target set. The second baseline period provided veridical feedback, was also 80 trials long, and was designed to wash-out any drift that may have developed in the no-feedback phase. Subjects were not asked to report their aim during these baseline trials as it was assumed that they would always be aiming at the target. A pause was included after the baseline trials so that the experimenter could explain the aiming procedure to the subjects. Subjects were explicitly instructed to indicate their intended reach direction, not the target or the cursor position, by tapping on the aiming ring with their left hand. Following these instructions, subjects completed 16 trials with veridical feedback to familiarize themselves with the touch screen and aiming procedure. The first day concluded with 608 perturbation trials, in which a 45° rotation between

movements of the hand and cursor feedback was abruptly introduced. Subjects were provided with a three-minute break at the midpoint of training on each day.

Subjects were carefully observed during the familiarization trials and for the first 50 training trials to ensure that they understood the re-aiming instructions; further direction was provided as necessary during this time. If a subject was observed to ‘re-aim’ repeatably on the target, instead of adjusting their aim, the subject was stopped and the experimenter insured that they understood the purpose of the aiming procedure.

On days two through five, the perturbation was present from the first trial. Subjects were given a quick refresher of the aiming and movement instructions at the beginning of each day and were told that they would “pick up right where they left off” the day before. Subjects completed 800 trials each day, with an approximately 3-minute break halfway through. A total of 3,808 training trials were completed. At the end of day five, a 32-trial washout period was completed without feedback. During this washout, subjects were instructed to discontinue any strategy they had developed and aim straight for the target.

In Experiment 3.1 (Rotation-Task), the perturbation was a 45° rotation in either the clockwise or counterclockwise direction (counterbalanced between subjects). In Experiment 3.2, the perturbation was a mirror about the vertical midline. In order to allow comparison between the two experiments, the four targets in Experiment 3.2 (Mirror-Task) were located 22.5° from the midline so that the solution to the perturbed reach was a 45° angle. The 16 targets in Experiment 3.1 were evenly spaced 22.5° apart on the aiming ring (Figure 3.2). Both experiments were programmed such that each of the targets appeared before any was again repeated.

Data and statistical analyses

The experiment presentation, data collection, and statistical analysis were all completed in Matlab (The Mathworks, 2016b). During both experiments, the digitizing tablet logged the trajectory of the right hand and the touchscreen monitor recorded the position tapped to indicate re-aiming location. To allow averaging across targets, hand trajectories were transformed to a common axis with the target at zero degrees. Additionally, the hand trajectories were transformed into heading angles by examining the average angle of the hand from a straight-line path to the target between 1 and 3 cm into movement. This procedure prevents the influence of visual feedback control on estimates of learning. An aiming angle for each trial was defined by the angle between the target and the location tapped on the touchscreen. The implicit process was calculated as the subtraction of the aim angle from the hand heading angle in each trial (Taylor, Krakauer, and Ivry, 2014). In all measures, a positive angle represents a counterclockwise divergence from the target. As we measured only two values, and computed the implicit process from those values, we conducted statistical analyses for only the aiming angles and the implicit process angles. We report the mean and standard deviation of hand angles for completeness only. Reaction times were calculated as the interval from target appearance to aim-report, except where noted. All data used in parametric statistical tests were tested for normality using Lilliefors test.

Modeling

To predict the time courses of explicit and implicit processes in the mirror task, we used a modified version of the two-state model (Smith et al., 2006). Here, we modeled explicit re-aiming as the fast learning process (X_f) and the implicit process as the slow learning process (X_s) over the course of 200 trials (for a detailed account of modeling explicit/implicit learning as fast/slow learning, see previous work by McDougle, Bond, and Taylor 2015). In addition, we assumed that explicit

re-aiming is updated based on target error, while the implicit process is updated based on the aim-to-cursor distance (Equations 1 and 2).

$$\textcircled{1} Xf(n + 1) = \alpha f \times Xf(n) + \beta f \times Ef(n)$$

$$\textcircled{2} Xs(n + 1) = \alpha s \times Xs(n) + \beta s \times Es(n)$$

Where Ef is the error between the target and the cursor locations, while Es is the error between explicit re-aiming (Xf)(n) and the cursor location. Note, we did not input actual values of explicit re-aiming in these simulations but instead simply treated them as a faster learning process. In addition, we did not fully collapse the target locations to a common axis in order to demonstrate how the mirror-reversal perturbation causes different signed errors depending on target location. As such, we separately simulated target locations in the first and third quadrants, and second and fourth quadrants. The values on αf , βf , αs , and βs were determined by hand tuning these parameters to the simple rotation case (Rotation-Task-Rotation-Correction Model, Figures 3.7b and 3.8b; $\alpha f = 1$, $\beta f = 0.2$, $\alpha s = 0.98$, and $\beta s = 0.01$). Note, we chose a value for αs that was slightly lower than one to simulate an asymptotic value of the implicit process that is similar to what has been observed experimentally. These same parameter values were then used for both the Mirror-Task-Rotation-Correction and the Mirror-Task-Mirror-Correction Model. In the Mirror-Task-Rotation-Correction Model, the error for the implicit process was defined in exactly the same way as in the Rotation-Task-Rotation-Correction Model.

To demonstrate the effect of the implicit process switching the error-correction calculation to be consistent with the mirror reversal, we flipped the sign of the implicit error term in the Mirror-Task-Mirror-Correction Model (Equation 3).

$$\textcircled{3} Xs(n + 1) = \alpha_s \times Xs(n) - \beta_s \times Es(n)$$

The modeling code described here is freely available online at [URL redacted for double-blind review]. The code is also available as Extended Data.

Results

Experiment 3.1

In Experiment One, twelve subjects participated in a visuomotor rotation task for one hour each day for five consecutive days. Subjects attempted to compensate for a visual perturbation of 45° while reaching to targets that appeared on a circle centered on the start-position. To dissociate explicit and implicit learning, subjects reported their intended end position (re-aiming) using a touch-screen monitor (Bond and Taylor, 2017; Hutter and Taylor, 2018, Figure 3.2).

As expected, participants were able to compensate for the visual perturbation during the first day of training (Figure 3.3). Nearly perfect performance, as measured by the hand heading angle, was achieved by the end of the first day ($44.8 \pm 2.0^\circ$) and was maintained in the fifth day of training ($44.7 \pm 0.8^\circ$). Interestingly, this high level of performance was supported by a relative mixing of explicit re-aiming and the implicit process both on the first day ($28.1 \pm 6.7^\circ$ re-aiming vs. $16.7 \pm 6.4^\circ$ adaptation) and on day five of the experiment ($30.7 \pm 16.2^\circ$ re-aiming vs. $13.9 \pm 16.3^\circ$ adaptation). Note, while it appears that explicit re-aiming accounts for nearly twice as much learning as the implicit process, performing inferential statistics between explicit and implicit learning is theoretically inappropriate since our implicit measure is derived from our explicit measure.




Figure 3.3. Time course of responses over five days of training to a 45° rotation. All responses were rotated to a common axis with the target at 0° and the solution at 45°. The vertical dashed lines represent first the onset of the rotational perturbation, then day breaks. Shaded areas represent between-subject standard error.

Analysis of within- and between-subject variance: While learning appears to be divided between explicit and implicit processes, visual inspection shows a notable increase in the between-subject variance for explicit re-aiming and the implicit process following the first day of training. This suggests that the relative contribution of explicit and the implicit process may differ dramatically between subjects. Indeed, if we sort the subjects by amount of the implicit process on Day 5 we find a full continuum of learning patterns (Figure 3.4). Our twelve subjects run the gamut from no the implicit process (Figure 3.4b), through a mixture of the implicit process and explicit re-aiming (Figure 3.4c), to complete the implicit process (Figure 3.4d). These results are consistent with the range of the implicit process previously found in numerous studies, including those using standard visuomotor rotation tasks without re-aiming reports (Taylor and Ivry, 2014), perturbations introduced in a pseudorandom walk (Stark-Inbar, 2017), and task-irrelevant-visual-error clamp (Morehead et al 2017; Kim et al 2018; Kim et al 2019).

Despite these large between-subject differences, participants appeared to be remarkably consistent in maintaining their personal relative contribution of explicit and implicit learning within and across days. To quantify this, we performed a regression comparing explicit re-aiming and the implicit process at the end of each day to the same measure at the beginning of the following day (Figure 3.5). We found a high correlation between the end of one day and the beginning of the next for both explicit re-aiming (*Pearson's* $r = 0.94$) and the implicit process ($r = 0.94$). To underscore that this correlation only reflects consistency at the individual level and not at the group level, we sought to quantify group consistency by shuffling the individual data. Here, we randomly assigned the prior days of each subject to the following days of a different subject 10,000 times. Averaging the correlation over these runs results in a correlation coefficient a fraction of the size seen for the true data (explicit re-aiming *average* $r = -0.07$; the implicit process *average* $r = -0.07$).




Figure 3.4. (A) A bar graph ordering subjects by the average amount of the implicit process as calculated from the last 200 trials of day five. (B) The subject (#2) reporting the most explicit re-aiming at the end of learning, assumed to have the least the implicit process (i.e. aiming directly at the solution with no perceptible the implicit process at the end of day five). (C) A subject (#6) taken from the center of the distribution. (D) The subject (#5) reporting the least explicit re-aiming (i.e. aiming at or very near to the target with all learning accounted for by the implicit process).




Figure 3.5. Average explicit re-aiming (A) and the implicit process (B) in last bin (16 trials) of day D plotted against first bin of day D+1 ($r = 0.94$). Each subject contributed 4 data points, and is represented by shape and color for visual clarity. The diagonal lines represent the unity line between responses on D and D+1. The marker color and shape represent an individual subject across all figures in this manuscript.

Analysis of adaptation measurements: When we use trial-by-trial aiming data to infer the implicit process levels, two subjects appear to compensate fully with the implicit process and five subjects appear to be using a mostly explicit strategy, with the five remaining subjects somewhere in-between. We compared this result with the traditional aftereffects measure, which was completed at the end of the 5th day of training. During washout trials where subjects were instructed to aim directly at the target without using any strategy and on which no feedback was given, not a single subject shows a full 45° aftereffect, or a 0° aftereffect – which would have been indicative of learning entirely via the implicit process and explicit re-aiming strategies respectively (Figure 3.6a).




Figure 3.6. (A) Average aftereffect for each subject averaged across the 32 trials of no-feedback washout (with subjects instructed to aim and move directly to the target). Error bars represent one standard error of the mean. (B) The time course of washout for the average across subjects, and individual trial values for the three example subjects defined in Figure 3. Shaded area represents between-subjects standard error. (C) Correlation between each individuals' adaptation as measured via subtraction (hand angle minus explicit re-aiming over the last 32 trials of training) vs washout (average hand angle in 32 washout trials). Error bars represent one standard error of the mean along each dimension. The marker color and shape represent an individual subject across all figures in this manuscript.

It is visually apparent that the aftereffects measure does not display the same range of adaptation that we saw using the aim subtraction method. To quantify this visual observation, we ran a correlation analysis comparing the implicit process in the final 36 trials of training (calculated with the subtraction measure) with the implicit process measured in the first 36 trials of the washout block (calculated by discrepancy from target location, with instruction to reach directly to target). We note that any correlation analyses on a limited sample size ($N=12$) should be viewed with caution. Additionally, it is important to note that we observed extremely variable trial-to-trial measurements within a subject. While the average adaptation, as reported in most studies of the

implicit process, corresponds well to the mean of trial-by-trial adaptation calculated through the aiming subtraction method, individual measures do not (Figure 3.6b).

Confidence is warranted in suggesting that some participants fully strategized their way out of the rotation problem. This is because they had to be aware of the correct solution in order to tell us that they intended to move 45° away from the target. Contrast this with the fully implicit subject, who is telling us that they don't perceive a rotation and so taps on or very near the target each trial. Such behavior could also be explained by a misinterpretation of the instructions or simple disregard for the instruction. Note, however, participants were required to report their intended aiming strategy on each trial, limiting the effort saved by reporting inaccurately. Additionally, if participants didn't understand the instructions, we would expect that they never re-aimed or fully re-aimed, and that this behavior would be consistent from the introduction of the rotation. Instead, we see gradual changes with training. Participants who displayed full the implicit process followed a consistent time course where the implicit process gradually increased with training while their explicit re-aiming gradually decreased. Likewise, participants who displayed nearly full explicit re-aiming increased their re-aiming angle with training while the implicit process gradually decreased. Importantly, in post-participation debriefing one of the participants who displayed full the implicit process indicated that they thought the perturbation was removed; the second such participant declined to be interviewed.

We conducted several post-hoc analyses in an attempt to determine the cause of the variability in the implicit process found in this experiment. Regression analyses with day-five the implicit process regressed by inter-trial interval, reaction time, movement time, or reach end-point variability, all produced null results ($p > .26$). It is likely, however, that this is the result of limited

power driven by a small sample size for this type of analysis. A study specifically designed – with much greater power – will be required to identify the source of individual variability.

Experiment 3.1 Conclusion: The average response to long-term adaptation in a visuomotor rotation tasks suggests that both the implicit process and explicit re-aiming continue to play a role in performance even after several days of training. On average, we found the implicit process to be constrained, which is consistent with the findings from studies examining the capacity of single-session implicit learning. However, we saw great variability between subjects, including evidence from two subjects that implicit recalibration could eventually replace explicit re-aiming.

We next investigate the usefulness of implicit recalibration processes in a *de novo* skill learning task, the mirror reversal task.

Experiment 3.2

In Experiment 3.2, twelve subjects participated in a visuomotor mirror reversal task for five consecutive days. Subjects compensated for a mirror perturbation originating at the midline while reaching to targets that appeared on a circle around the start position. Targets were positioned such that the correct response resulted in a 45° angle between hand and feedback. Subjects reported intended re-aiming using a touch-screen monitor before each trial. The procedure was exactly the same as previously used in Experiment 3.1.

As shown in Figure 3.1, useful recalibration in rotation paradigms has been described as that responding to the discrepancy between the cursor feedback and aiming location (Taylor and Ivry, 2011; Day, Roemmich, Taylor, and Bastian, 2016). For the implicit process to be useful in the mirror reversal task, the implicit error signal must be calculated in the opposite direction as it is for rotational perturbations (Figure 3.7). Additionally, the mirror reversal imposes a

counterclockwise error in quadrants one and three but a clockwise error in quadrants two and four – similar to a dual adaptation paradigm (Howard et al., 2012; Schween et al., 2019). This complicates the transformation space over targets and contributes to the learning difficulty.




Figure 3.7. Six scenarios with hand angle (magenta) responding only to implicit recalibration (red). Explicit strategy (blue) is held constant to simplify the figure. The implicit process (red) is driven by the error (black) between the re-aimed location (blue) and feedback (purple). Additive adaptation pulls the hand away from the re-aimed location and results in decreased error for a rotational perturbation, and increased error under a mirror reversal. The error sign is flipped in quadrants two and four under a mirror reversal, but not a rotation. Greater transparency of the vectors and endpoints indicates trials further into the past.

For clarity of visualization on graphics, and to underscore the difficulty of accumulating the implicit process during a mirror task, we combined responses for the targets in quadrants 1 and 3 (i.e., at 67.5° with those at 247.5°) and responses for the targets in quadrants 2 and 4 (i.e., at 112.5° with those at 292.5° ; Figure 3.7). Unlike in the rotation task, normalizing all target locations

to zero does not assist in visualizing the data. Statistical analyses were conducted on the full data set, combining all target locations by flipping the sign of quadrants two and four prior to averaging.

Unlike the visuomotor rotation in Experiment 3.1, participants on average slightly over-compensated for the mirror reversal perturbation on the first day ($48.5 \pm 3.8^\circ$), however, performance on the final day was nearly perfect ($45.7 \pm 1.3^\circ$, Figures 3.8a and 3.9a).




Figure 3.8. (A) Mirror reversal learning for all 12 subjects to targets in the first and third quadrants, normalized to 67.5° . The solution is 112.5° . Implicit recalibration is shown separately with positive values being appropriate for a rotational perturbation. (B) Rotation-Task-Rotation-Correction Model for target at 67.5° with parameters tuned to limit the implicit process to about 15° . (C) Mirror-Task-Rotation-Correction Model with parameters defined as in ‘B’. (D) Mirror-Task-Mirror-Correction Model with parameters defined as in previous. Implicit error term is reversed from the rotation correction.




Figure 3.9. (A) Mirror reversal learning for all 12 subjects to targets in the second and fourth quadrants, normalized to 112.5° . The solution is 67.5° . Implicit recalibration is shown separately with negative values being consistent with a rotational perturbation. (B) Rotation-Task-Rotation-Correction Model for target at 112.5° with parameters tuned to limit the implicit process to about 15° . (C) Mirror-Task-Rotation-Correction Model with parameters defined as in ‘B’. (D) Mirror-Task-Mirror-Correction Model with parameters defined as in previous. Implicit error term is reversed from the rotation correction.

On the first day, performance was again a mixture of the implicit process ($8.0 \pm 4.6^\circ$) and explicit re-aiming ($40.5 \pm 5.4^\circ$). However, while the implicit process initially contributed to learning (t -test of first day, $t = 6.050$, $p < .001$), it gradually decreased over time (paired t -test between adaptation on first and last day, $t = 2.202$, $p = .049$) and was not significantly different from zero on the final day ($2.2 \pm 7.8^\circ$, $t = 0.961$, $p = .357$). On average, participants fully compensated for the mirror reversal with explicit re-aiming ($43.6 \pm 7.6^\circ$) on the final day of training. Assuming a constant rate of change and that the implicit process is capped at

approximately 15° (as seen in Experiment 3.1), the implicit process would become more appropriate (i.e., switch signs) for a mirror-reversal after 10-14 days.

