Task goals influence online corrections and adaptation of reaching movements

John E. Schlerf and Richard B. Ivry

Task goals influence online corrections and adaptation of reaching movements

John E. Schlerf[1,2] **and Richard B. Ivry**[1,3]

[1]*Helen Wills Neuroscience Institute, University of California, Berkeley, California;* [2]*Department of Physical Medicine and Rehabilitation, Johns Hopkins University, Baltimore, Maryland; and* [3]*Department of Psychology, University of California, Berkeley, California*

Schlerf JE, Ivry RB. Task goals influence online corrections and adaptation of reaching movements. *J Neurophysiol* 106: 2622–2631, 2011. First published August 17, 2011; doi:10.1152/jn.00301.2010.— Everyday movements often have multiple solutions. Many of these solutions arise from biomechanical redundancies. Often, however, the goal does not require a unique movement. To examine how people exploit task-related redundancy, we observed as participants produced three-dimensional (3-D) reaching movements, moving to one of two rectangular targets that were diagonally oriented in the frontal (*x, y*) plane. On most trials, the movement was perturbed by a vertical, velocity-dependent force. Since participants were free to move in 3-D space, online corrections could involve movement along the perturbed, vertical dimension, as well as the nonperturbed, horizontal dimension. If the motor system exploits task redundancies, then corrections along the horizontal dimension should depend on the orientation of the target. Consistent with this prediction, participants modified both the horizontal and vertical coordinates of the trajectory over the course of learning, and the horizontal component was sensitive to the orientation of the target. Furthermore, participants produced online corrections with a horizontal component that brought the hand closer to the target. These results suggest that we not only correct for mismatches between expected and experienced forces but also exploit task-specific redundancies to efficiently improve performance.

force field; optimal control

REACHING TO GRASP AN OBJECT can usually be accomplished by a large set of kinematic patterns and final joint configurations, the so-called degrees of freedom problem. Despite the potential selection problem created by an excess of degrees of freedom, the motor system effortlessly chooses a motor command to achieve a desired end position. Optimal control theory has provided a formal framework for understanding the constraints underlying this process (Todorov and Jordan 2002). For example, certain combinations of changes about the muscles and/or joints are less costly (require less energy expenditure or increase end-state comfort) than others. We thus perform optimally when we select an action with the lowest cost.

The principles of optimal control can be used to understand how people make adjustments to ongoing movements (Diedrichsen et al. 2010). Assuming that there is a control cost to such corrections, online adjustments should be most evident when they help ensure the desired task outcome. Conversely, adjustments should be reduced for deviations that are irrelevant to the task outcome. Consistent with these expectations, greater variability is observed along a dimension that is irrelevant for task outcome (a redundant dimension) than along a dimension that is crucial for task outcome (Müller and Sternad 2004;

Todorov and Jordan 2002). This idea is also encompassed in the notion of an "uncontrolled manifold," which proposes that variability is allowed among the set of coordinates in task space (or manifold) for which the task outcome is equivalent (Cusumano and Cesari 2006; Scholz and Schöner 1999).

Redundancy does not solely arise because of the excessive degrees of freedom in the motor system. Many tasks entail goals that afford redundancies unrelated to the biomechanics of our limbs. When closing a door, for example, the same force can be applied anywhere along the vertical axis of the door with equivalent results. Similarly, to increase stability, we can grasp a stairway railing at multiple locations. In such situations, we should expect that an optimal planning system would allow greater variability along redundant dimensions defined in task space. Indeed, when learning a new skill, people have been shown to exploit such redundancies, producing greater variability along task-irrelevant dimensions compared with task-relevant dimensions (Cusumano and Cesari 2006; Müller and Sternad 2004). This aspect of optimal control has also been highlighted in recent studies examining aiming strategies when people reach for a target surrounded by asymmetric penalty zones. Under such conditions, participants aim for an optimal point that accommodates uncertainty related to their ability to control movements (Trommershäuser et al. 2005) as well as uncertainty inherent in the environment (Trommershäuser et al. 2003).

Although the work reviewed above emphasizes how task-based redundancy affects learning, less attention has been given to how task-based redundancy influences the online control of well-learned movements. This issue is important not only when considering movement execution but also for planning and learning. The stability of motor performance is sensitive to many factors. Fatigue, injury, clothing, gravity, and the current posture all affect how a limb responds to a neural command. Despite these multiple sources of variability, we manage to move with comparable proficiency across an extensive range of conditions.

