

Mirror Reversal and Visual Rotation Are Learned and Consolidated via Separate Mechanisms: Recalibrating or Learning *De Novo*?

Sebastian Telgen, Darius Parvin, and  Jörn Diedrichsen

Institute of Cognitive Neuroscience, University College London, London WC1N 3AR, United Kingdom

Motor learning tasks are often classified into adaptation tasks, which involve the recalibration of an existing control policy (the mapping that determines both feedforward and feedback commands), and skill-learning tasks, requiring the acquisition of new control policies. We show here that this distinction also applies to two different visuomotor transformations during reaching in humans: Mirror-reversal (left-right reversal over a mid-sagittal axis) of visual feedback versus rotation of visual feedback around the movement origin. During mirror-reversal learning, correct movement initiation (feedforward commands) and online corrections (feedback responses) were only generated at longer latencies. The earliest responses were directed into a nonmirrored direction, even after two training sessions. In contrast, for visual rotation learning, no dependency of directional error on reaction time emerged, and fast feedback responses to visual displacements of the cursor were immediately adapted. These results suggest that the motor system acquires a new control policy for mirror reversal, which initially requires extra processing time, while it recalibrates an existing control policy for visual rotations, exploiting established fast computational processes. Importantly, memory for visual rotation decayed between sessions, whereas memory for mirror reversals showed offline gains, leading to better performance at the beginning of the second session than in the end of the first. With shifts in time-accuracy tradeoff and offline gains, mirror-reversal learning shares common features with other skill-learning tasks. We suggest that different neuronal mechanisms underlie the recalibration of an existing versus acquisition of a new control policy and that offline gains between sessions are a characteristic of latter.

Key words: consolidation; mirror reversal; reaching movements; sensorimotor adaptation; skill learning; visual rotation

Introduction

Humans are experts in adjusting their movements to changing task demands (Helmholtz, 1866; McLaughlin, 1967; Gentilucci et al., 1995). Learning a new task requires a change in the functions that translate goals (and states) into motor commands. These functions have been synonymously referred to as visuomotor mappings, control policies, or inverse models (Sutton and Barto, 1998; Todorov and Jordan, 2002).

But are all new tasks learned the same way? Here we contrast the learning processes for two different visuomotor transformations: visual rotation (VR) and mirror reversal (MR). It has been suggested that MR and VR are learned using separate learning mechanisms (Werner and Bock, 2010). Here we hypothesize that VR can be learned by a gradual recalibration of the existing con-

trol policy, whereas MR requires the establishment of a novel mapping. This idea is motivated by how the motor system uses error to update future movements (Fig. 1). When confronted with VR, the correction calculated under the old policy will be directed approximately (for rotations $<90^\circ$) in the appropriate direction. The new policy therefore could be learned by updating the next motor command with the correction calculated following the outdated mapping (Kawato and Gomi, 1992). Repeated applications of this learning rule lead to the correct policy. During MR, however, the update inferred from the old mapping points in the wrong direction, and a novel policy would have to be acquired instead.

Krakauer and colleagues suggested that the difference between recalibration and acquisition is visible in speed-accuracy tradeoffs (Reis et al., 2009; Shmuelof et al., 2012). Because fast sequential movements require the rapid generation of feedforward and feedback commands, this likely relates to the speed of the underlying computational processes: When the system recalibrates a well-learned control policy, it should be able to use existing fast automatic processes and generate accurate responses even under time pressure. The establishment of a new control policy, however, should entail initially slower, and possibly more explicit, components (Hikosaka et al., 2002), requiring additional processing time. Only with long practice, it should become automatized and achieve equivalent performance at shorter time

Received Dec. 18, 2013; revised Aug. 17, 2014; accepted Aug. 29, 2014.

Author contributions: S.T. and J.D. designed research; S.T. and D.P. performed research; S.T. and J.D. analyzed data; S.T. and J.D. wrote the paper.

This work was supported by the Marie Curie Initial Training Network Cerebellum-C7 within the 7th framework program of the European Union. We thank Maurice Smith, Nobuhiro Hagura, John Krakauer, Alexandra Reichenbach, and George Prichard for helpful comments.


The authors declare no competing financial interests.

Correspondence should be addressed to Dr. Jörn Diedrichsen, Institute of Cognitive Neuroscience, University College London, Alexandra House, 17 Queen Square, London WC1N 3AR, United Kingdom. E-mail: j.diedrichsen@ucl.ac.uk.

DOI:10.1523/JNEUROSCI.5306-13.2014

Copyright © 2014 the authors 0270-6474/14/3413768-12\$15.00/0

Visuomotor rotation



Mirror reversal




Figure 1. Schematic drawing of recalibration during MR and VR. The hyphenated vertical line indicates the mirror reversal axis. In trial n , hand (red) movements toward the -20° target (Fig. 2 for coordinate frame) result in the cursor (blue) traveling to 20° , thus producing an error (hyphenated black arrow) of 40° . A fraction of this error vector is used to update the next motor command. On trial $n + 1$, the hand movement direction (solid red arrow) is therefore shifted from the previous movement direction (hyphenated red arrow). During VR (top), this leads to error reduction between cursor (solid blue arrow) and target compared with the previous movement. During MR (bottom), the same update results in an increased error.

intervals. Thus, we expected that the acquisition of a control policy would be accompanied by a shift in time-accuracy tradeoffs. We tested this idea by studying fast feedforward and feedback commands.

Finally, we also tested whether VR and MR learning differ in how the memory consolidates between sessions. Adaptation tasks typically show forgetting between sessions (Tong et al., 2002; Klassen et al., 2005; Krakauer et al., 2005; Trempe and Proteau, 2010), whereas skill-learning tasks, such as learning novel sequences of finger movements, show little forgetting (Reis et al., 2009), and sometimes even offline gains (Stickgold, 2005; Doyon et al., 2009; Brawn et al., 2010; Wright et al., 2010; Abe et al., 2011). Given that skill-learning tasks are also characterized by shifts in speed-accuracy tradeoff (Reis et al., 2009; Shmuelof et al., 2012), we hypothesized that MR learning may also show offline gains between sessions.

Materials and Methods

Participants. All participants ($N = 112$, 52 male) were right-handed according to the Edinburgh Handedness Inventory (median 84.6, interquartile range 25.3) (Oldfield, 1971) and were 18–30 years of age. None had a history of neurological illness and/or was taking medication. Participants were recruited through online advertising and received monetary compensation (£7/h) at the conclusion of the study. Informed consent was obtained before the study started, and all procedures were approved by the University College London Ethics Committee.

General procedure. Participants made 15 cm center-out reaching movements to targets displayed on a TFT LCD while holding a robotic handle with the right hand. The robotic device allowed unrestrained movement in the horizontal plane and was able to exert forces to the participant's hand. Movements were recorded at 200 Hz. Visual feedback was provided on a monitor (60 Hz refresh rate) that was viewed via a horizontal mirror placed over the participant's hand. The delay of the visual display (65 ms) was empirically measured using a photodiode and

taken into account in the analysis of the data. Because of the mirror, the arm and hand were not directly visible. The position of the right hand was represented on the mirror by a cursor (2 mm diameter).

At the beginning of each trial, the robot guided the participant's hand to the start location, a small rectangle, ~ 15 cm in front of the participant's chest. After the hand remained inside the start rectangle for >400 ms, a target (0.7×0.7 cm² square) appeared on the screen. To probe the time dependency of the forward command under the two visuomotor mappings, it was essential to enforce tight bounds on reaction time (RT), the time from target appearance to movement onset. Thus, participants were instructed that their first priority was to react quickly to the onset of the target. We played an unpleasant buzzing tone for slow reactions ($RT > 385$ ms) and an unpleasant high beep for anticipatory movements ($RT < 35$ ms).

A movement was considered started when the tangential velocity exceeded 3.5 cm/s and ended when it fell <3.5 cm/s. For offline analysis, the velocity threshold for the movement start was set to 2.5 cm/s. Participants were also instructed that their movements had to be fast and accurate to receive points. If the movement time (MT), the duration from movement onset to termination, was too long or if the peak velocity was too low (<40 cm/s), all items turned blue; if the peak velocity was too high (>100 cm/s), yellow. Green feedback indicated that the peak velocity was in the correct range but the movement was terminated outside of the tolerance zone around the target. Only when all criteria were met, did all items in the visual display turn red and a pleasant sound was played, signaling that the participants had gained a point. Participants were explicitly informed and then familiarized with these criteria over the first four practice blocks. The target zone in which the movement had to end was initially set to 1.2 cm, and the maximum MT to 1200 ms. These criteria were manually adjusted after each block to maintain a constant average success rate: If a participant achieved $>50\%$ of all points in the last block, both criteria were decreased by 0.1 cm and 100 ms, respectively, until they reached 0.7 cm or 800 ms. This adjustment ensured that the rate of reward stayed within a motivating range. Visually, the target always remained the same size (0.7 cm) because changes of target size might have caused participants to alter their strategy. For offline analysis, we included all trials, regardless of whether they satisfied the criteria described above (see Data analysis).

