

Adapting to inversion of the visual field: a new twist on an old problem

Timothy P. Lillicrap · Pablo Moreno-Briseño ·
Rosalinda Diaz · Douglas B. Tweed · Nikolaus F. Troje ·
Juan Fernandez-Ruiz

Received: 4 March 2013 / Accepted: 5 May 2013 / Published online: 23 May 2013
© Springer-Verlag Berlin Heidelberg 2013

Abstract While sensorimotor adaptation to prisms that displace the visual field takes minutes, adapting to an inversion of the visual field takes weeks. In spite of a long history of the study, the basis of this profound difference remains poorly understood. Here, we describe the computational issue that underpins this phenomenon and presents experiments designed to explore the mechanisms involved. We

show that displacements can be mastered without altering the updated rule used to adjust the motor commands. In contrast, inversions flip the sign of crucial variables called sensitivity derivatives—variables that capture how changes in motor commands affect task error and therefore require an update of the feedback learning rule itself. Models of sensorimotor learning that assume internal estimates of these variables are known and fixed predicted that when the sign of a sensitivity derivative is flipped, adaptations should become increasingly counterproductive. In contrast, models that relearn these derivatives predict that performance should initially worsen, but then improve smoothly and remain stable once the estimate of the new sensitivity derivative has been corrected. Here, we evaluated these predictions by looking at human performance on a set of pointing tasks with vision perturbed by displacing and inverting prisms. Our experimental data corroborate the classic observation that subjects reduce their motor errors under inverted vision. Subjects' accuracy initially worsened and then improved. However, improvement was jagged rather than smooth and performance remained unstable even after 8 days of continually inverted vision, suggesting that subjects improve via an unknown mechanism, perhaps a combination of cognitive and implicit strategies. These results offer a new perspective on classic work with inverted vision.

Timothy P. Lillicrap and Pablo Moreno-Briseño contributed equally to this work.

Electronic supplementary material The online version of this article (doi:10.1007/s00221-013-3565-6) contains supplementary material, which is available to authorized users.

T. P. Lillicrap (✉)
Department of Pharmacology, University of Oxford,
Mansfield Road, Oxford OX1 3QT, UK
e-mail: timothy.lillicrap@pharm.ox.ac.uk

P. Moreno-Briseño · R. Diaz · J. Fernandez-Ruiz (✉)
Laboratorio de Neuropsicología, Departamento de Fisiología,
Facultad de Medicina, Universidad Nacional Autónoma de
México, Coyoacán, Distrito Federal, 04510 México, México
e-mail: jfr@unam.mx

D. B. Tweed
Department of Physiology, University of Toronto, Toronto,
ON, Canada

N. F. Troje
Department of Psychology, Queen's University, Kingston,
ON, Canada

N. F. Troje
School of Computing, Queen's University, Kingston, ON, Canada

J. Fernandez-Ruiz
Facultad de Psicología, Universidad Veracruzana, Xalapa,
México

Keywords Visual inversion · Visuomotor learning ·
Feedback error · Visual rotation · Control theory ·
Sensitivity derivative

Introduction

While sensorimotor adaptation to prisms that displace the visual field by a fixed amount takes only a few

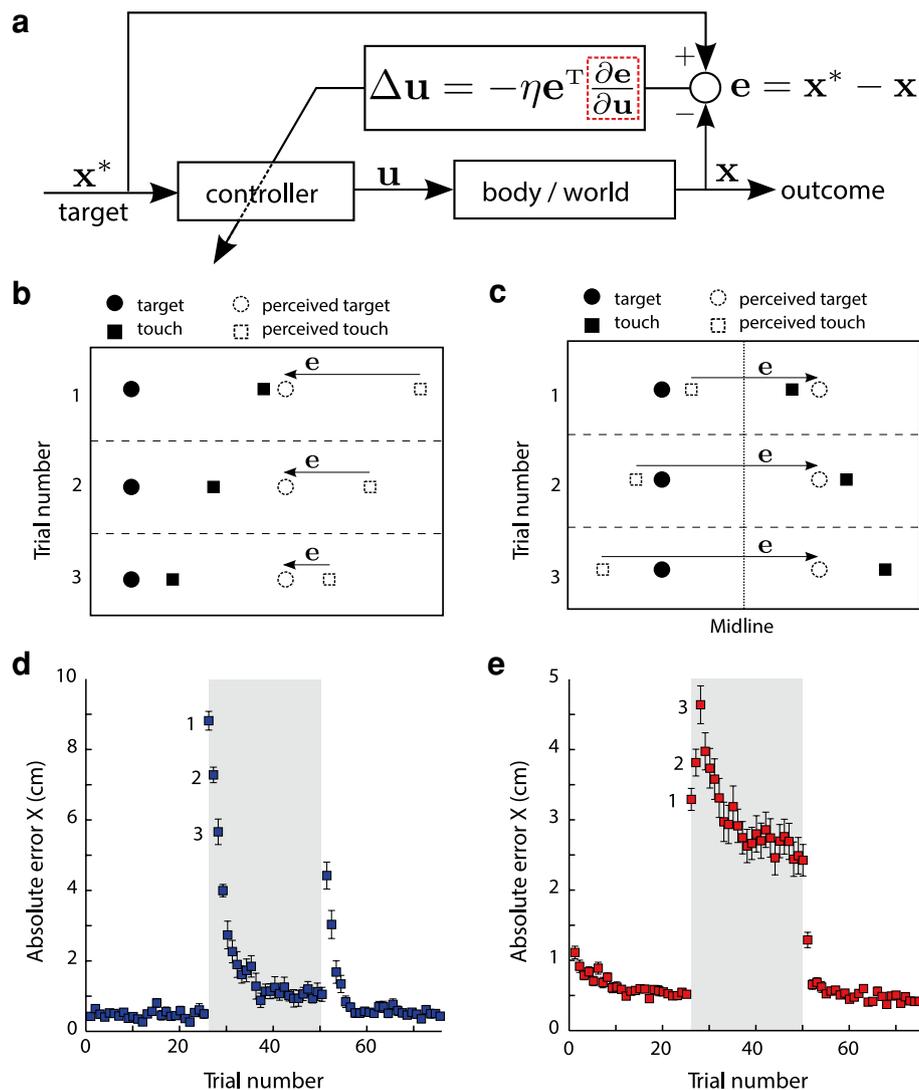


Fig. 1 Visuomotor pointing errors under two optical perturbations. **a** Schematic of visuomotor learning. A target, x^* , is presented to a controller which sends command signals, u , to the body and leads to an outcome, x , in the world. An error, e , is computed from the difference between x^* and x . When the body or world change, the controller needs to alter its mapping between targets and commands. To make appropriate command updates, Δu , the brain requires an estimate of the sensitivity derivatives, $\langle \partial e / \partial u \rangle$, which capture the relationship between errors and commands. While displacing prisms do not significantly alter this relationship, inverting prisms flip the sign of one of these derivatives, which causes learning to become maladaptive. In order for visuomotor learning to be successful again, the

sensitivity derivatives must themselves be relearned. **b, c** Diagram of error vector changes for the first three trials under horizontal refracting and inverting prism conditions, respectively. With the refracting prisms, the normal updated rule (i.e., correct in the direction of the error vector) causes the error to decrease, but with the inverting prisms, the same updated rule causes the error to increase. **d, e** Mean absolute pointing errors in the horizontal (X) direction across baseline, perturbed (*shaded area*), and post-perturbation phases (25 trials each). Subjects pointed to targets presented on a touch screen monitor in a dark room (see Experimental Procedures and Online Resource 1 Figure 1). During the perturbed phase, the horizontal refracting prism (**d**) or inverting prism (**e**) was introduced. *Error bars* are SEM

minutes, adapting to an inversion of the visual field takes weeks. In spite of a long history of the study, the basis of this profound difference remains poorly understood. To identify the underpinning cause of this discrepancy, we start with a simple question: How do we learn to move our arm to a visual goal? The basic procedure seems clear: We try to reach the target, then we calculate an

error (e.g., the vector from finger to goal), and we use this error to adjust our motor commands (Fig. 1a). But how do we know which adjustments will help, for example, if we increase a given motor command, will the error shrink or grow? To make appropriate choices, the brain needs to know how changes in command signals affect error.

