Behavioral/Systems/Cognitive

Use-Dependent and Error-Based Learning of Motor Behaviors

Jörn Diedrichsen,[1,2] **Olivier White,**[2] **Darren Newman,**[2] **and Níall Lally**[2]

[1]Institute of Cognitive Neuroscience, University College London, London WC1N 3AR, United Kingdom, and [2]Wolfson Centre for Cognitive and Clinical Neuroscience, School of Psychology, Bangor University, Gwynedd LL57 2AS, United Kingdom

Human motor behavior is constantly adapted through the process of error-based learning. When the motor system encounters an error, its estimate about the body and environment will change, and the next movement will be immediately modified to counteract the underlying perturbation. Here, we show that a second mechanism, use-dependent learning, simultaneously changes movements to become more similar to the last movement. In three experiments, participants made reaching movements toward a horizontally elongated target, such that errors in the initial movement direction did not have to be corrected. Along this task-redundant dimension, we were able to induce use-dependent learning by passively guiding movements in a direction angled by 8° from the previous direction. In a second study, we show that error-based and use-dependent learning can change motor behavior simultaneously in opposing directions by physically constraining the direction of active movements. After removal of the constraint, participants briefly exhibit an error-based aftereffect against the direction of the constraint, followed by a longer-lasting use-dependent aftereffect in the direction of the constraint. In the third experiment, we show that these two learning mechanisms together determine the solution the motor system adopts when learning a motor task.

Introduction

Motor behaviors such as reaching (Shadmehr and Mussa-Ivaldi, 1994; Donchin et al., 2003), eye movements (Srimal et al., 2008), walking (Morton and Bastian, 2006), and object manipulation (Witney et al., 2000) quickly adapt to the changing dynamics of the body and environment. Learning in this context is generally conceptualized as an estimation problem. For example, for reaching movements, the motor system may estimate the force that will act on the arm. If a perturbing force is experienced, the estimate will be updated using the prediction error, the difference between predicted and observed force. Crucially, during the next movement, the motor system attempts to cancel the perturbation by producing a force that counteracts the current estimated force. Error-based learning is a robust and well studied phenomenon that helps to keep motor behaviors finely calibrated in a changing environment.

It is, however, not always equally important to resist environmental changes. We show here that, when perturbations are applied along redundant dimensions of the movement (i.e., alter the movement without hindering successful task performance), a second learning mechanism can simultaneously lead to persistent changes in the direction of the perturbation. We refer to this mechanism as use-dependent learning, because it biases the next

movement to become similar to the last movement execution. The term use-dependent learning has been used in previous work (Classen et al., 1998; Bütefisch et al., 2000) to refer to neural or behavioral changes that are induced through the simple repetition of movements in the absence of systematic errors. Our usage here is compatible with this description, although we imply a slightly more specific meaning in terms of the underlying learning rule.

Redundancy is a fundamental feature of the human motor system that arises from the fact that there are more degrees of freedom available to control a movement than are strictly necessary to achieve the task goal (Bernstein, 1967). For example, when reaching to a target, multiple joint combinations can result in the same endpoint of the hand. Perturbations to the joint configuration that do not change the position of the hand do not hinder successful task performance and therefore do not need to be fully corrected (Todorov and Jordan, 2002). We show here that such perturbations, but not perturbations along task-relevant dimensions, induce use-dependent learning.

To study this phenomenon, we created a redundant variant of the standard reaching task. In this task, participants held on to a robotic device and made fast reaching movements toward a target. As opposed to the standard reaching task, in which the target defines both the vertical and horizontal goal position, the target in the redundant task was an elongated box, which could be hit at any position along its horizontal extent (see Fig. 1A). Visual feedback about the hand position was provided in form of a line that moved only in the task-relevant vertical direction, but not in the task-redundant horizontal direction.

In experiment 1, we induced use-dependent learning along the task-redundant dimension by passively guiding movements

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

in a slightly tilted direction. In experiment 2, we constrained active movements to demonstrate that error-based and use-dependent learning can occur simultaneously in opposing directions. Finally, experiment 3 provides an example of how the interaction of these two learning mechanisms determines the final solution during adaptation to a novel force field.

Materials and Methods

Apparatus and procedures. Participants were seated in front of a virtual-environment setup with their head on a chin rest. They made 12 cm reaching movements with their right hand while holding a robotic device (Phantom 3.0; SensAble Technologies), which recorded the position of the hand at 200 Hz. Movements were made in a nearly upward direction (on a plane tilted 20° away from the face) and involved elbow and shoulder joints. At the starting position, the elbow angle was ~90°. A simulated spring (200 N/m) restricted the movements to a plane. A stereo display was calibrated such that a cursor (6 mm sphere) could be presented at the correct visual three-dimensional location. The apparatus prevented participants from seeing their hands and arms at all times.

To start a trial, participant moved the cursor into a starting sphere, 6 cm to the right of the body midline at breast height. In a standard task (experiment 3), participants were instructed to push the cursor into a 6 mm target sphere by moving their hand rapidly upward. In the redundant task, the cursor was replaced with a 12 cm horizontal line, displayed at the vertical position of the hand. This line did not move horizontally. The task was to push this line into a 12×1 cm box, displayed 12 cm above the start.

