Implicit visuomotor adaptation remains limited after several days of training

Sarah A. Wilterson and Jordan A. Taylor

Department of Psychology
Princeton Neuroscience Institute
Princeton University

Abstract

Learning in sensorimotor adaptation tasks has been historically viewed as solely an implicit learning phenomenon. However, recent findings suggest that implicit adaptation is heavily constrained, calling into question its utility in motor learning, and the theoretical framework of sensorimotor adaptation paradigms. These inferences have been based mainly on results from single bouts of training, thus, it is possible that implicit adaptation processes supersede explicit compensation strategies, such as explicitly re-aiming their intended movement direction, over repeated practice sessions. We tested this by dissociating the contributions of explicit re-aiming strategies and implicit adaptation over five consecutive days of training. Implicit adaptation plateaued at a value far short of complete learning. We sought to determine if these constraints on implicit adaptation extend to another sensorimotor task, mirror reversal. Implicit adaptation was inappropriate for mirror reversal and became suppressed during training. These findings are consistent with a handful of recent studies suggesting that implicit adaptation processes, as studied in sensorimotor adaptation paradigms, cannot fully support long-term motor skill.

New and Noteworthy

In this set of studies, we find that implicit adaptation cannot fully account for learning in adaptation tasks, such as the visuomotor rotation and mirror reversal tasks, even following several days of training. In fact, implicit adaptation can be counterproductive to learning. These findings question the utility of implicit adaptation processes to motor learning more broadly.

Contact

Sarah Wilterson
wilterson@princeton.edu

Introduction

A fundamental function of the motor system is the adjustment of previously learned skills to accommodate changes in the environment or body. Many people might be familiar with such an accommodation from experience adjusting to new prescription eyeglasses. In the late 19th century, George Stratton (1896, 1897) noted that even radically altered visual input could be overcome by the motor system to allow for normal interaction with everyday objects. 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 (e.g. knocking over drinks) and visual oddities (e.g. stepping over a street sign that appears to be on the walkway), which eventually gave way to complete adaptation to the visual perturbation. Kohler (1951b) reported that after five days of continuous prism goggle use, the visual world suddenly became upright once more. Eventual removal of the prism goggles after complete compensation resulted in a strong aftereffect: a world that once again seemed upside-down. The presence of this slow adaptation followed by an aftereffect, coupled with the seminal findings of mirror reversal learning in patient H.M. (Corkin, 1968), led researchers to assume that the motor system's compensation to visual perturbations is the result of an implicit memory system. Indeed, it is unlikely that the subjects in Kohler's experiments were leveraging cognitive strategies to flip the world right-way-up throughout the entire time course of learning.

Building on these anecdotal observations, the field of motor learning has focused on carefully characterizing this presumably implicit process 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 Schlank, 1959; Cunningham, 1989; Imamizu et al., 1995; Pine and Krakauer, 1996; Krakauer, 2009). These experiments have revealed that learning follows an approximate power law function as the subjects adapt to the perturbed feedback. Importantly, aftereffects as first described by Helmholtz (1867), are observed when the perturbation is removed, which echo the anecdotal observations from Stratton and Kohler.

Much has been learned about motor learning using sensorimotor adaptation methods, such as learning rates (Scheidt, Dingwell, and Mussa-Ivaldi, 2001; Baddeley, Ingram, and Miall, 2003), response sensitivity functions (Fine and Thoroughman 2006, 2007; Wei and Körding, 2009), and consolidation (Brashers-Krug, Shadmehr, and Bizzi, 1996; Shadmehr and Brashers-Krug, 1997, Caithness et al, 2004), just to name a few. Importantly, these findings have been largely interpreted under the view that learning was entirely implicit. Indeed, the small, trial-by-trial adaptation effects were assumed to underlie the mechanisms of long-term motor learning (Jordan and Rumelhart, 1992).

A number of studies have challenged the view that learning is an entirely implicit process. Initially, these studies disassociated explicit and implicit learning 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 process of implicit adaptation is not necessary to compensate for visuomotor rotations. More recent work has assayed the operation of explicit strategies and implicit adaptation on a trial-by-trial basis. This was made possible by asking participants to report their intended aiming location (i.e. the direction that they intend to move in order to compensate for the perturbation), by verbally

specifying a location relative to the target (Taylor et al, 2014, Bond and Taylor 2015; McDougle et al 2015; Brudner et al 2016; McDougle et al 2017; Butcher et al 2017), or by tapping the surface of a touch screen (Bond and Taylor, 2017; Hutter and Taylor, 2018). In these experiments, 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 (implicit adaptation). These trial-by-trial assays have been extensively validated against traditional measures of implicit adaptation through direct comparison (Taylor and Ivry, 2011; Taylor et al, 2014), and convergent evidence (Poh and Taylor 2016; Butcher and Taylor 2018; Poh and Taylor 2019; McDougle and Taylor 2019). Given the advantages of disassociating explicit and implicit learning on a trial-by-trial basis, we adopted a re-aiming procedure for our experiments.

Even more challenging for a unitary implicit motor learning account: implicit adaptation appears to be highly stereotyped regardless of the particular task demands, such as the size of the visual perturbation (Bond and Taylor 2015). Implicit adaptation being a stereotyped process is supported by studies examining the adaptive response when feedback is not dependent on subjects' movements and is irrelevant to the task (Morehead, Taylor, Parvin, and Ivry, 2017; Kim, Morehead, et al. 2018). Morehead and colleagues revealed stereotyped asymptotic adaptation even when the incentive to strategize was completely removed. Surprisingly, asymptotic adaptation insufficient to account for the perturbation remained after many hours of practice in the perturbed environment (Morehead and Smith, 2017). Given these results, it is difficult to see how implicit adaptation could produce long-term motor learning in the absence of explicit re-aiming.

A limitation to many of these visuomotor rotation experiments is that they are conducted over a single block of massed learning, with the full experiment taking place on a single day. This is particularly problematic because 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 is insufficient to determine how implicit adaptation may contribute to long-term learning. The full effect of adaptation seen by Stratton (1896, 1897) and Kohler (1941, 1951a, 1951b) occurred following five full days of exposure. However, recent research has suggested that implicit adaptation remains stable over two days of training (de Brouwer et al, 2018; Huberdeau et al, 2019). This may be explained if it takes more than two days of exposure for fragile motor memories to transform into consolidated memories that are protected from disruption and more difficult to unlearn (Krakauer and Shadmehr, 2006; Criscimagna-Hemminger and Shadmehr, 2008; but see Caithness et al. 2004). Indeed, consolidation of motor memories has previously been suggested to require days between learning and trial exposures (Krakauer and Shadmehr, 2006). Implicit learning may require extensive consolidation periods in order to reach a sufficient magnitude to fully compensate for a large visuomotor rotation. However, we cannot evaluate this claim from previous studies of consolidation as these experiments did not distinguish between explicit re-aiming and implicit adaptation.

In Experiment One, we characterize the relative contribution of explicit re-aiming and implicit adaptation in a long-term visuomotor rotation training paradigm, one that took place over five separate days of training. Our goal was to determine if sufficient practice would result in explicit strategies giving way to full implicit adaptation.

An additional complication to long-term motor learning, not considered in Experiment One, is the directional nature of implicit adaptation. Several

studies have revealed that the key factor to the development of implicit adaptation is the direction of the perturbation, as opposed to the size of the perturbation (Fine and Thoroughman 2006, 2007; Wei and Kording, 2009; Semrau, Daitch, and Thoroughman, 2012; Morehead et al 2017; Butcher and Taylor, 2018). The implicit response to directional information appears to be automatic: the motor system automatically 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 automatic response to directional information has led researchers to question the role that implicit adaptation might play in visuomotor mirror reversal tasks (Gritsenko and Kalaska, 2010; Lillicrap et al., 2013; Telgen, Parvin, and Diedrichsen, 2014; Krakauer, 2019). It is worth noting that a mirror-reversal task could seem very similar to a mimicry or imitation task at face value; however, unlike social mimicry (e.g. Bek et al, 2016; for review see Gattis et al, 2002) and imitation tasks (Chiavarino et al, 2007), the mirror-reversal task does not require transforming the limb posture of another person or understanding social pressures. Instead, the mirror-reversal task is thought to probe the acquisition of a new sensorimotor mapping.