This robust performance indicates that the motor system is highly adaptive, adjusting the control signal to incorporate the context. Studies of sensorimotor adaptation have generally been limited to conditions in which the task goal is defined as a single point (with some tolerance) in a two-dimensional (2-D) workspace. For example, in studies involving force field perturbations (Izawa et al. 2008; Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994; Taylor and Thoroughman 2007; Thoroughman and Shadmehr 2000) or visuomotor transformations (Fishbach and Mussa-Ivaldi 2008; Mazzoni and Krakauer 2006; Sober and Sabes 2003; Tseng et al. 2007), the task goal is defined by a target location in 2-D space. In these

Address for reprint requests and other correspondence: J. E. Schlerf, 707 N Broadway, Rm. G-04, Baltimore, MD 21205 (e-mail: jschlerf@jhmi.edu).

2622   0022-3077/11 Copyright © 2011 the American Physiological Society   www.jn.org

contexts, adaptation requires the adjustment of an internal model to counteract the effect of the perturbation such that the movement terminates in the vicinity of the target location. Optimal control models have provided elegant accounts of learning under such conditions, particularly through explaining deviations from perfectly straight trajectories in visuomotor (Fishbach and Mussa-Ivaldi 2008) and force field tasks (Izawa et al. 2008). However, the focal nature of the targets in these studies precludes the analysis of whether adaptive processes exploit task-based redundancies. The goal of the current study was to address this issue.

Dimensional redundancy is not present in the typical force field study (e.g., when participants make planar movements in a viscous curl field). The force field involves two dimensions, and thus there is no irrelevant dimension. To introduce redundancy, we had participants reach in a 3-D workspace. We presented the target as a rectangular region, oriented diagonally on a virtual surface in the frontal plane. Contact at any point within the region was considered a successful reach. Use of this target instead of a single point made the extraneous third dimension relevant to task performance. To evaluate whether online feedback and adaptive processes incorporate information regarding task-based redundancies, we introduced a consistent force perturbation during the movement. This perturbation was restricted to the vertical dimension, thus displacing the hand from the target at an oblique angle. Any horizontal component of the correction is irrelevant to force adaptation (e.g., does not cancel out the perturbation) but, nonetheless, remains relevant to task performance.

We focused on how participants learned to respond to this perturbation. If learning involves generating an accurate model of the environmental perturbation, then the participants' behavior should be independent of the orientation of the target. That is, we would expect to observe an anticipatory trajectory that counteracts the perturbing effects of the force field. Alternatively, learning may incorporate task-based redundancy related to the rectangular targets. This hypothesis predicts that participants will not only adjust their trajectories to counteract the effects of the vertical perturbation but will also show systematic deviations along the horizontal axis to bring the hand closer to the target.

MATERIALS AND METHODS

Participants. Twenty-one right-handed, college-age individuals (10 male, 11 female, mean age 19.5 ± 1.8 yr) with normal or corrected-to-normal vision participated in the main experiment. All volunteers provided informed consent and were compensated for their time. All experiments were performed under a protocol approved by the Committee for the Protection of Human Subjects at University of California, Berkeley. We randomly assigned these participants to 2 groups, with 10 participants initially trained with a 45-degree target, and 11 initially trained with a 135-degree target.

Apparatus. Participants grasped the handle of a robotic manipulandum (PHANToM 3.0L; SensAble Technologies; http://www.sensable.com) capable of recording position and generating force along any of the three Cartesian axes. The robot was controlled by custom software written in Visual C++ using the OpenHaptics library. Control signals to the robot were updated at 1,000 Hz, and the output of the device was subsampled at 200 Hz for off-line analysis. As a safety measure, the force output was capped at 9.0 N. The handle was allowed to rotate freely around any axis (roll, pitch, and yaw), but the angle of handle rotation was not recorded. Participants viewed the environment

through a mirror. While the mirror precluded vision of the participant's arm, the cursor indicating hand position was presented to appear near the actual location of the hand, facilitating the subjective feeling of immersion in a 3-D environment (see Fig. 1A).