In control theory, this relationship between commands, u , and errors, e , is captured by the matrix, $\partial e/\partial u$, of *sensitivity derivatives* (Åström and Wittenmark 1995)—also called the *control jacobian* (Callier and Desoer 1991). This matrix quantifies how small changes in a command (e.g., moving the hand a little to the left) give rise to changes in the task error (e.g., the distance between the finger tip and a target). As we shall see, these derivatives lie at the heart of the profound differences between visuomotor adaptations to displaced versus inverted vision.

Visuomotor adaptation to perturbations that displace the visual field, for example, from left to right, is widely studied and well characterized (Harris 1965; Kohler 1963; Kornheiser 1976; Redding and Wallace 1990). Pointing, throwing, and reaching tasks have been used to assess adaptation, and in these tasks, human subjects adapt quickly and smoothly to displacements, typically within minutes (Fernandez-Ruiz et al. 2006; Kitazawa et al. 1997; Martin et al. 1996; Redding et al. 2005; Redding and Wallace 1990). When prisms are worn, which displace targets and responses to the right and thus initially produce a leftward error (Fig. 1b), subjects reduce their errors by correcting in the leftward direction on subsequent trials. Doing so, subjects make use of an implicit assumption about how error vectors ought to be used to update motor commands (Fig. 1a). The assumption, which holds in the case of displaced vision, is that the relationship between commands and errors (i.e., the sensitivity derivatives) has not been altered.

Comparatively, little is understood about visuomotor adaptation to inversions of the visual field—for example, a perturbation which *flips* the visual field from left to right about the midline (Fig. 1c). Studies have reported that, although subjects are initially severely impaired by inversions, they were eventually able to reacquire even complex sensorimotor skills, such as riding a bicycle (Harris 1965; Kohler 1963). However, most studies have been qualitative in nature (Rock 1966, 1973; Stratton 1896, 1897) or else have focused on perceptual rather than motor adaptations (Linden et al. 1999; Sekiyama et al. 2000). Thus, the reason for the profound difference in the time course of visuomotor adaptation, the manner in which adaptation unfolds, and the mechanisms involved are not well studied.

To help address these gaps in understanding, we investigated the computations required to adapt to inversions. We show that when inverting prisms are worn, the sign of one or more sensitivity derivatives is flipped, rendering the “standard” feedback rule used to adapt to displacements ineffective. For example, the catastrophic result of following the “standard” feedback rule (i.e., where leftward errors lead to leftward corrections) is shown in Fig. 1c. Counterintuitively, for visual inversions, leftward errors now require rightward corrections and vice versa. Thus, the feedback

rule itself, and not just the command signals, must be updated in order to adapt successfully to inversions.

We experimentally examined long- and short-term human adaptation to visual inversions. Visuomotor adaptation was probed using simple pointing tasks in which subjects were required to quickly touch targets presented on a monitor. We compared the results of these experiments with predictions made by two classes of existing models: those that update their feedback rule (i.e., those which reestimate sensitivity derivatives) and those which do not.

Superficially, our experimental data agree with the class of gradient-based models which update their feedback learning rule. However, closer examination of our results suggests that adaptation to inversions involves a complex mixture of implicit (i.e., gradient or reinforcement learning) and explicit or “cognitive” processes (e.g., Mazzoni and Krakauer 2006), which is not well modeled by the existing theory.

Methods

The present study tested the effects of two different kinds of visual perturbation—displacing and reversing—during simple motor tasks. All of the tasks required subjects to reach to and touch targets presented on a computer screen. We examined fixed target tasks under both visual perturbations, as well as a variant in which probe trials (using a second target) were used to evaluate generalization. To limit the effectiveness of straightforward cognitive strategies, we examined another variant of the task in which target positions were randomized; again, we looked at the effects of both displacing and reversing perturbations during this task. Finally, we studied the effect of long-term training in the reversing condition.

Visual distortions and task overview

In our *no-reversal* tasks, subjects reached to targets while wearing refracting prisms that displace the stimulus horizontally by 20° (Fig. 1b; details found below). This task resembles spear-fishing, where light from the target is refracted uniformly in one direction. This shift has no effect on sensitivity derivatives (Fig. 1a; see Online Resource 1), so for these tasks, control can improve based on the existing estimates of the derivatives.

In our *reversal* tasks, subjects reached to targets while wearing right-left inverting prisms (Figure 1c; details found below), which flip the visual image around the mid-line and, with it, one of the sensitivity derivatives (Fig. 1a; see Online Resource 1). If subjects tried to learn this task without updating their estimates of the derivatives, they would

never improve but would just get worse and worse (Abdelghani et al. 2008; Abdelghani and Tweed 2010).

To provide a more complete picture of the time course of adaptation to inversions, we performed both long- and short-term versions of the reversal tasks. While many studies have looked qualitatively at long-term adaptation to inverted vision, few of these have provided quantitative evaluation of visuomotor skills. In the long-term experiment, a single human subject, MW, continuously wore up-down reversing dove prisms set into custom-made goggles for 8 days. MW performed two variants of the reversal task after 12 h, 4 days, and 8 days of wearing the prisms. The first variant was essentially the same as the *reversal* task described above except that the reversal was up-down rather than right-left. The second task, called the *reversal-jump* task, was similar to the *reversal* task except that, after initiating a reaching movement, targets jumped a random distance up or down to a new location (Gritsenko and Kalaska 2010). MW was instructed to try to touch this new location. It is possible that only the feedforward component of reaching would undergo adaptation under inverted vision. The *reversal-jump* variant was included to examine the extent to which online feedback corrections were also adapted.

Subjects

All 128 subjects were right-handed volunteers between the ages of 18 and 30 years old. Seventy-eight of them participated in the short-term fixed target task, 25 in the short-term experiment with probe trials, 15 in the short-term random target task, and 10 in the long-term experiment. They were healthy at the time of the experiment, had normal or corrected to normal vision (MW had normal visual acuity), and had no history of mental or neurological injury. Subjects were informed about the general purpose and procedures of the experiment and were aware when a prism was inserted into and removed from the visual path, but were entirely naive about the perturbation induced by the prism. Subjects in the long-term study were the exception; because subject MW was aware of the nature of the perturbation, controls were also informed that the prism goggles would flip the visual field upside-down. The experiments were approved by Facultad de Medicina at UNAM and the Queen's University Human Research Ethics Boards.

Apparatus and task details

Short-term experiments, fixed target task variants: Subjects had their head stabilized using a chin support and their left eye was occluded throughout the experiments. The room in which the experiments were performed was kept dark. Subjects viewed a touch screen monitor with their right eye

through a small (2.6 cm in diameter) window in a wood panel. The touch screen monitor was positioned so that it could be seen and reached comfortably, approximately 40 cm from the chin rest. When they were ready, a green target (a filled circle 5 mm in diameter) appeared on a black background in a fixed location in their visual field. The subjects were instructed to press a key to make the target disappear and then point quickly and accurately to the target location. Subjects were asked to touch the target in a single movement (i.e., without online corrections) and were monitored to insure compliance. After touching the monitor, the target and a red dot (a filled circle 5 mm in diameter) appeared on the screen to give the subject visual feedback about the location of their response. After 250 ms, the feedback and target disappeared and a new green target appeared to signal the start of a new trial. There were three phases: (1) baseline, (2) perturbed, and (3) washout. Each phase lasted 25 trials. During the baseline and washout phases, subjects viewed the screen normally. During the perturbed phase, a horizontal displacing (20 diopter refracting) prism (for the *no-reversal* task) or a horizontal inverting (dove) prism (for the *reversal* task) was introduced in front of the right eye so that the screen was visible only via the prism. The real location of the visual target was changed during each phase so that it was perceived to be in the same location in the visual field.