Indeed, there is evidence that error-corrective feedback responses under mirror reversal must be overcome by explicit strategies or suppressed, even after a few days of training (Figure 1). Furthermore, Telgen and colleges (2014) found that consolidation and changes in a speed-accuracy tradeoff differ considerably between rotations and mirror reversals. It is worth noting, that a mirror reversal is exactly the task that HM was able to complete, which raises the question as to how this was accomplished if implicit adaptation – as we know it from previous visuomotor rotation studies – appears unable to solve a mirror reversal.

In Experiment Two, we sought to replicate and extend previous experiments and determine if implicit adaptation could become useful in a mirror task given a sufficient training period. Of particular interest was the eventual comprisal of a repeatably practiced, and mastered, compensation to a mirror reversal perturbation as either explicit re-aiming, implicit adaptation, or a combination thereof.

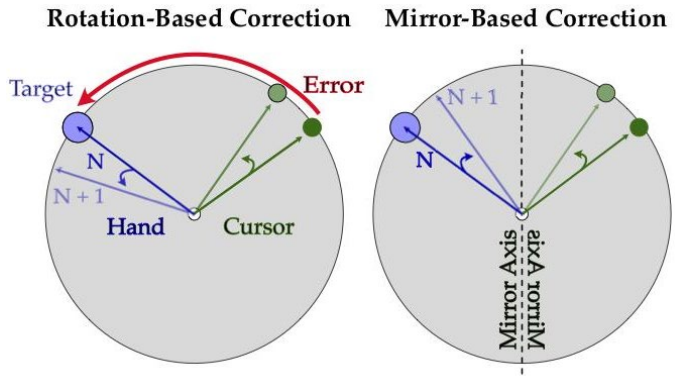

Figure 1. Appropriate correction for both rotation and mirror reversal perturbations.

Materials and Methods

Participants

Seventeen first year graduate students, undergraduate students, and community members were recruited to participate in Experiment 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 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 Princeton 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 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).

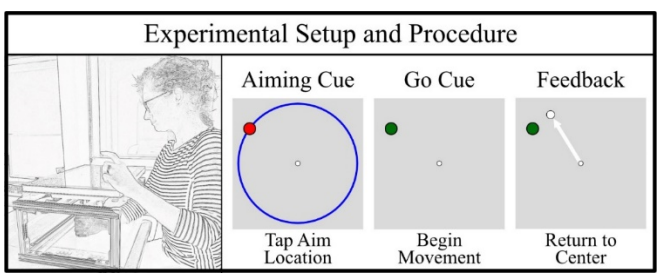

Figure 2. Experimental paradigm for assaying explicit re-aiming and implicit adaptation to an imposed visual perturbation. Participants indicate their intended movement direction (re-aiming) by tapping on the "aiming ring", which is shown here in blue. They then use the right hand to reach for the location that they think will bring the cursor to the target (the re-aimed location). Any learning not accounted for by explicit re-aiming is presumed to be implicit adaptation.

Procedure

Each trial began with subjects holding their right hand in the center of the workspace for 500 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 1).

Once the subject's right-hand exited the starting area, a small circular cursor (0.15-cm radius) provided online feedback of the current position of their hand. 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. 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. The next trial began once the subject reached the starting position.

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. Subjects became highly accurate at stopping this movement directly on the center of the workspace. This greatly increased the number of trials that could be completed in each one-hour session. Once the right hand moved out past 7 cm, endpoint feedback was displayed for 500 ms. 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 800 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 was abruptly introduced between movements of the hand and cursor feedback. Subjects were carefully observed for the first 50 perturbation trials to ensure that they understood the re-aiming instructions; further direction was provided as necessary. Subjects were provided with a three-minute break at the midpoint of training on each day.

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 1 (Rotation-Task), the perturbation was a 45° rotation in either the clockwise or counterclockwise direction (counterbalanced between subjects). In Experiment 2, the perturbation was a mirror about the vertical midline. In order to allow comparison between the two experiments, the four targets in Experiment 2 (Mirror-Task) were located 22.5° from the midline so that the solution to the perturbed reach was a 45° angle. The 12 targets in Experiment 1 were evenly spaced 22.5° apart on the aiming ring. 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 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. Implicit adaptation 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 implicit adaptation from those values, we conducted statistical analyses for only the aiming angles and implicit adaptation angles (Hutter and Taylor, 2018). 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 (Xf) and implicit adaptation as the slow learning process (Xs) over the course of 200 trials (for a detailed account of modeling explicit/implicit learning as fast/slow learning, see McDougle, Bond, and Taylor 2015). In addition, we assumed that the explicit re-aiming is updated based on target error, while implicit adaptation is updated based on the aim-to-cursor distance (McDougle, Bond, and Taylor, 2015; Equations 1 and 2).

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

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 7B and 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 implicit adaptation 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 implicit adaptation was defined in exactly the same way as in the Rotation-Task-Rotation-Correction Model.

To demonstrate the effect of implicit adaptation 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).

③ $Xs(n+1) = \alpha s \times Xs(n) - \beta s \times Es(n)$

Results

Experiment One

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 by 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 2).

As expected, participants were able to compensate for the visual perturbation during the first day of training (Figure 3). Nearly perfect performance, as measure by the hand heading angle, was achieved by the end of the first day (44.8 ± 2.0°) and was maintained in the fifth day of training (44.7 ± 0.8°). Interestingly, this high level of performance was supported by a relative mixing of explicit re-aiming and implicit adaptation both on the first day (28.1 ± 6.7° re-aiming vs. 16.7 ± 6.4° adaptation) and on day five of the experiment (30.7 ± 16.2° re-aiming vs. 13.9 ± 16.3° adaptation). Note, while it appears that explicit re-aiming accounts for nearly twice as

much learning as implicit adaptation, performing inferential statistics between explicit and implicit learning is theoretically inappropriate since our implicit measure is derived from our explicit measure.

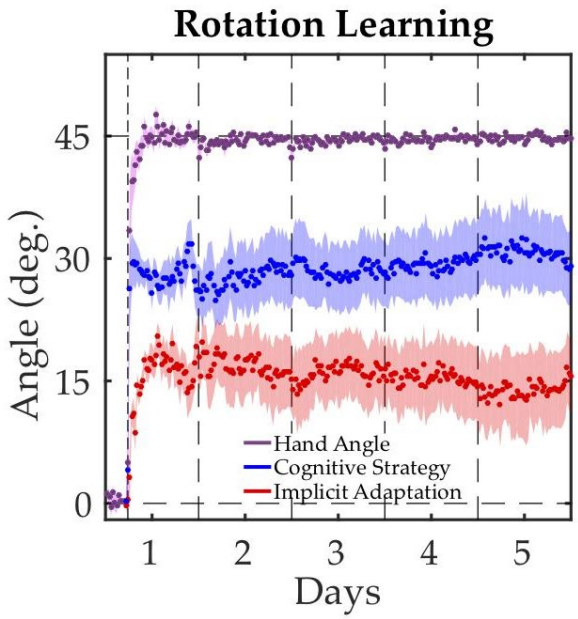

Figure 3. Time course of responses over five days of training to a 45° rotation. The solution is at 45°. The vertical dashed lines represent first the onset of the rotational perturbation, then day breaks. Shaded areas represent standard error.

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 implicit adaptation following the first day of training. This suggests that the relative contribution of explicit and implicit adaptation may differ dramatically between subjects. Indeed, if we examine the median subject and the two extremes of the distribution, we find a full continuum of learning patterns (Figure 4). Our twelve subjects run the gamut from full explicit re-aiming (Figure 4B), through a mixture of implicit adaptation and explicit re-aiming (Figure 4C), to complete implicit adaptation (Figure 4D).

These results are consistent with the range of implicit adaptation 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 individual 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 implicit adaptation at the end of each day to the same measure at the beginning of the following day (Figure 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 implicit adaptation (*Pearson's 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; implicit adaptation *average r* = -0.07).