Task. By moving the manipulandum, the participant controlled a 6-mm white spherical cursor that moved in the 3-D workspace. The cursor remained visible at all times during the experiment. At the onset of a trial, the hand was pulled gently by the robot to the start position, positioning the cursor within an 8-mm sphere located at the participants' vertical midline. In virtual space, the simulated start position was ~10 cm in front of the eyes and 25 cm below eye level. After the cursor was maintained within this sphere for 1,000 ms, one of two rectangular targets, 12 cm long and 1 cm wide, was presented along the back wall, which was 12 cm from the start location. The long axis of the rectangular targets was oriented at either 45 or 135 degrees from horizontal (angles increase counterclockwise). Participants were informed that the center of the bar was located nearest to the hand's starting position but that contact with any spot along the bar would be rewarded equally. Participants were required to make a single reaching movement, attempting to land within the target. At the termination of the reach (contact with the virtual back wall), an 8-mm sphere (the "feedback sphere") was presented on the surface to provide additional feedback of the movement endpoint. We also manipulated the color of the feedback sphere to train participants to move at a relatively constant speed given that the perturbing force was velocity dependent (see below). If the movement duration was <275 ms, the feedback sphere was red, informing the participant to slow down. If the movement duration was >325 ms, the feedback sphere was green, informing the participant to speed up. Movement durations between 275 and 325 ms fell within the desired speed criterion; on these trials, the feedback sphere was white. Thirty-five percent of all total reaches fell within this criterion. Since most of the other movements were performed at a speed close to this range [average movement time 310 (SD 68) ms] we did not exclude trials from the analysis based on the movement duration criterion.

To further motivate the participants, a running point tally was presented on the screen after each trial. If the cursor landed in the target within the appropriate movement duration window, 5 points were awarded. If the endpoint location was outside the target region, the score was decreased by 1 point. If the endpoint location was accurate but the movement duration was outside the desired range, the score remained unchanged.

Participants first practiced the task until they were comfortable moving within the virtual environment and could readily interpret the feedback. This typically involved 10–20 reaches. The main experiment began with a training block of 50 reaches in a null field (no perturbing forces). The orientation of the target was fixed for the entire block (45 or 135 degrees, counterbalanced across participants). The force field was then introduced in a second block of 50 trials (the exposure block), with the target orientation the same as in the training block. In this block, the movements were perturbed by a viscous curl field in which an upward vertical force was generated as a function of velocity into the workspace (z-axis). The force was generated as follows:

$$\begin{bmatrix} F_x \\ F_y \\ F_z \end{bmatrix} = B \begin{bmatrix} V_x \\ V_y \\ V_z \end{bmatrix}$$

$$B = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & \mu \\ 0 & 0 & 0 \end{bmatrix}$$