Short-term experiment with probe trials: This task was identical to the fixed target task described above, except that probe trials are interleaved to test generalization. Probe trials differed from the normal trials in two aspects. First, the target appeared in the horizontal mirror position of the normal trials, and second, the subjects did not see their hand during the reach, nor any error feedback at the end of the reach. To enforce these conditions, subjects wore computer-controlled occluding goggles that were programmed to be transparent for target presentation and opaque from the time that the subject's finger left the key, up until the end of the trial.

Short-term experiments, random target task variants: The random target task was the same as the fixed target tasks except that the target location on each trial was chosen randomly from a uniform distribution over a rectangular patch in the subject's visual field (15 cm × 15 cm). Also, to insure that no in-flight modification of pointing occurred, the subjects wore computer-controlled occluding goggles that blocked vision during the reach. The goggles were programmed to be transparent for target and feedback presentation, and opaque from the time that the subject's finger left the key until it hit the touch screen monitor. The number of trials for the three phases was different. In the *no-reversal* task: baseline = 30, perturbed = 110, and washout = 30. In the *reversal* task: baseline = 50, perturbed = 525, and washout = 50.

Long-term experiments, general information: MW wore the prism goggles for 8 full days. They were put on at 17:30 hours on August 3, 2009, and she took them off at 17:30 on August 11, 2009. Before she went to bed, we replaced the goggles with a blindfold. She also used a blindfold in the shower. Sometimes we had to clean the goggles from inside during the day. If this was necessary, she always held her eyes closed while the goggles were off. MW was tested on two sensorimotor tasks 4 times: after 12 h, 4 days, and 8 days of prism exposure, and then again in a post-test 45 min after removing the prisms. Control subjects ($n = 9$) were tested twice during the same day: once in a pre-condition without prisms and then once again with the same prisms as worn by MW. The pre- and post-tests were done with goggles that were identical to the inversion goggles except that they had no prisms fitted to them. Both the inversion goggles and the control goggles restricted the field of view, which was effectively 34° horizontal \times 25° vertical. Both pairs of goggles allowed vision through both eyes and tasks were performed with both eyes.

Long-term experiments, random targets no-jump task variant: Subjects had their head stabilized using a chin support. The room in which the experiments were performed was kept dark. A touch screen monitor was positioned so that it could be seen and reached comfortably, approximately 40 cm from the chin rest. When they were ready, subjects pressed a button with their right hand to make a white target (a filled square 5 mm in diameter) appear. Target locations were drawn randomly from a uniform distribution over a $15\text{ cm} \times 15\text{ cm}$ patch of the screen. As soon as the target appeared, subjects reached with their right hand to touch the target. If they made the reach within 700 ms of target onset, the subjects received a green feedback target (a filled square 5 mm in diameter) indicating where they had touched and that the trial was successful. If it took longer than 700 ms to touch the screen, subjects received a red feedback target (a filled square 5 mm in diameter) indicating where they had touched and that the trial was not fast enough; in this case, subjects were required to repeat the trial. Subjects were asked to touch the target in a single movement (i.e., without online corrections). Control subjects performed this task twice: once with the mocked-up goggles (80 successful trials) and once with the real prism goggles (80 successful trials). MW performed this task a total of 4 times (80 successful trials each time).

Long-term experiments, random targets jump task variant: The apparatus was the same as in the no-jumps variant described in the previous paragraph. Subjects pressed and held a button to make a white target (a filled square 5 mm in diameter) appear. When they were ready, subjects reached to the target. Target locations were drawn randomly from

a uniform distribution over a $15\text{ cm} \times 15\text{ cm}$ patch of the screen; 50 ms after their finger left the button (i.e., reach onset), the target jumped a random distance (drawn from a uniform distribution over the range 0–4 cm) up or down to a new location. Subjects were instructed to try to touch the new location of the target; that is, they needed to make an online correction. Subjects were required to touch the screen within 400 ms of reach onset. If they were successful, a green feedback target appeared to indicate the location they touched. Otherwise, a red feedback light appeared and they were required to repeat the trial. The strict criterion of 400 ms (as opposed to 700 ms above) was used in the jump task, because it was found during testing that if given 700 ms, people often attempt to stop and adjust their movement trajectory. By enforcing a maximum reach time of 400 ms, subjects were forced to make a single fluid movement (mid-flight adjustment was still possible, but they did not have time to stop and change course). Control subjects performed this task twice: once with the mocked-up goggles (100 successful trials) and once with the real prism goggles (300 successful trials). MW performed this task a total of 4 times (300 successful trials each time except during the post-test in which MW performed only 100 successful trials).

In all the experiments, subjects were monitored by the experimenter to insure that movements were executed as soon as targets were shown and were effectively “ballistic” (i.e., made in one smooth motion from the keyboard to the screen).

Data collection and analysis

A custom-built program was used to display the targets and feedback, control the occluding goggles, and record the location of subject responses. Data were analyzed and plotted using Matlab and the Statistics Toolbox. The Wilcoxon signed-rank test (a nonparametric alternative to the paired Student's t -test) was performed using Matlab's `signrank()` function, the unpaired t test with `t test 2()`, and regression with `regress()`.

Results

Short-term adaptation to visuomotor distortions

In the no-reversal task with a fixed target, subjects improved performance rapidly when the visual perturbation was introduced ($p = 5.9 \times 10^{-5}$; Wilcoxon signed-rank test between first and tenth perturbed trial), over the course of approximately 15 trials (Fig. 1d). Later, when they removed the prisms, they showed negative aftereffects for another 4 trials ($p = 9.2 \times 10^{-5}$, 0.0024, 0.13, 0.0064; Wilcoxon

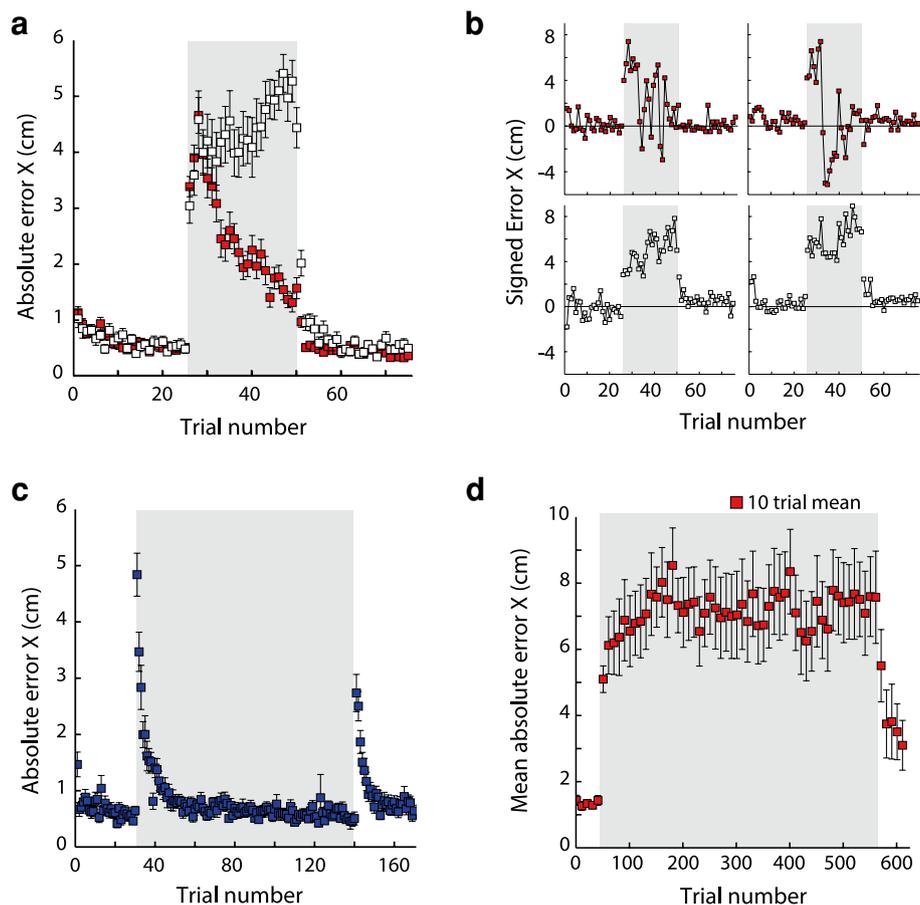


Fig. 2 A closer look at visuomotor adaptation to reversing prisms. **a** The performance of two groups of subjects in the inverted prism task with a fixed target. The majority of subjects (55 of 78) found a way to reduce the task error (*red squares*; error reduction was characterized by a mean absolute error of <3 cm across the last 5 perturbed trials), while 23 increased the error (*white squares*). **b** Four examples of single subject performance in the fixed target reversing condition. Note that while performance is improved over time, trial-to-trial

performance is erratic. **c** Average performance for subjects ($n = 15$) under the displacing prism condition with randomly presented targets. Performance is initially perturbed but quickly returns to near baseline over about 20 trials. **d** Average performance for subjects ($n = 15$) under the reversing prism condition with randomly presented targets. Under the perturbation, performance gradually worsens for approximately 150 trials and then levels off, neither improving nor degrading significantly (color figure online)

signed-rank test between the last baseline trial and the first, second, third, fourth washout trial, respectively; Figure 2a of Online Resource 1 shows the same data as Figure 1, but with signed error), as they unlearned their adaptation.