Analysis of within- and between-subject variance: As with the visuomotor rotation, there was significant variance in how subjects responded to the visuomotor mirror reversal (approximately -15° to 15° , Figure 3.10). Subjects were successful through a combination of the implicit process and explicit re-aiming, lost all the implicit process by the end of training, or reversed the direction of the implicit process.

As in Experiment 3.1, we conducted a linear regression analysis to determine how consistent subjects were across days (Figure 3.11). We again found remarkable consistency within each subject; expressed as a high correlation between the end of one day and the beginning of the next for both explicit re-aiming ($r = 0.865$) and the implicit process ($r = 0.770$). When this result was contrasted with a randomization between days D and D+1, the result is a far worse average correlation than the true data for both explicit re-aiming (*average* $r = -0.067$) and the implicit process (*average* $r = -0.071$). Again, while we note that correlations on small sample sizes should be viewed with caution, due to the extremely high correlational relationships we don't think there is reason to believe that the inclusion of more participants would change the conclusion in this case.




Figure 3.10: (A) A bar graph ordering subjects by the average amount of explicit re-aiming that they reported during the last 200 trials of day five. (B) The subject (#5) reporting to have reversed their implicit process so that it is appropriate for the mirror reversal at the end of day five. (C) A subject (#2) representative of the average taken from the middle of the distribution. (D) The subject (#12) reporting the most adaptation in the direction that is more appropriate for a rotation.




Figure 3.11. Average explicit re-aiming (A) and the implicit process (B) in last bin (16 trials) of day D plotted against first bin of day D+1 ($r = 0.77$ and $r = 0.87$, respectively). Each subject contributed 4 data points, and is represented by shape and color for visual clarity. The diagonal lines represent the unity line between responses on D and D+1.

Analysis of adaptation measurements: Mirror learning has been viewed as a skill learning process, as opposed to adaptation, because it does not result in the same response properties as a visuomotor rotation, such as rate of learning, offline consolidation, and shifts in a speed-accuracy tradeoff (Telgen, Parvin, and Diedrichsen, 2014). Given the complexity of forming an intuition about the implicit process in the mirror (Figure 3.7), we simulated the potential results using a modified two-state-space model (Smith et al., 2006). The goal was to form graphical representations of responses based on mirror-appropriate adaptation or adaptation to an alternating rotation. The two-state model characterizes learning in visuomotor rotation tasks in terms of two separate learning processes: one that learns slowly but retains memory and the other learns fast but quickly forgets. Recent work suggests that the slow process might be the implicit process while the fast process is explicit re-aiming (McDoughle, Bond, and Taylor, 2015). We adopted the

conventions of this model to predict learning based on the sign of the error signal. Note, our goal here is not to model mirror reversal learning *per se*, but to determine if the pattern of the implicit process we observed is more consistent with learning appropriate for a mirror reversal or for a visuomotor rotation.

Within the model, we assumed that the explicit re-aiming is updated based on target error, while the implicit process is updated based on the angle between the aim and the cursor (See methods for details). The learning and forgetting rates were determined by tuning these parameters to the simple rotation case in order to cap the implicit process at approximately 15° (Rotation-Task-Rotation-Correction Model, Figures 3.8b and 3.9b; Morehead and Smith, 2017). These same parameter values were then used for both of the Mirror-Models. In the Mirror-Task-Rotation-Correction Model, the error for the implicit process was defined in exactly the same way as in the Rotation-Task-Rotation-Correction Model (Figure 3.8c and 3.9c). However, to demonstrate the effect of the implicit process switching the error-correction calculation to be consistent with the mirror reversal, we flipped the sign of the implicit error term in the Mirror-Task-Mirror-Correction Model (Figure 3.8d and 3.9d).

When the mirror task is simulated with error calculated as appropriate for a rotation [$aim - cursor$], the result is a response pattern much like that of the data. In contrast, an error calculation that is appropriate for the mirror [$-(aim - cursor)$], produces results in a radically different pattern from the data (compare Figure 3.8.a to 3.8.c-d and Figure 3.9a to 3.9.c-d). The modeling results demonstrate that the implicit process under a mirror reversal first operates as if the perturbation was rotational in nature, then suppresses this adaptation.




Figure 3.12. (A) Average corrected aftereffects for each subject across the 32 trials of no-feedback, washout (with subjects instructed to aim and move directly to the target). Each subject is jittered left to right in order of least to most explicit re-aiming at the end of Day 5. (B) The time course of washout for the average across subjects and individual trial values for three representative subjects [#5 (Lowest), #8 (Median), #12 (Highest), re-aiming on last day). The shaded area represents the standard error of the mean for all subjects.

In the washout phase, at the end of day five, subjects were instructed to aim directly at the target without using any strategy to measure the size of the aftereffect. Similar to Experiment 3.1, we observed a very limited range of aftereffects compared with the implicit process calculated

throughout training. After correcting for baseline bias, aftereffects hover slightly, yet significantly, below zero (average corrected aftereffect = $-3.06 \pm 3.2^\circ$, t-value = -3.32, $p = .007$, Figure 3.1). As in Experiment 3.1, subjects were instructed to aim directly at the target without using any strategy in washout trials and no feedback was given. In addition to exhibiting the high within-subject variance seen in Experiment 3.1, the washout phase was confounded by the discrepancy between aiming to hit a target in training versus aiming to hit that same target in washout. Adaptation appears to be centered around the aim location, not the target location (Day et al 2016; McDougle and Taylor, 2019; Schween et al 2019). In our task, a correct re-aiming location in training is approximately the same location as another target on the opposite side of the mirror axis. The new aiming location in washout (i.e. directly at the target) is at the same location as that for the opposite target in training. Adaptation that resulted in positive aftereffects (relative to the adaptation expected in a visuomotor rotation task) during training would have resulted in negative aftereffects in washout. This suggests that we should read the average aftereffect as being in the direction to correct for a rotational error, not a mirror reversal.

Comparison of the Results of Experiments 3.1 and 3.2

Although Experiments 3.1 and 3.2 were conducted sequentially, and therefore should not be statistically compared, we took care to keep both experiments as similar as possible to allow qualitative comparison.

We saw a steadily declining average implicit recalibration for subjects training to a mirror reversal in Experiment 3.2. This is in contrast to the flat average of the implicit process function seen for the duration of training to a visuomotor rotation. A visual comparison of these time courses shows starkly different behavior between the two groups (Figure 3.13a). While subjects

completing the mirror reversal task did not, on average, reverse their implicit process sufficiently for it to work with their learning goal, they did suppress adaptation when compared to the subjects learning a rotation. This suppression is evident from the first day of training, where we see a slowing of initial adaptation rate. The source of this suppression remains an open question (see discussion).

We find that re-aiming reaction times, the time between target appearance and the touchscreen registering a re-aim tap, gradually decrease to approximately half over the course of five days of training (1.4 ± 0.2 s in first five bins of rotation training to 0.7 ± 0.07 s in last five bins of rotation training, Figure 3.13b). With no difference between the groups.




Figure 3.13. Comparison between subjects in Experiment 3.1 (Rotation) and Experiment 3. (Mirror Reversal). The horizontal dashed lines represent first the onset of the perturbation, then day break. The shaded areas represent the standard error of the mean. A) Time course of calculated the implicit process. B) Reaction time across training calculated as interval from target onset to re-aim registered. Because re-aiming was introduced right before perturbation onset, this time-course starts when the perturbation was initiated.

Discussion

The idea that adaptation of the motor system in response to altered feedback is a purely implicit process can be dated as far back as Stratton (1896). Our everyday experiences reflect this view. For example, compensating for the lateral and vertical shift of a computer mouse or a new pair of eyeglasses do not feel like explicit processes. However, it is also well-known implicit processes play a limited role in learning, particularly learning *de novo*. In two experiments, we described the timecourse of implicit adaptation to a rotation and mirror reversal perturbations.

Implicit Recalibration and Measurement Reliability

In these experiments, we measured the implicit process by subtraction from explicitly indicated re-aiming ((Taylor and Ivry, 2014; Bond and Taylor, 2017), and as an aftereffect (Hollerbach and Atkeson, 1987; Jordan and Rumelhart, 1992). Regardless of the metric used to measure the implicit process, we found that implicit processes do not fully account for learning in a rotation task, and were suppressed in the mirror reversal task. This group-level result may have been expected given recent evidence that the implicit process falls well short of full learning (Morehead and Smith, 2017; Morehead, Taylor, Parvin, and Ivry, 2017; Kim et al., 2018; Kim, Parvin, and Ivry, 2019). We hypothesized that this failure of the implicit process was due to a lack of sufficient opportunity for consolidation, an idea that failed to be supported by our five-day experiments.

While group-level data does not provide evidence that the implicit process is sufficient for complete learning in a visuomotor rotation task, certain individuals do show complete adaptation under the subtraction measure. Enormous variability, spanning the entire range from zero to full adaptation through implicit processes, was seen across individuals. This suggests that the

traditional method of averaging over participants is simplifying, and perhaps obscuring, experimental results. What information about the implicit process can be gleaned from these individual differences?

The mirror reversal results present an exciting future opportunity to determine the error signal that initiates suppression of the implicit process in the mirror reversal task. Large variability in individual behavior provides a foothold to understanding suppression of implicit processes. Here we consider a few possible explanations that could be a target for future research.

The mirror reversal, unlike a rotation task, has an error-signal with a constantly switching sign. Previous work has shown that inconsistency in error signals reduces overall learning (Castro et al., 2014) and the implicit process in particular (Hutter and Taylor, 2018), which is consistent with the motor system storing a history of errors (Herzfeld et al., 2014). Here, we would expect participants that experienced more error switching, through their unique history of missing the target, to suppress implicit adaptation more quickly and fully.

Second, the difference between rotations and mirror-reversals could be thought of as a reparameterization, as previously considered in rotations, compared with learning a completely new structure, suggested for the mirror reversal (Braun, Aertsen, and Wolpert, 2009). The reparameterization of the transformation matrix, as seen in adaptation to a rotation, might be the fundamental function of the implicit process. Indeed, the implicit process may be exceptional at matrix reparameterization. In contrast, learning a mirror reversal requires learning the full structure of the task, including the relationship between parameters in the transformation matrix. It appears that the implicit process tries to reparametrize the transformation matrix as if it were a rotation even if doing so is inappropriate for the task. Only after days of training does the implicit process appear to stop reparametrizing inappropriately. It could be that once a new structure is specified,

the structure of the mirror reversal, the implicit process becomes appropriate for this new structure. The speed of learning this structure would dictate how quickly implicit processes began reparametrizing appropriately.

A slightly lower-level of description of this problem, one rooted in control engineering, was put forth in a recent study by Hadjiosif and colleagues (2020), which compared adaptation under a visuomotor rotation and a mirror reversal as a way to distinguish between updating forward model updating and updating a control policy. They too find that the implicit process in a mirror reversal was directed in the wrong direction to counter a mirror reversal, which is inconsistent with the updating a forward model. Instead, implicit learning of a mirror reversal may require direct updating of a control policy. Learning a new control policy would again have the outcome of changing the structure of the transformation matrix.

These explanations are highly speculative, and additional research is necessary to pin down the reason for implicit processes being counterproductive in a mirror reversal task, and the subsequent suppression of implicit recalibration.

Proceduralization of Explicit Re-aiming

Although the implicit process is insufficient for learning, we submit that it is highly unlikely that subjects are performing a time-intensive, computationally-demanding, strategizing process at the end of five days of training. A more likely explanation is that the re-aiming strategy has become partially proceduralized, also referred to as caching or habitualization. We'd expect to find evidence for proceduralization in participants' reaction or preparation times (Logan, 1980; Cohen, Dunbar, and McClelland, 1990). Increased reaction time increases are commonly observed at the onset of visuomotor perturbations and gently decline with training (Saijo and Gomi, 2010;

Fernandez-Ruiz, Wong, et al., 2011; McDougle and Taylor, 2019). If preparation time is limited, performance is significantly hindered in visuomotor rotation tasks as participants may not have sufficient time to re-aim their movements to counter the rotation (Fernandez-Ruiz et al., 2011; Haith et al., 2015). However, one study found that performance under constrained preparation time can be restored following two days of training, potentially reflecting a proceduralization of the re-aiming strategy (Huberdeau, Krakauer, and Haith 2019).

Here, we find that re-aiming reaction times in both experiments gradually decrease to approximately half over the course of five days of training (1.4 ± 0.2 s in first five bins of rotation training to 0.7 ± 0.07 s in last five bins of rotation training, Figure 3.13b). We would like to interpret this result as indicative of proceduralization of the re-aiming strategy in both tasks. Note: the absolute value of reaction time during training is expected to differ from previous reports, as our measure of reaction time necessarily includes the time taken to report the re-aiming location with the non-dominant hand (e.g. Fernandez-Ruiz, Wong, et al. 2011; Haith, Huberdeau, and Krakauer 2015). Nonetheless, the fact that we observe a significant drop in reaction times over the course of training, and given the traditional view of interpreting reaction times as reflective of computational demands, we suspect that this is an indication that proceduralization has occurred.

The question turns to what might underlie proceduralization of explicit re-aiming. One possibility is that the process is as simple as forming a stimulus-response mapping, allowing subjects to automatically perform an action that was previously conducted through computation (Daw, Niv, and Dayan, 2005). Additionally, stimulus-response mappings may result in a cached response or habit, which could simulate an implicit process by being extremely fast and relatively robust to added cognitive loads (Daw, Niv, and Dayan, 2005; Dayan, 2009; Haith and Krakauer, 2013). McDougle and Taylor (2019) found evidence that initial rotation re-aiming strategies

eventually give-way to a theorized look-up table for each target location (a stimulus-response map).

Another possibility to account for the decline in reaction time is Logan's (1988) dual process theory. Unlike the formation of a stimulus-response mapping, dual process theory explicitly imposes a second process that is engaged at each instance of reaching to a target. This process competes with the strategy until it produces better reaction times, then the proceduralization takes over. As with a stimulus-response mapping, the correct response would be specific to individual targets and generalization is not expected. It is worth noting that a single-session rotation learning experiment found no evidence supporting a dual-process account (Provost et al, 2013).

In these two experiments, we found that, on average, the implicit process does not account for the majority of adaptive learning, even after many days of consolidation. However, we also saw large variation between individual subjects. This variation questions the validity of averaging over subjects to make claims about implicit processes, and raises questions as to why one measure captures this variability while the other does not. Finally, we found evidence in the mirror reversal task that the implicit process is slowly suppressed to compensate for alternate task demands.

Chapter IV

Cognitive Strategies in Mirror Reversal Tasks

Introduction

The cognitive strategies used in motor skill learning have received increasing attention in the past few years. As an example of the importance of cognitive strategies, consider learning to play a piano piece. This could be accomplished by learning the correct motor response for every note, a stimulus-response (S-R) association. However, building the S-R association is cumbersome, especially as the target set size exceeds working memory capacity (Miller, 1956). Performance would be very poor while the appropriate S-R associations were being constructed. This limitation can be avoided by using the common parametric strategy of anchoring at middle C and referencing other notes with respect to this anchor. The strategy rapidly improves performance while the skill is being developed.

We can study this phenomenon using simplified motor tasks that impose artificial errors on a motor skill, such as the visuomotor rotation task. In these tasks, learners often use cognitive strategies to speed learning and restore performance (Taylor et al., 2014; McDougle et al., 2019). These strategies appear to be deliberative, controlled processes as they can be highly variable and their use increases reaction times. McDougle and Taylor (2019) show that the explicit strategy employed in visuomotor rotation tasks is parametric in nature, bearing remarkable similarities to mental rotation (Kosslyn et al, 1978; Shepard and Metzler, 1971). However, when the target set size is low enough, this computation can be quickly replaced by an S-R association, such that the action vector is automatically retrieved at the solution location and the computational cost of large perturbations is eliminated. McDougle and Taylor suggest that a low set size is required to see the development of S-R association, in a relatively short training session, because a unique S-R

association must be created for every target position. Later in this paper, we will argue that the underlying computation of the parametric strategy can itself be cached into an S-R association, allowing S-R associations to work at infinite target set sizes with very little training.

Much like the visuomotor rotation tasks, the mirror-reversal task probes motor learning in an artificially alerted environment. At face value, these tasks create the same learning pressures and require similar strategies to overcome. A rule governs the relationship between the movement of the limb and the visual outcome (a rotation vs a mirror reversal, some examples in Figure 4.1) and learning this rule provides all of the information necessary to fully compensate for the perturbation. However, these two tasks are learned quite differently. Mirror reversal tasks do not produce the systematic error reliably seen in visuomotor rotation tasks when visual feedback is returned to a normal, unperturbed state (aftereffects; Telgen et al, 2014). Additionally, automatic adaptation in mirror reversal tasks works against learning (Telgen et al, 2014; Kasuga et al, 2015) and is eventually subsumed by explicit strategies (Wilterson and Taylor, 2019).

This paper replicates the results of McDougle et al (2019, Experiments 1 & 2) and extends these findings with evidence that, unlike visuomotor rotation tasks, the strategy employed in mirror reversal tasks is not dependent on angular perturbation size. Additionally, we show that complex environments which can be solved with a simple rule, like a large target set under a mirror reversal, afford more efficient learning than a simple environment requiring a complex solution, like a small target set under a visuomotor rotation.