Experiment 1: mirror reversal, feedforward control. The experiment consisted of two testing sessions, in which 15 participants were exposed to a mirror-reversed environment. The two experimental sessions took place between 4 and 10 PM on two consecutive days for all participants. Participants reached from a central starting location to one of six possible targets located at -20° , 0° , 20° , 160° , 180° , and -160° (Fig. 2).

Each session consisted of 16 blocks, each comprising 72 trials. The first session started with four training blocks to familiarize participants with the performance feedback (not included in the analysis) followed by four baseline blocks (blocks 1–4). Visual feedback was mirrored during the following eight blocks of the first session (blocks 5–12); e.g., to reach to the right target, one had to generate a reaching movement to the left. In the second session, visual feedback was mirrored during the first 12 blocks (blocks 13–24). In the last 4 blocks of the second session, visual feedback was returned to normal (blocks 25–28). Each block contained a total of 72 trials consisting of 12 reaches toward each of the six targets. The four lateral targets (-160° , -20° , 20° , and 160°) were chosen so that the required change in the motor command equaled 40° and would match the required change in the visual rotation condition (see below). To assess the state of the feedforward command in all experiments, we measured the initial movement direction, the angular hand position averaged from 100 to 150 ms after movement onset. This early measure is relatively uninfluenced by possible feedback corrections (Franklin and Wolpert, 2008).

In Experiments 1–4, participants were informed in the break between block 4 and 5 that a visuomotor transformation would be imposed, and the nature of the transformation (visual rotation or mirror reversal) was explained to them. We then stressed that their first priority should be to initiate their movement within RT limits, even if it meant that they missed the target. These restrictions largely prevented participants from consciously replanning their movement endpoint (Georgopoulos and




Figure 2. Target arrangements in Experiments 1 and 2. Gray circles represent target locations in Experiment 1, whereas white circles represent target locations in Experiment 2. Targets at 0° and 180° are half-gray/half-white because they were presented in both experiments. The hyphenated vertical line indicates the mirror reversal axis in Experiment 1. In Experiment 2, the rotations were applied relative to the start location.

Massey, 1987; Mazzoni and Krakauer, 2006; Neely and Heath, 2009; Taylor et al., 2010; Taylor and Ivry, 2011).

Experiment 2: visual rotation, feedforward control. Experiment 2 had generally the same structure as Experiment 1, with two testing sessions taking place on consecutive days. This time the participants ($N = 15$) were exposed to a 40° visual rotation instead of a mirror reversal of the cursor. As noted above, the required change in the motor command from the original to the new mapping in Experiment 1 was also 40°, such that the magnitude of the mapping change was equal in both experiments. Center-out reaching movements were executed toward eight circularly arranged targets (Fig. 2). Feedback regarding movement performance was given following the same criteria that were used for Experiment 1. Each session consisted of 16 blocks, and each block contained 72 trials, with each target appearing 9 times in random order. Again, the first 4 of the 16 blocks in the first session were training blocks and were excluded from all further analyses. This was followed by 4 baseline blocks, and 8 blocks in which a 40° visual rotation was imposed. The second session began with 12 VR blocks, followed by 4 blocks without rotation.

Experiment 3: mirror reversal, feedback control, and sleep dependency. Whereas Experiments 1 and 2 assessed learning of feedforward control, Experiment 3 was designed to also assess learning of fast feedback commands with mirror reversed visual feedback, by laterally displacing the cursor on a fraction of trials. Additionally, we tested the hypothesis that consolidation between sessions depended on sleep, motivated by the finding that sleep has been reported to benefit offline consolidation (Walker et al., 2002; Stickgold, 2005). Experiment 3 had generally the same structure as Experiments 1 and 2, using identical feedback procedures, number of trials per block, and the number of blocks per day. We tested feedback control only for the 0° target, as here no change in the feedforward command was required that could possibly confound the measurement. To increase the number of reaches to each target, we only tested targets at -20° , 0° , and 20° . Each block was divided into 9 miniblocks, and each miniblock consisted of 8 different trials (Table 1), designed to test either feedforward or feedback control. The trials within each miniblock were ordered randomly, with each trial type occurring once. To test changes in feedforward commands, reaching targets in trial Types 1 and 2 were presented at an angle of 20° or -20° from straight-ahead. As in Experiments 1 and 2, the angular hand position averaged from 100 to 150 ms after movement onset was measured for studying feedforward control. In the remaining 6 trials in each miniblock, partic-

Table 1. Trial types within every miniblock in Experiment 3^a

Trial type	1	2	3	4	5	6	7	8
Force channel	×	×	×	×	×	✓	✓	✓
Target location	-20	20	0	0	0	0	0	0
Cursor displacement	←	←	×	←	→	×	←	→

^aFor trial Types 1 and 2, each of the three cursor displacements (none, left, right) occurs only once for every 3 miniblocks. Crosses indicate the absence of cursor displacements or force channels, whereas ✓ marks and arrows indicate the presence of force channels and direction of cursor displacements, respectively.

Table 2. Experimental groups in Experiment 3 with testing sessions at different times of day^a

Blocks/groups	4 normal	8 MR	Break	12 MR	4 normal
ME	Morning		12 h	Evening	
EM	Evening		12 h	Morning	
EE	Evening		24 h	Evening	
MM	Morning		24 h	Morning	

^aBoth days consisted of 16 blocks, each containing 72 reaching movements. The first 4 blocks of day 1 were training blocks with normal visual feedback and are not listed in the table.

ipants reached to the straight-ahead target, and we tested fast feedback mechanisms. For trial types 4, 5, 7, and 8, we displaced the cursor by 1.5 cm to the left or right after the hand had traveled >1 cm from the origin. Cursor displacements elicit an automatic corrective response in the opposite direction with the aim of bringing the cursor back to the initial trajectory. This response has shorter latencies than voluntary response initiation (Franklin et al., 2008) and cannot be voluntarily suppressed.

To obtain a sensitive measure of the feedback response, we clamped the hand to a straight-line trajectory toward the target using a force channel for trial Types 6–8. These channels exerted a spring-like force of 6000 N/m. When a cursor was displaced, participants pushed into the channel wall attempting to correct for the displacement. The hand force was immediately counteracted by an equal amount of force from the robotic handle, which could then be used as a reliable measure of correction. On force channel trials, the cursor was displaced back to the original trajectory after the hand had moved >10 cm in the channel to allow the participants to reach the target. Because the automatic return of the cursor can cause attenuation of feedback responses (Franklin and Wolpert, 2008), we also added trials without channels (trial Types 4 and 5) in which the cursor was not returned. These trials therefore required a correction to reach the target. For the same reason, we also displaced, and did not return, the cursor on 2 of 3 trials in which the movement was directed at lateral targets (trial Types 1 and 2).

To determine whether performance changes between the sessions (forgetting or offline gains) depended on sleep, we assigned participants to one of four groups (Table 2). The first group (morning evening [ME]; 16 participants) had the first session in the morning and the second session 12 h later on the same day. The second group (evening morning [EM]; 15 participants) had the first session in the evening and the next session 12 h later after a night of sleep in the morning of the next day. To control for the effect of the time of day on performance, we included one control group that did both sessions in the evening (evening evening [EE]; 13 participants) and one that did both sessions in the morning (morning morning [MM]; 17 participants). For both groups, the sessions were separated by a 24 h break and a night of sleep. There were no significant age or gender differences between the four groups. Morning sessions took place between 7:30 and 10:30 AM and evening sessions between 7:30 and 10:30 PM. The role of sleep was only tested for MR, but not for VR, because no offline improvements were found for the latter.

Experiment 4: visual rotation, feedback control. Experiment 4 was designed to assess changes in fast feedback control during VR learning and was again similar in length and structure to Experiments 1–3. Movements were executed toward 8 targets. Instead of a 40° rotation, we imposed 60° or -60° rotations (balanced across 18 participants) to achieve sufficient power to detect changes in the direction of feedback corrections. On 48 of 72 trials, the cursor position was displaced by 1.5 cm once

the hand had traveled >1 cm from the origin. Because force channels are only suitable to measure feedback corrections orthogonal to the movement direction, we assessed fast feedback responses using the direction of the initial corrective response in free movements. This was measured by computing the difference in instantaneous velocity of the hand on trials with and without displacements. The cursor displacement was applied after the hand had traveled 1 cm from the start at an angle of -90° or 90° relative to the initial movement direction of the cursor, and therefore always at an angle of -30° or 150° relative to the movement direction of the hand (see Fig. 6B). An unadapted feedback response would yield an initial hand direction exactly opposing the visual displacement. For example, if the cursor was displaced -90° relative to the cursor direction (or -30° relative to the hand, hyphenated dark blue arrow), the correction should be directed toward 150° (see Fig. 6B, solid light blue arrow). A fully adapted feedback response would be rotated by 60° opposite to the imposed visual rotation, thus resulting in a 90° correction if the cursor was displaced -30° relative to the hand (see Fig. 6B, solid dark blue arrow).