Experiment 1. Fourteen right-handed volunteers (two males; mean age, 22.7 years) participated in experiment 1. Experimental and consent procedures were approved by the ethics committee of the School of Psychology at Bangor University. The first block consisted of 70 training trials in the redundant task. Movements with a peak velocity of 50 cm/s (the width of the acceptable window was adjusted to keep participants at 50% correct) were rewarded if movement times were <900 ms and spatial accuracy was within 3 mm of the target. Each of the next eight blocks started with 15 active movements, followed by 15 passive movements. For the passive movements, the robot moved the hand along a straight trajectory with a minimum-jerk profile. Participants were instructed to offer as little resistance as possible. These trials were indicated by a green starting sphere and no feedback about task success or failure was given. Participants were told that the purpose of the passive movements was to teach them a specific movement speed that they would be asked to replicate during the active movement phase. Unbeknown to the participants, on each block the trajectory for passive trials was angled 8° to the left or right, alternating between blocks, with the order counterbalanced between participants. To account for the slightly different baseline movement directions, we offset the direction of the passive movements for each pair of blocks by the average movement angle of the last pair. In the following 40 active movements, we then measured the influence of the passive movements onto the movement direction.

To show that the effects observed in experiment 1 were not attributable to visuomotor adaptation, we also conducted a control experiment with eight right-handed volunteers (six males; average age, 27 years). The study was identical with experiment 1, with the following changes. Each block contained only 30 instead of 40 active movements after the passive movements. To test for visuomotor adaptation, every fifth trial during the active movements was a probe trial in which a point target was presented at the currently preferred endpoint of the participants, which was estimated online using a running mean of the overall movement direction (z), $m_{n+1} = 0.7m_n + 0.3z_n$. Participants were instructed to move the hand at the required speed to the target as accurately as possible. During the movement, they only had information about the vertical position of their hand in form of the line cursor and received feedback of their full hand position only after movement end. For the first four blocks, participants performed the task with a line that indicated the vertical hand position only as in experiment 1, for the last four blocks with a cursor specifying both vertical and horizontal hand positions. Finally, to avoid an easy visual reference of the target and cursor mid-

point, we extended both target and cursor to be as wide as the screen and moved the starting position 4 cm closer to the body midline.

Experiment 2. Seventeen right-handed participants (four males; mean age, 23.4 years) were tested, 12 of whom had participated in experiment 1. Experiment 2 was identical with experiment 1, except that the 15 trials of passive movements were replaced with 15 active movements that were constrained by a force channel. Participants were not informed in advance of the channel trials. The channel was applied in a direction that deviated by ±8° from the average produced movement direction, as measured over the last two blocks. Any deviation from this trajectory resulted in a position-dependent force (1000 N/m) that pushed the hand back onto the straight-line trajectory.

Experiment 3. In experiment 3, we studied the adaptation process of reaching movements to a velocity-dependent force field in the standard and redundant task. Ten new right-handed participants (four males; mean age, 22.1 years) were recruited. For force field trials, the robot produced a force proportional to the instantiations velocity of the hand, calculated as follows:

$$\begin{bmatrix} f_x \\ f_y \end{bmatrix} = \begin{bmatrix} 0 & 4 \\ -4 & 0 \end{bmatrix} N/(m/s) \begin{bmatrix} v_x \\ v_y \end{bmatrix}, \qquad (1)$$

pushing the hand to the right during an upward movement. For counterclockwise curl fields, the direction of forces was mirrored. Trials were rewarded if participants stopped their movements within 6 mm of the target with a movement time <650–800 ms. The time criterion was adjusted individually to keep the success rate at 50%.

To dissociate the adaptation to a force field (error-based learning) from changes in the planned movement direction (use-dependent learning), we applied a force channel on randomly selected trials during adaptation. On these channel trials, the robot generated a position-dependent force (1000 N/m) that pushed the hand back onto a predetermined, straight trajectory. In the standard task, this trajectory was toward the target. In the redundant task, the channel was applied in three directions, +8.8, 0, and −8.8°, around the average movement direction (m), which was estimated as the running mean of the past overall movement directions (z), $m_{n+1} = 0.9m_n + 0.1z_n$.

Experiment 3 started with two practice blocks of 48 trials in the standard and the redundant task, followed by four conditions, tested in counterbalanced order, as follows: (1) standard task, 12 movements without a force field, followed by 84 movements with a clockwise force field; in 12 of the last 48 trials (randomly chosen), we applied a force channel; (2) same as (1), but with a counterclockwise force field; (3) redundant task, the first block was 48 trials without force field (including 12 channel trials); this was followed by 36 force field trials for initial adaptation and 96 trials including 24 channel trials; (4) same as (3), but with a counterclockwise force field.

To test the hypothesis that the observed adaptation in the redundant task was influenced by visuomotor adaptation, we performed control experiment 2. Eight new participants performed a version of experiment 3, in which we displayed the cursor at the true hand position additionally to the line in the redundant task.

Data analysis. The angle of the first 200 ms (y) after movement start ($v > 3$ cm/s) and the angle of the overall movement (z) were the dependent variables of primary interest. These two group-averaged time series were fitted simultaneously by applying Equations 5–8 iteratively. To remove possible left–right asymmetries, the clockwise and counterclockwise conditions were mirrored and averaged, such that we could set $w_0 = 0$. For a full probabilistic model, we assumed that the initial and overall movement directions were normally distributed around the predicted value with variance σ_y^2 and σ_z^2, respectively. The parameters $\Theta = \{A,B,C,D,E,F,\sigma_y,\sigma_z\}$ (for full model, see Results) were then estimated to maximize the log-likelihood of the data given the parameters. In experiment 1, parameters A–D were not fit since there was no force perturbation. We determined 95% confidence intervals of the parameters by drawing 100,000 samples from the joint posterior distribution over all parameters using Markov chain Monte Carlo sampling (Congdon, 2006).