In the reversal task with a fixed target, subjects on average got worse for 3 trials ($p = 6.7 \times 10^{-5}$, 7.67×10^{-7} ; Wilcoxon signed-rank test between the first and the second, third perturbed trials, respectively) and then improved ($p = 0.0024$; Wilcoxon signed-rank test between first and last perturbed trials; Fig. 1e), as we would expect if they “learned” in the wrong direction initially, because of an obsolete estimate of the sensitivity derivatives, and then started to learn appropriately after they corrected the estimate. But on closer analysis, our 78 subjects fell roughly into two groups (Fig. 2a). Twenty-three of them did not improve or got increasingly worse,

as if they could not relearn the sensitivity derivatives at all (see examples in Fig. 2b). When their prisms were removed, these non-learners showed large positive, rather than negative, aftereffects for 1–5 trials (Online Resource 1 Figure 2b). The other 55 subjects did worsen and then improve (improvement was characterized by a mean absolute error of <3 cm across the last 5 perturbed trials), but even their behavior diverged, in two important ways, from what we might expect if they were simply correcting their estimates of the sensitivity derivatives. First, their performance remained unstable rather than improving steadily (see examples in Fig. 2b); that is, even when they were able to reach to a location very near the target, their performance would quickly degrade. Second, they showed no consistent aftereffects (see also Online Resource 1 Figure 2c).

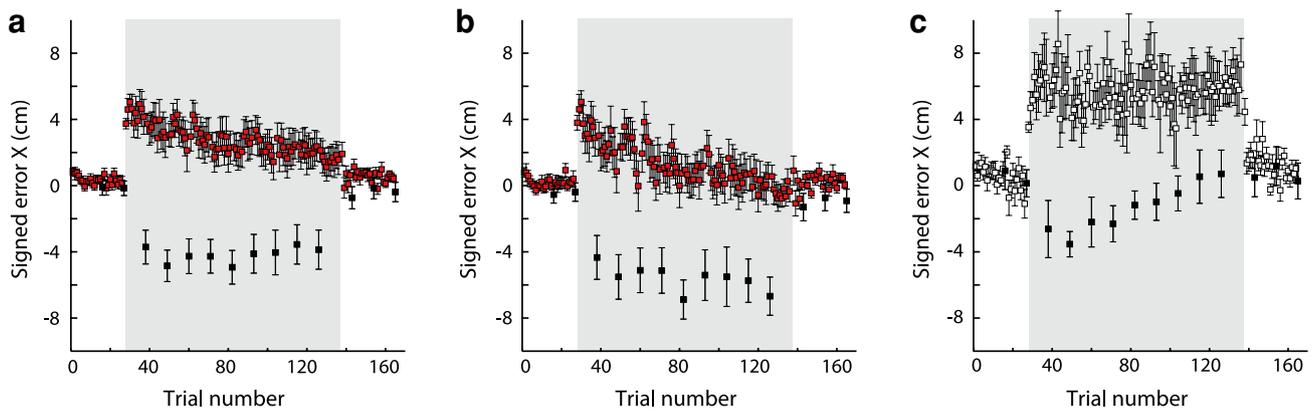


Fig. 3 Testing generalization using probe trials. **a** Performance in the inverted prism task with a fixed target during trials with visual feedback (*red*) and probe trials which used a target at a second location (at the mirror image position of the original target) and without error feedback (*black*). **b** Eighteen out of $n = 25$ subjects showed error reduction

In the reversal task with probe trials, average performance degraded for 3 trials of the perturbation phase ($p = 0.005$; Wilcoxon signed-rank test between first and third perturbed trials) and then improved ($p = 0.02$; Wilcoxon signed-rank test between the first 5 and final 5 perturbed trials). Analysis of the probes showed that during the reversal perturbation, the probes did not show any sign of error reduction ($p = 0.46$; Wilcoxon signed-rank test between the first and final probe trials during the perturbed phase) (Fig. 3a). However, after separating the subjects into two subgroups using the same criteria as in the preceding paragraph, further analyses showed two different trends. The group that showed error reduction during the non-probe trials ($n = 18$) exhibited a significant trend toward worsening errors during the probe trials ($p = 0.011$, $R^2 = 0.54$, slope = -0.23 ; linear regression on probe trial performance) followed by a lack of aftereffect ($p = 0.17$; Wilcoxon signed-rank test between last 5 baseline trials and the first 5 washout trials) (Fig. 3b). The group that did not show error reductions during the non-probe trials ($n = 7$) showed a significant trend toward decreasing errors during the probe trials ($p = 0.014$, $R^2 = 0.57$, slope = 0.34 ; linear regression on probe trial performance), as well as a significant aftereffect ($p = 0.03$; Wilcoxon signed-rank test between last 5 baseline trials and the first 5 washout trials) (Fig. 3c).

In the no-reversal task with random targets, subjects again improved rapidly, returning to near-baseline performance in about 30 trials (Fig. 2c). As expected, when the prisms were removed, they showed a robust negative aftereffect for about 15 trials.

In the reversal task with random targets, the average performance of subjects initially worsened for about 150 trials ($p < 10^{-5}$, $R^2 = 0.40$, slope = 0.82 ; linear regression

during the normal trials and a weak trend toward worsening error during the probe trials. **c** The remaining seven subjects showed an increase in error during the perturbed phase accompanied by a weak trend toward error reduction for the probe trials. The shaded area indicates the perturbed phase; error bars are SEM (color figure online)

of mean performance on trials 50–200) and then plateaued, showing no significant improvement even after an additional 300 trials of practice ($p = 0.9$; Wilcoxon signed-rank test on mean performance between trials 100–200 and 475–575; Fig. 2d). One out of the 15 subjects showed a small but significant reduction in error over the course of the perturbed trials ($p < 0.026$; Wilcoxon signed-rank test on performance during trials 100–200 and 475–575). On average, and across all subjects, a large absolute aftereffect was evident when the prisms were removed ($p < 10^{-5}$ for mean performance, $p < 0.03$ for all subjects; t test between baseline and washout trials).

Long-term adaptation experiment

During the long-term study, subject MW's improvement was clearly evident. Moments after donning the goggles, MW was rendered nearly helpless on even simple visuomotor tasks. After eight days of continuously wearing the goggles, MW was able to ride a bicycle in a controlled environment and take part in many day-to-day activities.

On the eighth day, MW showed marked improvement over control subjects in the reversal task ($p = 0.019$; t test between controls and MW; Fig. 4a). In the jump task, which tested her ability to make appropriate online corrections to a sudden change in target location, MW also showed significant improvement after 8 days ($p < 10^{-3}$; t test between controls and MW; Fig. 4b). However, in both the reversal and jump tasks, performance never reached that of the controls during the no-prisms condition ($p = 0.050$ and $p = 0.013$, respectively; t test between the controls during the no-prism condition and MW's performance on day 8). MW's performance on both tasks returned to control levels within 45 min of removing the prisms (Fig. 4a, b).

In addition, from debriefing MW on day 8, it was apparent that the improvement observed in both tasks was still highly reliant on attention and had not become automatic.