Experiment 5: control experiment for feedback response. Experiment 4 relies on the assumption that the feedback response is always opposite to the cursor displacement, independent of the direction of hand movement. That is, we assumed that the visuomotor system corrects equally for displacements parallel and orthogonal to the direction of movement. To test this assumption, three participants performed reaching movements over 16 blocks toward 8 different targets without a visual rotation. We then displaced the cursor by 1.5 cm at angles of -150° , -90° , -30° , 30° , 90° , and 150° relative to the initial hand and cursor movement direction (Fig. 6A, hyphenated colored arrows). If both orthogonal and parallel displacement components are corrected equally, the correction should always be exactly opposed to the displacement (Fig. 6A, solid colored arrows). In addition, each block contained two movements without displacement toward each target.

Data analysis. The data were analyzed using custom-written MATLAB (MathWorks) routines. For all five experiments, we excluded movements where the angle between the first and the second 100 ms segment after movement onset was $>60^\circ$, as a large difference between the two segments indicates that the movement was initially not directed at the target and only corrected online thereafter. Trials with peak movement velocities <40 or >100 cm/s or RTs <50 ms or RTs >730 ms were excluded in Experiments 1–3. For Experiment 3, we further excluded channel trials where force responses exceeded 5 Newton (N) at any point in time between 150 and 400 ms after the cursor displacement. Because the main variable of interest in Experiments 4 and 5 was the corrective velocity vector, we excluded for these experiments trials where the peak velocity deviated by >25 cm/s from the median in the respective block but included all trials independent of their reaction time. Combined, these criteria led to an exclusion of 5.4% of the trials in Experiment 1, 5.5% in Experiment 2, 4.5% in Experiment 3, 4.8% in Experiment 4, and 4.4% in Experiment 5.

In Experiment 1, tradeoffs between preparation time and accuracy of the feedforward command were quantified by the slope of the simple linear regression between RT and error. A tradeoff would show up as a negative relationship between these two variables. Assessing this relationship is complicated by the fact that both RT and error reduce over the course of learning, leading to a positive relationship that could obscure existing time-accuracy tradeoffs. To account for this effect, we first removed (within each subject and block) any linear trend across the block for RT and error independently. The movements toward the peripheral targets were then assigned to 1 of 5 bins according to this relative RT. This was done for each block, each participant, and each target separately. To obtain more stable estimates, we then combined the data across all four lateral targets by mirroring results toward the -20° and 160° onto the 20° and -160° targets. Furthermore, we averaged the data across 4 blocks for each participant. As a measure of the relationship between RT and error, we performed a simple linear regression analysis with the mean RT of each bin as the independent, and the mean signed error as the dependent variable, separately for each subject and block. The slope values were then compared using paired *t* tests. The time-accuracy tradeoff for visual ro-

tations in Experiment 2 was assessed using a similar analysis, while rotating the data to combine results across all 8 targets.

In Experiment 3, we compared the state of the feedforward command across days. Because of the possible RT dependency of the feedforward command and because mean RTs could change from session to session, we determined the expected initial error for a RT of 250 ms. For this, the relationship between RT and error was fitted for each participant, each block, and each target separately. Because this relationship was slightly nonlinear, we used Gaussian Process Regression (Rasmussen, 2006), which can accommodate any smooth relationship between two variables. The values of the length scale, variance, and noise variance hyperparameters were determined by fitting the data from all subjects together for each mirror reversed block and then taking the median values.

For Experiments 4 and 5, data were combined across all targets by rotating the movement data such that the movement direction 1 cm into the movement was located at 0° because the cursor displacements were always performed at an angle relative to this initial movement direction. We then used the difference between the average instantaneous velocity vector of trials with and without displacements to compute the velocity component that was due to the corrective response.

Results

Time-accuracy tradeoff in feedforward commands

We hypothesized that the learning of mirror reversal would be associated with a new time-dependent process that maps targets to actions, whereas visual rotation learning would be supported by the recalibration of an existing control policy and should therefore require no extra processing time.

We tested this idea by enforcing fast RTs in all reaching tasks. For MR learning (Experiment 1; Fig. 3A), RTs increased at the onset of MR by 145 ms (± 18 ms SE), $t_{(14)} = -8.232$, $p < 0.9 \times 10^{-7}$. RTs reached a plateau in the late MR blocks of the second session and approached the levels of the baseline performance. However, when the visual feedback switched back to the nonreversed mapping in block 25, RTs increased at first but subsequently decreased to 272 ms (± 5 ms) in the last block, yielding almost significantly shorter RTs than the last MR block ($t_{(14)} = 2.123$, $p = 0.052$). Thus, even after 2 d of training, movements in a MR environment required slightly more preparation time than in the normal environment.

For the equivalent VR experiment (Experiment 2, Fig. 3B), we expected RT to increase to a lesser degree, if at all. Average RT increased by 45 ms (± 8 ms) when the rotation was first introduced ($t_{(28)} = -2.918$, $p = 0.007$) (Fig. 3B). Thus, the increase of RT during VR learning was considerably smaller than the increase during MR learning ($t_{(28)} = -5.170$, $p = 1.74 \times 10^{-5}$). During the second day of training, none of the VR blocks differed significantly from baseline anymore (block 13: $t_{(14)} = -1.683$, $p = 0.114$). After the rotation had washed out (last block), the RTs were not significantly shorter than in the fourth block of training ($t_{(14)} = -1.256$; $p = 0.23$). Thus, we found that visual rotations induced less than one-third of the RT increase compared with mirror reversals.

Our main prediction, however, was that the difference between the two learning mechanisms should become visible in a time-accuracy tradeoff: the fact that, for a given adaptation state, trials with longer RTs show smaller errors. Because RTs as well as movement errors decreased over the course of the experiment, we first subtracted out any possible linear relationship between trial number and error and between trial number and reaction time for each participant and block separately in the MR and the VR conditions. We then plotted the initial movement direction of the hand (averaged from 100 to 150 ms after movement onset) as a function of RT for different groups of 4 blocks (Fig. 4). For MR learning (Experiment 1; Fig. 4A), baseline reaching angles were




Figure 3. Group-average reaction time across Experiments 1 and 2. White background represents reaching under normal visual feedback, whereas gray background represents reaching during mirror reversed or rotated visual feedback. The vertical line indicates the break between sessions. **A**, RT for -160° , -20° , 20° , and 160° targets during mirror reversal learning (Experiment 1). **B**, RT for reaching toward 8 targets during VR (Experiment 2). Error bars indicate between-subject SE.

offset from zero by $\sim 5^\circ$, indicating that participants showed a bias toward moving in the straight forward or backward direction (for how angles were combined across targets, see Fig. 3 legend), an effect likely caused by the unequal distribution of targets around the circle.

To determine whether there was a time-accuracy tradeoff, we calculated the regression slope between error and RT across bins (see Materials and Methods) (Fig. 4). In the MR experiment (blocks 1–4), there was a small, but significant, negative slope ($t_{(14)} = -4.477$, $p = 0.001$) during baseline. With the beginning of MR learning (blocks 5–8), the slope became significantly more negative compared with baseline ($t_{(14)} = 5.004$, $p = 1.93 \times 10^{-4}$). For long RTs, participants produced the correctly mirrored movements. However, for the fastest RT bin, movements started in the direction of the visually presented target, rather than in the opposite, correct direction; the error was significantly larger than 20° , where a 20° error signifies a movement toward the mirror reversal axis ($t_{(14)} = 3.812$, $p = 0.001$). As training proceeded, the relationship between RT and movement error retained similar slopes across all groups of 4 blocks (repeated-measures ANOVA with groups of 4 blocks as within-subject factor: $F_{(4,56)} = 0.588$, $p = 0.673$). Even in the end of training in Experiment 1, the difference in the RT-error relationship was still significant compared with baseline ($t_{(14)} = 3.995$, $p = .001$). However, the time-accuracy curve shifted sideways, such that higher accuracies could be achieved at shorter RTs. To quantify this observation, we calculated the RT necessary to reduce the error to 12° , as this time point allowed for assessment for all groups of 4 blocks of the experiment (Fig. 4A), by assuming an approximately linear relationship between error and RT in the range tested here and linearly predicting the reaction time for an error of 12° for each participant and quadruple of blocks. We found significant differences between blocks 5–8 and blocks 9–12 ($t_{(14)} = 2.405$, $p = 0.031$), blocks 13–16 ($t_{(14)} = 4.836$, $p = 2.64 \times 10^{-4}$), blocks 17–20 ($t_{(14)} = 3.769$, $p = 0.002$), and blocks 21–24 ($t_{(14)} = 3.860$, $p = 0.002$). Likewise, we found significant horizontal shifts between blocks 9–12 and blocks 13–16 ($t_{(14)} = 2.806$, $p = 0.014$), blocks 17–20 ($t_{(14)} = 3.405$, $p = 0.004$), and blocks 21–24 ($t_{(14)} = 3.353$, $p = 0.005$), meaning that each curve on day 2 was significantly shifted compared with each curve on day 1. In other words, MR training led to automatization of the new target-to-movement mapping, visible in a shift of the time-accuracy tradeoff.