Analysis of adaptation measurements: When we use trial-by-trial aiming data to infer implicit adaptation levels, two subjects appear to compensate fully with implicit adaptation and five subjects appear to be using a completely explicit strategy. 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. However not a single subject shows a full 45° aftereffect, or a 0° aftereffect – which would have been indicative of learning entirely via implicit adaptation and explicit re-aiming strategies respectively (Figure 6A).

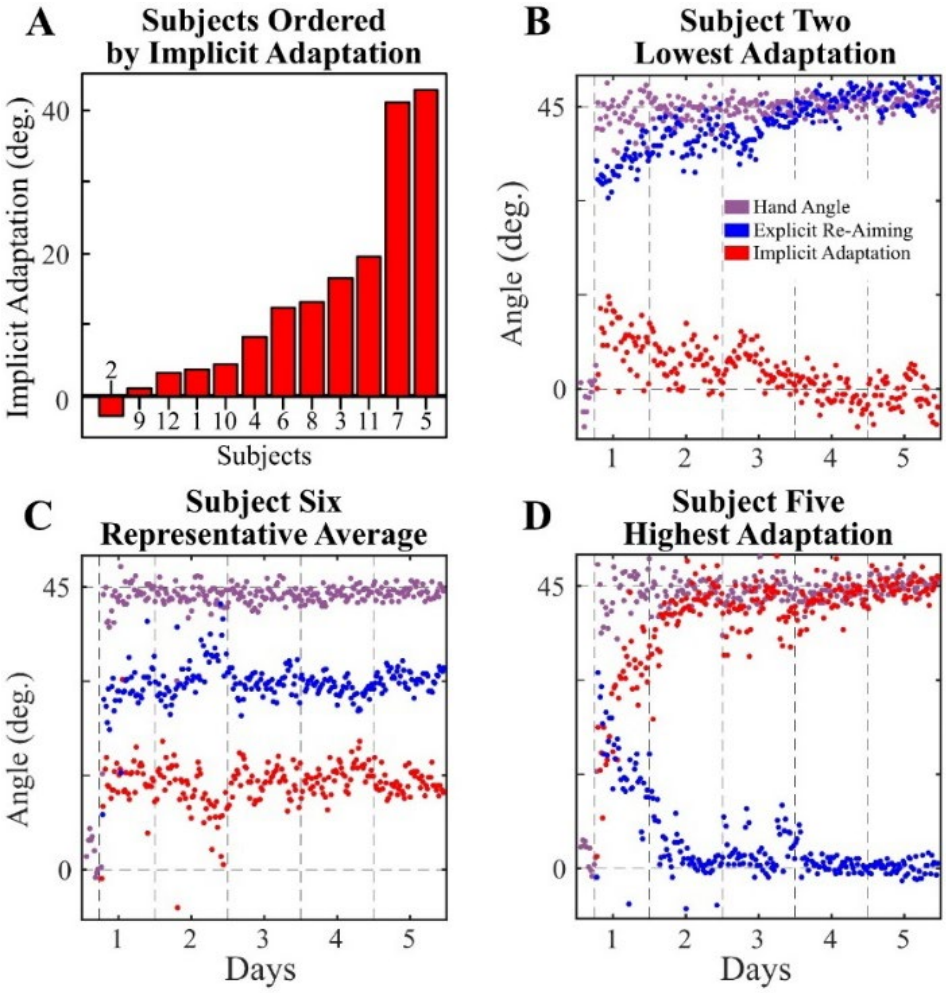

Figure 4.
(A) A bar graph ordering subjects by the average amount of implicit adaptation as calculated from the last 200 trials of day five.

(B) The subject (#2) reporting the most explicit re-aiming at the end of learning, assumed to have the least implicit adaptation (i.e. aiming at the solution with no perceptible implicit adaptation at the end of day five).

(C) A subject (#6) representative of the average taken from the middle of the distribution.

(D) The subject (#5) reporting the least explicit re-aiming (i.e. aiming at or very near to the target).

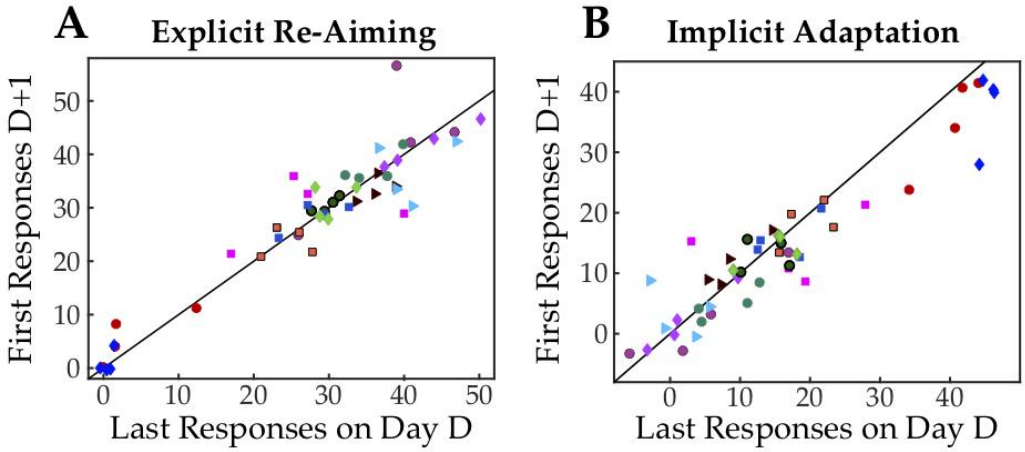

Figure 5. Average explicit re-aiming (A) and implicit adaptation (B) in last bin (16 trials) of day D plotted against first bin of day D+1 (r = 0.94). Each subject contributed 4 data points, and is represented by shape and color for visual clarity. The diagonal lines represent the unity line between responses on D and D+1.

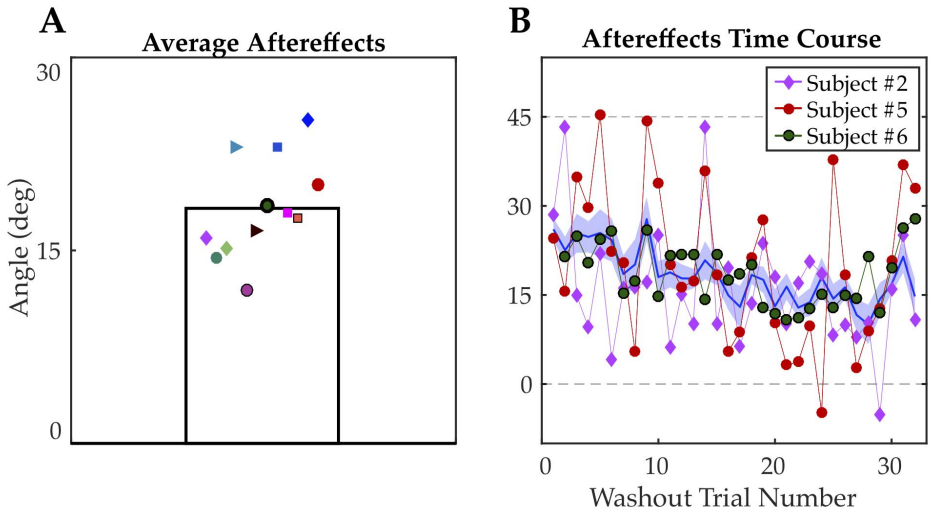

Figure 6. (A) Average aftereffect for each subject across the 32 trials of no-feedback washout (with subjects instructed to aim and move directly to the target). Each subject is labeled congruent with Figure 5 and jittered left to right in order of least to most explicit re-aiming at the end of Day 5 (see Figure 4A). (B) The time course of washout for the average across subjects and individual trial values for the three representative subjects defined in Figure 3. The individual subject markers are congruent with Figures 5 and 6A. Shaded area represents the standard error of the mean for all subjects.

While it is visually apparent that the aftereffects measure does not display the same range of adaptation that we saw using the aim subtraction method, 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. This calls into question the idea these aftereffects are an accurate representation of each individual subject's implicit adaptation. One possible explanation for this variability is the formation of a habit or proceduralization of learning (Hardwick et al 2019). Subjects may be attempting to suppress an explicit strategy that has become proceduralized over the five days of training. This suppression may fail, resulting in the inclusion of re-aiming in some trials of the aftereffect block (see Discussion). Regardless, this variability poses an additional challenge to our interpretation of the comparison between aftereffects and reported adaptation within a single subject.