Computational issues and model predictions

How should we understand these results? To begin, we represent the learning problem mathematically. We consider the parameterizable controller, f , written as, $u = f(x; \alpha)$, where u is a vector specifying the motor commands, x is a vector containing the sensory, context, and goal information used by the controller, and α is a matrix parameterizing the controller (e.g., α might contain the adjustable weights in a radial basis function network which maps sensory inputs to motor outputs via, $u_i = \sum_j \alpha_{ij} \varphi_j(x)$, where $\varphi_j(x)$, are the basis functions). The controller tries to move to a visual target, x^* ; so, we take its error vector to be $e = x^* - u$, and if we assume that the goal is to minimize the squared error, or loss function $L = (1/2)e^T e$, then the “standard” learning rule is as follows:

$$\Delta\alpha_{ij} = -\eta \left\langle \frac{\partial L}{\partial \alpha_{ij}} \right\rangle = -\eta e^T \left\langle \frac{\partial e}{\partial u} \right\rangle \frac{\partial u}{\partial \alpha_{ij}}$$

where $\eta > 0$ is a learning rate constant, e is the feedback error, $\langle \partial e / \partial u \rangle$ is an estimate of the sensitivity derivatives matrix, and $\partial u / \partial \alpha_{ij}$ are the derivatives of the output commands with respect to the parameter α_{ij} . In other words, the controller’s parameters are adjusted from one trial, t , to the next based on the gradient of the loss, which is calculated using sensitivity derivative estimates, that is, $\alpha_{ij}(t+1) = \alpha_{ij}(t) + \Delta\alpha_{ij}$. Notably, if the sign of the elements of the estimated sensitivity matrix is incorrect (e.g., when perturbations such as inverting prisms are suddenly introduced), then gradient updates will be in the wrong direction, and adaptation will be maladaptive (see (Abdelghani et al. 2008) and Online Resource 1 for details). In other words, a sudden flip in sensitivity will make learning unstable—even small errors will be amplified over subsequent trials.

In many theories of sensorimotor learning, there is no mechanism in place to adjust the estimates of the sensitivity derivatives themselves (Braun et al. 2009; Dean et al. 2002; Kawato and Gomi 1992; Porrill et al. 2004; Pouget and Snyder 2000; Thoroughman and Shadmehr 2000; Todorov and Jordan 2002). Fixed sensitivity estimates work very well under a wide range of conditions. Many tasks, like reaching to targets while wearing refracting prisms, do not alter the sensitivity matrix at all (Figure 1b; Online Resource 1). Other tasks alter the magnitudes, but not the signs of the matrix elements (e.g., reaching to targets viewed through compressing or expanding prisms, see Online Resource 1), and these also cause no problems because the standard learning rule will still yield parameter

updates of the right signs, so learning will proceed in the right direction—it might just take a little longer. But some tasks, like pointing with inverted vision, alter the signs of sensitivity derivatives, and this alteration is fatal to the fixed updated rules (Fig. 1c). In most cases, one can judge whether a visual distortion, $T(\cdot)$, will cause trouble for the feedback rule by looking for sign changes in the sensitivity matrix, $\partial e / \partial u$, after redefining the feedback error to take into account the transformation of the target and reached location, for example, $e = T(x^*) - T(u)$. See Online Resource 1 and Abdelghani et al. (2008) for an in-depth discussion of this criterion.

Two mechanisms for updating sensitivity derivative estimates have been proposed in the literature. It has been suggested that sensitivity estimates can be learned via a technique called the *distal teacher* (Jordan and Rumelhart 1992). Under this mechanism, a forward model (Mehta and Schaal 2002; Wolpert and Miall 1996) of the motor system and world is learned and motor command updates (i.e., Δu in Fig. 1a) are computed by back propagating sensory errors through this model. In this case, sensitivity derivatives are stored implicitly in the parameters of the forward model which itself is trained by minimizing the error between predicted and actual sensory consequences.

While distal teachers solve the problem of updating sensitivity estimates, they largely have been ignored since back propagation is considered by many to be biologically implausible (Kawato and Gomi 1992; Mazzoni and Krakauer 2006; Wolpert and Miall 1996). More recently, Abdelghani et al. (2008) suggested a mechanism called *implicit supervision*, which learns to explicitly represent sensitivity derivatives in its neural firing and thus does not require back propagation to deliver the correct motor errors for learning (Abdelghani et al. 2008). Implicit supervision circumvents the need for back propagation by building a model of *changes* in trial-to-trial errors (rather than a forward model of sensory consequences) and by employing an expansion recoding scheme (e.g., (Kawato and Gomi 1992; Rolls and Deco 2002) to provide the requisite basis function for learning new sensitivity derivatives. While these two mechanisms suggest fundamentally different neural implementations, both use gradient-based updates to adjust their sensitivity estimates and make qualitatively similar predictions about overt behavior on the inverted vision tasks.

To demonstrate the predictions made by the two classes of visuomotor learning (i.e., those which update their sensitivity derivatives and those which do not), we simulated exemplar models performing our reversal task with random targets. The first model used a fixed estimate of the sensitivity matrix and the second model used the mechanism of implicit supervision to update an estimate of the sensitivity matrix online (Abdelghani et al. 2008). Both models

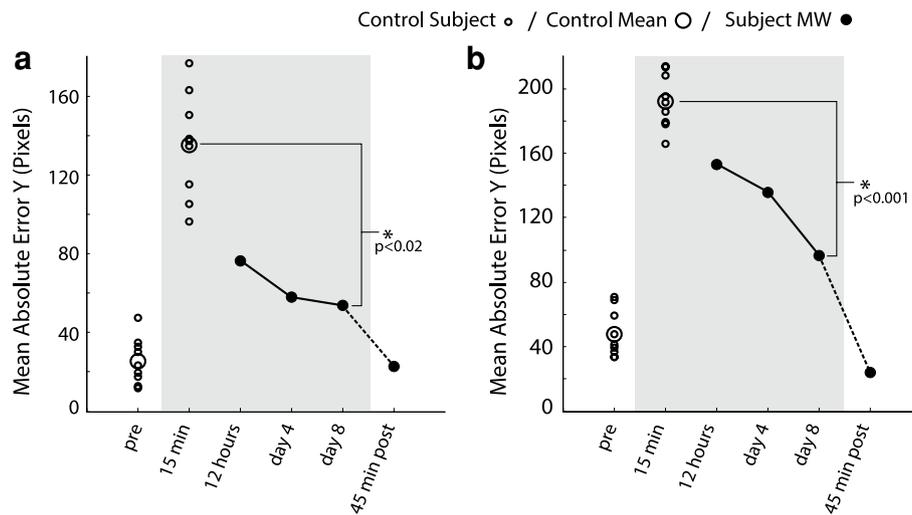


Fig. 4 Results of long-term prism adaptation experiments. Subject MW wore up-down reversing prisms continually for 8 days. MW was tested on two tasks after wearing the prisms for: 12 h, 4 days, 8 days, and then again 45 min after removing the prisms. **a** Mean performance of MW (*small filled circles*) and control subjects (who wore the same reversing prisms, but only for the duration of the task; *small open circles*; $n = 9$) in the reversal reaching task with random targets in the no prisms (mean of 80 targets) and prisms condition (mean of 80 targets). In the “no prisms” task subjects wore goggles with the same field of view as in the prisms variant, but with no prisms. MW shows marked improvement in performance after 12 h, and again after 4 days, but only a

small additional improvement after 8 days; 45 min after removing the goggles, MW’s performance returned to normal levels. **b** Mean performance of MW and control subjects ($n = 9$) on the jump task in which targets were displaced a random distance above or below the original target during the subject’s reach (no prisms = mean of 100 trials, prisms = mean of 300 trials). MW’s performance again shows marked improvements from the controls after 12 h, 4 days, and 8 days. However, even after wearing the reversing goggles continuously for 8 days, MW is still significantly worse than control subjects ($p = 0.013$; two sample t test); 45 min after removing the goggles, MW’s performance was as good or better than unperturbed control subjects

adapted the synaptic weight parameters of a simple radial basis function network (Moody and Darken 1989), which mapped 2D input targets to 2D reach locations (see Online Resource 1 for details). When the reversing perturbation is introduced, the performance of the standard model gets progressively worse during perturbed trials and only begins to improve once the perturbation is removed (Fig. 5a). In contrast, the implicit supervision variant initially worsens when the perturbation is introduced, but then, as soon as its sensitivity estimate has been sufficiently updated (i.e., once it has crossed zero, Fig. 5c), its performance improves smoothly toward baseline (Fig. 5b). Qualitatively similar behavior is given by models which use a distal teacher to update sensitivity derivatives.