In contrast, we hypothesized that VR learning (Experiment 2) is achieved by the recalibration of an existing control policy. Participants should therefore be able to exploit the automaticity of the old mapping even during learning and should thus not require additional time for processing. Hence, we predicted that, for VR learning, longer reaction times should not result in lower errors. This is indeed what we found (Fig. 4B). At baseline, there was a small, but significantly positive, relationship between error and RT ($t_{(14)} = 3.453$, $p = 0.004$). However, with the introduction of the VR, this relationship did not change (t test between the slopes of blocks 5–8 and blocks 1–4: $t_{(14)} = -1.442$, $p = 0.171$). Thus, although angular errors increased as soon as the visual display was rotated (blocks 5–8), longer RTs did not result in smaller errors. In subsequent blocks, the error reduced further, but no change in the dependency on RT was observed (t test between the slopes of blocks 21–24 and blocks 1–4: $t_{(14)} = 0.503$, $p = 0.623$).

Although the range of RTs between Experiments 1 and 2 were slightly different, the RT distribution overlapped considerably, especially for the later learning phases. To compare the MR and VR conditions directly, we recalculated the slopes between RT and reach angle for the fastest 4 bins during MR and the slowest 4 bins during VR learning, such that the average reaction time used for calculating the slopes in MR (292 ± 9 ms) and VR (279 ± 9 ms) were not significantly different ($t_{(28)} = 1.053$, $p = 0.301$). After subtracting the baseline slopes from all other phases we found that in all phases, there was a significant difference between the time-accuracy slope of the MR and VR learning conditions (blocks 5–8: $t_{(28)} = 4.429$, $p = 1.4 \times 10^{-4}$; blocks 9–12: $t_{(28)} = 5.101$, $p = 2.1 \times 10^{-5}$; blocks 13–16: $t_{(28)} = -4.781$, $p = 5.05 \times 10^{-5}$; blocks 17–20: $t_{(28)} = 3.420$, $p = 0.002$; blocks 21–24: $t_{(28)} = -4.401$, $p = 1.4 \times 10^{-4}$). Thus, over a comparable range of RTs, the MR learning group clearly showed a significantly stronger dependency of accuracy on RT than the VR learning group.

Adaptation of fast feedback responses

A second window of insight into how computations in the motor system unfold over time is to investigate fast feedback responses. If a new control policy requires more time to compute a motor command, then the feedback responses after learning should also be delayed, or possibly the early responses should be dominated by the old policy. If, however, an existing policy was recalibrated,

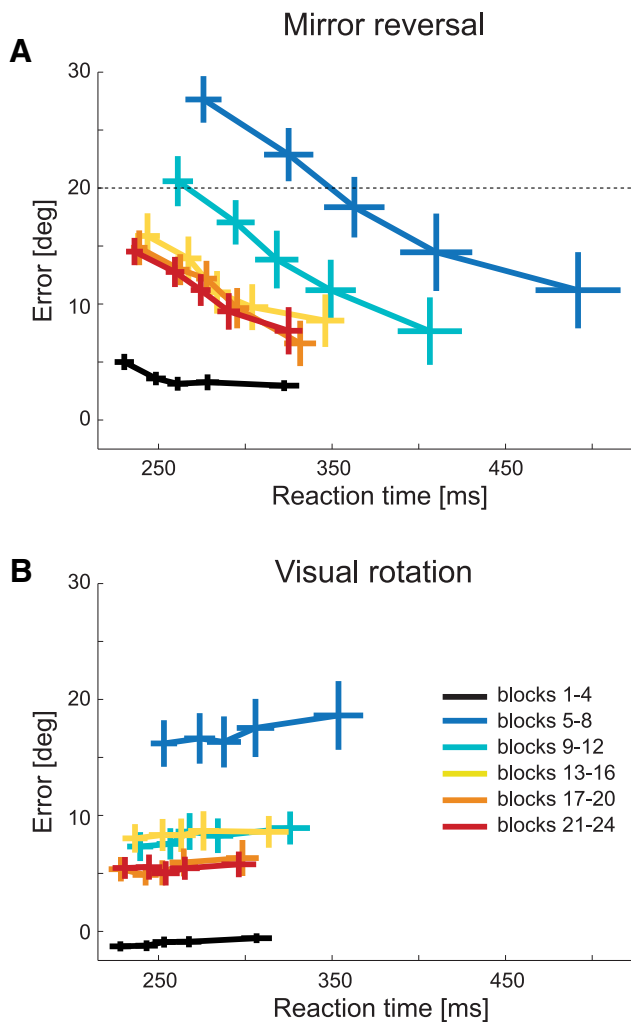


Figure 4. Relationship between RT and the directional error in Experiments 1 and 2. Blocks 1–4 were collected during baseline and blocks 5–24 during MR or VR. The trials were binned by RT for each target, participant, and block. Visual feedback was veridical during blocks 1–4 and mirror reversed or rotated during blocks 5–24. Blocks 1–12 were measured during the first session, blocks 13–24 during the second session. **A**, Mirror reversal: visual errors from movements toward the -160° and 20° target were flipped to allow averaging with errors from the -20° and 160° targets. Visual errors $<20^\circ$ indicate that the hand reached into the wrong (unmirrored) direction. Completely unadapted responses would yield an error of 40° . **B**, Visual rotation. A completely unadapted response would result in an error of 40° . Error bars indicate between-subject SE.

then both early and late components of the feedback response should adapt simultaneously.

To address this question for mirror reversal learning, Experiment 3 probed the reactions of the arm to sudden displacements of the cursor (Sarlegna et al., 2003). We then calculated the difference between force responses to left and rightward cursor jumps and halved it to inspect the temporal evolution of the feedback correction in different groups of 4 blocks of the experiment. Fig. 5 shows the results averaged across the four consolidation conditions. During unmirrored baseline movements, the corrective response began ~ 110 ms after the onset of the displacement and reached ~ 1 N after 250 ms. In the first 4 mirror-reversed blocks (blocks 5–8), it still reached ~ 0.8 N in the same direction but became less sustained thereafter; in the time window 250–350 ms, it was significantly lower than during baseline, $t_{(60)} = 8.35$, $p = 1.2 \times 10^{-11}$. This unreversed response would increase the visual error, rather than compensate for it (Fig. 1). In




Figure 5. Relationship between time and feedback response during mirror reversal learning (Experiment 3). Shown is the force measured in the channel produced in reaction to a 1.5 cm cursor displacement. Blocks 1–4 were collected during baseline and blocks 5–24 during MR. The hyphenated line indicates the reversed baseline response to serve as an illustration of what a perfectly mirror reversed feedback response would have looked like. Shaded area represents between-subject SE.

blocks 9–12, the force response further decreased but still did not reverse. Only during the second day (blocks 13–24) did we observe a reversal of the force response in the time window 250–350 ms (blocks 13–16, -0.14 ± 0.038 N, $t_{(60)} = -3.695$, $p = 4.8 \times 10^{-4}$). Yet, even in blocks 21–24, the initial incorrect force response was not fully abolished: in the time window between 130 and 200 ms, it remained significantly positive (0.13 ± 0.018 N, $t_{(60)} = 8.028$, $p = 4.3 \times 10^{-11}$).

In sum, feedback responses during MR learning provide a very similar picture as feedforward responses. While the system generates correct movements after additional processing time, the fast and automatic responses remained unadapted even after 2 training sessions (Gritsenko and Kalaska, 2010). The data clearly showed a progression of learning in which the correct response was progressively generated at shorter delays, suggesting that the new control policy, which was initially rather slow, became automatized.

Determining how feedback commands adapt during visual rotation is more challenging, as the adapted and unadapted responses are not opposite to each other but differ only by the imposed rotation angle. To amplify the contrast, we conducted another study (Experiment 4) in which participants adapted to either a 60° or a -60° rotation and probed feedback responses by displacing the cursor orthogonally to the cursor movement ($\pm 90^\circ$, Fig. 6B, hyphenated dark blue and red arrows). In the condition in which the cursor was rotated by 60° , the effective visual displacement was in a direction -30° and 150° relative to the hand movement. For a fully adapted feedback response, the hand should correct orthogonally to the hand trajectory as before (Fig. 6B, solid red and dark blue arrows). In contrast, if the feedback response is unadapted, the correction should be opposite to the visual displacement (i.e., 150° or -30° relative to the hand movement direction) (Fig. 6B, solid orange and light blue arrows).