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 laziness. 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 implicit adaptation followed a consistent time course where their implicit adaptation 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 their implicit adaptation

gradually decreased. Importantly, in post-participation debriefing one of the participants who displayed full implicit adaptation 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 implicit adaptation found in this experiment. Regression analyses with day-five implicit adaptation 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 One Conclusion The average response to long-term adaptation in a visuomotor rotation tasks suggests that both implicit adaptation and explicit re-aiming continue to play a role in performance even after several days of training. However, we found implicit adaptation to be fairly constrained, which is surprisingly consistent with the findings from studies examining the capacity of single-session implicit learning. We saw evidence in only two subjects that implicit adaptation could eventually replace explicit re-aiming. In Experiment Two, we sought to determine if implicit adaptation could appropriately compensate for a task with a different error structure, a mirror reversal.

Experiment Two

In Experiment Two, 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 One.

As shown in Figure 1, useful implicit adaptation 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 implicit adaptation 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. Additionally, the mirror reversal imposes a counterclockwise error in quadrants one and three but a clockwise error in quadrants two and four (Figure 7) - similar to a dual adaptation paradigm (Howard et al., 2012; Schween et al., 2019). This complicates the transformation space over targets generally used to view average behavior.

Given the opposing errors, we combined responses for the targets in quadrants 1 and 3 (i.e., at 22.5° with those at 247.5°; Figure 7) and responses for the targets in quadrants 2 and 4 (i.e., at 112.5° with those at 292.5°; Figure 7) for clarity of visualization on graphics. 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 One, participants on average slightly over-compensated for the mirror reversal perturbation on the first day (48.5 ± 3.8°), however, performance on the final day was nearly perfect (45.7 ± 1.3°, Figures 8A and 9A).

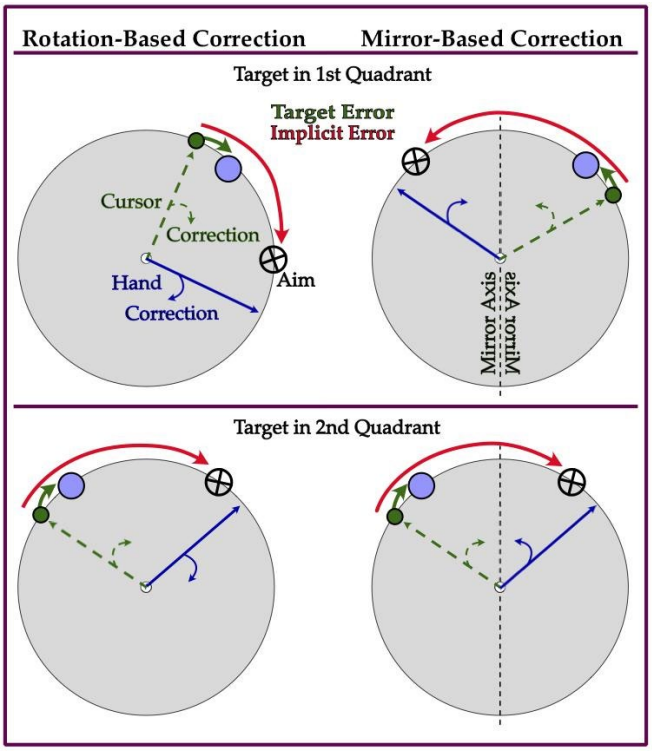

Figure 7. Diagram of appropriate corrections for both rotation and mirror-reversal perturbations given targets in the first and second quadrants. The error sign changes under a mirror perturbation.

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). To verify if the mirror reversal task resulted in mirror-appropriate adaptation, or if the error was processed as if it was an alternating rotation, we simulated the potential results using a modified two-state-space model (Smith et al., 2006). 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 implicit adaptation while the fast process is explicit re-aiming (McDougle, 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 implicit adaptation we observed is more consistent with learning appropriate for a mirror reversal or for a visuomotor rotation.

On the first day, performance was again a mixture of implicit adaptation (8.0 ± 4.6°) and explicit re-aiming (40.5 ± 5.4°). However, while implicit adaptation 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 ± 7.8°, $t = 0.961$, $p = .357$). On average, participants fully compensated for the mirror reversal with explicit re-aiming (43.6 ± 7.6°) on the final day of training.

Within the model, we assumed that the explicit re-aiming is updated based on target error,

while implicit adaptation is updated based on the aim-to-cursor distance (Taylor and Ivry 2011; McDougle, Bond, and Taylor, 2016; See methods for details). The learning and forgetting rates were determined by tuning these parameters to the simple rotation case in order to cap implicit adaptation at approximately 15° (Rotation-Task-Rotation-Correction Model, Figures 8B and 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 implicit adaptation was defined in exactly the same way as in the Rotation-Task-Rotation-Correction Model (Figure 8C and 9C). However, to demonstrate the effect of implicit adaptation 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 8D and 9D).

Within the model, we assumed that the explicit re-aiming is updated based on target error, while implicit adaptation is updated based on the aim-to-cursor distance (Taylor and Ivry 2011; McDougle, Bond, and Taylor, 2016; See methods for details). The learning and forgetting rates were determined by tuning these parameters to the simple rotation case in order to cap implicit adaptation at approximately 15° (Rotation-Task-Rotation-Correction Model, Figures 8B and 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 implicit adaptation was defined in exactly the same way as in the Rotation-Task-Rotation-Correction Model (Figure 8C and 9C). However, to demonstrate the effect of implicit adaptation 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 8D and 9D).

When the mirror task is simulated with error calculated as appropriate for a rotation $[aim - target]$, the result is a response pattern much like that of the data. In contrast, an error calculation that is appropriate for the mirror $[-(aim - target)]$, produces results in a radically different pattern from the data (compare Figure 8.A to 8.C-D and Figure 9.A to 9.C-D). The modeling results demonstrate that implicit adaptation under a mirror reversal first operates as if the perturbation was rotational in nature, then suppresses this adaptation.

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 1, we observed a very limited range of aftereffects compared with implicit adaptation calculated throughout training. After correcting for baseline bias, aftereffects hover slightly, yet significantly, below zero (average corrected aftereffect = -3.06 ± 3.2°, t-value =-3.32, p = .007, Figure 10). As in Experiment One, 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 One, 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.

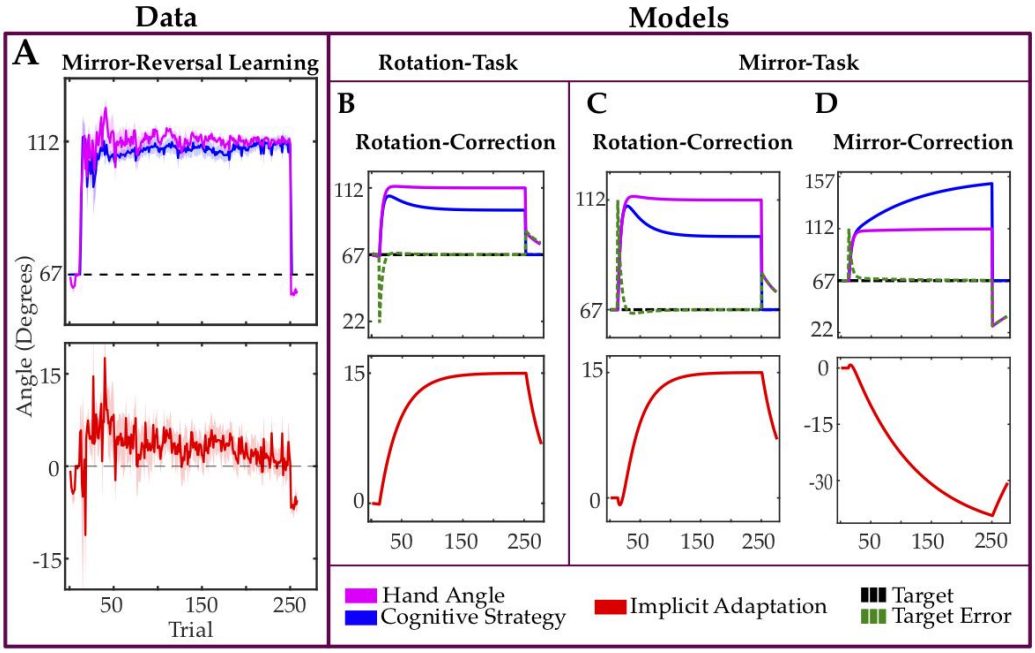

Figure 8. (A) Mirror reversal learning for all 12 subjects to targets in the first and third quadrants. (B) Rotation-Task-Rotation-Correction Model for target at 67.5° with parameters tuned to limit implicit adaptation to about 15°. (C) Mirror-Task-Rotation-Correction Model with parameters defined as in 'B'. (D) Mirror-Task-Mirror-Correction Model with parameters defined as in previous. Implicit error term is reversed from the rotation correction.