Materials and Methods

Participants

120 subjects were recruited from the research subject pool maintained by Princeton University and compensated with course credit or a small stipend (15 in Experiment 4.1, 50 in Experiment 4.2, 40 in Experiment 4.3, and 15 in Experiment 4.4). Experiments were conducted in accordance with the procedures of the institutional review board at Princeton University and all subjects provided written, informed consent. Subjects were verified to be right handed by the Edinburgh Handedness Inventory (Oldfield, 1971) and self-reported having normal or corrected-to-normal vision.

The original target sample sizes were chosen based on those reported in McDougale and Taylor (2019, Experiments 1 and 2), as the current experiments are a direct follow-up to their results. As in McDougale and Taylor's paper, we conservatively collected more subjects than we would if simply following a power analysis. This decision was made because 85% power analyses, using effect sizes calculated from McDougale and Taylor's data, indicate that we should collect samples of 5-10 subjects per condition. Such small samples are generally considered unsuitable for publication, so we choose higher target sample size. Note that all collected samples are greater than those required by this power analyses, with the exception of one condition in Experiment 4.3 as noted below.

We originally intended to recruit 160 subjects in total (15 in Experiment 4.1, 60 in Experiment 4.2, 60 in Experiment 4.3, and 15 in Experiment 4.4). However, we were unable to finish data collection due to the effective closure of behavioral experimentation at Princeton University resulting from the SARS-COV2 global pandemic from March 2020 to the time of writing (January 2021). This has substantial effects only on the rotation replication group in

Experiment 4.3, which has been replaced by previously published data (McDougle and Taylor, 2019) instead of the intended replication.

Apparatus

Subjects performed straight-line movements, sliding across a Wacom tablet with their right hand while holding a magnetic digitizing pen (Wacom Co., Kazo, Japan). The tablet was placed on a flat, horizontal surface directly in front of the subject (Figure 4.1a). Feedback stimuli were displayed on a 60 Hz, 17-in., Planar touch sensitive monitor (Planar Systems, Hillsboro, Oregon) and computed by a Dell OptiPlex 7040 machine (Dell, Round Rock, Texas) running Windows 7 (Microsoft Co., Redmond, Washington). Aiming locations were recorded by tapping the touch sensitive monitor, which was placed 25 cm above the Wacom tablet and obscured visual access to the right hand. Presentation of experimental stimuli was controlled by custom software coded in MATLAB (The MathWorks, Natick, Massachusetts), using Psychtoolbox extensions (Brainard, 1997, Kleiner, Brainard, and Pelli, 2007).

Procedure

Each trial began with the subject holding their right hand in the center of the workspace for 500 ms. Then, a circular orange target (0.25 cm radius) appeared and the subject was able to begin the reach with their right hand. While reaching, subjects were instructed to move as quickly and smoothly as possible through the target. Once the subject's right-hand exited the starting area, they did not receive visual feedback until their hand passed an invisible boundary. This boundary was created for each trial by calculating the radial distance from the start position to the target and using this value to draw a circle around the starting position. Endpoint feedback was displayed for

1 s, beginning 2 s after the right hand moved past the boundary. Delaying presentation of feedback is a reliable way to suppress implicit adaptation, allowing us to better isolate strategy use (Kitazawa, Kohno, & Uka, 1995; Brudner et al., 2016; Schween & Hegele, 2017; Parvin et al., 2018). If the position of the cursor completely overlapped with the target (< 1 deg. of angular deviation), the subject heard a pleasant “ding”. Otherwise, an unpleasant “buzz” sounded. After each movement, participants were guided back to the start position by a white ring. The ring was centered on the start position and its radius represented the distance of the subjects’ hand to the start position. The next trial began once the subject reached the starting position.

The goal of **Experiment 4.1** was to determine if a parametric strategy is used in a mirror reversal task. In the mirror reversal task, the visual perturbation was a mirror-reversal about the vertical midline. Targets could appear in any location on a ring, centered on the starting position and having a radius of 7 cm. To ensure that the full workspace was covered, the ring was divided into sixteen ten-degree slices (Figure 4.1b). Targets were pseudo randomly generated so that no target appeared in the exactly the same place as another, but an equal number of targets appeared in each slice. Although targets could appear anywhere on the ring, the trials were ordered such that each section was visited before any were repeated.

Because the questions of this experiment did not require a learning time course, participants were fully informed as to the nature of the perturbation and the solution. Participants completed 360 trials with the perturbation present for the full session.

In **Experiment 4.2**, we wanted to specifically disassociate the horizontal reaching distance from the perturbation angle in both the mirror reversal and rotation tasks. To accomplish this, we assigned targets to appear within a circle, centered on the starting location and with a radius of 8 cm. Along with the angular randomization described above, targets were pseudo randomly

generated such that radial distances, between zero and eight cm from the center, were equally distributed (Figure 4.1c). Targets could appear anywhere within the circle and each slice was visited once before any was repeated.

In the rotation task, the visual perturbation was an angular rotation varying in direction, clockwise and counterclockwise, and size, 0° to 180° . During the rotation task, participants completed trial pairs so that they could determine the size and direction of the rotation (McDougle and Taylor, 2019). For the first trial of the pair, participants were instructed to reach directly to the target and observe the error. On the second trial of the pair, they were instructed to try and correct for the error observed in the first trial. Within each pair, the target was different for the initial observation trial and the probe trial (Figure 4.1e). Participants were explicitly told these parameters at the onset of the experiment, were told that there was no overarching pattern throughout the task, and that each trial pair was entirely independent of the others. The trial pair system was unnecessary for the mirror task, as the rules of the mirror-reversal give full knowledge of the perturbation (Figure 4.1e).

In the mirror reversal task, the visual perturbation was again a mirror reversal about the vertical midline and participants were fully informed as to the nature of the perturbation and the solution. In both conditions, participants completed 360 trials with the perturbation present from the first trial.

Experiment 4.3 probes the use of parametric strategies and S-R associations through a reaction time task (following McDougle and Taylor, 2019). The essential goal of this experiment was to directly compare the mirror task to the rotation task results found in McDougle and Taylor Experiment 1. Participants were divided into two experimental conditions and two control conditions, with the replication conditions replaced by previously published data (see Participants

section above). The experimental groups completed the mirror task, as described above, with either 2 or 12 target locations. The control groups completed the task without a perturbation to establish a baseline reaction time with 2 or 12 targets. Participants in this experiment were *not* told of the nature of the perturbations so that we could establish a learning curve.

Targets were restricted to the ring with radius 7 cm, as described for Experiment 4.1, and target set size was restricted to 2 or 12 targets. Take 0° to be on the positive veridical axis. In conditions with 12 targets, targets appeared at 15° , 30° , and 60° in Q1; 120° , 150° , and 165° in Q2; 195° , 210° , and 240° in Q3; and 300° , 330° , and 345° in Q4 (see Figure 4.1d). The exact location of targets was chosen so that the angular perturbation averaged to 70° (a target located at 15° with the solution reflected about the midline requires a reach to -15° , resulting in a 30° perturbation; perturbation sizes were 30° , 60° , and 120°); this is important because we are comparing to a rotation task with a consistent angular perturbation of 75° . In conditions where only 2 targets are present, they were randomly and independently drawn for each participant from the set of 12 targets, with one target in Q1 and the other in Q3 or one target in Q2 and the other in Q4.

Participants in the experimental groups completed 36 baseline trials without perturbation, 300 trials with the perturbation turned on, and 36 washout trials without perturbation. In the washout phase, participants were instructed to move directly towards the target without using any strategy that they may have previously developed. Participants in the control group completed 252 trials with no perturbation.

In **Experiment 4.4**, we modified the rotation task used in Experiment 4.2 so that a mirror reversal strategy could be used to solve the rotation task. Recall that in Experiment 4, participants completed trial pairs so that they could determine the size and direction of the rotation. The direction and size of the rotation being determined from the first trial and compensated for in the

second trial. In this experiment, we keep the target in the same location for both trials (Figure 4.1f). A participant could theoretically flip the sign of the perturbation over the target to find the correct movement vector. The target here acts as a mirror axis, but one that changes every trial pair instead of staying constant as in the mirror reversal task. Participants were fully informed as to the nature of the perturbation and the solution. Participants completed 360 trials with the perturbation present from the first trial.

Data and statistical analysis

The experiment presentation, data collection, and statistical analysis were all completed in Matlab (The Mathworks, 2019a). During both experiments, the digitizing tablet logged the trajectory of the right hand. To allow averaging across targets, hand trajectories were transformed to a common axis with the target at zero degrees with a positive angle representing a counterclockwise divergence from the target. The distance measure was defined as the horizontal component of the distance from the starting position to the target. Reaction times were calculated as the interval from the time that the target appeared on screen to the time that the participant's right hand left the starting position. Any individual trial with a measure greater than three times the standard deviation of the mean for that measure in that subject was excluded from all analyses. All data used in parametric statistical tests were tested for normality using Lilliefors test.




Figure 4.1. A) Experiment setup. Right hand is on a tracking tablet with the hand covered by the horizontally placed screen. B) In Experiments 4.1 and 4.4, 360 unique targets were pseudorandomly pulled from a 7cm radius ring with the restriction that an equal number of targets appeared in each 10° segment. C) In Experiment 4.2, 360 unique targets were pseudorandomly pulled from the circle with a radius of 8cm with the restriction that an equal number of targets appeared in each 10° slice. D) Target locations for Experiment 4.3. In a mirror reversal task these target locations create rotational perturbations of 30°, 60°, and 120°. E) The target pair paradigm used in Experiment 4.2. Participants in the rotation condition reached straight toward the target on the first trial of the pair, and used the observed error to correct for the perturbation on the second trial of the pair. Participants in the mirror reversal condition did not need an observation trial to learn the solution, and so corrected for the perturbation on every trial. F) Modified trial pair task used in Experiment 4.4. The target stayed in the same location for both trials in the pair. In this task, the solution can be calculated either by rotating the correct distance from the target as previously seen in this rotation task, or by using a mirror-like strategy. Here, a mirror strategy requires flipping the error over the target instead of flipping the target over the midline.

Results

Experiment 4.1

Previous experimentation has shown a direct relationship between the size of a rotational perturbation and the time that it takes for a participant to produce a movement (McDougle and Taylor, 2019). This has been taken as evidence of a parametric strategy, where the participant is actually rotating some vector from the target to the corrected position. The farther the vector has to be rotated, the longer it takes for the computation to be completed. In this experiment, we sought to determine if the computation is a constant or if it is task-dependent. We used a mirror rotation task, which has similar visual features to a rotation task but with several possible solutions. We discuss two solution spaces here.

If a rotation strategy is used in the mirror reversal task, we expect to see reaction time (RT) increase as the movement angle increases (Figure 4.2a, 4.2b). In this task, subjects reached to pseudorandomly generated targets positioned on a ring around the starting position. Subjects used shooting movements to reach past the target and received delayed, endpoint feedback. In the mirror reversal task, unlike a rotation task, target position dictates the angular perturbation (e.g. a target at -45° is hit by reaching to $+45^\circ$, a 90° movement angle). Given that targets were positioned all around the ring, subjects made movements resulting in angular perturbations of up to $\pm 180^\circ$.

We first analyzed our data using only movement angles that fell between 15° and 90° (matching the movement angles described in McDougle and Taylor, 2019). A t-test conducted on regression coefficients found no effect of movement angle (average slope = -0.5 ms/ $^\circ$, $t=-1.8$, $p=0.1$). Using the full data set with movement angles between 0° and 180° , we do find an effect of movement angle (average slope = 0.2 ms/ $^\circ$, $t=2.3$, $p=0.4$). This effect is driven by very large

rotation sizes ($>150^\circ$; Figure 4.2b). Taken together, we think that there is little evidence for use of a parametric rotation strategy in this task. We further examine this point in the discussion.

We next considered that the mirror task may be solved by a different parametric computation. The solution to a mirror reversal can be computed by traveling the distance, along the horizontal dimension, from the target to its opposite position on the ring (Figure 4.2a, 4.2c). If this computation is completed parametrically, we would expect to see RT increase as the horizontal distance from the target location to the solution is increases. When conducting this analysis, we ignore rotational information in favor of using horizontal distance as the predictive variable. Due to high collinearity between rotational and horizontal distance in this task, we do not report a multivariate analysis here (correlation coefficient = 0.98). This problem is addressed in Experiment 4.2.

T-tests conducted on the coefficients of the regression found a negative effect of movement distance (average slope = -0.69 ms/mm, $t=-3.4$, $p<.01$). This effect is in the opposite direction predicted by a parametric computation. In fact, the significant effect that we see here suggests a computation that is faster to farther targets. One possible explanation is that targets that are very close to zero time to parse, related to perceptual accuracy and not a parametric computation. There appears to be no measurable use of a parametric strategy in this task.

These results open the possibility that the underlying computation of the parametric strategy is itself being cached into an S-R association. As previously considered, the S-R relationship is between a target appearing in a specific location and the predefined vector given by use of the parametric strategy at that location. Sufficient repetitions of the parametric strategy at that target location (the stimulus) allow the system to drop the computation in favor of using only the resultant vector (the response). Because this process relies on repetition, using a large number

of targets prevents acquisition of the S-R mapping (at least on the timescales used in these experiments). The system is consequently left with the slower parametric strategy, the speed of which is dependent on the ‘amount’ of computation needed, which is further given by the size of the perturbation. Here, we see very little evidence of a parametric computation. Within the range of rotation sizes previously found to elicit a parametric computation, the response is performed with approximately equal RT regardless of the perturbation size. This suggests that the mirror task allows S-R associations to form in the presence of a very large target set size (360 unique targets).

Before further pursuing the idea that a parametric strategy is itself being cached into an S-R association, we wanted to ensure that our findings were not a statistical abnormality or the result of a confound between horizontal distance and rotational distance.




Figure 4.2. A) Rotational distance is measured as the angle between the hand and cursor vectors. Horizontal distance is measured from the mirror (vertical) axis to the hand location. B & C) Mean of median RTs from 15 subjects plotted as a function of movement angle (binned by 10°) and horizontal distance (binned by 5 mm). Error bars represent between-subject standard error. The green points represent the rotation sizes between 15° and 90° (as used in McDougle and Taylor, 2019, Experiment 1).

Experiment 4.2

In Experiment 4.2, we replicated the effect of rotation size on RT in a rotation task, as found in of Experiment 4.1 and McDougle and Taylor, 2019 (Experiment 2), and the null effect found in Experiment 4.1. Additionally, we began to dissociate rotational distance from horizontal distance by using targets throughout the workspace (correlation coefficient = 0.03 in the rotation task and 0.7 in the mirror task; Figure 4.3).

Again, beginning with movement angles that fell between 15° and 90° , separate bivariate regressions were conducted for the rotation and mirror task conditions. In the rotation task, t-tests conducted on the regression coefficients replicate a positive relationship between movement angle and RT (average slope = $2.1 \text{ ms}/^{\circ}$, $t=3.8$, $p<.01$) and horizontal distance produced no measurable effect on RT (average slope = $0.9 \text{ ms}/\text{mm}$, $t=1.6$, $p=.12$). In the mirror task, we find a negative relationship between movement angle and RT (average slope = $-0.7 \text{ ms}/^{\circ}$, $t=-3.3$, $p<.01$) and no effect of horizontal distance (average slope = $-0.3 \text{ ms}/\text{mm}$, $t=-1.2$, $p=.24$). *Post hoc* t-tests with Bonferroni correction (assuming four tests to account for those reported below) reveal a significant difference between the slope of the coefficients on movement angle between subjects in the rotation and mirror conditions ($t=4.0$, corrected $p<.01$). Consistent with horizontal distance having no effect on RT in either condition, the difference between the slope of the coefficients on distance was nonsignificant ($t=1.6$, corrected $p=.39$).




Figure 4.3. Using targets within the circle, instead of just on the perimeter ring, allows us to disassociate rotation size from horizontal distance. In every case, the mirror reversal can be solved by using horizontal distance to calculate the solution location. Subjects completed straight-line shooting movements and were not required to stop on the target.

Next, we ran bivariate regressions on the full data set, with movement angles between 0° and 180° . In the rotation task, we again find a positive relationship between movement angle and RT (average slope = $1.2 \text{ ms}/^\circ$, $t=4.0$, $p<.01$) and no effect attributable to horizontal distance (average slope = 0.2 ms/mm , $t=0.7$, $p=.50$). In the mirror task, we find no effect of movement angle (average slope = $-0.0 \text{ ms}/^\circ$, $t=-0.0$, $p<.98$), but a negative effect of horizontal distance (average slope = -0.6 ms/mm , $t=-2.3$, $p=.03$). As with the restricted movement angles, *post hoc* *t*-tests between subjects in the rotation and mirror conditions with Bonferroni correction reveal a significant difference between the slope of the coefficients on rotation ($t=3.1$, corrected $p=.01$), but not horizontal distance ($t=1.9$, corrected $p=.24$).




Figure 4.4. The first row shows the means of median RTs from 30 subjects in the rotation task condition plotted by (A) rotation size (binned by 10°) and (B) horizontal distance to target (binned by 5 mm). The second row shows the means of median RTs from 20 subjects in the mirror task condition plotted by (C) rotation and (D) horizontal distance to target. Error bars represent between-subject standard error. The green points represent rotation sizes between 15° and 90°

Experiment 4.3

One possible reason that rotation magnitude does not elicit an RT cost in the mirror reversal task is that the mirror strategy is efficiently cached, even with high target set sizes.