The latter prediction, however, relies on the assumption that participants would correct their hand movement opposite to the visual cursor displacements, even if the displacement were not orthogonal to the movement direction. Because it is possible that the motor system reacts less to the component of the visual displacement in the direction of the movement, we tested our as-

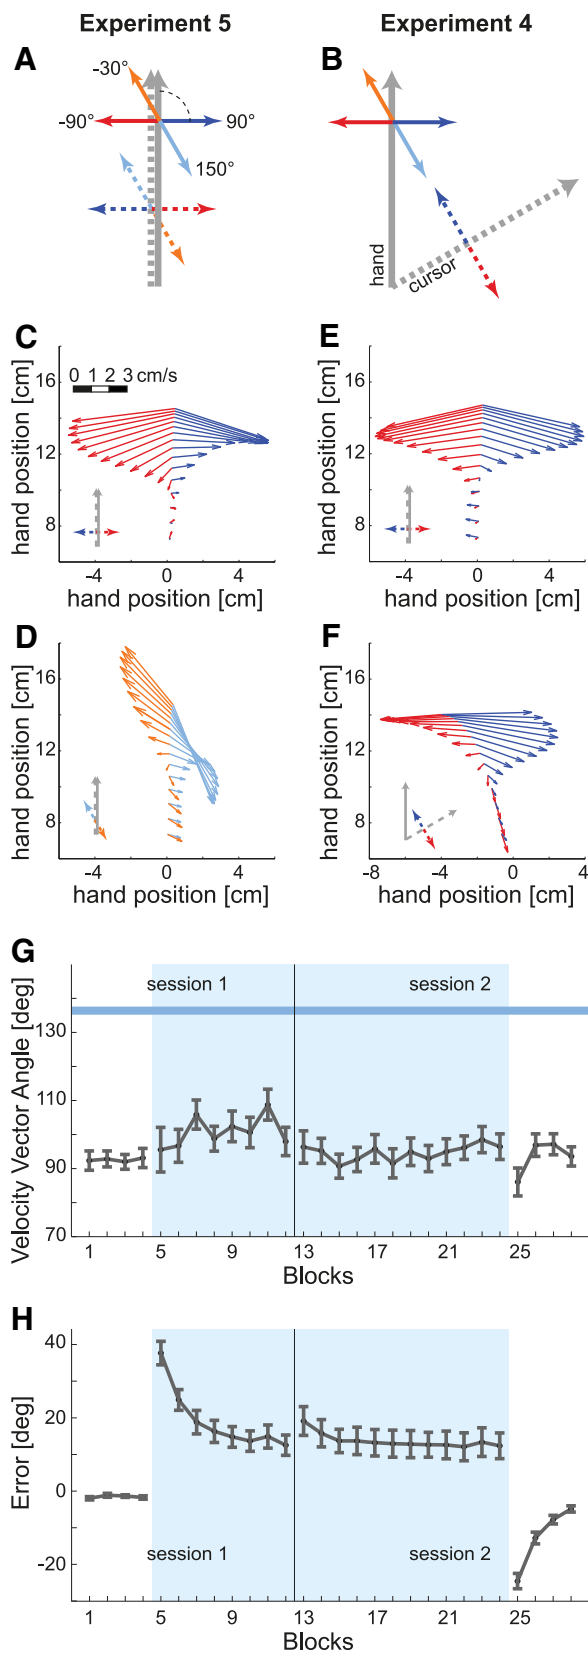


Figure 6. Feedback responses in Experiments 4 and 5. **A**, In Experiment 5, the cursor (hyphenated gray line) and the hand (solid gray line) moved in the same direction. The cursor was displaced (hyphenated colored arrows) at an angle of -90° (dark blue), -30° (light blue), 90° (red), or 150° (orange) relative to the movement direction. Displacements also occurred in 30° and -150° directions (data not shown). The hand movements that cancel out the cursor displacements are shown as solid arrows of the same color. **B**, In Experiment 4, the cursor

(hyphenated gray line) was rotated by 60° or -60° (only the 60° is shown in the schematic) from the hand movement (solid gray line). Displacements were -90° (blue hyphenated) or 90° (red hyphenated) relative to the movement direction of the cursor. The solid red and dark blue arrows indicate the required hand movement directions that cancel out the corresponding displacement (hyphenated arrow with the same color). The orange and the light blue arrows represent what an unadapted response would look like. **C**, Quiver plot of feedback responses in Experiment 5 to -90° (dark blue) and 90° (red) cursor displacements. The vector origin represents the average hand position at time points from 75 to 375 ms after the cursor displacement (20 ms resolution), and the vector the difference in instantaneous hand velocity between trials with and without displacement. **D**, Feedback responses to -30° (light blue) and 150° (orange) cursor displacements in Experiment 5. **E**, Response to -90° (dark blue) and 90° (red) cursor displacements during baseline reaching (i.e., before cursor rotation in Experiment 4) and **F** with rotated cursor (blocks 5–8). Results are shown averaged over the 60° and -60° rotation groups, by right-left flipping the results for the -60° group. **G**, Mean angular direction of feedback correction (\pm SE) 250–350 ms after the displacement plotted over all blocks of Experiment 4. Responses are combined across cursor displacements and rotation groups. Light blue background represents blocks with visual rotation. Blue line and shading represent prediction of fully unadapted feedback response, based on mean and SE of responses to oblique cursor displacement in Experiment 5. **H**, Mean angular error of the feedforward command (\pm SE) averaged from 100 to 150 ms after movement onset while adapting to the 60° rotation in Experiment 4 for comparison.

assumption in an additional experiment. In Experiment 5, we displaced the cursor by 1.5 cm at an angle of $\pm 30^\circ$, $\pm 90^\circ$, and $\pm 150^\circ$ relative to hand and cursor movement (Fig. 6A). Even for the oblique angles, the initial correction should be exactly opposite to the cursor displacements.

We used the difference between the instantaneous velocity vectors between trials with and without displacements at different time points after the displacement as a measure of the corrective response. We found that, for the 90° displacements under the natural mapping, the velocity difference vectors were slightly tilted downward, meaning that the hand not only corrected in the appropriate direction, but also decelerated along the main direction of movement (Fig. 6C). To summarize the effects across displacement directions offline, we rotated the correction vector for the -90° displacements by 180° , effectively canceling out any decelerating effect.

For oblique displacements, we found that the corrections were approximately opposite to the displacement (Fig. 6D). To analyze the responses together, we inverted the horizontal component of the responses to the 150° and 30° displacements, and the vertical component of the responses to the $\pm 150^\circ$ displacement, such that all corrections would superimpose with the correction for the -30° displacements (which requires a 150° correction for full cancellation). The angle of the resulting correction was 136.4° ($\pm 9.1^\circ$), slightly less than the ideal response of 150° , indicating that participants reacted to displacements in movement direction slightly less than to displacements orthogonal to it. Thus, based on these results, we would expect that a fully unadapted feedback response to an anticlockwise (-90°) cursor displacement under a 60° cursor rotation should be 136.4° .

In Experiment 4, we averaged the results of the 60° and -60° rotation groups, by flipping the trajectories for the group that underwent the -60° rotation. The average feedback responses during VR learning (Fig. 6F) did not resemble the feedback responses observed in the control experiment (Fig. 6D). Rather, the corrections were oriented -90° and 90° relative to the movement direction. In other words, the feedback response in VR appeared to be immediately oriented in the correct direction (Fig. 6G). Although we cannot directly compare the forces measured in Experiment 3 with the velocity vectors measured in Experiment

←

(hyphenated gray line) was rotated by 60° or -60° (only the 60° is shown in the schematic) from the hand movement (solid gray line). Displacements were -90° (blue hyphenated) or 90° (red hyphenated) relative to the movement direction of the cursor. The solid red and dark blue arrows indicate the required hand movement directions that cancel out the corresponding displacement (hyphenated arrow with the same color). The orange and the light blue arrows represent what an unadapted response would look like. **C**, Quiver plot of feedback responses in Experiment 5 to -90° (dark blue) and 90° (red) cursor displacements. The vector origin represents the average hand position at time points from 75 to 375 ms after the cursor displacement (20 ms resolution), and the vector the difference in instantaneous hand velocity between trials with and without displacement. **D**, Feedback responses to -30° (light blue) and 150° (orange) cursor displacements in Experiment 5. **E**, Response to -90° (dark blue) and 90° (red) cursor displacements during baseline reaching (i.e., before cursor rotation in Experiment 4) and **F** with rotated cursor (blocks 5–8). Results are shown averaged over the 60° and -60° rotation groups, by right-left flipping the results for the -60° group. **G**, Mean angular direction of feedback correction (\pm SE) 250–350 ms after the displacement plotted over all blocks of Experiment 4. Responses are combined across cursor displacements and rotation groups. Light blue background represents blocks with visual rotation. Blue line and shading represent prediction of fully unadapted feedback response, based on mean and SE of responses to oblique cursor displacement in Experiment 5. **H**, Mean angular error of the feedforward command (\pm SE) averaged from 100 to 150 ms after movement onset while adapting to the 60° rotation in Experiment 4 for comparison.

4, these results contrast starkly with the slow and incomplete adaptation of fast feedback responses during MR learning.

Our results therefore suggest a fundamental difference in the way in which MR and VR are learned. MR learning initially requires extra processing time to compute accurate feedforward and feedback commands, indicating that it may involve the establishment of a new control policy. Although the new motor commands could be generated more quickly after 2 d of training, it remained dependent on processing time. In contrast, VR learning did not show such dependency, even early in learning—consistent with the idea that here a fully automatized control policy was recalibrated.

Offline gains in performance between sessions

With the shifting time-accuracy tradeoff, MR learning shares an important feature with other motor learning tasks (Beilock et al., 2008). It has been recently suggested that such shifts should be considered the defining feature of “skill learning” (Reis et al., 2009; Shmuelof et al., 2012). Another characteristic of many tasks that are considered “skill” tasks concerns consolidation between sessions: For example, for learning of sequential movements, performance levels typically deteriorate very little overnight (Rickard et al., 2008) and sometimes even appear to show offline gains (Stickgold, 2005; Wright et al., 2010; Abe et al., 2011). In contrast, adaptation tasks that require a recalibration of an existing control policy nearly universally show some decay of the motor memory during an intervening interval (Tong et al., 2002; Klassen et al., 2005; Krakauer et al., 2005; Trempe and Proteau, 2010). If this different temporal dynamic of consolidation can be attributed to the suggested distinction of automatization of a new control policy versus recalibration of an existing control policy, then MR learning should show offline gains in the break between the two sessions, whereas VR learning should show offline forgetting.