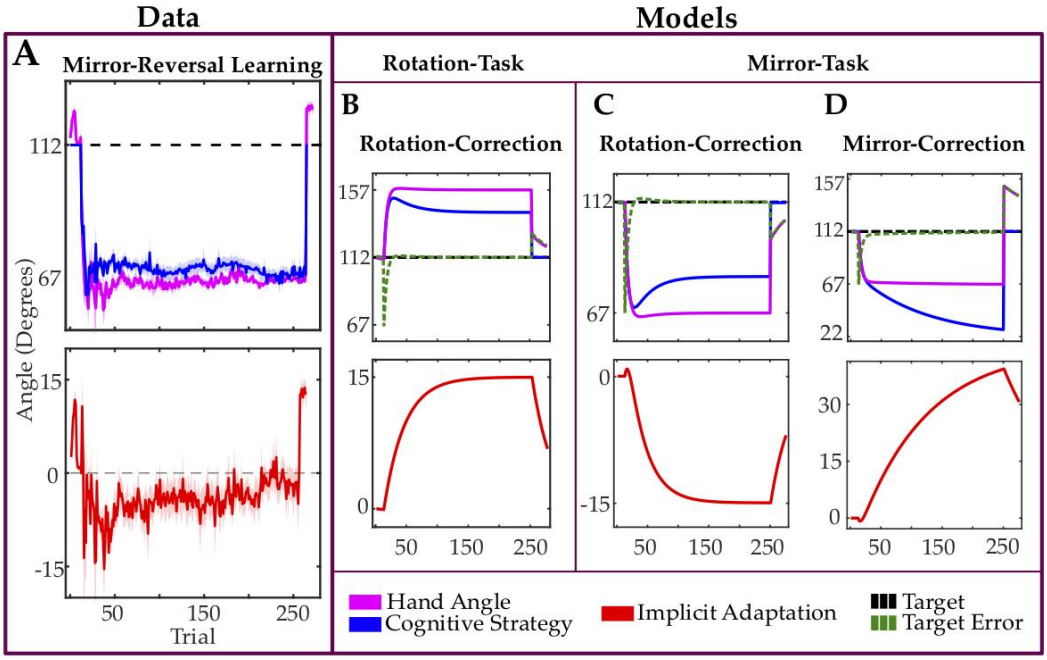

Figure 9. (A) Mirror reversal learning for all 12 subjects to targets in the second and fourth quadrants. (B) Rotation-Task-Rotation-Correction Model for target at 112.5° with parameters tuned to limit implicit adaptation to about 15°. (C) Mirror-Task-Rotation-Correction Model with parameters defined as in 'B'. (D) Mirror-Task-Mirror-Correction Model with parameters defined as in previous. Implicit error term is reversed from the rotation correction.

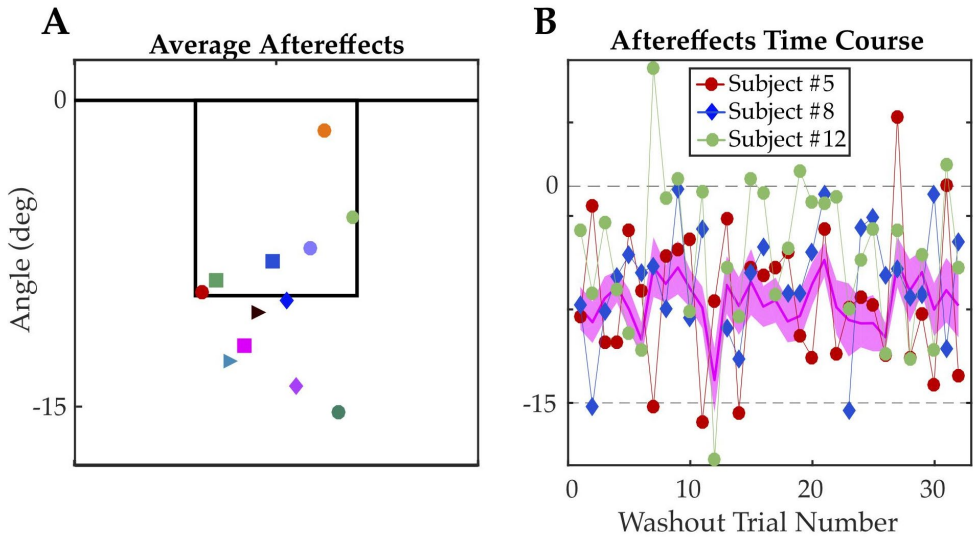

Figure 10. (A) Average corrected aftereffects for each subject across the 32 trials of no-feedback, washout (with subjects instructed to aim and move directly to the target). Each subject is jittered left to right in order of least to most explicit re-aiming at the end of Day 5. (B) The time course of 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.

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 (shown in Figure 10A) as being in the appropriate direction to correct for a rotational error.

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 11). Subjects were successful through a combination of implicit adaptation and explicit re-aiming, lost all implicit adaptation by the end of training, or reversed the direction of implicit adaptation. Although reversing the direction of implicit adaptation is the only way for the implicit adaptation process to be useful when compensating for a mirror reversal, we saw little evidence for such a reversal across individuals (particularly once corrected for baseline bias).

As in Experiment One, we conducted a linear regression analysis to determine how consistent subjects were across days (Figure 12). 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 (*Pearson's r* = 0.865) and implicit adaptation (*Pearson's 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 implicit adaptation (*average r* = -0.071). Again, while we note that correlations on small sample sizes should be viewed with caution, we don't think the inclusion of more participants would dramatically change the correlation.