McDougle and Taylor (2019; experiment 1) found that formation of S-R associations in a rotation task was modulated by target set size (McDougle and Taylor, 2019, Figure 2). When compensating for a 75° rotation, subjects in both low set size (two target) and high set size (twelve target) conditions showed an increase in RT, consistent with using a parametric strategy to mentally rotate the motor plan on each trial. However, as training progressed, subjects in the low set size condition reduced their RT significantly below that of subjects in the high set size condition. This is evidence that subjects in the low set size condition form an S-R association, while those in the high set size condition do not (or do so less effectively). That a small target set size speeds the formation of S-R associations implies that each target is a unique stimulus that has to be paired with its appropriate response. Under this framework, the S-R association is created when the solution to each target location is cached through repeated experience at that target location.

We hypothesized that RT in the mirror task did not scale with rotation magnitude in Experiment 4.2 because the solution was already cached. However, because there was such a large target set size (360 unique targets), this would require the solution to be cached without experience at the target location. Thus, we propose that the parametric computation itself was cached into an S-R association - leaving target set size inconsequential. The parametric computation used to solve the mirror reversal could be calculated either through a mental rotation strategy or by using the horizontal distance to manipulate the automatic, target-facing movement vector. Caching this

solution means that the automatic movement vector is initiated in the correct direction regardless of target location.

In this experiment, we tested the difference between the effect of set size on computation speed in the mirror reversal and rotation tasks. A control group, where subjects completed a shorted version of the task without any perturbation, was included to find the RT floor at 2 and 12 targets. We expected that the baseline RT over training would differ between set size conditions due to Hick's Law, a log-linear increase in RT attributable to increased set of objects under consideration (Hick, 1952). All participants reached to targets set on a 7cm ring and received endpoint feedback that was delayed by 2s. Data for the Rotation Condition is taken from McDougle and Taylor (2019; see participants section).

Because the Control and Mirror conditions were collected at a different time from the Rotation condition, we were concerned that there could be a systematic difference in the baseline RT of subjects. Indeed, a significant one-way ANOVA ($f=8.8$, $p<0.01$) and subsequent *post hoc* t-tests show that the average RTs of the Control and Mirror conditions in the final baseline block were similar (489 ± 136 ms and 509 ± 139 ms, respectively; $p=.95$), while the average RT for the final baseline block in the Rotation condition (696 ± 262 ms) differed significantly from both the Control condition ($p<.01$) and Mirror condition ($p<.01$). To conduct a conservative test of our hypothesis, we normalized the RTs by adding *Baseline Rotation* – (*Baseline Control* + *Baseline Mirror*) / 2 = 197 ms to the RTs of the Control and Mirror Conditions.

Our first analysis focused on the hypothesis that, unlike the Rotation Condition where subjects in the 12T condition maintain high RTs consistent with a parametric strategy, RTs from the mirror condition would not differ from the target set control condition (Figure 4.5). Consistent with this hypothesis, we see a neat stacking of RTs early in learning. The Rotation conditions are

higher than the Mirror conditions which are in turn higher than the control conditions. Additionally, the 2T groups within these conditions are faster than the 12T groups (Figure 4.5a). Consistent with the idea that target set size is less important to computation in the mirror conditions, all groups converged by the end of training, except for the 12T Rotation condition.

To quantify this, we examined the means of median RTs over 6 blocks between early learning (first 6 blocks) and late learning (last 6 blocks). We defined a block as the interval over which each subject was exposed to each target exactly once, 2 trials in the 2T group and 12 trials in the 12T group, to control for the number of exposures at each target location. RTs (Figure 4.2b) were submitted to a three-way ANOVA with factors of Time (early vs. late learning), Set Size (2T vs. 12T), and Task Type (Mirror vs. Rotation).

We found a significant main effect of Time ($f=16.1$, $p<.01$), reflecting decreased RT in all groups over training. We also observed a significant effect of Task Type, suggesting that subjects employed different computational strategies in the mirror reversal and rotation tasks ($f=12.67$, $p<.01$). We did not find a main effect of Set Size ($f=2.9$, $p=0.09$), but, found a significant two-way Time $\times$ Set Size interaction ($f=4.5$, $p=0.036$) and Time $\times$ Task Type ($f=4.16$, $p=0.018$). *Post hoc* t-tests (Bonferroni corrected) revealed an effect of set size on late RT in the rotation groups (2T vs. 12T rotation: $t=-4.0$; $p<.01$), but not in the mirror groups (2T vs. 12T mirror: $t=-1.3$; $p=0.2$).

One potential explanation for the additional time needed to complete the rotation task is that subjects were moving slower in order to be more accurate. However, there was more error in the rotation conditions than in the mirror conditions (Figure 4.5c). The same analysis described above conducted with the absolute value of target error yielded a null ANOVA (Figure 4.5d, all $p>.15$).




Figure 4.5. Subjects completed a reaching task under three conditions: 20 subjects completed a control task (no perturbation), 20 subjects completed a mirror reversal task with an average angular perturbation of 70°, and 40 subjects completed a rotation task with a constant 75° perturbation. (A) Time courses for the means of reaction times, the time from target appearance until movement began, averaged across twelve trials. Vertical lines represent transition from baseline to training and training to washout (B) Mean RT for the first and last six blocks of training. A block is defined as 2 trials for the 2T groups and 12 trials for the 12T groups. (C) Time courses for the means of target error averaged across twelve trials. Target error is defined as the difference between end-point cursor location and target location. (D) Mean of the absolute value of error for the first and last six blocks of training. All shaded regions and error bars represent one between-subjects s.e.m. Data for the rotation condition was taken from McDougale and Taylor (2019).

Experiment 4.4

In Experiments 4.2 and 4.3, we found that mirror reversal tasks are not subject to the same computational costs as rotation tasks, even with high set sizes. However, it could be that there is a visual, or other low-level, advantage to flipping across the x-axis that allows for an extremely efficient use of strategy. In this experiment, we control for visual symmetry by manipulating the rotation task. The rotation task, also used in Experiment 4.2, uses trial-pairs (Figure 4.1e). The first trial provides information as to the rotation magnitude and direction, and the subject corrects for this perturbation on the second trial of the pair. Here we allow for a mirror-like strategy by using the same target location for each trial of a pair. Essentially, the mirror-axis changes on each trial-pair to the location of the target (Figure 4.1f). Subjects could ‘flip’ the perturbation over the target, but the ‘horizontal’ component of the target location is less readily available.

Given the lack of evidence for an effect of horizontal distance on RT and our inability to completely remove its covariance with rotation size, we have excluded horizontal distance as a predictive variable in this experiment. For comparative purposes, we included the graphical description in Figure 4.6b. As with previous experiments, we first analyzed our data with a regression using only movement angles that fell between 15° and 90°. A t-test conducted on the regression coefficients showed a significant effect of movement angle (average slope = 1.6 ms/°, $t=2.4$, $p=.03$). Using the full data set with movement angles between 0° and 180°, we do not find an effect of movement angle (average slope = 0.9 ms/°, $t=1.8$, $p=.09$; Figure 4.66a).




Figure 4.6. Mean of median RTs from 15 subjects plotted as a function of movement angle (binned by 10°) and horizontal distance (binned by 5 mm). Error bars represent between-subject standard error. The green points show rotation sizes between 15° and 90°.

Discussion

It has previously been suggested that a parametric rotation strategy could be cached into an automatic S-R response (Provost et al., 2013; McDougale & Taylor, 2019; Wilterson & Taylor, 2019). However, evidence has been limited due to the problem that large target set sizes continue to require a parametric strategy (incidental evidence for a cached parametric rotation strategy with 16 targets is given in Wilterson & Taylor, 2019). If the parametric computation is truly cached, the set size should not limit strategy use any more than set size limits reaching straight towards the target in the absence of a perturbation. In this paper, we presented evidence that the parametric strategy used to solved a mirror reversal task can be quickly cached in a single training session and used on an unlimited number of targets.

In Experiments 4.1 and 4.2, we show evidence that the mirror reversal task does not use a parametric strategy related to rotation size or distance along the horizontal axis, regardless of target set size. In those cases where small horizontal distances drive a negative correlation, it is in the opposite direction as predicted by a parametric computation. Based on the training time course that we see in Experiment 4.3, it seems likely that a parametric strategy is used at the beginning of training – revealed by initially increased RT – but is quickly replaced with S-R associations. The S-R association cannot be between individual targets and their responses, because subjects had a single exposure to each target in Experiments 4.1 and 4.2. While this single-exposure problem could be circumvented if the S-R association was based on a generalized ‘cloud’ of targets (e.g. if all targets within 10° of each other were treated as the same target), we would expect this to result in increased error as the response would not be tailored to the individual target. If the S-R mapping is not one between target and solution, we suggest propose that it is a cached parametric strategy.

In Experiment 4.4, we allow for use of a mirror-like strategy in a rotation task by using the target location as a mirror axis. The solution to any individual trial of this task is considerably more complicated than in the mirror reversal task, because the reference point (i.e. the mirror axis) is constantly changing. Rotation size does have an effect on the RT in this task, however it is not as large nor as consistent as in the normal rotation task. We take this as preliminary evidence that the parametric strategy was beginning to become cached, but had not completed the transition. Because it is not learned so quickly that it is difficult to measure, but still shows signs of being cached in a single session, this task is a good candidate for further exploration of the time course of turning a parametric strategy into an S-R mapping.

It is not immediately apparent why the parametric strategy used in a mirror reversal task would be easier to cash than that used in a rotation task. Especially a rotation task with a consistent

rotation size, as in Experiment 4.3. One proposal is that the mirror reversal does not stress working memory in the same way that a rotation task does. Working memory has previously been cited as a possible reason that small target set sizes allow S-R associations. Working memory could be useful in holding the solution linked with each target long enough for an S-R association to be formed. Not having the solution readily available would increase error, slowing the acquisition speed of a classically reinforced S-R association. A task with a more readily available reference point, as in the mirror reversal task, may assist in reducing error without extensive reliance on working memory, speeding the S-R acquisition.

The associations as discussed above were mostly conceived as pairing an individual target and its solution, not a cached parametric strategy. The first is only useful at the individual target that the response has been paired to and requires learning the unique response for each target. So, the acquisition of individual target-response associations may be irrelevant, competitive, or may provide a scaffolding to cash the full parametric strategy. Further investigation would be needed to disassociate these possibilities.

Chapter V

Model-Free and Model-Based Learning in *de Novo* Motor Skills

Introduction

The effective control of movement is itself a challenging task, and the challenge only increases as our goals become progressively more variable or complex. Returning to the example at the top of this dissertation, building the action-outcome mapping needed to quickly and reliably produce the action sequences necessary for keyboarding takes significant time - on the order of months to years. However, adapting one's finger movements to account for minor environmental perturbations (e.g. the scaling between a desktop and laptop keyboard) is remarkably fast - on the order of minutes to hours. This recalibration is driven by an implicit process, which is slow to learn, limited in scope, and highly stereotyped. Strategy use increases the rate and accuracy of learning, supplementing implicit processes in adaptation tasks (Taylor et al., 2014; Bond and Taylor 2015; McDougle et al, 2015; McDougle et al 2017; Morehead et al 2017; Kim et al. 2018; Mazzoni and Krakauer 2006; Taylor et al 2011; Butcher and Taylor 2018; Bond and Taylor 2017; Hutter and Taylor 2018; Wilterson and Taylor 2019; Haith et al 2015; Morehead et al 2015).

The field of motor learning has focused on the recalibration of skilled movements, mainly reaching, grasping, and walking, following an unexpected disturbance to the action-outcome mapping (Jordan and Rumelhart 1992; Shadmehr, Krakauer, and Smith 2010; Shadmehr and Mussa-Ivaldi 1994; Krakauer et al., 2000; Wolpert and Ghahramani 2000; Shadmehr, Krakauer, and Smith 2010; Krakauer 2009; Petitot, O'Reilly 2018). Few studies have sought to determine how the action-outcome mapping is learned *de novo*. So, we have limited understanding of how action-outcome mapping, and the ability to adapt those mappings, initially develop.

There is some experimental evidence that model-free and model-based learning underlie *de novo* learning of an action-outcome mapping. Indeed, the model-based stage of learning is often characterized as slow and effortful, requiring a computationally intensive tree search of action choices (Daw, Niv, and Dayan, 2005), echoing the attention-demanding effort of early skill learning (Adams, 1982; Dickinson, 1985; Killcross and Coutureau, 2003). On the other hand, once learning becomes model-free, it is no longer necessary to search the decision tree for the best solution. This results in characteristically rigid performance that is quickly and automatically accessible without need for effortful processing or attention (Daw, Niv, and Dayan, 2005). We recognize this as a close description of a learned skill, which can be completed rapidly without effortful control.

Imposing time pressure on subjects when they are learning an unintuitive keyboard-to-cursor mapping results in suppression of their ability to generalize the learning to new target locations (Fermin et al, 2010). However, these participants were able to generalize to new target locations when provided sufficient time to prepare. This suggests that normal use of action-outcome mapping relies on a time-consuming computation, such as a model-based process. A model-based process could afford generalization by rolling out the consequences of various choices using learned action-outcome associations. However, imposing a time-limit on the motor controller may have forced reliance on a faster, but less flexible, model-free process.

We propose that skill learning is the product of an interaction between model-based planning, which provides rapid improvement, and low-level reinforcement-learning, which eventually captures the action-outcome mapping. A simulation of the outcome of potential outcomes leads to the selection of an action, but this simulation requires considerable computational recourses and processing time. Eventually, this cost can be reduced by forming a

cached representation of the action-outcome mapping through reinforcement. The action-outcome mapping can then be acted on to produce adaptation.

The strategic control and caching processes would proceed at the same time during learning. We would like to characterize the interaction between the two, which may be cooperative, competitive, or independent. The Cooperative prediction is that learning of the action-outcome mapping is improved by the strategic selection of correct actions, which are then reinforced. The Competitive prediction is that strategic planning blocks the reinforcement learning system, perhaps by reducing prediction errors. Finally, strategic planning and action-outcome caching may work in parallel without interacting.

To test these hypotheses, we needed to develop a model task for studying *de novo* learning. The mirror reversal task, as discussed in Chapter II, is arguably the most novel mapping available for study in a reaching paradigm, but subjects compensate for the mirror in less than one hour. There is no transformation of a reaching movement which doesn't have a relationship to the highly learned body control mapping established over decades of moving around in the natural environment. Given this, recent papers have suggested that there is no sensorimotor adaptation task which can probe *de novo* learning of an action-outcome mapping (Krakauer et al., 2019; Simpanto and Cisek 2020). However, all movement learning paradigms are on a continuous scale from the learned mapping to a completely *de novo* mapping.

We introduce a *de novo* learning paradigm which uses key-presses mapped onto directions (inspired by Fermin, et al., 2010). The key-press paradigm allows us to diminish the preconceived mappings between movement and outcome found in adaptation tasks. A straight line of letter keys (imagine using QWERTY to move six directions) has no readily available mapping to call upon. There is no intuitive reason to expect W to move up and T to the left. The exception is, perhaps,

that you may expect the two outermost keys, Q and Y in this example, to move left and right respectively. Nonetheless, there is very little in the way of an intuitive mapping to overcome, so we approach a task where learning is actually conducted *de novo*. Importantly, this task also allows for spatial rotation of the action-outcome mapping, so that we can integrate *de novo* learning with prior work on adaptive learning.

In Experiment 5.1, we establish that learning a novel action-outcome pairing in the key-press task is separate from sequence learning. This was continued in Experiment 5.2, where we show that the new action-outcome mapping can be used to pre-plan an action sequence. In Experiment 5.3, we test the three interaction hypotheses by directly manipulating access to computational time in a more difficult version of the task. We minimized use of model-based learning by severely restricting available reaction time to near the limit of a stimulus response (200 – 300 ms). Over twelve days of training, we periodically probed the action-outcome model to determine the effect of this intervention.

Experiment 5.1

In Experiment 5.1, we sought to answer the question: Is learning an action-outcome mapping obligatory? In this experiment, we used instruction to encourage a sequence-learning strategy in one group of subjects. All of the information was present to learn the action-outcome mapping, but the task could be completed using only the sequence. Learning was challenged with a generalization test, a skill which requires an action-outcome mapping and cannot be accomplished if only a sequence was learned. We found that the group given the sequence did not learn the action-outcome mapping as compared to a control group.

Materials and Methods

Participants

Fifty participants [32 female, age 36.6 ± 10 years] were recruited through the Amazon Mechanical Turk online platform. 41 subjects self-reported being right-handed. The experimental protocol was approved by the Princeton University Institutional Review Board and all subjects provided informed consent. Subjects received monetary compensation for their participation.

Apparatus

This was conducted virtually using subjects' personal computers and peripherals. Participants used the letter keys 'A-S-D' on their keyboard to manipulate a cursor around a 5x5 grid. The experiment was controlled by custom software coded in HTML, CSS, and JavaScript, and was hosted on Amazon Mechanical Turk.

Procedure

This experiment was separately replicated three times with minor changes (e.g. letter keys were changed from A-W-D to A-S-D), twice online and once in-lab. All replications produced the same results, which are represented here by data from the final replication.