Whereas the perturbation was velocity dependent, the viscosity term μ was position dependent to ensure that all trials were identical during the initial phase of the movement. Viscosity was zero for the initial 30 mm of movement and was then quickly ramped up to 7.5
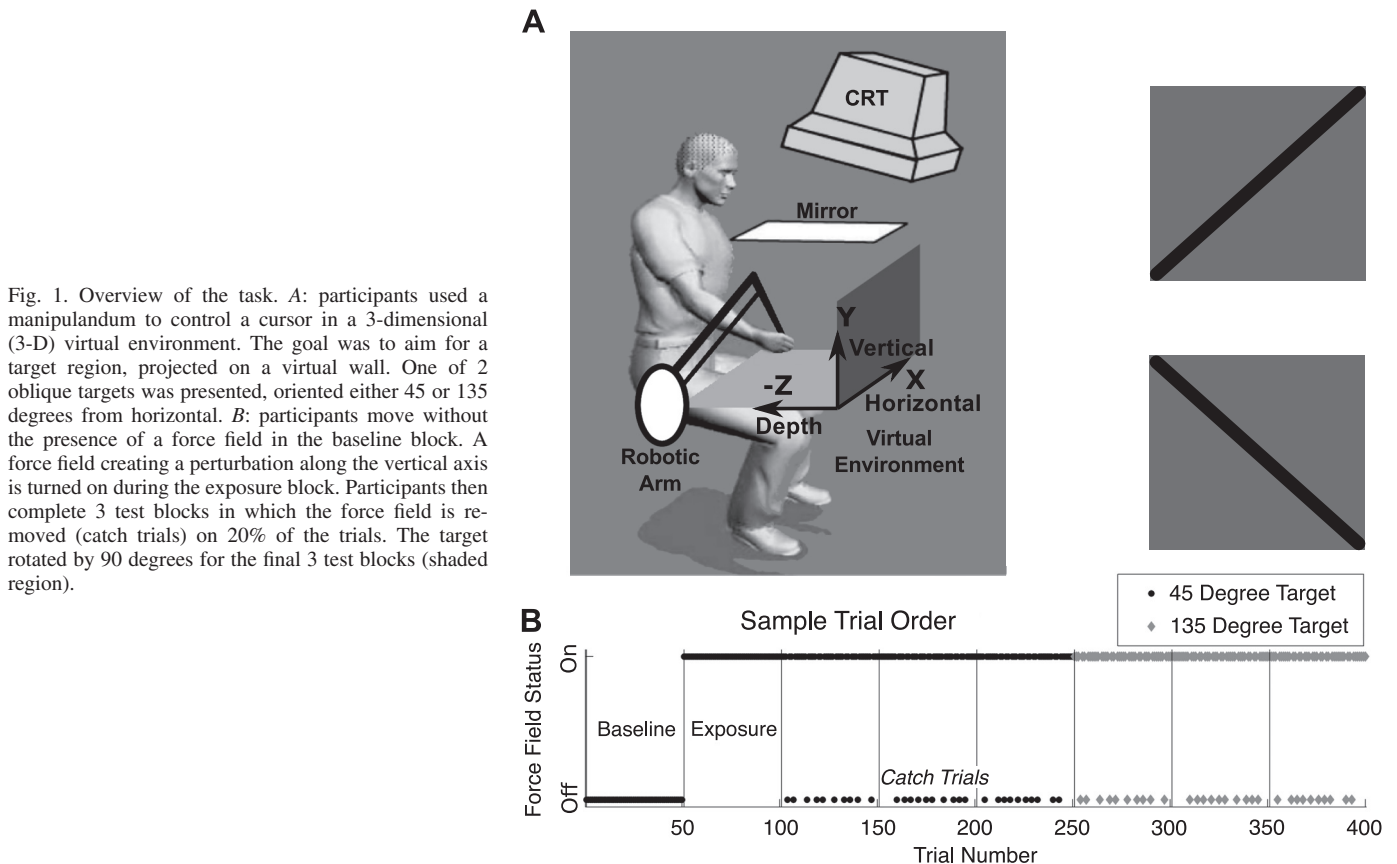
Fig. 1. Overview of the task. *A*: participants used a manipulandum to control a cursor in a 3-dimensional (3-D) virtual environment. The goal was to aim for a target region, projected on a virtual wall. One of 2 oblique targets was presented, oriented either 45 or 135 degrees from horizontal. *B*: participants move without the presence of a force field in the baseline block. A force field creating a perturbation along the vertical axis is turned on during the exposure block. Participants then complete 3 test blocks in which the force field is removed (catch trials) on 20% of the trials. The target rotated by 90 degrees for the final 3 test blocks (shaded region).

Ns/m over the next 30 mm of movement. This created a window during which any response to the perturbation was purely a result of feedforward processes. We used this window to assess aiming.

After the learning block, the participant completed 6 test blocks of 50 trials each (Fig. 1*B*). For the first three test blocks, the orientation of the target surface was always the same as in the training and learning blocks. For the last three test blocks, the target was rotated by 90 degrees. Within each test block, the force field was present for 40 (80%) of the trials and turned off on 10 of the trials (20% catch trials). The catch trials were randomly determined with the constraint that a catch trial was always preceded by at least two force field trials. A short break was provided between each block.

Reaching with the left hand. An additional group of 18 participants was recruited to repeat the experiment using the left arm instead of the right. This group allowed us to conduct a replication study of our main effects of interest, as well as to explore potential biomechanical interactions between the arm and manipulandum that may have influenced the results. Experimental procedures were identical to those used by the right-hand reaching group, with 8 of the participants initially reaching to a 45-degree target and 10 initially reaching to a 135-degree target.

Nonredundant control experiment. Ten additional subjects were recruited for a control experiment in which the target was a circle of 1 cm in diameter. All other factors were the same as in the main experiment. This control condition allowed us to establish how participants respond to a perturbation when reaching to a specific location rather than an oriented object.

Data analysis. Data analysis was performed in Matlab (The Math-Works; http://www.mathworks.com). Velocity profiles were manually reviewed to identify and remove trials in which participants stopped moving forward before hitting the wall or took a "wind-up" by initially moving the cursor away from the wall before moving forward (criterion of 20 mm from start position). Approximately 5% of trials were discarded based on these criteria.

Analysis of the vertical dimension assesses the response along the axis of the force perturbation. Since the force field had no horizontal component, the horizontal dimension was used to examine goal-oriented adjustments to the force. As such, analysis of the response along this axis provided a more direct test of the effect of target orientation on adaptation.

Catch trials, during which the perturbation was unexpectedly removed, were of special interest. The endpoints of these trials were compared with the endpoints of the immediately preceding trials to probe the feedforward prediction of the forces and their online changes when that prediction fails. These analyses were conducted using a 2 × 2 repeated-measures ANOVA with within-subject factors of trial type (precatch or catch trial) and target orientation.