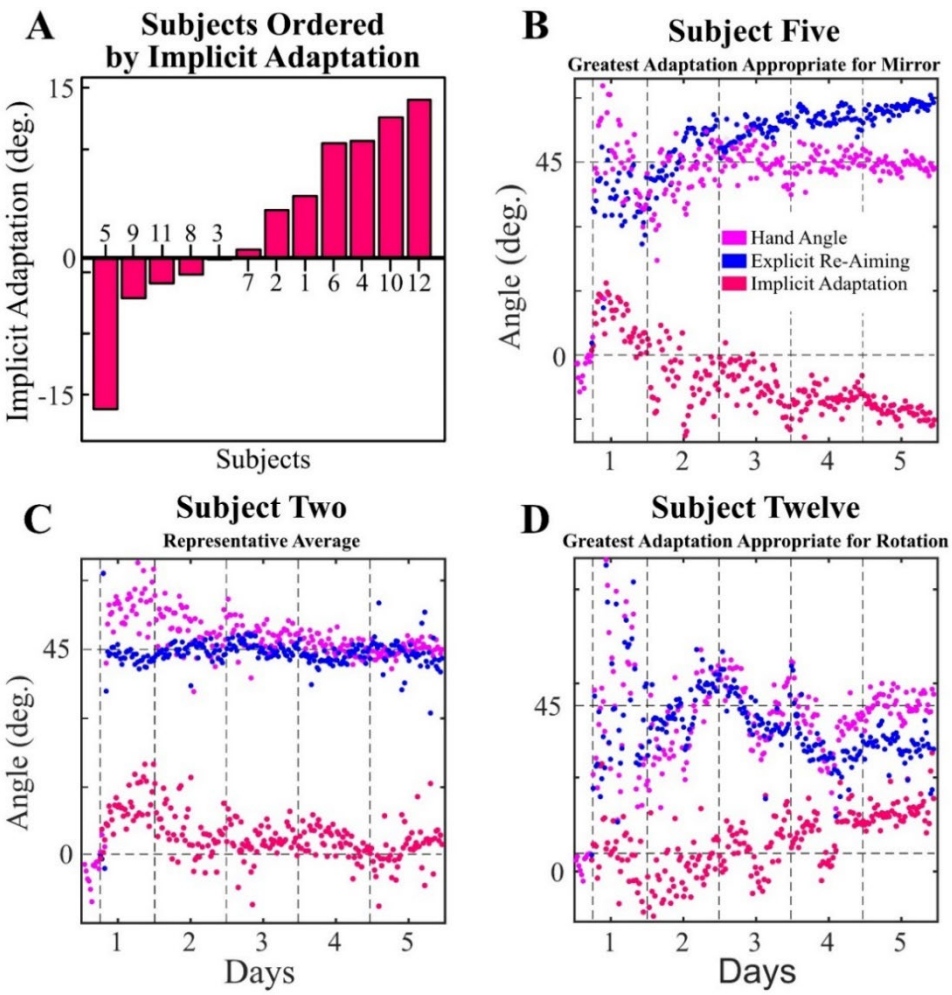

Figure 11.
(A) A bar graph ordering subjects by the average amount of explicit re-aiming that they reported during the last 200 trials of day five.

(B) The subject (#5) reporting to have reversed their implicit adaptation 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.

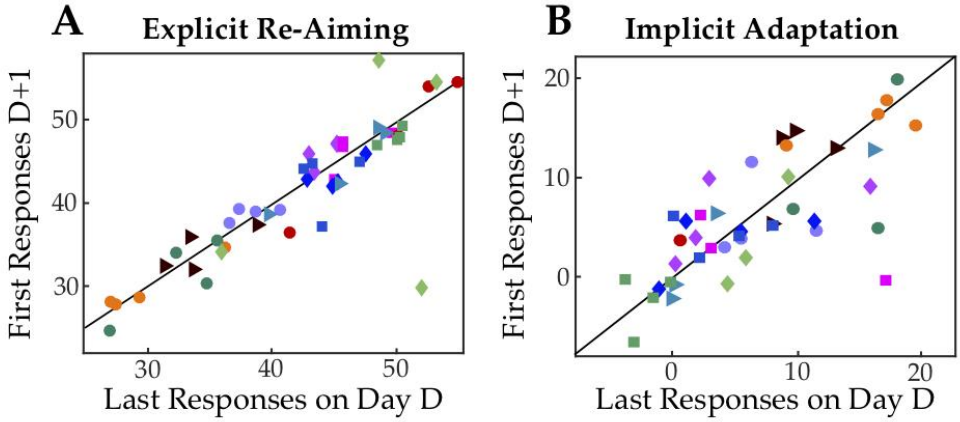

Figure 12. Average explicit re-aiming (A) and implicit adaptation (B) in last bin (16 trials) of day D plotted against first bin of day D+1 (r = 0.77 and r = 0.87, respectively). Each subject contributed 4 data points, and is represented by shape and color for visual clarity. The diagonal lines represent the unity line between responses on D and D+1.

Comparison of the Results of Experiments One and Two: Although Experiments One and Two were designed sequentially, and therefore should not be statistically compared, we took care to keep both experiments as similar as possible to allow qualitative comparison. To this end, the mirror targets were chosen to provide the same solution angle as a 45° rotation.

We saw a steadily declining average implicit adaptation function for subjects training to a mirror reversal in Experiment Two. This is in contrast to the flat average implicit adaptation function seen for the duration of training to a 45° visuomotor rotation. A visual comparison of these time courses shows starkly different behavior between the two groups (Figure 13).

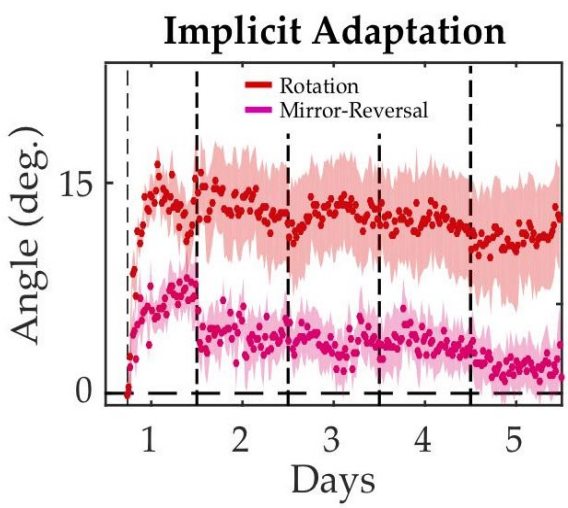

Figure 13. Time course of calculated implicit adaptation over five days of training to a 45° rotation and mirror reversal. The horizontal dashed lines represent first the onset of the rotational perturbation, then day breaks. Shaded areas represent standard error of the mean.

While subjects completing the mirror reversal task did not, on average, reverse their implicit adaptation 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, as observed by Telgen and colleagues (2014), remains an open question.

Discussion

Adaptation of the motor system in response to altered feedback as 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 does not feel like an explicit process. In experimental manipulations, implicit adaptation has been measured either by subtraction from explicitly indicated re-aiming or as an aftereffect, the latter of which has been interpreted as evidence for implicit learning underling skill learning (Hollerbach and Atkeson, 1987; Jordan and Rumelhart, 1992). Regardless of the metric used to measure implicit adaptation, across two experiments we found that implicit adaptation does not fully account for a highly practiced adaptation.

This result may have been expected given recent evidence that implicit adaptation 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 implicit adaptation was due to a lack of sufficient opportunity for consolidation, an idea that failed to be supported by our five-day experiments. The failure of what we have been measuring as implicit adaptation in the sensorimotor learning field is underscored by the mirror reversal task. Even after five days, implicit adaptation was suppressed or remained counter to the subject's goal. This supports previous studies which found counter-productive adaptation in mirror reversal tasks (Telgen, Parvin, and Diedrichsen, 2014). The inflexibility of implicit adaptation, even after multiple days of training, restricts its usefulness for many motor learning tasks. Humans exhibit great flexibility and adaptability under numerous motor learning conditions, as often illustrated by elite

athletes, and implicit adaptation appears insufficient to account for such performances.

Interestingly, implicit adaptation was suppressed over time in the mirror reversal task. One possible explanation is the presence of an error-signal that has a constantly switching sign. Previous work has shown that inconsistency in error signals reduces overall learning (Castro et al., 2014) and implicit adaptation in particular (Hutter and Taylor, 2018), which is consistent with the motor system storing a history of errors (Herzfeld et al., 2014). However, we acknowledge that this is speculative and future experiments will be necessary to pin down what features of the task may be responsible for modulating implicit adaptation processes.

Although implicit adaptation 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 habitualization, which may be reflected in participants' reaction or preparation times (Logan, 1980; Cohen, Dunbar, and McClelland, 1990). 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). One study found that performance under constrained preparation time can only 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 14). While we would like to interpret this result as indicative of a proceduralization of a re-aiming strategy in both tasks, the reaction times are still quite long compared to what might be expected under no perturbation; reaction time was 0.2 ± 0.02 s in the baseline phase. However, it is important to note that a direct comparison of reaction times in baseline and during perturbation training may be misleading given the comparative brevity of the baseline (80 trials without rotation), and the added requirement of reporting intended aiming location during perturbation training. Direct comparison with previous reports of reaction time during training is also problematic given the differences in task demands (i.e., reporting of re-aiming strategies). In these previous reports, reaction times less than approximately 300-400 ms resulted in lowered performance on a visuomotor rotation task (Fernandez-Ruiz, Wong, et al. 2011; Haith, Huberdeau, and Krakauer 2015), which is approximately half of our measured reaction times by the end of training on day five. 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 some degree of proceduralization has occurred.

Another suggestion of proceduralization can be seen in our aftereffects data (Figures 6B and 10B). The great variability within each subject may be the result of attempting, and partially failing, to suppress a proceduralized strategy. Hardwick and colleagues (2019), showed that a habit is difficult to suppress following several days of training. In our washout data, each trial may be the result of successfully suppressed strategy, a pure measure of implicit adaptation, or a mixture of unsuppressed strategy and implicit adaptation.