Force channel model. To decompose the responses in the force channel in experiment 3, we analyzed 11 average force profiles for each partici-
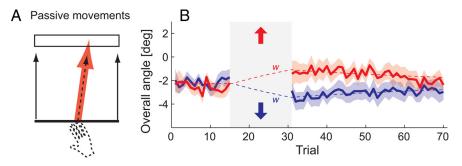
A Passive movements

B

Figure 1. Experiment 1: passively guided movements induce a change along a task-redundant dimension through use-dependent learning. ***A***, Participants moved a line indicating the vertical hand position to a goal box. The horizontal hand position was not relevant to the task. ***B***, Overall movement angle (direction of movement measured from beginning to end). During passive trials (gray box), the robot moved the participants' hand on a trajectory angled 8° rightward (red) or leftward (blue). The manipulation changed the angle of the subsequent free movements. The dashed lines indicate model fit (planned movement angle, w).

pant using a constrained linear model. The first two time series were derived from the average channel response in the standard task after clockwise and counterclockwise adaptation (see Fig. 4A). These were decomposed into an average response ($\mathbf{x}_0$, caused by slight curvature of the movement) and a component caused by the adaptation to a force field ($\mathbf{x}_1$) as follows:

$$\begin{aligned} \mathbf{y}_1 &= b_1\mathbf{x}_1 + \mathbf{x}_0 \\ \mathbf{y}_2 &= b_2\mathbf{x}_1 + \mathbf{x}_0. \end{aligned} \quad (2)$$

The next three time series came from the redundant task without force field adaptation and a channel tilted by $+8.8$, 0, or $-8.8°$ from the recent movement direction (see Fig. 4B). These were modeled as a mix of the intercept ($\mathbf{x}_0$) and a component ($\mathbf{x}_2$) attributable to the deviation of the channel angle from the planned movement angle (w_1) as follows:

$$\begin{aligned} \mathbf{y}_3 &= (-8.8 - w_1)\mathbf{x}_2 + \mathbf{x}_0 \\ \mathbf{y}_4 &= (0 - w_1)\mathbf{x}_2 + \mathbf{x}_0 \\ \mathbf{y}_5 &= (8.8 - w_1)\mathbf{x}_2 + \mathbf{x}_0. \end{aligned} \quad (3)$$

Finally, the six channel responses in the redundant task after force field adaptation (see Fig. 4C) were modeled as a linear combination of force field adaptation and responses attributable to deviation of the channel from the respective planned angles (w_1, w_2) as follows:

$$\begin{aligned} \mathbf{y}_6 &= b_3\mathbf{x}_1 + (-8.8 - w_2)\mathbf{x}_2 + \mathbf{x}_0 \\ \mathbf{y}_7 &= b_3\mathbf{x}_1 + (0 - w_2)\mathbf{x}_2 + \mathbf{x}_0 \\ \mathbf{y}_8 &= b_3\mathbf{x}_1 + (8.8 - w_2)\mathbf{x}_2 + \mathbf{x}_0 \\ \mathbf{y}_9 &= b_4\mathbf{x}_1 + (-8.8 - w_3)\mathbf{x}_2 + \mathbf{x}_0 \\ \mathbf{y}_{10} &= b_4\mathbf{x}_1 + (0 - w_3)\mathbf{x}_2 + \mathbf{x}_0 \\ \mathbf{y}_{11} &= b_4\mathbf{x}_1 + (8.8 - w_3)\mathbf{x}_2 + \mathbf{x}_0. \end{aligned} \quad (4)$$

The regressors ($\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2$) and the regression coefficients ($b_1, \ldots, b_4$, $w_1, \ldots, w_3$) were unknown. To solve this dual estimation problem, we first estimated the likely shape of the regressors $\mathbf{x}_0, \ldots, \mathbf{x}_2$, assuming $b_i = \pm 1$ (depending on the direction of the force field) and $w_j = 0$. We then solved the linear equations with an ordinary-least-square approach, yielding maximum-likelihood estimates of the seven regression coefficients under the assumption of Gaussian noise.

By comparing the size of the predicted response $b_i\mathbf{x}_1$ with the force profile during exposure to the force field (see Fig. 4A), we normalized the size of force field adaptation for redundant and standard tasks. The value w_i was added to the mean direction of the channel for each condition to obtain the planned movement direction.

Results

Experiment 1: passive movements induced use-dependent learning

In experiment 1, participants performed the redundant task. After 15 initial active movements, the robotic device passively guided the participants' hands for 15 movements toward the tar-

get (Fig. 1A). The robot moved the hand on a trajectory that deviated either 8° to the left or 8° to the right of the center, alternating between blocks of trials. Participants were not aware of this manipulation: After the experiment, we told them that the robot deviated their hand laterally and asked them to guess the direction of the deviation on the last block. Only 5 of 14 participants could answer this question correctly, a number not significantly different from chance ($p = 0.202$). In contrast, the subsequent active movements showed a prolonged aftereffect (Fig. 1B). The direction of the following 40 movements was significantly biased toward the direction of the preceding passive movements in that block ($t_{(13)} = 4.681$; $p < 0.001$).