Offline gains in skill learning experiments are often reported to depend on sleep (Walker et al., 2002; Cohen et al., 2005; Robertson et al., 2005; Stickgold, 2005; Rickard et al., 2008). For MR learning in Experiment 3, we therefore randomly assigned the participants to one of four groups. The ME group had the first session in the morning and the second session in the evening of the same day, and therefore did not have a night of sleep between the two sessions. The EM group had the first session in the evening and the next session in the morning of the next day. Both of these groups had a break of 12 h between their two sessions. To test whether potential differences depended on the time of day of the first or second session, rather than on the presence or absence of sleep, we included two additional groups, which performed the experiment either on the mornings (MM) or on the evenings (EE) of 2 consecutive days. If consolidation really depended on sleep but not time of day, then only the ME group (the only group without sleep) should show worse consolidation than any of the other three groups, whereas the other three groups should not differ from each other.

Because error depended on RT and because RT may differ from one session to the next, we quantified the skill level as the movement error that the participant would show for a fixed RT. The slightly nonlinear relationship between error and RT was fitted using Gaussian Process Regression (see Materials and Methods), and we then simply read off the movement error for an RT of 250 ms. Errors from movements toward the 20° target were inverted, so that the RT-corrected directional error for both peripheral targets could be averaged.

We found that MR learning did not show forgetting between sessions, but rather offline gains in performance (Fig. 7). Across

all groups, there was a significant improvement in feedforward performance from the last block of the first session to the first block of the second session ($t_{(60)} = -4.72, p = 1.4 \times 10^{-5}$). Tested individually, the EM group ($t_{(14)} = -2.678, p = 0.018$), the EE group ($t_{(12)} = -3.174, p = 0.008$), and the MM group ($t_{(16)} = -2.138, p = 0.048$) all significantly improved overnight. The only group that did not show significant improvements was the ME group ($t_{(15)} = -1.872, p = 0.081$), which did not have a night of sleep between the two sessions. However, there was no significant direct difference between the group without sleep and the groups with a night of sleep between the two sessions in terms of their change in movement error from Session 1 to Session 2 ($t_{(59)} = -1.471, p = 0.147$).

Offline gains were even more clearly visible in the feedback corrections (Fig. 8). For this analysis, we averaged the feedback response (Fig. 5) over the interval from 250 to 350 ms after the displacement, as this time period showed the most profound learning-related changes. Again, all participants combined showed very strong offline gains ($t_{(60)} = -4.637, p = 1.9 \times 10^{-5}$). We also plotted this measure as a function of block for all four groups separately. The EM group ($t_{(14)} = 2.265, p = 0.04$), the EE group ($t_{(12)} = 3.011, p = 0.011$), and the MM group ($t_{(16)} = 2.656, p = 0.017$) showed significant increases in performance from one session to the next. The only group that did not show improvements was the ME group ($t_{(15)} = 1.189, p = 0.253$) (i.e., the group that did not have a night of sleep between the two sessions). The groups with sleep had only marginally stronger offline gains than the group without sleep ($t_{(59)} = 1.837, p = 0.071$), indicating that offline improvements may have been enhanced by sleep. There was no significant effect of time of day of the first ($t_{(59)} = 1.220, p = 0.227$) or the second session ($t_{(59)} = 0.650, p = 0.518$), nor was there an effect of the duration of the break between the sessions ($t_{(59)} = 1.314, p = 0.194$). Together, these results clearly demonstrate the existence of offline gains during MR learning. In respect to the sleep dependency of this effect, our results remain inconclusive. Even though there are some trends in the data that indicate that an intermitted night of sleep may amplify this effect, the direct comparison of the groups failed to reach significance.

In contrast to MR, VR learning showed clear forgetting between sessions, in line with many other adaptation tasks (Tong et al., 2002; Klassen et al., 2005; Krakauer et al., 2005; Trempe and Proteau, 2010). Although we did not find a significant relationship between RT and angular error, we used, for the sake of consistency, the same method for RT correction as for the MR data. Within the first day, the initial error reduced from 24.4° ($\pm 2.1^\circ$) to 8.7° ($\pm 1.5^\circ$) (Fig. 7E). When participants returned on the second day, their error had increased again to 14.7.2° ($\pm 3^\circ$). Angular errors in the first block of the second session were significantly larger than angular errors in the last block of the first session ($t_{(28)} = -2.192, p = 0.049$; Fig. 7F). Thus, our results confirm previous literature showing that adaptation is forgotten between sessions, and provide evidence for a clear dissociation from MR learning, for which offline gains are observed.

Discussion

We directly contrasted learning of two different visuomotor transformations. For MR learning, we found a clear RT dependency of initial movement error, with faster responses leading to larger errors than slow responses. We hypothesized that MR learning involves the acquisition of a new sensorimotor mapping, which initially takes more time than the old mapping to perform the necessary computations. Therefore, under strict time con-




Figure 7. Consolidation of the feedforward command in Experiments 2 and 3. Average angular errors 100–150 ms after movement onset are plotted over the different blocks of the experiment (A–D) for the four mirror reversal groups (Experiment 3) and (E) the visual rotation group (Experiment 2). The error is corrected for the influence of time-accuracy tradeoff by calculating the average error at RT = 250 ms (see Materials and Methods). Colored background represents blocks with mirror reversal or visual rotation. The vertical hyphenated line separates the two sessions. All mirror reversal groups performed as well or better in the first block of the second session than in the last block of the first session. F, Bar graph of the difference in error between the first block in the second session (block 13) and the last block in the first session (block 12) split up by the visual rotation and the four mirror reversal groups: ME, EM, EE, MM; and the VR group. * $p < 0.05$ (significant t test against 0). Error bars indicate between-subject SE.

straints, the response was still dictated by the old mapping. With 2 d of training, we found that the new mapping became increasingly automatic, achieving the same movement error at shorter RTs. It did not, however, achieve the same automaticity as the baseline mapping.

For VR learning, movement error did not decrease with increasing RT. We propose that this form of motor learning relies on the recalibration of an already existing mapping and therefore can exploit the established automaticity of the underlying computational processes. Thus, in this view, the appearance of a time-accuracy tradeoff at the beginning of learning with subsequent shifts of this relationship is a cardinal sign that the motor system acquires and automatizes a new mapping from goals to motor commands (Shmuelof et al., 2012).

Intriguingly, we found a parallel dissociation between MR and VR learning during fast feedback responses to displacements of the visual cursor. For MR learning, the corrective response was initially directed into the wrong direction, even after 2 d of training (Day and Lyon, 2000; Gritsenko and Kalaska, 2010) and reversed only in the late phases of the response. Thus, feedforward and feedback control both require additional processing time in the beginning of learning and then are increasingly automatized.

In contrast, the feedback command during VR learning appeared to be fully adapted immediately. It has been suggested that feedback responses during large VR must adapt rapidly within a single trial because the hand would otherwise circle around the target (Braun et al., 2009a). Another explanation might be that the feedback command does not need to adapt at all because it




Figure 8. Consolidation of the feedback command in Experiment 3. The average feedback command 250–350 ms after the displacement is plotted over different blocks of the experiment. Colored background represents mirror reversal of the visual feedback. **A–D**, Feedback commands of the four mirror reversal groups. **E**, Bar graph of the force differences between the first block in the second session (block 13) and the last block in the first session (block 12) split up by the four groups: ME, EM, EE, and MM. * $p < 0.05$ (significant t test against zero). Error bars indicate between-subject SE.

always bases its reactions on the relative angle between the displacement and the visually observed trajectory. Whatever the exact mechanism, the presence of time-accuracy tradeoffs in MR, and their absence during VR, provides clear evidence that the two visual transformations are learned via separate processes.

A previous study found a relationship between RT and how quickly participants learned a 60° visual rotation (Fernandez-Ruiz et al., 2011). However, in this study, RTs were unconstrained and on average 400–600 ms. We argued that unconstrained RTs may have invited strategic replanning of the endpoint (Mazzoni and Krakauer, 2006; Taylor et al., 2010; Taylor and Ivry, 2011), a process more related to an explicit mental rotation of the desired movement direction (Georgopoulos and Massey, 1987; Neely and Heath, 2009) than to visuomotor adaptation. Indeed, when RTs were constrained to <350 ms, as in our study, no evidence for a time-accuracy tradeoff in VR learning was found. These results therefore argue that even visual rotations are not always learned purely through recalibration of an existing control policy: without speed constraints, additional time-consuming processes

(strategic remapping) can help to improve performance more quickly.