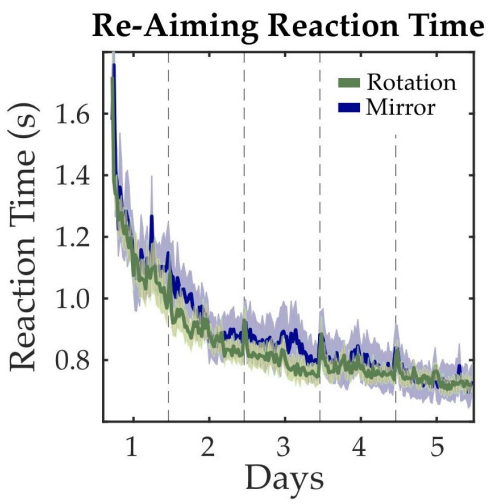

Figure 14. Reaction time across training for subjects in both Experiment Two (Rotation) and Experiment Two (Mirror-Reversal). The shaded areas represent the standard error of the mean.

This leaves open the question of 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). Stimulus-response relationships potentially explain previous assumptions that adaptive learning is necessarily implicit; pushing back on the idea that computations cannot be accomplished sufficiently quickly to allow learning in a forced reaction time experiment (Haith and Krakauer, 2013; Haith, Huberdeau, and Krakauer, 2015; Haith, Pakpoor, and Krakauer, 2016).

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).

Alternatively, 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 implicit adaptation. Indeed, implicit adaptation may be exceptional at matrix reparameterization. We acknowledge that we don't know why implicit adaptation is limited in reparameterization for large rotational changes (Wei and Kording, 2009, 2010; Hutter and Taylor, 2018), however this is an interesting avenue for future research. 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 implicit adaptation 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 implicit adaptation appear to stop reparametrizing inappropriately. It could be that once a new structure is specified, implicit adaptation becomes appropriate for this new structure. Learning this new structure, thereby reformatting adaptation processes, might be what we think of as motor skill learning.

A slightly different 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 implicit adaptation 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 have the outcome of changing the structure of the transformation matrix, again what we may think of as motor skill learning.

Regardless of the specific mechanistic explanation, we hold that classic implicit adaptation is not the driving force behind motor learning, but that another implicit system is largely responsible. The failure of adaptation is underscored by the mirror reversal task. At first glance, the two tasks don't appear radically different. One might expect implicit adaptation to be able to solve both. However, there are significant differences in the underlying computational demands, which the implicit adaptation system appears unable to rectify. As illustrated in Figure 1, the relationship between the observed error and the motor command are in opposite directions between the visuomotor rotation and mirror reversal tasks. Lillicrap and colleagues (2013) suggested that this distinction is captured by the sign of sensory derivatives (or control Jacobian). They found that models capable of updating the estimate of the sensory derivatives predicted the average human behavior while wearing inverting prisms: Errors initially increase, but once the sensitivity estimate is updated performance improves. However, the model could not predict individual subjects, which often showed erratic behavior. Based on our current findings, we suspect that this erratic behavior is related to explicit re-aiming strategies attempting to overcome both the inverting prisms and counterproductive implicit adaptation. Indeed, one participant in their study wore inverting prisms for eight consecutive days, but even by the end of the eighth day she reported that good performance required concentrated attention (Lillicrap et al., 2013). Thus, it appears that implicit adaptation may be highly constrained and operates regardless of task demands (Mazzoni and Krakauer 2006; Taylor and Ivry 2011; Morehead et al., 2017). Explicit strategies, on the other hand, appear readily able to solve either task, potentially by re-estimating the sensory derivatives between error signals and motor commands.

Implicit adaptation does not account for the majority of adaptive learning, even after allowing for many days of consolidation. We found no evidence that implicit adaptation is able to compensate for alternate task demands, as seen in a mirror reversal task. We believe that the majority of the learning is produced by explicit strategies and that the re-aiming strategies in our task are partially proceduralized through training. It is an open question to what this proceduralization process might be, and how implicit adaptation interacts with these novel compensatory structures. What is clear, however, is that the implicit adaptation process, as evidenced by aftereffects, does not appear able to account for implicit learning even in sensorimotor adaptation tasks – let alone motor skill learning.

Acknowledgements

This work was supported by the National Institute of Neurological Disorders and Stroke (Grant R01 NS-084948 to J. A. Taylor). We thank Peter Butcher for providing extensive coding assistance in preparing the main experimental task.

References

Baddeley, Ingram, & Miall (2003). System identification applied to a visuomotor task: near-optimal human performance in a noisy changing task. *Journal of Neuroscience*, 23(7), 3066-3075.

Bek, Poliakoff, Marshall, Trueman, & Gowen (2016). Enhancing voluntary imitation through attention and motor imagery. *Experimental brain research*, 234(7), 1819-1828.

Benson, Anguera, & Seidler, (2011). A spatial explicit strategy reduces error but interferes with sensorimotor adaptation. *Journal of neurophysiology*, 105(6), 2843-2851.

Bond, & Taylor, (2015). Flexible explicit but rigid implicit learning in a visuomotor adaptation task. *Journal of neurophysiology*, *113*(10), 3836-3849.

Bond and Taylor (2017). Structural learning in a visuomotor adaptation task is explicitly accessible. *eNeuro*, 4(4)

Brainard & Vision (1997). The psychophysics toolbox. *Spatial vision*, 10, 433-436.

Brashers-Krug, Shadmehr, & Bizzi (1996). Consolidation in human motor memory. *Nature*, *382*(6588), 252.

Brudner, Kethidi, Graeupner, Ivry, & Taylor (2016). Delayed feedback during sensorimotor learning selectively disrupts adaptation but not strategy use. *Journal of neurophysiology*, 115(3), 1499-1511.

Braun, Aertsen, Wolpert, & Mehring (2009). Motor task variation induces structural learning. *Current Biology*, 19(4), 352-357.

Butcher and Taylor (2018). Decomposition of a sensory-prediction error signal for visuomotor adaptation. *Journal of Experimental Psychology: Human Perception and Performance*, 44(2):175-194.

Butcher, Ivry, Kuo, Rydz, Krakauer, & Taylor (2017). The cerebellum does more than sensory prediction error-based learning in sensorimotor adaptation tasks. *Journal of neurophysiology*, 118(3), 1622-1636.

Caithness, Osu, Bays, Chase, Klassen, Kawato... & Flanagan (2004). Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. *Journal of Neuroscience*, 24(40), 8662-8671.

Castro, Hadjiosif, Hemphill, & Smith, (2014). Environmental consistency determines the rate of motor adaptation. *Current Biology*, 24(10), 1050-1061.

Chiavarino, Apperly, & Humphreys, (2007). Exploring the functional and anatomical bases of mirror-image and anatomical imitation: the role of the frontal lobes. *Neuropsychologia*, 45(4), 784-795.

Cohen, Dunbar, & McClelland, (1990). On the control of automatic processes: a parallel distributed processing account of the Stroop effect. *Psychological review*, 97(3), 332.

Corkin (1968). Acquisition of motor skill after bilateral medial temporal-lobe excision. *Neuropsychologia*, 6(3), 255-265.

Criscimagna-Hemminger & Shadmehr (2008). Consolidation patterns of human motor memory. *Journal of Neuroscience*, *28*(39), 9610-9618.

Cunningham (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, Niv, & Dayan (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature neuroscience*, 8(12), 1704.

Day, Roemmich, Taylor, & Bastian (2016). Visuomotor learning generalizes around the intended movement. *eNeuro*, 3(2).

Dayan (2009). Dopamine, reinforcement learning, and addiction. *Pharmacopsychiatry*, 42(S 01), S56-S65.

de Brouwer, Albaghdadi, Flanagan, & Gallivan (2018). Using gaze behavior to parcellate the explicit and implicit contributions to visuomotor learning. *Journal of neurophysiology*, 120(4), 1602-1615.

Fernandez-Ruiz, Wong, Armstrong, & Flanagan (2011). Relation between reaction time and reach errors during visuomotor adaptation. *Behavioural brain research*, 219(1), 8-14.