Discussion

We have linked the fundamental difference between visuomotor adaptations to displaced versus inverted vision to the concept of sensitivity derivatives. Unlike displacements, inversions require the brain to update its estimate of these derivatives in order to correct the feedback rule used to update the motor commands. We also looked at long- and short-term adaptation to inversions in human

subjects. Paralleling qualitative observations (Harris 1965; Kohler 1963; Rock 1966, 1973; Stratton 1896, 1897), our long-term experiment demonstrates quantitatively that subjects show significant, though incomplete, improvement on inverted visuomotor tasks. Our results also suggest a complex interplay between the gradient-based learning that appears to underly most “implicit” learning and cognitive strategies.

Many current theories of sensorimotor learning do not mention sensitivity derivatives. Some explicitly assumes that sensitivity estimates are known and fixed (i.e., known innately), citing the biological implausibility of distal teachers which rely on back propagation (Dean et al. 2002; Kawato and Gomi 1992; Porrill et al. 2004). Others implicitly assume that sensitivity derivatives are a fixed property of the system or at least offer no mechanism by which they could be learned (Pouget and Snyder 2000; Thoroughman and Shadmehr 2000; Todorov and Jordan 2002). Theories in this class differ greatly among themselves, but we will group them under the heading of “standard” models, meaning any scheme where learning is based on fixed estimates of sensitivity derivatives. Opposed to these, the *distal teacher* technique (Jordan and Rumelhart 1992; Wolpert and Kawato 1998; Wolpert and Miall 1996) and the recent theory of *implicit supervision* (Abdelghani et al. 2008)

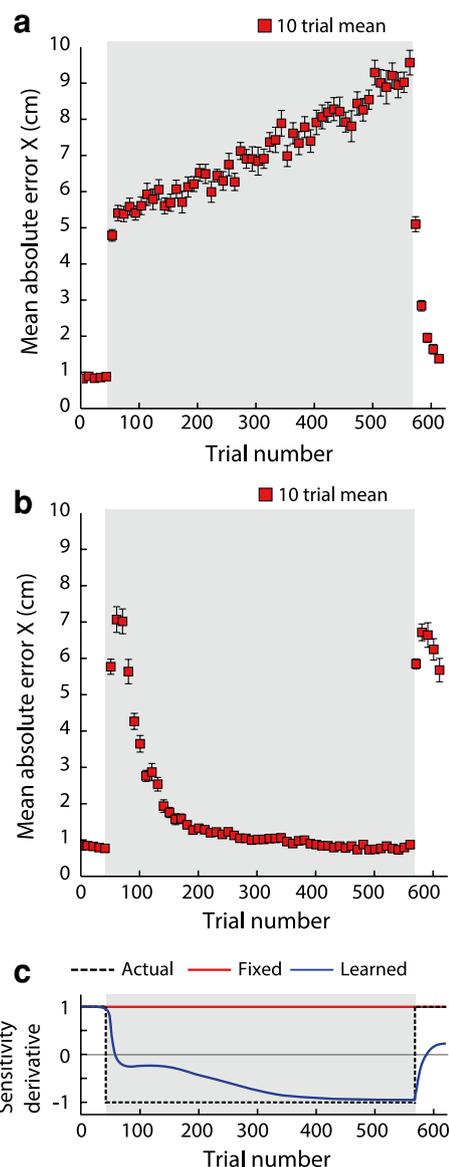
Fig. 5 Modeled results. **a** Simulated data produced by the model which does not update its feedback rule under the reversing prism condition when targets are drawn randomly from a uniform distribution; that is, task is the same as for Fig. 2d. Plot shows errors produced when the model does not learn the partial derivative (i.e., the “standard” gradient update). Note that performance continually degrades. The model predicts exponential degradation, so we have set the learning rate to a very small constant here for visualization purposes. **b** Simulated data when the model uses implicit supervision to learn the sensitivity derivatives online. The model initially produces an increase in error (until trial ~30), but quickly decreases the error as soon as its estimate of the partial has been corrected (i.e., crossed 0). Note the pronounced aftereffect. Error bars are SEM. **c** Mean estimates of the sensitivity derivative for the two model variants. When the prisms are introduced, the actual sensitivity (dashed line) flips from 1 to -1 . The standard model has its estimate (red line) fixed throughout the experiment. On the other hand, the model variant with implicit supervision learns the new sensitivity. Normal learning is resumed as soon as implicit supervision has learned the new sign of the sensitivity derivative (about 30 trials after the prisms are introduced). During washout, implicit supervision has to relearn the original sensitivity before normal learning may resume. The standard model begins to do useful learning as soon as the washout trials begin because the fixed sensitivity remains appropriate for the case where no perturbation is applied (color figure online)

incorporate mechanisms by which the sensitivity matrix itself can be learned.

The critical test between these two classes of theory is to see how subjects perform on tasks where the sensitivity derivatives are suddenly reversed in sign—not just changed in magnitude but carried across zero. As we show in the Results and Online Resource 1, the *reversal* task (pointing to targets under inverted vision) is just such a test. In such tasks, standard models imply that performance should become increasingly worse and never improve (Fig. 5a). Models which relearn sensitivity derivatives predict that performance should initially worsen, until the brain relearns the reversed sensitivity derivatives, and then get better in a smooth and stable fashion (Fig. 5b; Abdelghani et al. 2008).

Neither theory fits the empirical data well. Both classes of model straightforwardly account for the results of *no-reversal* tasks. However, standard models are explicitly ruled out by the results of the fixed target reversal experiments: The theory predicts that performance should degrade monotonically, but most subjects eventually improve.

Models which relearn sensitivity derivatives predict the *average* human data in the fixed target task: Subjects initially worsen and then improve. But the predictions do not match the traces for individual subjects. The theory predicts that, as soon as the new sensitivity derivatives are learned, improvement should be smooth and steady (Fig. 5b, c). In contrast, our individual subjects show erratic, rather than steady, improvement (Fig. 2b). Also, the model predicts a powerful aftereffect where again performance initially



degrades and then improves (Fig. 5b), whereas our subjects who improved in this task showed no consistent aftereffect (Online Resource 1 and Fig. 2c).

In the reversal random target experiment, results also mismatch theory. The theory again predicts initial worsening followed by steady improvement (Fig. 5b), but the expected timescale depends on learning constants and other factors. Subjects show initial degradation followed by a plateau in learning, though they might have improved given more practice (i.e., >500 trials). As well, during washout, subjects show a large aftereffect, but not of the kind predicted by implicit supervision, which predicts that performance should initially degrade before returning to the baseline (Fig. 5b).

Results from the long-term experiment are similarly difficult to reconcile with existing theories. By the eighth day,

MW did show significant improvement in both the reversal and reversal-jump tasks, ruling out standard models and suggesting that she may have learned the new sensitivity derivative. However, the tasks remained difficult for MW; adaptation was unstable and never complete, even on the final day of the study. Thus, while MW's performance eventually improved, as predicted by implicit supervision, her behavior never acquired the proficiency attained by control subjects without prisms or that expected if she had simply updated her sensitivity estimates and then proceeded to adapt as she would under displacing prisms. In addition, MW reported that her improved performance was contingent on actively maintaining attention and had not become automatic as with performance under displacing prisms. The fact that performance remained unstable even after 8 days and returned to baseline almost immediately after removing the prisms suggests that the original implicit rule (with an incorrectly signed estimate of the sensitivity derivative) is still active and dragging performance away from a good solution; that is, the implicit rule appears to continue to “fight” cognitive and/or other strategies even after long-term exposures.