Why does the brain have to learn a new control policy for mirror reversals whereas it appears to recalibrate an existing control policy for visual rotations? At a computational level of description (Marr and Poggio, 1976), MR and VR learning seems to be comparably difficult. Both can be described with a simple change in the function that transforms visual inputs into arm movements. However, what is difficult for the brain has to be viewed in the context of its prior experience. In ambiguous situations, the motor system appears to interpret visuomotor errors as being caused by VR (Turnham et al., 2011), possibly reflecting inherent assumptions about the structure of the environment. These priors can be changed through repeated exposure to different environments, a process termed structural learning (Braun et al., 2009b). Viewed in this framework, MR learning would be slow as it violates the learned structure of possible visuomotor transformations, requiring the slow acquisition of a new structure. A related explanation is based on the assumption that a

visuomotor mapping is adapted by adding some part of the corrective response under the old mapping to the old motor command (Kawato and Gomi, 1992). VRs up to 90° could be learned like this, whereas for MR the initial corrective response would point in the wrong direction (Fig. 1), again requiring the establishment of a new control policy. This hypothesis would make the yet to be tested prediction that rotations >90° should also show time-accuracy tradeoffs. Indeed, it has been suggested that such large rotations are learned by different mechanisms (Abeele and Bock, 2001).

Rather than providing a clear computational-level explanation, the main empirical contribution of the paper is to show that MR and VR learning clearly differs in their time-accuracy tradeoff, both in feedforward and feedback control. We hypothesize that these tradeoffs are tightly related to the tradeoff between movement speed and accuracy, as faster movements impose tighter time constraints on feedback processes. Consistent with our interpretation, shifts in such speed-accuracy tradeoffs have been interpreted as a sign of the establishment of a new control policy (Haith and Krakauer, 2013). Following this definition, the learning of new trajectories (Shmuelof et al., 2012), finger sequences (Karni et al., 1995), or finger configurations (Waters-Metenier et al., 2014) should have some similarity to MR learning.

Our second main finding is that the presence of a time-accuracy tradeoff is associated with how the learned behavior consolidates between sessions. For VR learning for which no time-accuracy tradeoff was found, forgetting occurred between sessions. This is in line with other studies of adaptation (Kassardjian et al., 2005; Krakauer et al., 2005; Galea et al., 2011). For MR learning, we found clear evidence for offline gains, both in the feedforward and the feedback command. So far, offline gains have mainly been reported for motor learning of sequential movements (Robertson et al., 2004). Our study provides to our knowledge the first reported instance of offline improvement for learning of visuomotor transformations during reaching movements.

There has been an extensive debate on whether true offline gains in sequential finger movements depend on sleep (Stickgold, 2005; Wright et al., 2010; Abe et al., 2011). Our results do not allow for a definite conclusion in the MR learning task: For both feedback and feedforward commands, we found trends indicating that offline gains are brought about by sleep; however, a direct comparison of the different MR groups did not reach statistical significance. Thus, our failure to find evidence of sleep dependency may be partly due to a lack of power, and the relationship between sleep and memory in this context may warrant further study.

The presence of a time-accuracy tradeoff and offline gains suggests that the learning mechanisms that underlie MR and 40° VR have different physiological underpinnings. Specifically, one may speculate that the establishment of a new control policy relies on corticostriatal circuits. Indeed, Gutierrez-Garralda et al. (2013) showed that basal ganglia patients exhibit normal learning in a dart throwing task when the visual scene is horizontally displaced but impaired performance when the visual scene is mirror reversed (but see Stebbins et al., 1997; Laforce and Doyon, 2001). The basal ganglia have been associated with action selection (Gerardin et al., 2004) and the acquisition of new control policies (Doya, 2000; Middleton and Strick, 2000; Hikosaka et al., 2002; Boyd et al., 2009; Doyon et al., 2009). In addition, Parkinson's and Huntington's disease patients are impaired in learning sequential finger movements and learning of other novel tasks (Gerardin et al., 2004; Boyd et al., 2009; Penhune and Steele, 2012). In

contrast, the adaptation of eye movements (Takagi et al., 1998, 2000), arm movements (Martin et al., 1996; Tseng et al., 2007), and gait (Reisman et al., 2007) heavily depends on the integrity of the cerebellum, whereas basal ganglia-associated disorders affect adaptation to a lesser degree (Fernandez-Ruiz et al., 2003; Marinelli et al., 2009; Gutierrez-Garralda et al., 2013).

A strict dissociation between the cerebellum as the substrate for adaptation/recalibration and the basal ganglia as the substrate for control policy acquisition has recently been called into question with increasing evidence that the cerebellum is involved in both adaptation and "skill learning" (Penhune and Steele, 2012). Cerebellar patients are impaired in dart throwing tasks with horizontally shifted as well as with mirror reversed visual feedback (Sanes et al., 1990; Vaca-Palomares et al., 2013).

To date, it has been very difficult to determine whether any differences found between adaptation and skill-learning tasks can be truly attributed to the underlying learning mechanism or the differences between the tasks that are used to measure them. Here we demonstrate that the two mechanisms are differently engaged in the learning of two different visuomotor mappings during reaching movements. The current paradigm may therefore be ideally suited for studying the neural correlates of acquisition and recalibration of control policies using functional imaging or neurophysiologic recordings within a single task.

References

- Abe M, Schambra H, Wassermann EM, Luckenbaugh D, Schweighofer N, Cohen LG (2011) Reward improves long-term retention of a motor memory through induction of offline memory gains. *Curr Biol* 21:557–562. [CrossRef Medline](#)
- Abeele S, Bock O (2001) Sensorimotor adaptation to rotated visual input: different mechanisms for small versus large rotations. *Exp Brain Res* 140:407–410. [CrossRef Medline](#)
- Beilock SL, Bertenthal BI, Hoerger M, Carr TH (2008) When does haste make waste? Speed-accuracy tradeoff, skill level, and the tools of the trade. *J Exp Psychol Appl* 14:340–352. [CrossRef Medline](#)
- Boyd LA, Edwards JD, Siengsukon CS, Vidoni ED, Wessel BD, Linsdell MA (2009) Motor sequence chunking is impaired by basal ganglia stroke. *Neurobiol Learn Mem* 92:35–44. [CrossRef Medline](#)
- Braun DA, Aertsen A, Wolpert DM, Mehring C (2009a) Learning optimal adaptation strategies in unpredictable motor tasks. *J Neurosci* 29:6472–6478. [CrossRef Medline](#)
- Braun DA, Aertsen A, Wolpert DM, Mehring C (2009b) Motor task variation induces structural learning. *Curr Biol* 19:352–357. [CrossRef Medline](#)
- Brawn TP, Fenn KM, Nusbaum HC, Margoliash D (2010) Consolidating the effects of waking and sleep on motor-sequence learning. *J Neurosci* 30:13977–13982. [CrossRef Medline](#)
- Cohen DA, Pascual-Leone A, Press DZ, Robertson EM (2005) Off-line learning of motor skill memory: a double dissociation of goal and movement. *Proc Natl Acad Sci U S A* 102:18237–18241. [CrossRef Medline](#)
- Day BL, Lyon IN (2000) Voluntary modification of automatic arm movements evoked by motion of a visual target. *Exp Brain Res* 130:159–168. [CrossRef Medline](#)
- Doya K (2000) Complementary roles of basal ganglia and cerebellum in learning and motor control. *Curr Opin Neurobiol* 10:732–739. [CrossRef Medline](#)
- Doyon J, Bellec P, Amsel R, Penhune V, Monchi O, Carrier J, Lehérycy S, Benali H (2009) Contributions of the basal ganglia and functionally related brain structures to motor learning. *Behav Brain Res* 199:61–75. [CrossRef Medline](#)
- Fernandez-Ruiz J, Diaz R, Hall-Haro C, Vergara P, Mischner J, Nuñez L, Drucker-Colin R, Ochoa A, Alonso ME (2003) Normal prism adaptation but reduced after-effect in basal ganglia disorders using a throwing task. *Eur J Neurosci* 18:689–694. [CrossRef Medline](#)
- Fernandez-Ruiz J, Wong W, Armstrong IT, Flanagan JR (2011) Relation between reaction time and reach errors during visuomotor adaptation. *Behav Brain Res* 219:8–14. [CrossRef Medline](#)
- Franklin DW, Wolpert DM (2008) Specificity of reflex adaptation for task-relevant variability. *J Neurosci* 28:14165–14175. [CrossRef Medline](#)