After agreeing to participate in the experiments, subjects were presented with the following instruction and cover task:

The goal of the game is to land your helicopter safely on the landing pad in as few moves as possible. To play, use the A, W, and D keys to move the helicopter across the board. Be careful, those keys might not do what you think they should.

Don't take too long thinking! You have a short amount of time to make it to the landing pad before you run out of fuel.

Each trial began with the helicopter and target visible on a 5x5 grid (Figure 5.1a). Subjects were able to begin moving across the board immediately and had six seconds to reach the target with unlimited moves (following Fermin et al, 2010). If the time expired and the subject had not reached the target, the feedback “Try again! 0 points” was displayed. If the subject reached the target, the helicopter remained on the target location for the duration of the six seconds. If the target was hit with the minimal number of moves, the feedback read “Perfect! 100 points”. For every move over the minimum, 5 points were deducted from the perfect 100. Feedback for a trial in which one additional move was used to reach the target would read “There is a faster route. 95 points”. A running total of points appeared under the trial feedback. All feedback was displayed for 2s, followed by a 3s hold period. The hold period was used to provide additional instructions, which varied by condition and are described below.

Movement of the helicopter was determined by one of three key mappings. Each key mapping was chosen to be unintuitive, and require use of both diagonals and cardinal directions. An example mapping can be seen in Figure 5.1a

The experiment began with twenty training trials, each requiring the subject to move between a single start-end pair (the Trained pair). Trained pairs were randomly selected from eight start-end pairs that constituted a set for a particular key mapping. Each start-end pair forced the subject to use all three keys to reach the target in the minimum number of moves, and the minimum number of moves was always six. Following training, without break or other indication of a change, each subject completed a generalization block which included all eight pairs of the set. The seven untrained pairs (Generalization pairs) provided a measure of subjects' ability to use the

action-outcome mapping. If subjects were using a model-free strategy, they would have trouble accessing the action-outcome mapping instead of the sequence used in training. One instance of the Trained pair was included in the generalization block to control for block effects.

To bias learning towards a model-based (learning the action-outcome mapping) or model-free (sequence learning) strategy, we manipulated the instructions given to participants during training. The Given Sequence group (GS) was told an exact sequence of button-presses that would reach the target in the minimal number of moves. For example, the Trained pair seen in Figure 5.1a would be accompanied by the message: “The answer is: A-S-S-A-D-A.”. This message would appear before the start of the first trial, and during the hold period in-between every training trial. Participants in the No Instruction group (NI), were not given this additional information. Importantly, both the GS and NI groups had total control over the helicopter. Both groups had full access to the action-outcome mapping relationship through on-line feedback, but the GS group did not have to actively use this information once they’d memorized the sequence. The NI group had to find the appropriate sequence themselves before they could use a route memorization strategy. Sequences were not provided during the generalization block.

The entire experiment lasted less than five minutes.

Data and Analysis

Data was used from each subject that completed the experiment and pressed at least one button on at least 25% of trials. As is common on MTurk, many participants completed the experiment with minimal engagement. The measure of performance was a combination of ability to hit the target and number of key presses used. Hitting the target with the minimum number of moves (six) was recorded as 100. Each additional move over the minimum decreased the score by

5. If the target was not reached within the allotted time, the score was recorded as 0. This rating scheme was determined before data collection began.

Each subject contributed two values to the Analysis of variance, one for the Trained location trial and another for the Generalization trials. Post hoc t-tests were corrected using the Tukey-Kramer method. Data was analyzed using custom software coded in MATLAB 2019a (The MathWorks, Natick, Massachusetts).

Results

The key measure in this experiment is the ability to use the action-outcome mapping to move to new locations within the generalization block. We conducted a 2-way ANOVA with variables Condition (GS vs NI) x Trial Type (Trained pair vs Generalization pair). There were main effects of both Condition ($p=.02$) and Trial Type ($p<.001$). More important to our hypothesis, there was a significant Condition x Trial Type interaction ($p=.04$). Despite a disadvantage in training, the NI group was significantly better at reaching novel targets than the GS group (*post hoc t-test* $p=.01$).

Discussion

Even though both groups controlled the movement of the helicopter and saw the result of each key-press, the group given a solution did not learn the underlying key mapping. This suggests that learning the action-outcome mapping is not obligatory. Speculatively, this could mean that learning *de novo* is costly and avoided when possible.

There were two major potential confounds to this experiment, which we eliminated through control experiments. The first that we addressed was that the difference in learning could be due

to variability in training, and that the GS group would have been obligated to learn the action-outcome mapping had they received similar variability. We tested this by yoking a GS group to a NI group, such that the GS group no longer trained to the same start-end locations, but were given end locations and sequences that matched the corresponding NI subject's, often wondering and errant, path. This made no difference in the experimental conclusion.

Second, the GS group held the given sequence in working memory throughout the task, which may block learning of the action-outcome mapping. We controlled for this by giving the NI group a six-number sequence to memorize and repeat at the end of the experiment. Again, there was no effect of the control.




Figure 5.1. A) An example of a key-mapping is located above the grids. An example of a trained start/end pair is shown in grey and a generalization start/end pair in red. B) Performance for fifty subjects (25 in each condition) during training. Perfect performance would be reaching the target within the allotted time and with the minimal number of moves. Shaded area represents standard error. The Instruction group was given the solution (button-press sequence) during training. C) Performance on the last five training trials, the start/end location that was trained, and the novel, generalization, start/end locations.

Experiment 5.2

In Experiment 5.2, we validated that the action-outcome mapping learned in the button-press task could be used to plan a sequence of movements before initiating the action. This is an important role for action-outcome mappings, eventually providing the ability to conduct a skill based on intention, instead of selecting individual actions. For example, right now I'm typing this sentence by thinking of words and phrases, not each letter individually.

Materials and Methods

Participants

Twenty participants [10 female, age 19.4 ± 1.3 years] were recruited from the subject pool maintained by the Psychology department at Princeton University. Sixteen subjects self-reported being right handed and all reported having normal or corrected-to-normal vision. The experimental protocol was approved by the Princeton University Institutional Review Board and all subjects provided informed consent. Participants received course credit for completing the experiment.

Apparatus

Stimuli were displayed 60 Hz Dell monitor and computed by a Dell OptiPlex 7050'a machine (Dell, Round Rock, Texas) running Windows 10 (Microsoft Co., Redmond, Washington). Participants made their responses using a standard desktop keyboard. Subjects were allowed to adjust their chairs and monitor distance for comfort. The game was controlled by custom software coded and run in MATLAB 2016a (The MathWorks, Natick, Massachusetts), using Psychtoolbox extensions (Brainard, 1997, Kleiner, Brainard, and Pelli, 2007).

Procedure

In Experiment 5.2, we tested the usability of an action-outcome controller to plan and execute a sequence without cursor feedback. We hypothesized that this ability would depend on the complexity of the mapping and the length of the sequence.

Much like Experiment 5.1, participants were trained to move a cartoon helicopter across a 5x5 grid to target locations marked by a landing pad and the movement of the helicopter was determined by an unintuitive key-mapping. Each subject completed the full experiment once with a 3-key mapping and once with a more difficult 6-key mapping. The presentation of the 3- and 6-key mappings was counterbalanced across participants. When beginning the experiment, participants received the following instructions:

The goal of the game is to land your helicopter safely on the landing pad in as few moves as possible. To play, use the D, F, and J [W, E, R and U, I, O] keys to move the helicopter. Be careful! Those keys might not do what you think they should.

IMPORTANT: In "test mode", your helicopter will disappear as soon as you start moving. It is ok if you do not know how to land the helicopter. Even if you have no idea which direction you will move, make your best guess.

It is better to guess than to do nothing

During training, the additional degrees of freedom afforded by the 6-key mapping prevented us from using the grid boundaries to force a 6-move sequence (as we did in Experiment 5.1), so we blocked off three to six spaces on each trial. The blocked spaces were marked with an “X” and subjects could not move into those spaces. Strategic placement of these blocked spaces forced subjects to use each key at least once, and made the shortest path always six moves. Blocked spaces

were also used in the 3-key mapping for consistency. Like in Experiment 5.1, subjects trained with a single start-end location.

Test trials were interleaved throughout training. A single, no-feedback, generalization trial would be presented after every five training trials. Before each generalization trial, the subject was warned with the message “Test Alert”. The start position was always the center of the grid, and the end location was location pseudo-randomly chosen to be 1, 2, 3, or 4 moves from the starting position. This training-test pattern was repeated for 140 training trials, with subjects receiving one generalization trial of each length once every 20 training trials. Generalization trials did not have blocked spaces.

Data and Analysis

Key measures of this task were ability to hit the target (hit rate) and reaction time from target appearance until first button-press (RT) on generalization trials. The target was not hit when the time limit (6 s) expired or the maximum number of moves (15) was reached. We excluded trials in which the time limit expired before a move was recorded.

Results

We conducted two separate ANOVAs on hit rate and RT with Condition (3-key vs 6-key mapping) and Distance (1:4) as predictive variables. We found main effects of both Condition and Distance on hit rate ($p=.002$ and $p=.006$, respectively) and RT ($p=.009$ and $p=.01$, respectively). There were no significant interactions (all $p>.7$). Increasing the complexity of the mapping had a detrimental effect on planning, and there was a considerable disadvantage to having to plan more movements without cursor feedback (Figure 5.2).




Figure 5.2. Accuracy (A) and reaction time (B) on test trials. 3-key trials are graphed in blue and 6-key trials in red. Error bars represent standard error and circles represent individual subjects.

Discussion

In this experiment, we found that the action-outcome mapping was used to pre-plan a series of actions. This planning was dependent on the complexity of the mapping. The complexity may be seen as a proxy for skill, in that the simpler mapping could be learned more quickly. Additionally, there was a strong effect of how far in advance an action had to be planned. Perhaps analogous to my ability to plan out typing one sentence, but not a whole paragraph.

We conducted an additional experiment to control for the presence of test trials during training, which may bias learning toward an action-outcome mapping that could be used to meet the test trial demands. This follow-up experiment moved all of the test trials to the end of training, and provided the same results as discussed above.

Experiment 5.3

In Experiment 5.3 we test how the action-outcome mapping is learned, as well as the result of that learning. We increased the difficulty of the task, using eight degrees of freedom and increasing the size of the grid to allow for fifty-two training locations. This increased difficulty slowed the learning curve, allowing us to assess acquisition of the skill over a month of training. We directly inhibited use of a computationally intensive model-based strategy by limiting preparation time. Over the course of training, the action-outcome mapping was tested using generalization, no-feedback planning, speed / accuracy tradeoff, and rapid feedback control tests. We found that inhibiting model-based learning had a major impact on performance, but did not suppress learning of the action-outcome mapping by any of our measures. This provides indirect support for the hypothesis that model-based and model-free learning are working independently.

While we have assumed that the action-outcome mapping is the target of adaptation in an adaptation task, this has not been shown experimentally. At the end of training, we correlated the strength of individuals' learned action-outcome mapping with their ability to adapt to a visual perturbation. We found that the action-outcome mapping is strongly correlated with adaptive learning.

Materials and Methods

Participants

43 participants [27 female, age 24.5 ± 4 years] were recruited to participate in this experiment. Of these, 40 completed all 14 sessions (information on missing data in *Data and Statistical Analysis* below). All subjects self-reported being right handed and having normal or corrected-to-normal vision. The experimental protocol was approved by the Princeton

University Institutional Review Board and all subjects provided informed consent. Subjects received monetary compensation for their participation.

Apparatus

This month-long experiment was conducted virtually using subjects' personal computers and peripherals (25 on MacOS and 18 on a Windows operating system). Participants used the arrow and letter keys on their keyboard to manipulate a cursor around a 7x9 grid. The experiment was controlled by custom software coded in HTML, CSS, and JavaScript, and was hosted on Amazon Mechanical Turk.

Procedure

The 14-session experiment took place over four weeks, with 4 sessions in weeks 1-3 and 2 sessions in week 4. Three of the training sessions were followed by tests designed to probe acquisition of the action-outcome model. The final session assessed adaptability of the learned action-outcome model.

In order to assess the role of time pressure on learning, subjects were assigned to two independent groups. The High Pressure (HP) group learned to initiate their first movement extremely quickly after seeing the target, while also learning the novel key-mapping. This group trained in a range, determined by pilot testing, to cause accuracy that is just above chance after a single session of learning (forced reaction time of 200 – 300 ms). The Low Pressure (LP) group was given more time to move, and perhaps plan, on each trial (forced RT of 400 – 500 ms). The LP group was not given an extensive amount of preparation time – you could imagine allowing a

full second or removing the timing requirements altogether – because we didn’t want to excessively change the task demands between groups.

Session 1: Participants began the first day with a series of familiarization tasks, designed to introduce the general task structure and timing restrictions. Participants used the arrow keys to guide their cursor (cartoon helicopter) across a 7x9 grid to a goal location (landing pad, Figure 1). The arrow keys always moved in veridical directions (“up” goes up). When the cursor was moved into the target space, a pleasant ‘bling’ sounded, an American flag appeared in the target location, and a new trial started. In the first three trials, there was no timing requirement. After successfully hitting a target three times, participants were shown a cover story for the timing instructions.

“For each mission to be successful, you must take off at exactly the right time. Too early and you’ll be in the line of fire while your backup is still on the runway. Too late and command will demote you for missing the meetup and failing the team. In order to ensure that everyone is synchronized, command will provide a countdown.”

Participants were then given an opportunity to hear the countdown (“Three-Two-One-BEEP”) Each sound in the countdown occurred in 750 ms intervals and the BEEP lasted 160 ms, for a total hold time of 2410 ms. Participants were told “Your job is to anticipate the BEEP, and begin moving as the BEEP sounds off”. After this point, each trial was initiated by the count-down. Timing performance was defined by an RT tolerance window around the end of the hold time wherein a subject could move without penalty. A too-slow trial resulted in hearing a male voice say “demoted” and a red icon appearing in the location of the helicopter. A too-fast trial resulted in hearing the sound of an explosion while a cartoon explosion was displayed in the location of the helicopter (Figure 5.3b).




Figure 5.3: (A) The 7x9 grid on which the game is played. Each trial begins with the cursor on the center (white) square. The red spaces show an example set of reserved generalization spaces that are not used in training. Generalization spaces were always equally distributed between 2, 3, and 4 moves from the start location. The speed-accuracy test uses the 8 spaces immediately surrounding the start position (blue). (B) Trial progression. A new trial begins with the countdown, which allows the subject to synchronize their first movement with the go cue. The helicopter cursor appears with the start of the countdown. The forced RT is set by target appearance (e.g. the target appears 600 ms before the end of the countdown, forcing the subject to have an RT of 600 ms). The RT tolerance allows for minor variation in the timing of movement initiation, the game is impossible without this allowance. Too-fast and too-slow trials occur when movement is initiated before or after the RT tolerant window, respectively. A trial in which the subject initiates a movement too quickly is marked with the explosion icon replacing the helicopter, and the sound of an explosion. A too-slow trial is marked with the demoted icon and a male voice saying “demoted”. If the target is hit, a pleasant “bling” is played and a flag appears over the target. (C) Each subject is trained on a *de novo* mapping of the letter keys Q-W-E-R and U-I-O-P. Also shown is the clockwise rotation of the example mapping (rotated mappings were used in Session 13).

During the next ten trials, the target appeared 800 ms before the end of the countdown (800 ms forced RT) and each trial was followed by feedback displayed on the screen. If a participant was too-slow, they saw the in-trial feedback as described above and then “TOO SLOW! You missed the meet window and have been recalled”. If the participant moved too quickly, the message “TOO FAST! You left before your backup and your position was overrun” was displayed. The message “Nicely done” appeared if the participant moved on-time. During these trials, we purposefully made it extremely difficult to move on-time (± 100 ms RT tolerance from end of hold time) so that participants would receive feedback and learn what the different sounds and icons indicated for their performance.

Next, participants completed 50 additional introductory trials with the arrow keys, veridical feedback, forced RT of 800 ms, and an RT tolerance of ± 150 ms. These trials did not have written feedback. An inner-press-interval (IPI) restriction was introduced such that the HP group had to press successive buttons within 300 ms and the LP group within 400 ms. Additionally, we introduced a movement number restriction. Each trial was capped at 20 movements. If a subject reached this limit, they received the same feedback as if they’d been too-slow on the first movement. The IPI and movement number restrictions were held constant throughout the introductory and the following training trials. The purpose of these introductory trials was to familiarize subjects with the countdown and IPI timing restrictions before introducing novel key-mappings and more difficult forced RTs. A break period was allowed after the introductory trials.

The training trials were introduced by instructing participants that the controls had changed and that they would use Q-W-E-R with the right hand and U-I-O-P with the left hand to maneuver the helicopter. The arrow keys were disabled. Subjects were also told to prioritize timing over moving toward the target during the game. The structure of the task was the same as described

above, except that subjects were randomly assigned to one of four *de novo* mappings (Figure 5.3c). The helicopter now moved in eight directions with an unintuitive mapping between keys and movement. Throughout this manuscript, we will refer to this as the Trained Mapping. During training, the HP group trained with forced RTs of 200, 225, 250, 275, and 300 ms, while the LP group was given 400, 425, 450, 475, and 500 ms between target appearance and the end of the countdown.

After 300 training trials, subjects completed a series of test trials to probe the action-outcome model. Except for the Generalization Test, which occurred after training without acknowledgement, a break and instructions were given between each test.