We hypothesize that this effect is attributable to use-dependent learning, a process by which the movement plan associated with a certain task is made more similar to the last executed movement. In this example, these changes occur in terms of the movement direction, a task-redundant dimension attributable to the structure of the target. Use-dependent learning can be understood as a simple averaging process between the last planned direction (w_n) and the last actual direction of movement (y_n) as follows:

$$w_{n+1} = Ew_n + Fy_n + (1 - E - F)w_0. \quad (5)$$

The last component of the update rule allows for a slow drift back to the baseline direction (w_0). When fitting this model to the group average data, the use-dependent learning rate F was estimated to be 0.011 [95% confidence interval, (0.008–0.014)]. The estimate for the retention parameter E was 0.966 (0.953–0.978). Although statistically less stable, the model can also be fitted to individual data. Eleven of the 14 participants showed a positive use-dependent learning rate F (sign test, $p = 0.028$), with a median of 0.013.

Before concluding that the passive movements induced use-dependent learning, we needed to consider the alternative hypothesis that the aftereffect was caused by a change in the visual-proprioceptive or visuomotor calibration (Sober and Sabes, 2003), possibly induced by an implicit error signal between the felt hand position and the assumed visual position of the hand (i.e., the middle of the line). In control experiment 1, we tested this hypothesis in two ways. First, we replicated experiment 1, but inserted probe trials, in which participants were instructed to move their unseen hand to a specific location in the target zone, indicated by a 6 mm target sphere. If passive movements changed the visuomotor calibration, we should find systematic errors in the direction of the passive movements. Contrary to this hypothesis, no such difference was found ($t_{(7)} = -0.619$; $p = 0.55$) (see supplemental Fig. 1B, available at www.jneurosci.org as supplemental material). Second, we again replicated experiment 1 but this time removed any systematic discrepancy between proprioceptive and visual estimated hand position by continually displaying a cursor at the veridical hand position. Participants now consciously perceived during passive movements that their hand was deviated to the side but still showed a significant aftereffect ($t_{(7)} = -3.096$; $p = 0.017$) (see supplemental Fig. 1A, available at www.jneurosci.org as supplemental material). These data clearly demonstrate that the change in preferred direction of movement
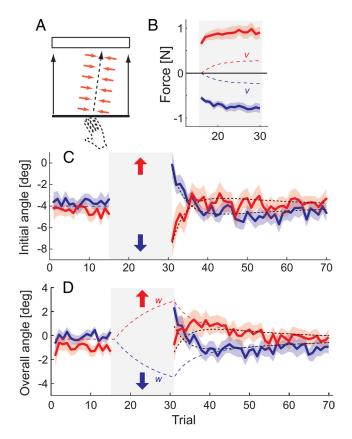
Figure 2. Experiment 2 shows simultaneous action of use-dependent and error-based learning. **A**, After a pretest of 15 movements, a tilted trajectory ($\pm 8°$) was imposed using a force channel (gray box). **B**, The lateral force exerted against the channel (measured at 200 ms after movement start) increases during the channel trials, indicating force field adaptation (v). **C**, The initial angle of the subsequent free movements (measured at 200 ms) shows a error-based aftereffect against the direction of the channel, which washes out quickly. **D**, The overall movement angle shows a persistent aftereffect of use-dependent learning after washout of the force field effect. The dashed black lines indicate model fit.

in the redundant task was not attributable to a change in visuo-motor calibration.

Experiment 2: use-dependent and error-based learning can act simultaneously

The second experiment demonstrates that use-dependent and error-based learning can act simultaneously and independently in opposing directions. In this experiment, participants moved actively during the induction phase. To constrain the movements to a certain direction, we applied a force channel to the hand: The robotic device simulated a stiff spring along a predetermined straight path that was rotated 8° to the left or right from the current movement direction (Fig. 2A). On the first channel trial, participants pushed against the robotic arm with an average force of 0.6 N (measured at 200 ms) because the hand deviated from the planned movement trajectory. The force increased to 0.72 N on the next trial ($t_{(16)} = 2.267$; $p = 0.038$) and reached a level of 0.84 on the 15th trial (Fig. 2B). After the force channel was removed, participants showed an error-based aftereffect: their initial movement directions deviated in the direction opposite to the channel they had experienced ($t_{(16)} = 2.95$; $p = 0.009$) (Fig. 2C). The error-based aftereffect quickly dissipated over four to five movements. Surprisingly, the initial and overall movement direction then continued to change into the opposite pattern (Fig. 2D). For the remainder of the block in the overall movement

direction deviated in the direction of the experienced channel ($t_{(16)} = -2.453$; $p = 0.026$).