- Galea JM, Vazquez A, Pasricha N, de Xivry JJ, Celnik P (2011) Dissociating the roles of the cerebellum and motor cortex during adaptive learning: the motor cortex retains what the cerebellum learns. *Cereb Cortex* 21:1761–1770. [CrossRef Medline](#)
- Gentilucci M, Daprati E, Toni I, Chieffi S, Saetti MC (1995) Unconscious updating of grasp motor program. *Exp Brain Res* 105:291–303. [CrossRef Medline](#)
- Georgopoulos AP, Massey JT (1987) Cognitive spatial-motor processes: 1. The making of movements at various angles from a stimulus direction. *Exp Brain Res* 65:361–370. [Medline](#)
- Gerardin E, Pochon JB, Poline JB, Tremblay L, Van de Moortele PF, Levy R, Dubois B, Le Bihan D, Lehericy S (2004) Distinct striatal regions support movement selection, preparation and execution. *Neuroreport* 15:2327–2331. [CrossRef Medline](#)
- Gritsenko V, Kalaska JF (2010) Rapid online correction is selectively suppressed during movement with a visuomotor transformation. *J Neurophysiol* 104:3084–3104. [CrossRef Medline](#)
- Gutierrez-Garralda JM, Moreno-Briseño P, Boll MC, Morgado-Valle C, Campos-Romo A, Diaz R, Fernandez-Ruiz J (2013) The effect of Parkinson's disease and Huntington's disease on human visuomotor learning. *Eur J Neurosci* 38:2933–2940. [CrossRef Medline](#)
- Haith AM, Krakauer JW (2013) Model-based and model-free mechanisms of human motor learning. In: *Progress in motor control* (Richardson MJ, Riley MA, Shockey K, eds), pp 1–21. *Advances in experimental medicine and biology*. New York: Springer.
- Helmholtz H (1866) *Handbuch der physiologischen Optik*. Leipzig: Voss.
- Hikosaka O, Nakamura K, Sakai K, Nakahara H (2002) Central mechanisms of motor skill learning. *Curr Opin Neurobiol* 12:217–222. [CrossRef Medline](#)
- Karni A, Meyer G, Jezzard P, Adams MM, Turner R, Ungerleider LG (1995) Functional MRI evidence for adult motor cortex plasticity during motor skill learning. *Nature* 377:155–158. [CrossRef Medline](#)
- Kassarjian CD, Tan YF, Chung JY, Heskin R, Peterson MJ, Broussard DM (2005) The site of a motor memory shifts with consolidation. *J Neurosci* 25:7979–7985. [CrossRef Medline](#)
- Kawato M, Gomi H (1992) The cerebellum and VOR/OKR learning models. *Trends Neurosci* 15:445–453. [CrossRef Medline](#)
- Klassen J, Tong C, Flanagan JR (2005) Learning and recall of incremental kinematic and dynamic sensorimotor transformations. *Exp Brain Res* 164:250–259. [CrossRef Medline](#)
- Krakauer JW, Ghez C, Ghilardi MF (2005) Adaptation to visuomotor transformations: consolidation, interference, and forgetting. *J Neurosci* 25:473–478. [CrossRef Medline](#)
- Laforce R Jr, Doyon J (2001) Distinct contribution of the striatum and cerebellum to motor learning. *Brain Cogn* 45:189–211. [CrossRef Medline](#)
- Marinelli L, Crupi D, Di Rocco A, Bove M, Eidelberg D, Abbruzzese G, Ghilardi MF (2009) Learning and consolidation of visuo-motor adaptation in Parkinson's disease. *Parkinsonism Relat Disord* 15:6–11. [CrossRef Medline](#)
- Marr D, Poggio T (1976) From understanding computation to understanding neural circuitry. *Neurosci Res Bull* 15:470–488.
- Martin TA, Keating JG, Goodkin HP, Bastian AJ, Thach WT (1996) Throwing while looking through prisms: I. Focal olivocerebellar lesions impair adaptation. *Brain* 119:1183–1198. [CrossRef Medline](#)
- Mazzoni P, Krakauer JW (2006) An implicit plan overrides an explicit strategy during visuomotor adaptation. *J Neurosci* 26:3642–3645. [CrossRef Medline](#)
- McLaughlin SC (1967) Parametric adjustment in saccadic eye movements. *Percept Psychophys* 2:359–362. [CrossRef](#)
- Middleton FA, Strick PL (2000) Basal ganglia and cerebellar loops: motor and cognitive circuits. *Brain Res Rev* 31:236–250. [CrossRef Medline](#)
- Neely KA, Heath M (2009) Visuomotor mental rotation: reaction time is not a function of the angle of rotation. *Neurosci Lett* 463:194–198. [CrossRef Medline](#)
- Oldfield RC (1971) The assessment and analysis of handedness: the Edinburgh inventory. *Neuropsychologia* 9:97–113. [CrossRef Medline](#)
- Penhune VB, Steele CJ (2012) Parallel contributions of cerebellar, striatal and M1 mechanisms to motor sequence learning. *Behav Brain Res* 226:579–591. [CrossRef Medline](#)
- Rasmussen CE (2006) *Gaussian processes for machine learning*. Cambridge, MA: Massachusetts Institute of Technology.
- Reis J, Schambra HM, Cohen LG, Buch ER, Fritsch B, Zarahn E, Celnik PA, Krakauer JW (2009) Noninvasive cortical stimulation enhances motor skill acquisition over multiple days through an effect on consolidation. *Proc Natl Acad Sci U S A* 106:1590–1595. [CrossRef Medline](#)
- Reisman DS, Wityk R, Silver K, Bastian AJ (2007) Locomotor adaptation on a split-belt treadmill can improve walking symmetry post-stroke. *Brain* 130:1861–1872. [CrossRef Medline](#)
- Rickard TC, Cai DJ, Rieth CA, Jones J, Ard MC (2008) Sleep does not enhance motor sequence learning. *J Exp Psychol Learn Mem Cogn* 34:834–842. [CrossRef Medline](#)
- Robertson EM, Pascual-Leone A, Miall RC (2004) Current concepts in procedural consolidation. *Nat Rev Neurosci* 5:576–582. [CrossRef Medline](#)
- Robertson EM, Press DZ, Pascual-Leone A (2005) Off-line learning and the primary motor cortex. *J Neurosci* 25:6372–6378. [CrossRef Medline](#)
- Sanes JN, Dimitrov B, Hallett M (1990) Motor learning in patients with cerebellar dysfunction. *Brain* 113:103–120. [CrossRef Medline](#)
- Sarlegna F, Blouin J, Bresciani JP, Bourdin C, Vercher JL, Gauthier GM (2003) Target and hand position information in the online control of goal-directed arm movements. *Exp Brain Res* 151:524–535. [CrossRef Medline](#)
- Shmuelof L, Krakauer JW, Mazzoni P (2012) How is a motor skill learned? Change and invariance at the levels of task success and trajectory control. *J Neurophysiol* 108:578–594. [CrossRef Medline](#)
- Stebbins ED, Glenn T, Singh J, Willingham DB, Goetz CG (1997) Intact mirror-tracing and impaired rotary-pursuit skill learning in patients with Huntington's disease: evidence for dissociable memory systems in skill learning. *Neuropsychology* 11:272–281. [CrossRef Medline](#)
- Stickgold R (2005) Sleep-dependent memory consolidation. *Nature* 437:1272–1278. [CrossRef Medline](#)
- Sutton RS, Barto AG (1998) *Introduction to reinforcement learning*, Ed 1. Cambridge, MA: Massachusetts Institute of Technology.
- Takagi M, Zee DS, Tamargo RJ (1998) Effects of lesions of the oculomotor vermis on eye movements in primate: saccades. *J Neurophysiol* 80:1911–1931. [Medline](#)
- Takagi M, Zee DS, Tamargo RJ (2000) Effects of lesions of the oculomotor cerebellar vermis on eye movements in primate: smooth pursuit. *J Neurophysiol* 83:2047–2062. [Medline](#)
- Taylor JA, Ivry RB (2011) Flexible cognitive strategies during motor learning. *PLoS Comput Biol* 7:e1001096. [CrossRef Medline](#)
- Taylor JA, Klemfuss NM, Ivry RB (2010) An explicit strategy prevails when the cerebellum fails to compute movement errors. *Cerebellum* 9:580–586. [CrossRef Medline](#)
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235. [CrossRef Medline](#)
- Tong C, Wolpert DM, Flanagan JR (2002) Kinematics and dynamics are not represented independently in motor working memory: evidence from an interference study. *J Neurosci* 22:1108–1113. [Medline](#)
- Trempe M, Proteau L (2010) Distinct consolidation outcomes in a visuomotor adaptation task: off-line leaning and persistent after-effect. *Brain Cogn* 73:135–145. [CrossRef Medline](#)
- Tseng YW, Diedrichsen J, Krakauer JW, Shadmehr R, Bastian AJ (2007) Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J Neurophysiol* 98:54–62. [CrossRef Medline](#)
- Turnham EJ, Braun DA, Wolpert DM (2011) Inferring visuomotor priors for sensorimotor learning. *PLoS Comput Biol* 7:e1001112. [CrossRef Medline](#)
- Vaca-Palomares I, Díaz R, Rodríguez-Labrada R, Medrano-Montero J, Vázquez-Mojena Y, Velázquez-Pérez L, Fernandez-Ruiz J (2013) Spinocerebellar ataxia type 2 neurodegeneration differentially affects error-based and strategic-based visuomotor learning. *Cerebellum* 12:848–855. [CrossRef Medline](#)
- Walker MP, Brakefield T, Morgan A, Hobson JA, Stickgold R (2002) Practice with sleep makes perfect: sleep-dependent motor skill learning. *Neuron* 35:205–211. [CrossRef Medline](#)
- Waters-Metenier S, Husain M, Wiestler T, Diedrichsen J (2014) Bihemispheric transcranial direct current stimulation enhances effector-independent representations of motor synergy and sequence learning. *J Neurosci* 34:1037–1050. [CrossRef Medline](#)
- Werner S, Bock O (2010) Mechanisms for visuomotor adaptation to left-right reversed vision. *Hum Mov Sci* 29:172–178. [CrossRef Medline](#)
- Wright DL, Rhee JH, Vaculin A (2010) Offline improvement during motor sequence learning is not restricted to developing motor chunks. *J Mot Behav* 42:317–324. [CrossRef Medline](#)