Fine & Thoroughman (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 & Thoroughman (2007). Trial-by-trial transformation of error into sensorimotor adaptation changes with environmental dynamics. *Journal of neurophysiology*, 98(3), 1392-1404.

Gritsenko & Kalaska (2010). Rapid online correction is selectively suppressed during movement with a visuomotor transformation. *Journal of neurophysiology*, 104(6), 3084-3104.

Hadjiosif, Krakauer, & Haith (2020). Implicit adaptation is driven by direct policy updates rather than forward-model-based learning. *bioRxiv*.

Haith & Krakauer, (2013). Model-based and model-free mechanisms of human motor learning. In *Progress in motor control* (pp. 1-21). Springer, New York, NY.

Haith, Huberdeau, & Krakauer (2015). The influence of movement preparation time on the expression of visuomotor learning and savings. *Journal of neuroscience*, 35(13), 5109-5117.

Haith, Pakpoor, & Krakauer (2016). Independence of movement preparation and movement initiation. *Journal of Neuroscience*, 36(10), 3007-3015.

Hardwick, Forrence, Krakauer, & Haith (2019). Time-dependent competition between goal-directed and habitual response preparation. *Nature Human Behaviour*, *3*(12), 1252-1262.

Hegele & Heuer (2010). Implicit and explicit components of dual adaptation to visuomotor rotations. *Consciousness and cognition*, 19(4), 906-917.

Held & Schlank (1959). Adaptation to disarranged eye-hand coördination in the distance-dimension. *The American journal of psychology*, 603-605.

Helmholtz (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, Vaswani, Marko, & Shadmehr (2014). A memory of errors in sensorimotor learning. Science, 345(6202), 1349-1353.

Heuer & Hegele (2008). Adaptation to visuomotor rotations in younger and older adults. *Psychology and aging*, 23(1), 190.

Hollerbach & Atkeson (1987). Deducing planning variables from experimental arm trajectories: Pitfalls and possibilities. *Biological Cybernetics*, 56(5-6), 279-292.

Howard, Ingram, Franklin, & Wolpert (2012). Gone in 0.6 seconds: the encoding of motor memories depends on recent sensorimotor states. *Journal of Neuroscience*, 32(37), 12756-12768.

Huberdeau, Krakauer, & Haith (2019). Practice induces a qualitative change in the memory representation for visuomotor learning. *J Neurophysiol,* 2019.

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

Imamizu, Uno, & Kawato (1995). Internal representations of the motor apparatus: implications from generalization in visuomotor learning. *Journal of Experimental Psychology: Human Perception and Performance*, 21(5), 1174.

Jordan & Rumelhart (1992). Forward models: Supervised learning with a distal teacher. *Cognitive science*, 16(3), 307-354.

Kim, Morehead, Parvin, Moazzezi, & Ivry (2018). Invariant errors reveal limitations in motor correction rather than constraints on error sensitivity. *Communications Biology*, *1*(1), 19.

Kim, Parvin, & Ivry (2019). The influence of task outcome on implicit motor learning. *eLife*, 8, e39882.

Kleiner, Brainard, Pelli, Ingling, Murray, & Broussard (2007). What's new in Psychtoolbox-3. *Perception*, 36(14), 1.

Kohler (1941). *Der Einfluß der Erfahrung in der optischen Wahrnehmung. Beleuchtet von Versuchen langdauernden Tragens bildverzerrender Prismen. Phil* (Doctoral dissertation, Dissertation, Innsbruck).

Kohler (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 (1951b). *Warum sehen wir aufrecht - obwohl die Bilder im Inneren des Auges verkehrt stehen? Die Pyramide*, 28e33.

Krakauer (2009). Motor learning and consolidation: the case of visuomotor rotation. In *Progress in motor control* (pp. 405-421). Springer, Boston, MA.

Krakauer, Hadjiosif, Xu, Wong, & Haith (2019). Motor learning. *Comprehensive Physiology*, *9*, 613-663.

Krakauer & Shadmehr (2006). Consolidation of motor memory. *Trends in neurosciences*, *29*(1), 58-64.

Lillicrap, Moreno-Briseño, Diaz, Tweed, Troje, & Fernandez-Ruiz (2013). Adapting to inversion of the visual field: a new twist on an old problem. *Experimental brain research*, 228(3), 327-339.

Logan (1980). Attention and automaticity in Stroop and priming tasks: Theory and data. *Cognitive psychology*, 12(4), 523-553.

Logan (1988). Toward an instance theory of automatization. *Psychological review*, 95(4), 492.

Mazzoni & Krakauer (2006). An implicit plan overrides an explicit strategy during visuomotor adaptation. *Journal of neuroscience*, 26(14), 3642-3645.

McDougle & Taylor (2019). Dissociable cognitive strategies for sensorimotor learning. *Nature communications*, *10*(1), 40.

McDougle, Bond, & Taylor (2015). Explicit and implicit processes constitute the fast and slow processes of sensorimotor learning. *Journal of Neuroscience*, 35(26), 9568-9579.

McDougle, Bond, & Taylor, (2017). Implications of plan-based generalization in sensorimotor adaptation. *Journal of neurophysiology*, 118(1), 383-393.

Morehead, Taylor, Parvin, & Ivry (2017). Characteristics of implicit sensorimotor adaptation revealed by task-irrelevant clamped feedback. *Journal of cognitive neuroscience*, *29*(6), 1061-1074.

Morehead, Smith (2017, November). The magnitude of implicit sensorimotor adaptation is limited by continuous forgetting. Abstract. Advances in Motor Learning & Motor Control, Washington DC.

Oldfield (1971). The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia*, *9*(1), 97-113.

Pine & Krakauer (1996). Learning of scaling. *Neuroreport*, 7, 2357-2361.

Poh & Taylor (2019). Generalization via superposition: combined effects of mixed reference frame representations for explicit and implicit learning in a visuomotor adaptation task. *Journal of neurophysiology*, 121(5), 1953-1966.

Provost, Johnson, Karayanidis, Brown, & Heathcote (2013). Two routes to expertise in mental rotation. *Cognitive science*, 37(7), 1321-1342.

Redding & Wallace (1996). Adaptive spatial alignment and strategic perceptual-motor control. *Journal of Experimental Psychology: Human Perception and Performance*, 22(2), 379.

Sachse, Beermann, Martini, Maran, Domeier, & Furtner (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, & Gomi (2010). Multiple motor learning strategies in visuomotor rotation. *PLoS One*, 5(2), e9399.

Schaefer, Shelly, & Thoroughman (2012). Beside the point: motor adaptation without feedback-based error correction in task-irrelevant conditions. *Journal of Neurophysiology*, 107(4), 1247-1256.

Scheidt, Dingwell, & Mussa-Ivaldi (2001). Learning to move amid uncertainty. *Journal of neurophysiology*, 86(2), 971-985.

Schween, McDougle, Hegele, & Taylor (2019). Explicit strategies in force field adaptation. *BioRxiv*, 694430.

Semrau, Daitch, & Thoroughman (2012). Environmental experience within and across testing days determines the strength of human visuomotor adaptation. *Experimental brain research,* 216(3), 409-418.

Shadmehr & Brashers-Krug (1997). Functional stages in the formation of human long-term motor memory. *Journal of Neuroscience*, 17(1), 409-419.

Smith, Ghazizadeh, & Shadmehr (2006). Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS biology*, 4(6), e179.

Stratton (1896). Some preliminary experiments on vision without inversion of the retinal image. *Psychological review*, *3*(6), 611.

Stratton (1897). Vision without inversion of the retinal image. *Psychological review*, *4*(4), 341.

Taylor, Krakauer, & Ivry (2014). Explicit and implicit contributions to learning in a sensorimotor adaptation task. *Journal of Neuroscience*, *34*(8), 3023-3032.

Taylor & Ivry (2011). Flexible strategies during motor learning. *PLoS Computational Biology*, 7(3)

Telgen, Parvin, & Diedrichsen (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.

Wei & Kording (2009). Relevance of error: what drives motor adaptation?. *Journal of neurophysiology*, *101*(2), 655-664.

Wei & Körding (2010). Uncertainty of feedback and state estimation determines the speed of motor adaptation. Frontiers in computational neuroscience, 4, 11.