The generalization trials tested subjects' ability to move to new spaces on the 7x9 grid. During training, 9 spaces were reserved for this purpose (Figure 5.3). The generalization spaces were chosen to be evenly divided between spaces that required 2, 3, and 4 movements. During the Generalization block, subjects moved to each trained location (52 spaces) once and each generalization location twice. Trained and generalization trials were intermixed, for a total of 70 trials. Forced RTs were set at 800 ms, with an RT tolerance of 1 s, and 5s IPI allowance.

Following the Generalization Test, subjects completed a Speed-Accuracy Tradeoff Test, the primary test of subjects' ability to use the action-outcome mapping. Targets in the Speed-Accuracy Tradeoff Test appeared in the 8 positions immediately surrounding the start location. Forced RTs were between 0 and 500 ms in 25ms steps, with an RT tolerance of 150 ms. No IPI was set as only a single movement was allowed on each trial and a new trial started following feedback for that one movement. Participants were informed before beginning the test that only one movement was allowed and that they should prioritize moving on-time over accuracy – even if that meant guessing. Each of the 21 RT times was presented 10 times, for a total of 210 trials.

Next, participants completed a No-Feedback Test to determine their ability to plan out a series of movements. They attempted to reach each of the 52-training locations without being able to see their cursor during intermediate steps. Once the first movement was initiated, the helicopter disappeared and remained hidden until the subject reached the target or ran out of moves. For this test, the number of movements was capped at 10. This reduced frustration as a subject was unlikely to make it to the target if they had not done so within 10 movements. Forced RTs were once again set at 800 ms, with an RT tolerance of 1 s, and 5s IPI allowance.

Finally, participants were tested on their rapid feedback control through a Target-Jump Test. On half of the trials in this test, the target ‘jumped’ to a new location after the first movement was initiated, requiring the subject to rapidly re-calculate any planned movements. The target jumped pseudorandomly to any space on the board that was at least 2 spaces from the original target and at least 2 spaces from the subject’s current position. These restrictions were imposed to prevent subjects from ‘accidentally’ reaching the target because it was in the previously planned path. There were 80 trials in this test, with 40 target-jump trials randomly distributed throughout. Forced RTs were set at 800 ms, with an RT tolerance of 1 s, and 5s IPI allowance.

Sessions 2 – 12: Subjects completed 300 training trials as described above. The beginning of each day began with a written reminder of the goal of the game and that subjects should prioritize timing. As in Session 1, the HP group trained with forced RTs 200:25:300, while the LP group was given 400:25:500 ms. The RT tolerance was 150 ms and IPIs were restricted to under 300 ms for the HP group and under 400 ms for the LP group. In **Session 6**, a Speed-Accuracy Tradeoff Test followed training. In **Session 12**, following training, subjects completed all of the test trials as in Session 1 (Generalization, No-Feedback, Speed-Accuracy, and Target-Jump).

Session 13: The goal of the last session was to determine if the *de novo* motor controller had been sufficiently learned to allow for adaptation. Half of subjects were given a rotated version of their Trained Mapping. The Rotated Mapping was modeled after a standard visuomotor rotation task, and so the movement direction of each key was shifted 45° (e.g. the key that moved UP in the Trained Mapping moves UP+RIGHT in the Rotated Mapping, Figure 5.3c). The direction, clockwise vs counterclockwise, of the rotation was counterbalanced across subjects. The other half of subjects were randomly assigned to a new Random Mapping, unrelated to the Trained Mapping. The assigned Random Mapping was one of the mappings used by other subjects in training. Subjects were informed that the key-mapping had changed, but not the nature of that change. Each subject completed 300 training trials with the new mapping. Then, a Speed-Accuracy Test was completed, followed by 50 more training trials as a refresh. The session ended with a 250-trial washout block, with the key-mapping returned to the Trained Mapping. Subjects were informed that the key-mapping was returned to the Trained Mapping. The washout block was intended to provide a measure the effect of Rotation vs Random on relearning.

Data and Analysis

Due to the nature of a multi-day experiment in which there is limited ability for the experimenters to assist participants with technical issues, there is data missing for several sessions/participants. All collected data was used regardless of missing data from that participant. Additionally, we used data from subjects that were late in completing sessions (Table 5.1).

Known reasons for missing data are listed below. Some data is missing without an identifiable cause (subjects were unable to troubleshoot conclusively).

1. A computer crashed due to another application initiating an update.
2. The browser was closed after the session was completed, but before the data was submitted.
3. Subject withdrew from the experiment.

Session *first of week	Number of Subjects	Time Between Sessions (days)
1*	22 HP / 21 LP	-
2	22 HP / 21 LP	1.7 ± 1.2
3	22 HP / 21 LP	1.4 ± 0.7
4	22 HP / 21 LP	1.8 ± 1.0
5*	22 HP / 21 LP	3.0 ± 1.6
6	22 HP / 20 LP	1.5 ± 0.9
7	22 HP / 21 LP	1.4 ± 0.7
8	22 HP / 21 LP	1.2 ± 0.4
9*	22 HP / 21 LP	2.9 ± 1.7
10	22 HP / 21 LP	1.6 ± 0.8
11	21 HP / 19 LP	1.3 ± 0.6
12	21 HP / 20 LP	1.5 ± 0.9
13*	18 HP / 20 LP	4.3 ± 2.1
14	19 HP / 21 LP	1.7 ± 1.3

Table 5.1: The second column contains the number of subjects that we received data from on each day. Note that some days have fewer subjects than the following, this is due to technical issues that prevented us from receiving the data for that day. The third column shows the average and standard deviation of the time between each session. As anticipated, there is much more variability between weeks than within.

On-time trials were defined as those in which the subject initiated their movement within the RT tolerance window. This measure was not dependent on subject's ability to hit the target after beginning their movement.

In training and test, we focused on three metrics for accuracy: ability to hit target, number of moves over the optimal (MoO), and reduction of distance to the target on the first move (RD). Target hits were defined simply as trials where the subject successfully navigated to the target. This inherently requires that the subject move within the RT tolerance window. On each trial, the number of moves over the optimal was found by taking the number of moves used to reach the target and subtracting the minimum number of moves. Trials where the subject ran out of moves were included in this measure as twenty (the maximal number of moves) minus path length. Trials in which the subject did not move on-time were excluded, as there wasn't an opportunity for the subject to use more than one movement. The ability to reduce the distance to the target was defined as any move that would allow the subject to hit the target in the minimum number of moves, regardless of whether or not they went on to do so. For targets that were greater than one step away from the starting position, there were two or three first moves that all could result in hitting the target using the minimum number of moves. These were all equally treated as a successful reduction of distance.

In the Generalization, No Feedback, and Target Jump tests: Reaction times (RT) were defined as the time elapsed from presentation of the target until the first button-press was initiated. In cleaning the data, we removed any RT above 1.5 times the tolerance window ($RT > 1.5\text{ s}$) and negative RTs (movements which were initiated before the target appeared, $<0.5\%$ total). Trials that were identified to be removed due to RT were also removed for all other metrics.

Finally, we included a *post hoc* measurement of participants' reliance on cardinal directions vs diagonals. The impetus for this analysis was post experiment surveys wherein several subjects indicated a strong preference for using only cardinal directions. For each trial, we took the number of moves made along the cardinal directions and divided by the total number of moves.

In analyses of variance (ANOVA) on training data, we averaged responses for each subject over Day 1 and Day 12. When analyzing data collected on Day 13, we used the first and last twenty-five trials. Each subject contributed one value to each bin of the ANOVA. *Post hoc* t-tests were corrected using the Tukey-Kramer method. Data was analyzed using custom software coded in MATLAB 2019a (The MathWorks, Natick, Massachusetts).

Results

Training

The training task had two major components, timing and movement. One goal of this experiment was to determine what aspects of the task subjects would prioritize in the absence of specific coaching from the experimenter. To this end, we considered three types of movement efficacy: proportional hit rate, number of moves over optimal (MoO), and reduction of distance to the target with the first move (RD; Figure 5.4a-c). Additionally, the ability to time movements correctly was measured as the proportion of trials where the subject made the first move within the RT tolerance window (Figure 5.4d). To quantify training performance, we conducted ANOVAs with the variables Condition (HP vs LP) and Time (average of the 1st day vs average of the 12th day). This ANOVA structure was used on each of the four outcome variables.




Figure 5.4. (A) Training timecourse for hit rate. Higher is better. (B) Number of moves over optimal (MoO). Lower is better. (C) Reducing distance to target on the first move (RD). Higher is better. (D) Timing timecourse showing the proportion of too-fast and too-slow trials, as well as the proportion of on time trials. All variables are shown in 50 trial bins. Shading represents one standard error of the mean.

The LP group moved to the target more efficiently than the HP group, shown in a main effect of Condition for all four measures: hit rate ($p=.003$), MoO ($p=.001$), RD ($p<.001$), and timing accuracy ($p=.01$). There was also a main effect of Time for improving hit rate ($p<.001$), decreasing MoO ($p<.001$), and honing timing ($p<.001$), but not ability to RD ($p=.06$). Condition x Time interactions for hit rate ($p=.002$) and timing ($p=.006$) show that the HP group was significantly worse performance on those measures on Day 1, but had caught up to the LP group by Day 12. The Condition x Time interaction for RD ($p=.009$) shows that the HP group did not learn to reduce the distance to the target on the first key-press, but the LP group did.

Speed/Accuracy Test

The speed/accuracy tradeoff test is our main measure of ability to use the action-outcome mapping (Figure 5.5). Subjects were presented with targets one move away from the start location and they had to select the single correct key that would move them into that space. Forced reaction times varied between zero and 500 ms. Moving at very low preparation times requires that the action-outcome mapping be computationally efficient. So, high accuracy at low speeds indicates a robust skill. Using a two-way ANOVA, Condition (HP vs LP) x Time (Day 1 vs Day 12), we see that there is no difference in slope between the Conditions ($p=.2$). Even though the groups appear to have prioritized different aspects of the task in training, their action-outcome mapping is equally effective within our power limitations. There is a main effect of time ($p<.001$), but not a Time x Condition interaction ($p=.77$). *Post hoc* t-tests show that the bulk of learning is concentrated between Day 1 and Day 6 (HP $t=3.1$, $p=.006$; LP $t=4.3$, $p<.001$). There was not significant learning between Day 6 and Day 12 (HP $t=1.8$, $p=.08$; LP $t=0.1$, $p=.91$).




Figure 5.5. (A) Speed/accuracy curves for each of the 3 test days in 50 ms bins. (B) Subtraction of speed/accuracy curves showing learning between days 1 and 6, but not days 6 and 12. Each subject contributed a single data point to each RT. Shading represents one standard error of the mean.

No Feedback Test

The No Feedback Test required subjects to move “in the dark”, providing a measure of their ability to plan out a route in advance of beginning the sequence. Accuracy was measured with hit rate and MoO (Figure 5.6a-b), and, as this test allowed variability in the RT tolerance ($\pm 1s$), we were able to use RT and inner-press-interval (IPI) as measurements of preparation time (Figure 5.6b-c). We conducted a separate ANOVA for each of the four outcome variables (hit rate, MoO, RT, and IPI) using Condition (HP vs LP), Time (Day 1 vs Day 12), and Distance to Target (1-4 moves away) as predictor variables. If suppressing computation time had an effect on the ability to use the action-outcome controller to plan an action sequence, we would expect a Condition x Trial Type interaction, which is not found here.

By Day 12, the LP group was much faster in initiating the movement sequence (RT), with no corresponding detriment to actually hitting the target or increase in the number of moves it took to do so. That is, there was a significant difference between the HP and LP conditions in RT ($p < .001$), but not a corresponding fluctuation in any of the other variables (hit rate $p = .50$; MoO $p = .28$; IPI $p = .08$). This suggests that the LP group was more efficient in planning an action sequence to reach the target. All variables showed a main effect of Time (all $p < .001$), such that performance was better on Day 12. There was also a main effect of Distance for hit rate, MoO, and IPI (all $p < .001$), but, unexpectedly, not RT ($p = .99$). There was a significant Condition x Time interaction ($p < .001$) on RT. The interaction shows that two Conditions had equivalent RT, a stand-in for planning time, on Day 1, but the LP group significantly improved their planning time over the HP group by Day 12.




Figure 5.6. Performance on the No Feedback test. Subjects moved without cursor feedback to each of the 52 training locations. (A) Percent of trials where the subject reached the target. Higher is better. (B) Number of moves taken over the optimal (MoO). Lower is better. (C) Reaction time from target appearance to first move. Lower is better. (D) Average time from one button press to the subsequent (IPI). Lower is better. For all, shading represents one standard error of the mean.

Target Jump Test

Target jump trials provide a measure of rapid feedback control, the ability to use online feedback to correct an errant path. In Jump trials, the target switched to a new grid location after the first move was initiated. We compare the responses immediately following the jump (moves 2-5), to the same moves of trials where the target did not change locations (Stay trials). In this test, we use two outcome variables, reduction of distance to the target (RD), as a measure of move-by-move accuracy, and IPI, as a measure of processing speed (Figure 5.7).

We analyzed the data using separate ANOVAs with predictive variables Condition (HP vs LP), Trial Type (Jump vs Stay), Move Number (2-5), and Time (Day 1 vs Day 12). There are main effects of all four independent variables to both measures. The HP group has higher IPI ($p < .001$) and is worse at RD ($p = .002$). Jump Trials require more processing time ($p < .001$) and have lower RD ($p < .001$) than stay trials. IPI decreases ($p < .001$) and RD first increases, then decreases with the number of moves ($p < .001$). Over training, IPI decreases ($p < .001$) and RD increases ($p < .001$). Additionally, both RD and IPI show Trial Type x Move Number interaction (both $p < .002$). RD is also dependent on a Trial Type x Time interaction ($p < .001$). The IPI ANOVA shows additional interactions between Condition x Time ($p < .001$), Move Number x Time ($p = .002$), and the three-way interaction Condition x Move Number x Time ($p = .01$).

If suppressing computation time had an effect on rapid feedback control, we would expect a Condition x Trial Type interaction, which was not found here. We did find that jumping the trial location caused significant interruptions, however, it is the timing of these interruptions that is interesting. It appears that the second key press, which occurs immediately after the target jump, is not where corrections are occurring. The IPI is low at this point the RD suffers on jump trials.

This suggests that the second key-press is locked in, such that the press occurs before a correction can be made. On the third key-press, the IPI slows as subjects correct for the jump, increasing RD.




Figure 5.7. Performance on the Target Jump test. Note that the x-axis on these graphs represent moves within a trial (key-presses), not target distance. (A and B) Reduction of distance to the target as a result of each move on Day 1 and Day 12. Higher is better. (C and D) Time between key presses on Day 1 and 12. Shading represents one standard error of the mean.

Generalization Test

The generalization test is a measure of subjects' ability to move to a new special location, that were not used on training (Figure 5.3a). Accuracy was measured with hit rate and MoO (Figure 5.8), and we used RT and IPI as measures of processing speed (Figure 5.9). We analyzed the data using separate ANOVAs with predictive variables Condition (HP vs LP), Trial Type (Generalization vs Trained), Target Distance (2-4), and Time (Day 1 vs Day 12). If the training condition held an effect on generalization, we would expect to see Condition x Trial Type interactions, which are not present here (all $p > .7$).

The HP group had a significantly lower hit rate ($p = .005$), used more moves over the optimal ($p < .001$), and was slower to initiate movement ($p < .001$), but was similar to the LP group in IPI ($p = .08$). The Trial type had no effect on performance (all $p > .4$). Distance to Target significantly increased the number of MoO used to reach the target ($p = .007$) and IPI ($p = .002$). Finally, training Time had a significant effect on all variables (all $p < .001$). The only interactions were Condition x Time in target hit rate ($p = .001$) and IPI ($p = .002$).

The HP group had overall worse performance than the LP group and moving to a farther target decreased performance. However, there was no detrimental effect of going to new locations over trained ones.




Figure 5.8. Hit rate and number of moves over optimal as a function of Condition, Trial Type, Target Distance, and Time. Shading represents one standard error of the mean.




Figure 5.9. Reaction time and inner press interval as a function of Condition, Trial Type, Target Distance, and Time. Shading represents one standard error of the mean.

Adaptation and De Novo Mappings

On Day 13 of the experiment, each subject was exposed to a new key-mapping: either a rotation of the Learned mapping or a new, random mapping. New mapping assignments were counterbalanced across HP and LP conditions. If the Trained action-outcome mapping was adaptable, and the Rotated mapping had an appropriate statistical relationship to allow adaptation, we expected to see an advantage for the Rotated mapping over the Random mapping. The Rotated mapping was not compared to the first day of exposure for the Trained mapping because of the many general task demands, such as timing, that were learned over the course of training. Having learned these domain general skills would artificially inflate the learning timecourses of the Rotated mapping over the Trained mapping. Thus, we used the Random mapping as a control.

Conducted a separate ANOVA for hit rate, RD, MoO, and timing accuracy with Condition (HP vs LP), Task Type (Random vs Rotation), and Time (First 25 Trials vs Last 25 Trials; Figure 5.10). There was no effect of Condition in any of these tests (all $p > .1$). Participants given the Rotation Task were better able to reduce distance to target on the first move ($p < .001$) and used fewer movements over the optimal ($p = .002$); hit rate ($p = .12$) and timing ($p = .06$) were not significantly affected by Task Type. Over training, performance generally improved with higher hit rate, increased RD, decreased MoO, and better timing accuracy (all $p < .001$). There were no significant interactions for any measure (all $p > .06$).