Similar patterns can be seen in earlier experiments and everyday observations, if we examine them with sensitivity derivatives in mind. For example, there are many human tasks that require adaptation to reversed derivatives; for example, the sensitivity matrix is reversed when a dentist drills teeth viewed in a mirror, or when a doctor performs surgery by way of an endoscopic camera (Cresswell et al. 1999). The derivatives are also flipped when we switch to a new up-down throttle mapping for a flight simulator or video game. In all of these cases, the “normal” relation between error and command is reversed; for example, feedback errors ordinarily requiring a leftward correction now require a rightward one, and vice versa (Fig. 1b). Yet, people manage to cope, contradicting standard models and in part supporting implicit supervision (Fig. 5).

Other experiments presented animals with adaptation problems by surgically altering their sensors or limbs (Brinkman et al. 1983; Harris 1965; Illert et al. 1986; Kohler 1963; Rock 1966, 1973; Singer et al. 1982; Sperry 1941, 1943a, b, 1947; Stratton 1896, 1897), and again, the hard problems are the ones that change the signs of sensitivity derivatives. These studies do not mention sensitivity derivatives, but they did show that some perturbations are much more difficult to adapt to than others, either permanently disabling the subject or requiring very long adaptation. For example, Sperry transposed antagonistic nerves and muscles in rats (Sperry 1941, 1943b), while others surgically rotated the eyeballs of newts, frogs, and cats by 180° (Singer et al. 1982; Sperry 1943a). In all of these studies, the animals were permanently impaired: For example, for ever after, the newts and frogs would spring to the left

to catch food presented on their right. These results contradict implicit supervision and support the standard model. To reconcile them with the prism results, we can point to the species difference and suggest that simpler brains are unable to learn sensitivity derivatives. In monkeys, the story is less clear. Monkeys whose nerves for antagonistic arm muscles were transposed showed some recovery of function (Sperry 1947; Sugita 1996): Over the course of 3 years, most learned to inhibit maladaptive movements and reach their goals. But to do it, they used what Sperry called “tricks,” and they still occasionally lapsed into inappropriate, reversed movements, even for well-practiced actions.

How can we explain the monkeys' lapses and our own subjects' erratic progress and short-lived aftereffects? There are several possible explanations. Unstable performance might reflect a conflict between two controllers (see for example Wolpert and Kawato 1998), the old and the new, or it may mean that estimates of sensitivity derivatives, while they are being revised, are labile because they have not yet been consolidated. We might also attribute them to the interaction between straightforward gradient learning and “cognitive” or “explicit” strategies (Sperry 1947), though these latter cannot yet be cashed out into a concrete computational theory. Reliance upon attention and short-lived aftereffects are both hallmarks of cognitive strategies, and adaptation instability induced by an interaction between implicit learning and cognitive strategies has been observed previously (Cunningham 1989; Brooks et al. 1995; Martin et al. 1996; Mazzoni and Krakauer 2006; Sailer et al. 2005; Anguera et al. 2010). Thus, understanding inversion adaptation as a mixture of simple gradient learning and “cognitive” strategies seems promising (Fernandez-Ruiz et al. 2011).

Welch (1974) was among the first to advance the idea that motor learning might be composed of deliberate and automatic parts, and more recently, several authors have explicitly explored the contribution of cognitive strategies (Taylor and Ivry 2011; Block and Bastian 2012; Seidler et al. 2012). This view—that adaptation to inversions involves a combination of gradient-based and cognitive strategies—also presents two distinct challenges. First, a mixture implies that the extent to which observable behavior represents gradient learning versus a cognitive strategy is unclear—making empirical data difficult to decipher. Second, there are few computational models that attempt to mechanistically capture “cognitive” motor strategies and their interactions with gradient-based learning (see Online Resource 1). In Mazzoni and Krakauer (2006), which observed an interaction between gradient (or “implicit”) learning and a cognitive strategy, this second issue is mitigated by the fact that the cognitive strategy employed was known—that is, subjects were asked to employ a specific

strategy. However, for visual inversions and muscle reversals, the cognitive strategies employed by subjects are not explicit and indeed may vary from moment to moment. We might guess at some of the cognitive strategies employed by subjects. For example, subjects may resort to trial and error exploration of the mapping between their actions and perceived consequences after gradient learning is recognized to fail, or they may have a discrete insight about the perturbation and subsequently reach to the location opposite to where they view the target. But subjects may also be employing other less straightforward strategies, and behavior may be a mix of several. Future work on understanding inversion adaptation may do well to focus on the algorithmic properties we need to cover the facts, namely that the mechanism invoked by reversals is unstable, apparently lapsing or overcorrecting from one trial to the next, and leaves little or no trace when sensitivity derivatives are restored to normal.

Our task with probe trials tested the extent to which people could discover and use the “obvious” cognitive strategy (i.e., aiming at the mirror position of the presented target). Interestingly, subjects who successfully adapted in this task worsened on the generalization probes, while the non-adapters group improved on the probes. The fact that the adapters group did not improve on the probe trials indicates that they settled on a restricted strategy that does not generalize. It is difficult to explain the fact that adapters and non-adapters displayed opposite probe performance trends, though it seems plausible that both trends result from spatial generalization of maladaptive implicit learning [e.g., Mazzoni and Krakauer (2006)], which in the case of the non-adapters coincidentally drives their probe trial responses closer to the target.

The present study has focused on a simple measure of performance—end-point error. As such, there are several limitations which should be acknowledged. Notably, neither in-flight trajectories of reach nor the reaction times of participants (i.e., the time between target presentation and movement initiation) were examined. These factors are examined by Gritsenko and Kalaska (2010) who studied rapid online corrections to a reversing or “mirror” visuomotor transformation. They found that subjects adapted to the transformation (notably with fixed target positions) but that when the targets were displaced laterally immediately after movement onset, there was a suppression of rapid feedback corrections. Corrections under the mirror transformation occurred at a longer latency that was similar to the onset time of voluntary corrections in subjects who had not previously experienced the mirror transformation. Their results suggest relatively immutable fast feedback control which can at best be suppressed—thus performance gains in the reversal tasks take place almost exclusively via a slower “voluntary” process.

Werner and Bock (2010) examined adaptation to left-right reversed vision and present additional evidence supporting the idea that adaptation involves a mixture of implicit and cognitive processes. Interestingly, their data show that target location (with respect to the axis of rotation) appears to alter the relative contribution of these two components. For example, for targets in the direction perpendicular to the axis of reversal, performance improvement is discrete and relatively unstable, suggesting a primarily cognitive solution. Behavior for targets in the direction parallel to the reversal axis exhibits the phenomena observed here; performance is seen to worsen and then plateau, likely due to the application of a maladaptive implicit learning. Adaptation to targets between the two extremes exhibits features of both cognitive and implicit processes.

Our attention in the paper had focused on two approaches to accomplishing the inversion task: cognitive strategies and gradient learning. A third mechanism, which may be at work, is reinforcement learning (RL). The idea being that if gradient-based learning has failed and no good cognitive strategy is known, then trying random perturbations of previous actions and subsequently using actions similar to those which gave better outcomes is a viable approach. The line between gradient-based learning and RL is somewhat blurred. Reinforcement learning can be viewed as gradient learning where the gradient is unknown and is thus estimated from trial-to-trial using correlations between the loss function and the learned parameters. After it becomes clear that the sensitivity information is wrong, RL learning might act as a natural fallback. Future modeling work could examine the patterns of behavior expected from a mixture of a reinforcement and gradient-based learning with incorrect sensitivity information.

In conclusion, we have identified sensitivity derivatives as a crucial concept for understanding visuomotor learning and have shown that, from an algorithmic point of view, there is a sharp distinction between learning tasks in which the derivatives change only in magnitude and those where they are reversed in sign. This idea accounts for the basic difference in adaptation time course between displaced versus inverted vision. Our work also demonstrates that a complete understanding of inversion adaptation will require both empirical and computational innovations.