These effects can be understood as arising from the parallel activity of error-based and use-dependent learning. According to well established models of error-based learning (Thoroughman and Shadmehr, 2000; Donchin et al., 2003), participants have an estimate of force that will act on their hand during the movement (v_n). When experiencing the unexpected force (f_n) that deviated their hand during the channel movement, the estimate will be updated based on the prediction error, the different between felt and expected force as follows:

$$v_{n+1} = Av_n + B(f_n - v_n), \qquad (6)$$

where the learning parameter B and the retention parameter A dictate the speed of the learning and forgetting process. The increase in estimated force could be observed directly, because participants attempted to counteract the channel force more and more over the 15 movements (Fig. 2B). After removal of the channel ($f_n = 0$), this compensation led to a deviation of the initial movement direction (y_n) in a direction opposite to the channel as follows:

$$y_n = D(f_n - v_n) + w_n. \qquad (7)$$

The parameter D relates inversely to the stiffness of the arm. The initial error in the movement ($y_n - w_n$) was then partially corrected to produce the overall movement direction (z_n). This explains why the force field aftereffect was most pronounced in the initial movement direction as follows:

$$z_n = y_n + C(y_n - w_n). \qquad (8)$$

The novel aspect of our model is the addition of use-dependent learning (Eq. 5), which, simultaneously with error-based learning, changed the planned movement during the channel movements in the direction of the channel (w) (Fig. 2D). Because retention of error-based learning is poorer than for use-dependent learning, the error-based aftereffect washed out rather quickly, revealing a use-dependent aftereffect in the opposite direction. Consistent with this idea, the simultaneous fit of Equations 5–8 could capture well all observed effect (Fig. 2, dashed line). The estimates for the learning and retention parameters for group average data indicated a faster process for error-based, $B = 0.10$ ($0.07–0.14$), $A = 0.83$ ($0.74–0.87$), than for use-dependent learning, $F = 0.038$ ($0.012–0.087$), and $E = 0.91$ ($0.87–0.97$). We confirmed this result by applying the model to individual data, yielding positive use-dependent learning rates for 16 of 17 participants ($p < 0.001$; median, 0.059), and positive error-based learning rates for 15 of 17 participants ($p < 0.001$; median, 0.091).

Finally, we asked whether the strength of error-based and use-dependent learning was correlated across participants. Consistent with the notion that these processes are independent, the force produced during the last five movements in the channel—a measure of error-based learning—correlated with the force field aftereffect (trials 31–34) ($r = 0.589$; $p = 0.013$), but not with the use-dependent aftereffect (trial 38–70) ($r = -0.078$; $p = 0.765$). Thus, together these results provide strong evidence for the parallel and independent action of these two learning mechanisms.

Experiment 3: adaptation solution is shaped by both use-dependent and error-based learning

In the third experiment, we demonstrate that parallel action of use-dependent and error-based learning determines the final so-
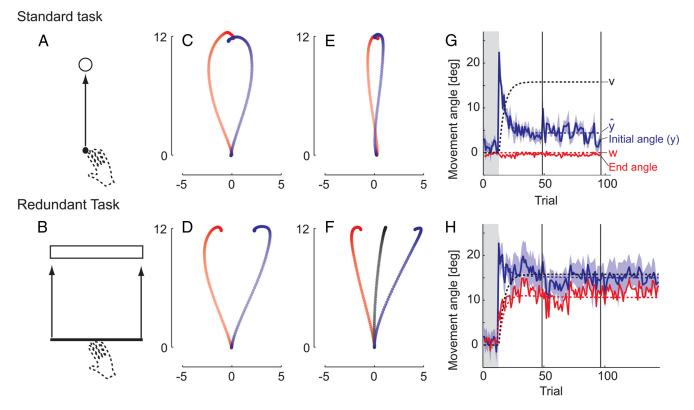
Standard task

Redundant Task

Figure 3. Experiment 3. Shown are force field adaptation in the standard reaching task (**A**) and redundant reaching task (**B**). Initially, the clockwise (blue) and counterclockwise (red) force field caused large errors in the trajectory that is fully corrected in the standard (**C**), but not in the redundant task (**D**). After adaptation, movements in the standard task are nearly straight (**E**), whereas they deviate into the direction of the force field in the redundant task (**F**). **G**, **H**, Time evolution of adaptation and fit of the use-dependent learning model. The initial movement angle (*y*) and fit ($\hat{y}$) are shown in blue. In both tasks, there is an equivalent increase in the expected force (*v*, black dashed line; here scaled by *D*). Only in the redundant task is there a concurrent change in the planned direction (*w*, red dashed line), leading to a shift of both the initial (blue) and overall (red) movement angle.

lution during adaptation to a viscous force field (Shadmehr and Mussa-Ivaldi, 1994). During a standard (Fig. 3*A*) or redundant (Fig. 3*B*) reaching task, the robot repeatedly applied a clockwise or counterclockwise force, perpendicular and proportional to hand velocity. In contrast to the force channel used in experiment 2, the force field dependent on the velocity rather than on the position of the hand, thereby inducing large deviation of the initial direction of the movement. Participants corrected this initial error fully in the standard task [Fig. 3*C*, *C* = 1.07 (1.0–1.11)], but significantly less in the redundant task [Fig. 3*D*, *C* = 0.75 (0.68–0.80)], a difference that is predicted qualitatively by optimal control theory (Todorov and Jordan, 2002; Diedrichsen, 2007).

Given this difference in online correction, how would the system adapt to prolonged application of a force field? After adaptation in the standard reaching task, the hand moved again in a straight-ahead direction (Fig. 3*E*, −6.3 and 1.95°). In the redundant task, however, the initial movement direction was shifted toward the direction of the force field (−12.9 and 17.7°) (Fig. 3*F*).