Figure 5.10. (A-D) Hit rate, RD, MoO, Timing Accuracy and Cardinal Directions as a function of Condition, Task Type, and Time. As there were no significant effects of HP vs LP, those Conditions were combined for visual clarity. Shading represents one standard error of the mean.

The RD measure should be interpreted with care, as it is confounded to favor the Rotated Mapping. As previously noted, there are many targets which have equivalent first moves; that is, two or more first moves which would allow the target to eventually be reached in the optimal number of moves. For example, say that the equivalent moves are directly to the right and up/right.

Moving to either of these spaces would constitute a success under the RD measure. In the Rotated mapping, you could mean to move directly to the right, but accidentally use the Trained mapping. This action would take you up/right under a counterclockwise rotation, which would be counted as a success. However, in the Random mapping, inadvertently using the Trained mapping is highly unlikely to result in moving to a space that is counted as correct under the RD measure. This is only true for training, *not* the Speed/Accuracy test, which has a unique solution to each trial.

Following training, all subjects participated in a Speed/Accuracy tradeoff test. This test is our main measure of ability to use the action-outcome mapping (Figure 5.11). We found a strong advantage for the Rotated mapping over the Random mapping. The slopes from this test were evaluated using an ANOVA with variables Condition (HP vs LP) and Task Type (Random vs Rotation). There was a main effect of Task Type ($p=.01$), but no effect of Condition ($p=.26$) or Task Type x Condition interaction ($p=.64$).




Figure 5.11. Speed/accuracy curves for Random and Rotated mappings in 50ms bins.

Action-outcome mapping strength predicts adaptability

One of our underlying assumptions when beginning this experiment was that the action-outcome mapping is the target of adaptation. However, this has not been shown experimentally. Traditional paradigms do not provide sufficient variability in the Learned mapping to use as a predictive factor for adaptation. However, this experiment does provide sufficient variability between subjects, and so we sought to use the action-outcome mapping on Day 12 to predict adaptation to a Rotated mapping. We used our cleanest measure of having learned the action-outcome mapping, the Speed/Accuracy test.

We *a priori* planned to combine the HP and LP groups for this analysis, as slitting our subjects into four groups (HP/LP x Rotation/ Random) wouldn't provide sufficient power to use a linear model. This was further justified as we found no significant differences between conditions on the Speed/Accuracy test. We used a linear model to test the relationship between the Speed/Accuracy test conducted on the Trained Mapping on Day 12 and on the Rotated mappings on Day 13 (Figure 5.12). We found a significant correlation between performance on Trained and Rotated mappings ($p < .001$), but not between the Trained and Random mappings ($p = .99$). This supports our hypothesis that a more established action-outcome mapping allows for better adaptation of that mapping, but does not assist with learning of a completely novel mapping.




Figure 5.12. Correlation plots with slope from Speed/Accuracy tradeoff test of Trained mapping on Day 12 predicting test slopes of (A) Rotated and (B) De Novo mappings. Solid line is the line of best fit. Dashed lines represent the confidence intervals. Subjects from the HP group are coded in red and subjects from the LP group are coded in blue.

Exploratory analyses

While debriefing subjects, several (3 HP, 5 LP) spontaneously shared that they choose to navigate using only cardinal directions, ignoring their ability to move diagonally. This was not an expected behavior and so was not part of our initial hypotheses. However, in retrospect, we believe this dimension-reduction is a normal response to a difficult task. We have consequently quantified the use of cardinal directions (Figure 5.13). In training, we find that cardinal directions are favored to diagonals from very early in training. There is a main effect of Condition ($p=.002$) and Time ($p=.002$), with no interaction ($p=.35$).

When a new mapping is introduced on Day 13, we see that the tendency for using cardinal directions is amplified in the Random group and suppressed in the Rotation group. An ANOVA on use of cardinal directions showed a main effect of Task Type ($p<.001$) and Time ($p<.001$), with no effect of Condition ($p=.3$). There were no significant interactions (all $p>.4$). This would suggest that the Rotation group “sticks” to the keys they used in training, the cardinal directions that are now diagonals, while the Random group finds the new cardinal directions. Both Random and Rotation groups are arriving at similar training performance while using very different strategies.




Figure 5.13. Proportional use of cardinal directions in (A) training and (B) Day 13 new mapping. Note that the two groups are HP vs LP in training, and Random vs Rotation on Day 13.




Figure 5.14. Percent of hit trials on the Day 13 speed accuracy test when (A) the solution was a cardinal direction or (B) a diagonal.

Discussion

In this experiment, we found a large effect of time pressure on performance in task. The HP group took about ten days to catch up to the LP group in ability to move on-time and hit the target. Additionally, the HP group was never as efficient as the LP group. As anticipated, model-based learning appears to bootstrap performance, allowing faster, more efficient learning. However, suppressing model-based learning does not appear to have affected the development of an action-outcome mapping. Both groups performed equally well on all of our test types and had equivalent ability to adapt the learned mapping. This suggests, assuming that we had sufficient power to detect differences that may exist, that model-based learning is running independent of model-free learning. Even though model-based learning was suppressed, model-free learning was able to work in the background to learn the action-outcome mapping.

This experiment provided a model of several ways that the key-press paradigm can be used to study learning *de novo*: preparation time manipulation, generalization, no feedback pre-planning, rapid feedback control, and adaptation. There is a wealth of hypotheses that could be tested using this method. While we are excited about the promising opportunities for research using this paradigm to study the progression of *de novo* learning, this iteration had several limitations that should be considered in future work. First, we could have simply missed the point at which the HP and LP groups differed. We did not have test blocks every day and so we don't know if there was a time, perhaps somewhere between days two and five, where the two groups diverged.

Additionally, we were not expecting subjects to stick on cardinal directions. We assumed that even if people began by using a limited subset of the available degrees of freedom, they would eventually get bored and explore the rest. However, we found that when left to their own devices,

people were inclined to stick with the first solution that they discovered and not learn new ones. Moving forward, if not specifically interested in this dimension-reduction strategy, we suggest a ‘hallway’ type design instead of the grid, this way there would only be one possible path to reach the target space. The experimenter could then force use of all of the supplied degrees of freedom. Additionally, this would allow for short trials, as we used in this experiment, or continuous movement designs.

Perhaps causing or exacerbating the use of only cardinal directions by suppressing boredom, many subjects admitted to listening to podcasts or watching television while completing the experiment. This is an unfortunate side-effect of an online experiment. It is unclear how, or if, forcing subjects to fully attend to the experiment would change performance.

Chapter VI

Conclusion

The overarching goal of this dissertation was to expand our knowledge of the mechanisms underlying motor skill learning. We began by investigating the interaction between implicit adaptation and explicit strategies. Implicit adaptation adjusts the accuracy of learned skills when a perturbation is applied. The process reacts proportionately to small perturbations, systematically adapting the skill to the new environment (Chapter II). While slow, gradual adaptation is sufficient for most learned tasks, sometimes a leap in learning is necessary to accommodate a radical environmental change. Explicit strategies, in this case re-aiming, are generally deployed whenever the perturbation is large enough to be noticed by the learner. However, explicit re-aiming is voluntary and learners do not employ strategies when they are not useful. In contrast, implicit adaptation proceeds even when counterproductive (Chapter III). Thus, strategies are particularly important when the adjustments made by implicit recalibration are inappropriate for the situation, such as in the mirror-reversal task. The mirror-reversal task has been thought to require a second mechanism to develop a new *de novo* skills (Chapters III and IV). *De novo* skills require development of a new action-outcome mapping, the specification of what action to take in order to achieve a particular goal or series of goals. To further investigate, we developed a model task which isolates *de novo* skill learning (Chapter V). These experiments provided evidence that new skills are learned through a reinforcement learning-like process.

Based on the evidence presented in this manuscript, I would argue that there are two processes. One is responsible for building the skill *de novo* (DN), the other maintains that skill once built (recalibration). A main confusion between these systems stems from “adaptation” tasks,

which have been considered synonymously with recalibration. However, they can produce both recalibration and *de novo* learning (Chapter III).

We see this combination of recalibration and *de novo* learning in the two visuomotor tasks used in Chapters III – IV. The goal was to hit a target along the perimeter of a circle using a center-out reaching movements. These tasks manipulate the established skill, and its associated action-outcome mapping, of moving the hand in a straight line. By applying perturbations to visual feedback, we can simulate motor system drift and environmental changes that lead to movement errors. These perturbations produce adaptive learning. Implicit adaptation, which is most effective at compensating for small error sizes, and explicit re-aiming, which is used to correct a wide range of errors, work together to produce the appropriate response for any particular visual error. Specifically, implicit adaptation allows for small updates to the internal model, necessary to avoid accumulation of small errors and to avoid drift in the system, and explicit strategies are available for fast learning of large large perturbations. This suggests an important role for both strategies and implicit recalibration in adaptive learning.

However, the adaptation tasks should not be considered pure measures of recalibration of a learned skill. The rules of the hand-reaching action-outcome mapping require that feedback follows the hand. That is, moving clockwise brings the cursor clockwise even if the feedback is trailing or leading the hand. If this action-outcome mapping is maintained and the perturbation is small, adaptation is able to fully compensate and restore performance. We see DN learning, inability of recalibration to restore performance, when the perturbation is large or breaks the rules of the established action-outcome mapping. The clearest version of this break is when the mirror-reversal task deploys feedback in the opposite direction of the hand movement. This new action-outcome mapping requires DN learning, and we previously showed that the mirror-reversal task

must be solved using a strategy. Importantly, this implies that any DN learning taking place in the mirror-reversal task must be contained within the cognitive strategy, as the cognitive strategy is defined as all of the learning that is not implicit recalibration.

Because recalibration is inappropriate in the mirror-reversal task, it has been considered a *de novo* learning task. However, the absence or inappropriate operation of recalibration does not necessarily mean that the learning is DN in a general sense. In the mirror-reversal, up is still up and down is still down. To better isolate the DN learning process, we created a task with an action-outcome mapping that is fully arbitrary, the key-press experiments introduced in Chapter V. The goal of the task was to learn how to move in 3-8 arbitrary directions using 3-8 arbitrary keyboard keys. Unlike the mirror-reversal task, where the hand and arm mechanics are still inextricably tied to the direction of the cursor, the key-press task has no *a priori* mapping. While you, obviously, know where your fingers are located on the keyboard, there is no particular reason to believe that the e-key will move your cursor left instead of right. Consequently, adaptation in the key-press task is theoretically impossible until an action-outcome mapping has been learned. If the logic of the task is true, and there is no *a priori* mapping, any learning must be DN learning until some action-outcome mapping is established. Mechanisms that produce learning in the key-press task must be processes that support DN learning.

In the key-press experiments, we used model-based and model-free learning as a framework for thinking about skill learning. Of particular interest, we found strategy-use that coincided with situations conducive to model-based learning. Model-based learning requires computational time that is dependent on complexity, the number of possible choices to choose from and the range of situations in which those choices must be made. In the key-press task, providing additional time for computation produces faster learning, presumably as a result of

model-based-like processing. This additional time also produces a dimension-reduction strategy. Learners use the extra time to reduce their possibility space from eight directions to four, preferentially using the four cardinal directions more than learners that were not given additional computational time.

While DN learning holds characteristic properties of both cognitive strategy use and model-based learning, it is unclear from this data how they relate to each other. For example, a strategy could be a type of model-based learning, or it could be an entirely separate computation. Perhaps this computation reduces cognitive load and/or provides space for use of working memory, which then coincidentally produces an environment favorable to model-based learning.

A first suggestion is that a strategy is a very coarse example of model-based learning. Both model-based learning and strategic compensation are hypothesis-driven approaches “I think X will happen if I do Y”. Contrast to experience-driven model-free learning “last time I did X, Y happened”. The model-based / strategy learner makes a hypothesis about the action-outcome mapping, and simulates the consequences of actions taken under that hypothesis. When the action-outcome mapping is learned, these hypothesis making devices should always be correct, but are no longer necessary. Additionally, both cognitive strategies and model-based learning are generalizable computations. A mirror-reversal strategy can be transferred to a new section of a workspace and model-based learning allows for new routes through a reinforcement task. In the key-press experiment, we see that the strategy of using cardinal directions is carried through the generalization and no-feedback tests.

One issue with this hypothesis is: if a strategy is a kind of model-based learning process, what does it mean when we say that a strategy has been “cached” as discussed in Chapter IV? There, we suggested that a strategy itself may be cached in such a way that the computation is

skipped. The difference between this strategy caching and normal model-free learning, the that the strategy would be generalizable and not necessarily reliant on previous experience with individual targets. It is the strategy, not the action, that is the target of the model-free learning. However, if a strategy is an instantiation of model-based learning, as suggested above, it cannot itself be the target of model-free learning.

One solution is that the ‘caching’ of the strategy is the point at which the strategy is replaced by the action-outcome mapping. Under this hypothesis, the strategy has considerable control over the shape of the action-outcome mapping because it has been supplying the inputs to the model-free process. This would be a modified version of the interaction hypothesis proposed in Chapter V, where learning of the action-outcome mapping is improved by the strategic selection of correct actions which can be reinforced. This interaction between strategy and model-free learning jump-starts DN learning by narrowing the action space. Model-free learning is able to operate more quickly because it is fed correct actions to reinforce. In this case, it would appear as though the strategy was still being used because the action-outcome mapping was scaffolded by that strategy. The reason that the parametric computation can be skipped is not because the strategy itself is cached, but because an action-outcome mapping has replaced it.

In discussing the use of strategies in these tasks, I have ignored the fact that dimensionality reduction and simple rule-based strategies are relatively simple in comparison to the strategies that may be developed in more naturalistic settings. These reductions could have an illusory effect, making the skill learning process look more like reinforcement learning than is accurate. Future research may explore the range of explicit strategies available to learners, and determine if more complicated strategies hold the same parallels to reinforcement learning.

It is still unclear if there is significant and useful overlap between reinforcement learning and *de novo* skill learning. However, we are moving closer to understanding these phenomena with the development of new model tasks that better define and isolate the processes involved.

REFERENCES

- Abeelee, S., & Bock, O. (2001). Sensorimotor adaptation to rotated visual input: different mechanisms for small versus large rotations. *Experimental Brain Research*, 140(4), 407-410.
- Adams, C. D. (1982). Variations in the sensitivity of instrumental responding to reinforcer devaluation. *The Quarterly Journal of Experimental Psychology Section B*, 34(2b), 77-98.
- Baddeley, R. J., Ingram, H. A., & Miall, R. C. (2003). System identification applied to a visuomotor task: near-optimal human performance in a noisy changing task. *Journal of Neuroscience*, 23(7), 3066-3075.
- Benson, B. L., Anguera, J. A., & Seidler, R. D. (2011). A spatial explicit strategy reduces error but interferes with sensorimotor adaptation. *Journal of neurophysiology*, 105(6), 2843-2851.
- Bock, O., Abeelee, S., & Eversheim, U. (2003). Human adaptation to rotated vision: interplay of a continuous and a discrete process. *Experimental brain research*, 152(4), 528-532.
- Bond, K. M., & Taylor, J. A. (2015). Flexible explicit but rigid implicit learning in a visuomotor adaptation task. *Journal of neurophysiology*, 113(10), 3836-3849.
- Bond, K. M., & Taylor, J. A. (2017). Structural Learning in a Visuomotor Adaptation Task Is Explicitly Accessible. *eNeuro*, 4(4), ENEURO-0122.
- Brainard, D. H., & Vision, S. (1997). The psychophysics toolbox. *Spatial vision*, 10, 433-436.
- Brashers-Krug, T., Shadmehr, R., & Bizzi, E. (1996). Consolidation in human motor memory. *Nature*, 382(6588), 252.
- Braun, D. A., Aertsen, A., Wolpert, D. M., & Mehring, C. (2009). Motor task variation induces structural learning. *Current Biology*, 19(4), 352-357.

- Brudner, S. N., Kethidi, N., Graeupner, D., Ivry, R. B., & Taylor, J. A. (2016). Delayed feedback during sensorimotor learning selectively disrupts adaptation but not strategy use. *Journal of neurophysiology*, 115(3), 1499-1511.
- Butcher, P. A., & Taylor, J. A. (2018). Decomposition of a sensory prediction error signal for visuomotor adaptation. *Journal of experimental psychology: human perception and performance*, 44(2), 176.
- Butcher, P. A., Ivry, R. B., Kuo, S. H., Rydz, D., Krakauer, J. W., & Taylor, J. A. (2017). The cerebellum does more than sensory prediction error-based learning in sensorimotor adaptation tasks. *Journal of neurophysiology*, 118(3), 1622-1636.
- Caithness, G., Osu, R., Bays, P., Chase, H., Klassen, J., Kawato, M., ... & Flanagan, J. R. (2004). Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. *Journal of Neuroscience*, 24(40), 8662-8671.
- Castro, L. N. G., Hadjiosif, A. M., Hemphill, M. A., & Smith, M. A. (2014). Environmental consistency determines the rate of motor adaptation. *Current Biology*, 24(10), 1050-1061.
- Cheng, S., & Sabes, P. N. (2006). Modeling sensorimotor learning with linear dynamical systems. *Neural computation*, 18(4), 760-793.
- Choi, Y., Shin, E. Y., & Kim, S. (2020). Spatiotemporal dissociation of fMRI activity in the caudate nucleus underlies human de novo motor skill learning. *Proceedings of the National Academy of Sciences*, 117(38), 23886-23897.
- Choi, Y., Shin, E. Y., & Kim, S. (2020). Spatiotemporal dissociation of fMRI activity in the caudate nucleus underlies human de novo motor skill learning. *Proceedings of the National Academy of Sciences*, 117(38), 23886-23897.

- Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: a parallel distributed processing account of the Stroop effect. *Psychological review*, 97(3), 332.
- Corkin, S. (1968). Acquisition of motor skill after bilateral medial temporal-lobe excision. *Neuropsychologia*, 6(3), 255-265.
- Criscimagna-Hemminger, S. E., & Shadmehr, R. (2008). Consolidation patterns of human motor memory. *Journal of Neuroscience*, 28(39), 9610-9618.
- Crossman, E. R. F. W. (1959). A theory of the acquisition of speed-skill. *Ergonomics*, 2(2), 153-166.
- Cunningham, H. A. (1989). Aiming error under transformed spatial mappings suggests a structure for visual-motor maps. *Journal of experimental psychology: Human perception and performance*, 15(3), 493.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, 8(12), 1704.
- Day, K. A., Roemmich, R. T., Taylor, J. A., & Bastian, A. J. (2016). Visuomotor learning generalizes around the intended movement. *eNeuro*, 3(2).
- Dayan, P. (2009). Dopamine, reinforcement learning, and addiction. *Pharmacopsychiatry*, 42(S 01), S56-S65.
- de Brouwer, A. J., Albaghdadi, M., Flanagan, J. R., & Gallivan, J. P. (2018). Using gaze behavior to parcellate the explicit and implicit contributions to visuomotor learning. *Journal of neurophysiology*, 120(4), 1602-1615.
- Dickinson, A. (1985). Actions and habits: the development of behavioural autonomy. *Philosophical Transactions of the Royal Society B*, 308(1135), 67-78.

- Fermin, A., Yoshida, T., Ito, M., Yoshimoto, J., & Doya, K. (2010). Evidence for model-based action planning in a sequential finger movement task. *Journal of motor behavior*, 42(6), 371-379.
- Fernandez-Ruiz, J., Wong, W., Armstrong, I. T., & Flanagan, J. R. (2011). Relation between reaction time and reach errors during visuomotor adaptation. *Behavioural brain research*, 219(1), 8-14.
- Fine, M. S., & Thoroughman, K. A. (2006). Motor adaptation to single force pulses: sensitive to direction but insensitive to within-movement pulse placement and magnitude. *Journal of neurophysiology*, 96(2), 710-720.
- Fine, M. S., & Thoroughman, K. A. (2007). Trial-by-trial transformation of error into sensorimotor adaptation changes with environmental dynamics. *Journal of neurophysiology*, 98(3), 1392-1404.
- Fitts, P. M., & Posner, M. I. (1967). Human performance. Wadsworth Publishing Co. Belmont, CA.
- Ghilardi, M. F., Gordon, J., & Ghez, C. (1995). Learning a visuomotor transformation in a local area of work space produces directional biases in other areas. *Journal of Neurophysiology*, 73, 2535–2539.
- Gritsenko, V., & Kalaska, J. F. (2010). Rapid online correction is selectively suppressed during movement with a visuomotor transformation. *Journal of neurophysiology*, 104(6), 3084-3104.
- Hadjiosif, A. M., Krakauer, J. W., & Haith, A. M. (2020). Implicit adaptation is driven by direct policy updates rather than forward-model-based learning. *bioRxiv*.

- Haith, A. M., & Krakauer, J. W. (2013). Model-based and model-free mechanisms of human motor learning. *In Progress in motor control* (pp. 1-21). Springer, New York, NY.
- Haith, A. M., Huberdeau, D. M., & Krakauer, J. W. (2015). The influence of movement preparation time on the expression of visuomotor learning and savings. *Journal of neuroscience*, 35(13), 5109-5117.
- Haith, A. M., Pakpoor, J., & Krakauer, J. W. (2016). Independence of movement preparation and movement initiation. *Journal of Neuroscience*, 36(10), 3007-3015.
- Hardwick, R. M., Forrence, A. D., Krakauer, J. W., & Haith, A. M. (2019). Time-dependent competition between goal-directed and habitual response preparation. *Nature Human Behaviour*, 3(12), 1252-1262.
- Hegele, M., & Heuer, H. (2010). Implicit and explicit components of dual adaptation to visuomotor rotations. *Consciousness and cognition*, 19(4), 906-917.
- Held, R., & Schlank, M. (1959). Adaptation to disarranged eye-hand coördination in the distance-dimension. *The American journal of psychology*, 603-605.
- Helmholtz, H. L. F. (1867). *Handbuch der Physiologischen Optik* (published as "Helmholtz's treatise on physiological optics, Translated from the Third German Edition," ed. James PC Southall; 1925). Washington, DC: The Optical Society of America.
- Herzfeld, D. J., Vaswani, P. A., Marko, M. K., & Shadmehr, R. (2014). A memory of errors in sensorimotor learning. *Science*, 345(6202), 1349-1353.
- Heuer, H., & Hegele, M. (2008). Adaptation to visuomotor rotations in younger and older adults. *Psychology and aging*, 23(1), 190.
- Heuer, H., & Hegele, M. (2008). Constraints on visuo-motor adaptation depend on the type of visual feedback during practice. *Experimental Brain Research*, 185(1), 101-110.

- Holland, P. J., Codol, O., & Galea, J. M. (2018). The contribution of explicit processes to reinforcement-based motor learning. *Journal of neurophysiology*.
- Hollerbach, J. M., & Atkeson, C. G. (1987). Deducing planning variables from experimental arm trajectories: Pitfalls and possibilities. *Biological Cybernetics*, 56(5-6), 279-292.
- Howard, I. S., Ingram, J. N., Franklin, D. W., & Wolpert, D. M. (2012). Gone in 0.6 seconds: the encoding of motor memories depends on recent sensorimotor states. *Journal of Neuroscience*, 32(37), 12756-12768.
- Huberdeau, D. M., Krakauer, J. W., & Haith, A. M. (2019). Practice induces a qualitative change in the memory representation for visuomotor learning. *Journal of neurophysiology*, 2019.
- Hutter, S. A., & Taylor, J. A. (2018). Relative sensitivity of explicit re-aiming and implicit motor adaptation. *Journal of neurophysiology*. 120(5): 2640-2648.
- Imamizu, H., Uno, Y., & Kawato, M. (1995). Internal representations of the motor apparatus: implications from generalization in visuomotor learning. *Journal of Experimental Psychology: Human Perception and Performance*, 21(5), 1174.
- Jiang, W., Yuan, X., Yin, C., & Wei, K. (2018). Visuomotor learning is dependent on direction-specific error saliency. *Journal of neurophysiology*, 120(1), 162-170.
- Jordan, M. I., & Rumelhart, D. E. (1992). Forward models: Supervised learning with a distal teacher. *Cognitive science*, 16(3), 307-354.
- Kasuga, S., Kurata, M., Liu, M., & Ushiba, J. (2015). Alteration of a motor learning rule under mirror-reversal transformation does not depend on the amplitude of visual error. *Neuroscience research*, 94, 62-69.
- Killcross, S., & Coutureau, E. (2003). Coordination of actions and habits in the medial prefrontal cortex of rats. *Cerebral cortex*, 13(4), 400-408.

- Kim, H. E., Morehead, J. R., Parvin, D. E., Moazzezi, R., & Ivry, R. B. (2018). Invariant errors reveal limitations in motor correction rather than constraints on error sensitivity. *Communications Biology*, 1(1), 19.
- Kim, H. E., Parvin, D. E., & Ivry, R. B. (2019). The influence of task outcome on implicit motor learning. *eLife*, 8, e39882.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, 36(14), 1.
- Kohler, I. (1941). Der Einfluß der Erfahrung in der optischen Wahrnehmung. Beleuchtet von Versuchen langdauernden Tragens bildverzerrender Prismen. Phil (Doctoral dissertation, Dissertation, Innsbruck).
- Kohler, I. (1951a). U"ber Aufbau und Wandlungen der Wahrnehmungswelt. Insbesondere u"ber 'bedingte Empfindungen'. Wien: Rohrer: Osterreichische Akademie der € Wissenschaften, Philosophisch-historische Klasse. Sitzungsberichte 227.
- Kohler, I. (1951b). Warum sehen wir aufrecht - obwohl die Bilder im Inneren des Auges verkehrt stehen? *Die Pyramide*, 28e33.
- Kosslyn, S. M., Ball, T. M. & Reiser, B. J. (1978) Visual images preserve metric spatial information: evidence from studies of image scanning. *Journal of Experimental Psychology Human Perception and Performance* 4, 47–60.
- Krakauer, J. W. (2009). Motor learning and consolidation: the case of visuomotor rotation. In *Progress in motor control* (pp. 405-421). Springer, Boston, MA.
- Krakauer, J. W., & Shadmehr, R. (2006). Consolidation of motor memory. *Trends in neurosciences*, 29(1), 58-64.

- Krakauer, J. W., Hadjiosif, A. M., Xu, J., Wong, A. L., & Haith, A. M. (2019). Motor learning. *Comprehensive Physiology*, 9, 613-663.
- Lillicrap, T. P., Moreno-Briseño, P., Diaz, R., Tweed, D. B., Troje, N. F., & Fernandez-Ruiz, J. (2013). Adapting to inversion of the visual field: a new twist on an old problem. *Experimental brain research*, 228(3), 327-339.
- Liu, X., & Scheidt, R. A. (2008). Contributions of online visual feedback to the learning and generalization of novel finger coordination patterns. *Journal of neurophysiology*, 99(5), 2546-2557.
- Logan, G. D. (1980). Attention and automaticity in Stroop and priming tasks: Theory and data. *Cognitive psychology*, 12(4), 523-553.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological review*, 95(4), 492.
- Marko, M. K., Haith, A. M., Harran, M. D., & Shadmehr, R. (2012). Sensitivity to prediction error in reach adaptation. *Journal of neurophysiology*, 108(6), 1752-1763.
- Martin, T. A., Keating, J. G., Goodkin, H. P., Bastian, A. J., & Thach, W. T. (1996). Throwing while looking through prisms: II. Specificity and storage of multiple gaze—throw calibrations. *Brain*, 119(4), 1199-1211.
- Mazzoni, P., & Krakauer, J. W. (2006). An implicit plan overrides an explicit strategy during visuomotor adaptation. *Journal of neuroscience*, 26(14), 3642-3645.
- McDougle, S. D., & Taylor, J. A. (2019). Dissociable cognitive strategies for sensorimotor learning. *Nature communications*, 10(1), 40.

- McDougle, S. D., Bond, K. M., & Taylor, J. A. (2015). Explicit and implicit processes constitute the fast and slow processes of sensorimotor learning. *Journal of Neuroscience*, 35(26), 9568-9579.
- McDougle, S. D., Bond, K. M., & Taylor, J. A. (2017). Implications of plan-based generalization in sensorimotor adaptation. *Journal of neurophysiology*, 118(1), 383-393.
- Miall, R. C., & Wolpert, D. M. (1996). Forward models for physiological motor control. *Neural networks*, 9(8), 1265-1279.
- Miller, G. A. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review* 63, 81–97 (1956).
- Morehead, J. R., Taylor, J. A., Parvin, D. E., & Ivry, R. B. (2017). Characteristics of implicit sensorimotor adaptation revealed by task-irrelevant clamped feedback. *Journal of cognitive neuroscience*, 29(6), 1061-1074.
- Morehead, J.R., Smith, M. (2017, November). The magnitude of implicit sensorimotor adaptation is limited by continuous forgetting. Abstract. Advances in Motor Learning & Motor Control, Washington DC.
- Mosier, K. M., Scheidt, R. A., Acosta, S., & Mussa-Ivaldi, F. A. (2005). Remapping hand movements in a novel geometrical environment. *Journal of neurophysiology*, 94(6), 4362-4372.
- Oldfield, R. C. (1971). The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia*, 9(1), 97-113.
- Pine, Z. M., & Krakauer, J. W. (1996). Learning of scaling. *Neuroreport*, 7, 2357-2361.
- Pine, Z. M., Krakauer, J. W., Gordon, J., & Claude, G. C. (1996). Learning of scaling factors and reference axes for reaching. *NeuroReport*, 7, 2357-2361.

- Poh, E., Carroll, T. J., & Taylor, J. A. (2016). Effect of coordinate frame compatibility on the transfer of implicit and explicit learning across limbs. *Journal of neurophysiology*, 116(3), 1239-1249.
- Pouget, A., & Snyder, L. H. (2000). Computational approaches to sensorimotor transformations. *Nature neuroscience*, 3(11s), 1192.
- Provost, A., Johnson, B., Karayanidis, F., Brown, S. D., & Heathcote, A. (2013). Two routes to expertise in mental rotation. *Cognitive science*, 37(7), 1321-1342.
- Ranganathan, R., Wieser, J., Mosier, K. M., Mussa-Ivaldi, F. A., & Scheidt, R. A. (2014). Learning redundant motor tasks with and without overlapping dimensions: facilitation and interference effects. *Journal of Neuroscience*, 34(24), 8289-8299.
- Reichenthal, M., Avraham, G., Karniel, A., & Shmuelof, L. (2016). Target size matters: target errors contribute to the generalization of implicit visuomotor learning. *Journal of neurophysiology*, 116(2), 411-424.
- S. M. Radhakrishnan, S. N. Baker, A. Jackson, Learning a novel myoelectric-controlled interface task. *Journal of Neurophysiology*, 100, 2397–2408 (2008)
- S. M. Radhakrishnan, S. N. Baker, A. Jackson, Learning a novel myoelectric-controlled interface task. *Journal of Neurophysiology*, 100, 2397–2408 (2008)
- Sachse, P., Beermann, U., Martini, M., Maran, T., Domeier, M., & Furtner, M. R. (2017). “The world is upside down”—The Innsbruck goggle experiments of Theodor Erismann (1883–1961) and Ivo Kohler (1915–1985). *Cortex*, 92, 222-232.
- Saijo, N., & Gomi, H. (2010). Multiple motor learning strategies in visuomotor rotation. *PLoS One*, 5(2), e9399.

- Schaefer, S. Y., Shelly, I. L., & Thoroughman, K. A. (2012). Beside the point: motor adaptation without feedback-based error correction in task-irrelevant conditions. *Journal of Neurophysiology*, 107(4), 1247-1256.
- Scheidt, R. A., Dingwell, J. B., & Mussa-Ivaldi, F. A. (2001). Learning to move amid uncertainty. *Journal of neurophysiology*, 86(2), 971-985.
- Schween, R., McDougale, S. D., Hegele, M., & Taylor, J. A. (2019). Explicit strategies in force field adaptation. *BioRxiv*, 694430.
- Semrau, J. A., Daitch, A. L., & Thoroughman, K. A. (2012). Environmental experience within and across testing days determines the strength of human visuomotor adaptation. *Experimental brain research*, 216(3), 409-418.
- Shadmehr, R., & Brashers-Krug, T. (1997). Functional stages in the formation of human long-term motor memory. *Journal of Neuroscience*, 17(1), 409-419.
- Shepard, R. N. & Metzler, J. (1971) Mental rotation of three-dimensional objects. *Science* 171, 701–703.
- Smith, M. A., Ghazizadeh, A., & Shadmehr, R. (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS biology*, 4(6), e179.
- Stratton, G. M. (1896). Some preliminary experiments on vision without inversion of the retinal image. *Psychological review*, 3(6), 611.
- Stratton, G. M. (1897). Vision without inversion of the retinal image. *Psychological review*, 4(4), 341.
- Taylor, J. A., & Ivry, R. B. (2011). Flexible cognitive strategies during motor learning. *PLoS computational biology*, 7(3), e1001096.

- Taylor, J. A., Krakauer, J. W., & Ivry, R. B. (2014). Explicit and implicit contributions to learning in a sensorimotor adaptation task. *Journal of Neuroscience*, 34(8), 3023-3032.
- Telgen, S., Parvin, D., & Diedrichsen, J. (2014). Mirror reversal and visual rotation are learned and consolidated via separate mechanisms: recalibrating or learning de novo?. *Journal of Neuroscience*, 34(41), 13768-13779.
- Thoroughman, K. A., & Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature*, 407(6805), 742.
- Thoroughman, K. A., & Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature*, 407(6805), 742.
- Wei, K., & Körding, K. (2009). Relevance of error: what drives motor adaptation?. *Journal of neurophysiology*, 101(2), 655-664.
- Wei, K., & Körding, K. (2010). Uncertainty of feedback and state estimation determines the speed of motor adaptation. *Frontiers in computational neuroscience*, 4, 11.
- Werner, S., & Bock, O. (2010). Mechanisms for visuomotor adaptation to left–right reversed vision. *Human movement science*, 29(2), 172-178.
- Wilterson, S.A., & Taylor, J.A. (2019). Implicit visuomotor adaptation remains limited after several days of training. *BioRxiv*: <https://doi.org/10.1101/711598>

PRIOR PUBLICATION

Portions of this dissertation have previously been published in:

Hutter, S. A., & Taylor, J. A. (2018). Relative sensitivity of explicit re-aiming and implicit motor adaptation. *Journal of neurophysiology*. 120(5): 2640-2648.

Wilterson, S.A., & Taylor, J.A. (2019). Implicit visuomotor adaptation remains limited after several days of training. *BioRxiv*: <https://doi.org/10.1101/711598>