Acknowledgments We would like to thank Mohamed N. Abdelghani for conversations relating to the theoretical aspects of this paper. T.P.L. was supported by an Ontario Graduate Scholarship and by the Sun Microsystems of Canada Scholarship in Computational Sciences and Engineering. J.F.R. was supported by UNAM grant PAPIIT IN202810 and CONACYT grant 102314. P.M-B. was supported by a CONACyT award. D.B.T. was supported by CIHR. N.F.T. was supported by NSERC.

References

- Abdelghani MN, Tweed DB (2010) Learning course adjustments during arm movements with reversed sensitivity derivatives. *BMC Neurosci* 11:150
- Abdelghani MN, Lillcrap TP, Tweed DB (2008) Sensitivity derivatives for flexible sensorimotor learning. *Neural Comput* 20(8):2085–2111
- Anguera JA, Reuter-Lorenz PA, Willingham DT, Seidler RD (2010) Contributions of spatial working memory to visuomotor learning. *J Cogn Neurosci* 22(9):1917–1930
- Åström KJ, Wittenmark B (1995) Adaptive control, 2nd edn. Addison-Wesley, Reading
- Block HJ, Bastian AJ (2012) Cerebellar involvement in motor but not sensory adaptation. *Neuropsychologia* 50:1766–1775
- Braun DA, Aertsen A, Wolpert DM, Mehring C (2009) Motor task variation induces structural learning. *Curr Biol* 19(4):352–357
- Brinkman C, Porter R, Norman J (1983) Plasticity of motor behavior in monkeys with crossed forelimb nerves. *Science* 220(4595):438–440
- Brooks V, Hilperath F, Brooks M, Ross HG, Freund HJ (1995) Learning “what” and “how” in a human motor task. *Learn Mem* 2(5):225–242
- Callier FM, Desoer CA (1991) Linear system theory. Springer texts in electrical engineering. Springer-Verlag, New York
- Cresswell AB, Macmillan AI, Hanna GB, Cuschieri A (1999) Methods for improving performance under reverse alignment conditions during endoscopic surgery. *Surg Endosc* 13(6):591–594
- Cunningham HA (1989) Aiming error under transformed spatial mappings suggests a structure for visual-motor maps. *J Exp Psychol* 15(3):493–506
- Dean P, Porrill J, Stone JV (2002) Decorrelation control by the cerebellum achieves oculomotor plant compensation in simulated vestibulo-ocular reflex. *Proc R Soc Lond B* 269:1895–1904
- Fernandez-Ruiz J, Diaz R, Moreno-Briseno P, Campos-Romo A, Ojeda R (2006) Rapid topographical plasticity of the visuomotor spatial transformation. *J Neurosci* 26(7):1986–1990
- Fernandez-Ruiz J, Wong W, Armstrong IT, Flanagan JR (2011) Relation between reaction time and reach errors during visuomotor adaptation. *Behav Brain Res* 219(1):8–14
- Gritsenko V, Kalaska JF (2010) Rapid online correction is selectively suppressed during movement with a visuomotor transformation. *J Neurophysiol* 104:3084–3104
- Harris CS (1965) Perceptual adaptation to inverted, reversed, and displaced vision. *Psychol Rev* 72(6):419–444
- Illert M, Trauner M, Weller E, Wiedemann E (1986) Forearm muscles of man can reverse their function after tendon transfers: an electromyographic study. *Neurosci Lett* 67(2):129–134
- Jordan MI, Rumelhart DE (1992) Forward models: supervised learning with a distal teacher. *Cogn Sci* 16(3):307–354
- Kawato M, Gomi H (1992) The cerebellum and VOR/OKR learning models. *Trends Neurosci* 15(11):445–453
- Kitazawa S, Kimura T, Uka T (1997) Prism adaptation of reaching movements: specificity for the velocity of reaching. *J Neurosci* 17(4):1481–1492
- Kohler I (1963) The formation and transformation of the perceptual world. *Psychological Issues*, vol 3, No. 4, Monograph 12, International Universities Press, New York, pp 1–173
- Kornheiser AS (1976) Adaptation to laterally displaced vision: a review. *Psychol Bull* 83(5):783–816
- Linden DE, Kallenbach U, Heinecke A, Singer W, Goebel R (1999) The myth of upright vision. A psychophysical and functional imaging study of adaptation to inverting spectacles. *Perception* 28(4):469–481
- Martin TA, Keating JG, Goodkin HP, Bastian AJ, Thach WT (1996) Throwing while looking through prisms. II. Specificity and storage of multiple gaze-throw calibrations. *Brain* 119(4):1199–1211
- Mazzoni P, Krakauer JW (2006) An implicit plan overrides an explicit strategy during visuomotor adaptation. *J Neurosci* 26(14):3642–3645
- Mehta B, Schaal S (2002) Forward models in visuomotor control. *J Neurophysiol* 88(2):942–953
- Moody J, Darken CJ (1989) Fast learning in networks of locally-tuned processing units. *Neural Comput* 1(2):281–294
- Porrill J, Dean P, Stone JV (2004) Recurrent cerebellar architecture solves the motor-error problem. *Proc Biol Sci R Soc Lond B* 271(1541):789–796
- Pouget A, Snyder LH (2000) Computational approaches to sensorimotor transformations. *Nat Neurosci* 3:1192–1198
- Redding GM, Wallace B (1990) Effects on prism adaptation of duration and timing of visual feedback during pointing. *J Mot Behav* 22(2):209–224
- Redding GM, Rossetti Y, Wallace B (2005) Applications of prism adaptation: a tutorial in theory and method. *Neurosci Biobehav R* 29(3):431–444
- Rock I (1966) The nature of perceptual adaptation. Basic Books, New York
- Rock I (1973) Orientation and form. Academic Press, New York
- Rolls ET, Deco G (2002) Computational neuroscience of vision. Oxford University Press, Oxford
- Sailer U, Flanagan JR, Johansson RS (2005) Eye-hand coordination during learning of a novel visuomotor task. *J Neurosci* 25(39):8833–8842
- Seidler RD, Bo J, Anguera JA (2012) Neurocognitive contributions to motor skill learning: the role of working memory. *J Mot Behav* 44(6):445–453
- Sekiya K, Miyauchi S, Imaruoka T, Egusa H, Tashiro T (2000) Body image as a visuomotor transformation device revealed in adaptation to reversed vision. *Nature* 407(6802):374–376
- Singer W, Treter F, Yinon U (1982) Central gating of developmental plasticity in kitten visual cortex. *J Physiol* 324:221–237
- Sperry RW (1941) The effect of crossing nerves to antagonistic muscles in the hind limb of the rat. *J Comp Neurol* 75(1):1–19
- Sperry RW (1943a) Effect of 180 degree rotation of the retinal field on visuomotor coordination. *J Exp Zool* 92(3):263–279
- Sperry RW (1943b) Functional results of crossing sensory nerves in the rat. *J Comp Neurol* 78(1):59–90
- Sperry RW (1947) Effect of crossing nerves to antagonistic limb muscles in the monkey. *Arch Neurol Psychiatry* 58(4):452–473
- Stratton GM (1896) Some preliminary experiments on vision without inversion of the retinal image. *Psychol Rev* 3(6):611
- Stratton GM (1897) Vision without inversion of the retinal image. *Psychol Rev* 4(4):341
- Sugita Y (1996) Global plasticity in adult visual cortex following reversal of visual input. *Nature* 380(6574):523–526
- Taylor JA, Ivry RB (2011) Flexible cognitive strategies during motor learning. *PLoS Comput Biol* 7(3):e1001096
- Thoroughman KA, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. *Nature* 407(6805):742–747
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5(11):1226–1235
- Welch RB, Choe CS, Heinrich DR (1974) Evidence for a three-component model of prism adaptation. *J Exp Psychol* 103(4):700–705
- Werner S, Bock O (2010) Mechanisms for visuomotor adaptation to left–right reversed vision. *Hum Mov Sci* 29:172–178
- Wolpert DM, Kawato M (1998) Multiple paired forward and inverse models for motor control. *Neural Netw* 11(7–8):1317–1329
- Wolpert DM, Miall RC (1996) Forward models for physiological motor control. *Neural Netw* 9(8):1265–1279