We considered two possible explanations for this effect. First, participants may have simply adapted less to the force field in the redundant task, because lateral deviations here were task irrelevant. Our model of use-dependent and error-based learning offers a different explanation. The model (Eqs. 5–8) predicts that the degree of force field adaptation would be identical across standard and redundant tasks because error-based learning is driven only by the difference between expected and experienced forces, but not by the relevance of the error for the task (Eq. 6). Simultaneously, use-dependent learning (Eq. 5) would shift the planned movement in the direction of the force field. As outlined

in Introduction, we hypothesize that use-dependent learning only occurs for perturbations along task-redundant dimensions. Therefore, the planned movement direction should only change in the redundant task, but not in the standard task, in which the perturbation interfered with successful task performance.

Thus, the model predicts equivalent force field adaptation in standard and redundant task, combined with a shift of the planned movement direction (not just endpoint) in the direction of the force field in the redundant task. To test these predictions, we applied a force channel on randomly selected trials before and after adaptation. We exploited the fact the force profiles produced against the wall of the channel had a different temporal shape when participants previously had adapted to a velocity-dependent force field, compared with when they had changed their planned movement direction. To demonstrate this fact, we first measured the response in channel trials in the standard task after force field adaptation. Here, participants counteracted the expected force by pressing against the wall of the channel, matching size and velocity dependence of the force field (Fig. 4*A*) (Scheidt et al., 2000; Wagner and Smith, 2008). We compared this response with the response induced by shifts in planned direction. We applied force channels in the redundant task before force field adaptation in three directions around the mean direction of movement. Now, participants pressed against the channel toward the desired endpoint, even after the movement had stopped and the task goal had been reached (Fig. 4*B*).

Finally, we applied channels in three different directions around the average direction of in the redundant task after force field adaptation (Fig. 4*C*). Because of their different temporal shapes, we now could decompose these channel responses into a

component attributable to force field adaptation and a component caused by shifts in the planned movement direction. The observed force profiles were modeled using a constrained linear model (see Materials and Methods). The fits (Fig. 4A–C, dashed lines) explained 94.7% (SD = 2.1%) of the variance for each individual subject's force profiles.

Parameter estimates confirmed that the extent of force field adaptation (Fig. 4D) was statistically equivalent between the redundant (74.5%) and the standard task (76.8%) ($t_{(9)} = 0.34$; $p = 0.742$). The analysis also revealed that the planned movement direction changed significantly in the direction of the force field in the redundant task (Fig. 4E). A counterclockwise force field induced a −9.7° change ($t_{(9)} = 4.51$; $p = 0.001$), whereas a clockwise force field induced a +11.9° change ($t_{(9)} = 3.82$; $p = 0.004$). These findings provide strong evidence against the hypothesis that force field adaptation is reduced in the redundant task. Rather, our analysis clearly shows that the adaptation solution between standard and redundant task are only different, because the planned movement direction shifted into the direction of the force field during the process of adaptation, a behavior only predicted by use-dependent learning.

The final adaptation solution is determined by the relative speeds of these two processes. When fit to the group average data for both the standard (Fig. 3G) and redundant task (Fig. 3H), the learning parameters were estimated to be $A = 0.95$, $B = 0.17$, $E = 0.83$, and $F = 0.11$. This again was confirmed using model fits to individual data with a median use-dependent learning rate of $F = 0.11$, and the data from only one participant resulting in a negative estimate (sign test, $p = 0.011$). The two-process model also predicts that this solution will be stable. While w drifts back slowly toward the old preferred solution (w_0), the less-than-full adaptation ($v < f$) and lack of corrections cause the movements to systematically deviate from the planned direction ($y = w + \delta$). Therefore, the solution $w - w_0 = F/(1 - E - F)\delta$ will be stable.

Finally, we considered two alternative explanations for the observed effect. First, the solution for the redundant task may have been chosen, because it constitutes the optimal solution under the task requirements (Izawa et al., 2008). Although we cannot exclude the possibility that successful task performance can partly account for the high use-dependent learning rate found in experiment 3 compared with experiments 1 and 2 (see Discussion), we believe that it is unlikely that an optimization process alone can account for the observed adaptation solution. Energetically, the solution chosen in the redundant task was clearly more effortful than the solution chosen in the standard task. To achieve the same movement time, participants had to move faster (59.4 vs 52.8 cm/s; $t_{(9)} = 4.78$; $p < 0.001$), and—because the force field was velocity dependent and the force field adaptation identical between tasks—had to counteract a stronger
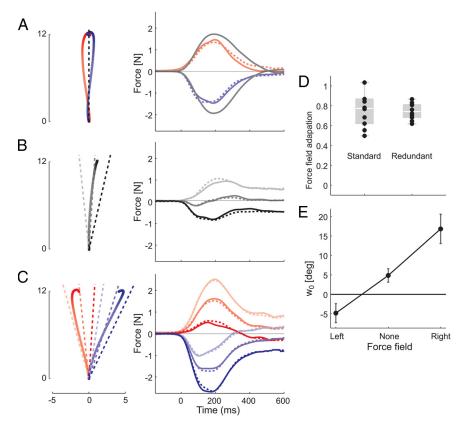

Figure 4. Estimating amount of force field adaptation and shift in planned direction in experiment 3. **A**, After adaptation to a clockwise (blue) or counterclockwise (red) force field in the standard task, a channel was applied in the target direction (black dashed line). Participants produced anticipatory forces against the wall of the channel (solid red and blue line), similar to the force field experienced on nonchannel trials (gray line). Model fits (see Materials and Methods) are shown in dashed lines in all panels. **B**, Before adaptation in the redundant task, a channel was applied in three directions (light, −8.8°; medium, 0°; dark, −8.8°) around the average movement direction. Even after the end of the movement ($t > 400$ ms), participants pressed against the channel, providing a measure of their planned movement direction. **C**, After force field adaptation in the redundant task, a channel was applied in three directions around the respective average movement direction. Resultant force profiles were modeled as a mixture of force field adaptation and deviations of planned direction (dashed lines). **D, E**, Parameter estimates indicate equivalent force field learning in standard and redundant tasks (**D**), and a shift of the planned direction in the direction of the force field in the redundant task (**E**). Error bars indicate SEM.

force field. Based on an analysis of inverse dynamics (see supplemental material, available at www.jneurosci.org), adapted movements for the redundant task entailed a 22% higher control cost than for the standard task. Furthermore, there were no significant differences in the accuracy in the vertical dimension for standard and redundant reaching task ($t_{(9)} = 0.44$; $p = 0.67$). Although it is still possible that lateral movements were beneficial for accuracy considering the higher movement speeds in the redundant task, an optimization approach would need to explain why participants drifted back toward straight-ahead movements in experiments 1 and 2, rather than continuing with the supposedly better lateral movements.

Second, we reconsidered the hypothesis that the shifts in the redundant task were not caused by use-dependent learning, but by changes in the alignment between vision and proprioception, as had been previously hypothesized (Scheidt et al., 2005). In control experiment 2 ($n = 8$), we replicated experiment 3, while continually displaying a cursor at the veridical hand position. Although we thereby prevented any conflict between the seen and felt hand position, results were comparable with experiment 3 (see supplemental Figs. 2, 3, available at www.jneurosci.org as supplemental material).

Discussion

In sum, our experiments reveal the existence of two fundamentally different learning mechanisms in response to a perturbation. First, the motor system estimates and counteracts the perturbation through a mechanism that is driven by prediction errors (Thoroughman and Shadmehr, 2000; Donchin et al., 2003). Second, we now show that at the same time the nervous system associates the current goal with the last executed movement, a form of Hebbian learning. In contrast to error-based learning, which attempts to cancel a perturbation, use-dependent learning changes movements in the direction of a perturbation. This process can be shown when movements are passively guided to one side (experiment 1), or when active movements are systematically perturbed to one side (experiments 2, 3). Although use-dependent learning has been previously proposed as a mechanism that underlies the acquisition of motor behaviors through simple repetition (Classen et al., 1998), we show for the first time that error-based and use-dependent learning simultaneously contribute to the learning of the same motor behavior.

Error-based learning appears to be insensitive to task goals. We show in experiment 3 that the amount of force field adaptation was equivalent in the standard task, in which the perturbation hindered task performance, and in the redundant task, in which the perturbation did not have to be corrected fully. Congruent with the notion of task insensitivity, it has been shown that participants even show adaptation if it hurts goal achievement (Mazzoni and Krakauer, 2006).

In contrast to error-based learning, use-dependent learning appears to be task sensitive. In the standard task, in which the perturbation obstructed the achievement of the goal, no use-dependent learning was observed. Only by applying perturbation along a task-redundant dimension (i.e., the lateral position of the hand), were we able to show its existence. How might such a task sensitivity arise?

One possibility is that use-dependent learning is modulated by reward signals. Standard reinforcement learning, however, cannot easily account for the observed effects. Most learning rules change the motor commands along a gradient estimated using the temporal difference error (Sutton and Barto, 1998), the difference between received and expected reward ($r_n - Er$). Because the temporal difference error will be on average zero, this learning rule would only change the behavior, if the laterally perturbed movements were more rewarding than the original straight movement. We do not believe that this is the case, because the participants, left to their own devices, drifted back toward the obviously preferable straight movements.

However, reinforcement learning can also be achieved by making the learning rate for use-dependent learning F proportional to the normalized reward obtained from the last movement $r_n/(r_n + Er)$. Following this rule, the next movement will be biased toward the last movement execution, even if the reward for the last movement was identical with the expected reward. Thus, if the sampling of new experiences is biased in a certain direction, the learned behavior will shift into that direction, a behavior clearly demonstrated in our experiments.

Although both learning mechanism will converge on the best solution, the updated motor plan for reward-weighted averaging will always lie between the last plan and the last executed movement. In contrast, a rule that weights recent experience by the temporal difference error, attempts to extrapolate outside of the recent experience: one-half of the time it will move in the opposite direction as the last movement. This difference may make reward-weighted use-dependent learning more stable for nonlinear and high-dimensional problems (Peters and Schaal, 2007; Hoffmann et al., 2008).

It should be noted, however, that the hypothesized reward modulation could not simply be based on the overall movement outcome (i.e., the number of points received). Both in the standard and the redundant task, early movements in the force field were rewarded in 33 versus 48% of the trials, and this difference is insufficient to explain the big difference in use-dependent learning rate. Instead, the motor system may determine the contribution of each movement component (initial movement direction, online correction) to the overall reward and base subsequent learning on these signals. In the standard task, the task success could be fully attributed to the online correction, and no use-dependent learning of the initial movement direction would occur. In the redundant task, the initial movement direction could be credited with some of the movement success (as part of the correction was task irrelevant), allowing use-dependent learning to occur. Although this may provide a plausible explanation, much more work is needed to address the factors that determine the strength of use-dependent learning and its connection to reinforcement learning.

Overall, however, use-dependent learning provides a novel account of a number of perplexing phenomena in the motor learning literature. For example, the curvature of a movement, induced by the presence of an obstacle, influences the curvature of the next movement, even when the obstacle has been removed (Jax and Rosenbaum, 2007). Similarly, when learning to move along a curved force channel, participants exhibit force field adaptation against the channel, as well as changes in the planned curvature of the movement (Chib et al., 2006). Finally, use-dependent learning can explain why the motor system arrives at seemingly nonoptimal solutions when adapting to novel dynamics in redundant tasks, as observed in experiment 3 (Scheidt et al., 2005; Diedrichsen, 2007).

Although use-dependent learning can be revealed in experimental tasks with redundancy, it is reasonable to assume that human motor learning always involves both use- and error-based learning. For example, during visuomotor rotation task, error-based learning will produce fast, trial-by-trial adaptation, independent of task goals (Mazzoni and Krakauer, 2006) or movement corrections (Tseng et al., 2007). Once the movement direction changes, use-dependent learning will associate the changed motor plan with the goal, leading to additional consolidation.

Although error-based learning depends on the integrity of the cerebellum (Martin et al., 1996; Diedrichsen et al., 2005; Smith and Shadmehr, 2005), we hypothesize that use-dependent learning may only depend on local changes in cortical motor areas such as primary motor cortex (Classen et al., 1998; Bütefisch et al., 2000). Our current results provide the requisite behavioral tools to investigate such neural dissociations. Moreover, given that perturbations along task-redundant dimensions can bring about long-lasting changes of motor behavior, it may constitute a promising technique for the robotic-assisted physical rehabilitation of stroke patients (Huang and Krakauer, 2009).

References

Bernstein NA (1967) The co-ordination and regulation of movement. Oxford: Pergamon.

Bütefisch CM, Davis BC, Wise SP, Sawaki L, Kopylev L, Classen J, Cohen LG (2000) Mechanisms of use-dependent plasticity in the human motor cortex. Proc Natl Acad Sci U S A 97:3661–3665.

Chib VS, Patton JL, Lynch KM, Mussa-Ivaldi FA (2006) Haptic identification of surfaces as fields of force. J Neurophysiol 95:1068–1077.

Classen J, Liepert J, Wise SP, Hallett M, Cohen LG (1998) Rapid plasticity of human cortical movement representation induced by practice. J Neurophysiol 79:1117–1123.

Congdon P (2006) Bayesian statistical modeling, Ed 2. London: Wiley.

Diedrichsen J (2007) Optimal task-dependent changes of bimanual feedback control and adaptation. Curr Biol 17:1675–1679.

Diedrichsen J, Verstynen T, Lehman SL, Ivry RB (2005) Cerebellar involvement in anticipating the consequences of self-produced actions during bimanual movements. J Neurophysiol 93:801–812.

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

Hoffmann H, Theodorou E, Schaal S (2008) Optimization strategies in human reinforcement learning. In: Advances in computational motor control VII. Washington, D.C.: Society for Neuroscience.

Huang VS, Krakauer JW (2009) Robotic neurorehabilitation: a computational motor learning perspective. J Neuroeng Rehabil 6:5.

Izawa J, Rane T, Donchin O, Shadmehr R (2008) Motor adaptation as a process of reoptimization. J Neurosci 28:2883–2891.

Jax SA, Rosenbaum DA (2007) Hand path priming in manual obstacle avoidance: evidence that the dorsal stream does not only control visually guided actions in real time. J Exp Psychol Hum Percept Perform 33:425–441.

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.

Mazzoni P, Krakauer JW (2006) An implicit plan overrides an explicit strategy during visuomotor adaptation. J Neurosci 26:3642–3645.

Morton SM, Bastian AJ (2006) Cerebellar contributions to locomotor adaptations during splitbelt treadmill walking. J Neurosci 26:9107–9116.

Peters J, Schaal S (2007) Using reward-weighted regression for reinforcement learning of task space control. Paper presented at the 2007 IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning, Honolulu, HI, April.

Scheidt RA, Reinkensmeyer DJ, Conditt MA, Rymer WZ, Mussa-Ivaldi FA (2000) Persistence of motor adaptation during constrained, multi-joint, arm movements. J Neurophysiol 84:853–862.

Scheidt RA, Conditt MA, Secco EL, Mussa-Ivaldi FA (2005) Interaction of visual and proprioceptive feedback during adaptation of human reaching movements. J Neurophysiol 93:3200–3213.

Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. J Neurosci 14:3208–3224.

Smith MA, Shadmehr R (2005) Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. J Neurophysiol 93:2809–2821.

Sober SJ, Sabes PN (2003) Multisensory integration during motor planning. J Neurosci 23:6982–6992.

Srimal R, Diedrichsen J, Ryklin EB, Curtis CE (2008) Obligatory adaptation of saccade gains. J Neurophysiol 99:1554–1558.

Sutton RS, Barto AG (1998) Reinforcement learning. Cambridge, MA: MIT

Thoroughman KA, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. Nature 407:742–747.

Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. Nat Neurosci 5:1226–1235.

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.

Wagner MJ, Smith MA (2008) Shared internal models for feedforward and feedback control. J Neurosci 28:10663–10673.

Witney AG, Goodbody SJ, Wolpert DM (2000) Learning and decay of prediction in object manipulation. J Neurophysiol 84:334–343.