The absence of the force field on catch trials likely caused single-trial learning (Thoroughman and Shadmehr 2000). To probe this phenomenon, we used the trials immediately following the catch trials (postcatch), focusing on the early phase of the trajectories. For this analysis, the position of the hand immediately before the onset of the force field (30 mm into the reach) was recorded and analyzed in a 2 × 2 repeated-measures ANOVA with factors of trial type (catch or postcatch) and target orientation. To reduce the amount of variability in these early measures, the preforce position on each trial was adjusted by subtracting the value on the immediately preceding precatch trial.

For trajectory analyses, the reaches were standardized such that there was one value per millimeter along the z-axis (by binning existing values and interpolating missing values). This procedure allowed for spatially accurate averaging across the horizontal (x) and vertical (y) dimensions. A permutation analysis was used to statistically evaluate the divergence of two trajectories. For each point along the z-axis (the reach direction), we computed the difference between the average position (across subjects) when reaching toward the 45- and 135-degree targets. A null distribution was then synthesized by taking the observed trajectories and assigning them at random to

either the 45-degree or the 135-degree groups (independent of the actual target) and computing the difference. This procedure was repeated 10,000 times. Assuming a null hypothesis of no true difference in trajectory, the actual differences should have fallen somewhere in the middle of the synthesized distribution. We defined significant divergence by requiring the actual difference to fall within a tail of the null distribution (taking an α of 0.05) for more than 10 mm.

RESULTS

Endpoint distribution during baseline. Participants found this task challenging. During the baseline block, before any force perturbation was present, participants landed in the target region on an average of 14.8 (SD 4.1) of the 25 trials. Despite not being explicitly instructed to do so, reach endpoints were grouped near the center of the target (see Fig. 2, *A–C*, for sample subjects' performance). Note that this was the closest point on the wall from the starting location (and participants were aware of this). The distribution of endpoints was largely elliptical, with the major axis skewed vertically (see Fig. 2*A* for a sample subject). We assessed the angle of the distribution in two ways. First, we used a principal component analysis to examine the direction of maximal spread (Fig. 2, *D* and *E*). When subjects reached to the 45-degree target, the principle axis was oriented, on average, 82.5 (SD 15.3) degrees; when subjects reached to the 135° target, the principle axis was oriented at 112.3 (SD 12.2) degrees. The angles of the principle axes were influenced by the orientation of the target, an effect confirmed by a Wilcoxon rank sum test (rank sum 63, $P < 0.002$).

Next, we looked for the direction that contained the most trial-by-trial fluctuation. We assessed this by projecting the endpoint data (originally in *x* and *y*) onto a new unit vector and rotating this new vector until we found the direction in which the autocorrelation at lag 1 was furthest from zero (Fig. 2, *D* and *E*). Because this measurement will be symmetric across 180 degrees, we computed this twice: once with the unit vector constrained between 0 and 180 degrees, and a second time with the unit vector constrained between 90 and 270 degrees. The second measurement had lower variance across subjects, indicating that this was closer to the true distribution; thus we report summary statistics for this direction. The direction of greatest trial-by-trial fluctuation was nearly horizontal for both targets (black lines in Fig. 2, *D* and *E*). For the 45-degree target, the angle of maximal correction was 179.2 (SD 43.3) degrees, whereas for the 135-degree target, the angle of maximal correction was 192.3 (SD 31.6) degrees. No significant difference was observed between target angles (rank sum 100, $P = 0.5$). This result indicates that the anisotropy observed in the endpoints is not solely a result of trial-by-trial adjustments. It may instead be due to uncorrected variability along the long dimension of the target, consistent with what would be expected by optimal feedback control theory.

In the nonredundant control experiment, movements were directed toward a circular target. The endpoints were distributed along the vertical dimension (Fig. 2*C*), with the principal axis of the endpoints at 80.3 (SD 22.3) degrees (Fig. 2*F*). More importantly, the axis along which participants were correcting in the control experiment was also near vertical [mean 103.1 (SD 25.2) degrees], in contrast to that observed with the elongated targets in the main experiment. We assume that the anisotropy in endpoints for the control task may reflect the combined effects of feedback corrections within a trial as well as adjustments between trials. With the circular target, both of these components would influence the anisotropy along the same axis.

Learning. When exposed to the force field in the exposure block, the endpoints of the movements were immediately perturbed
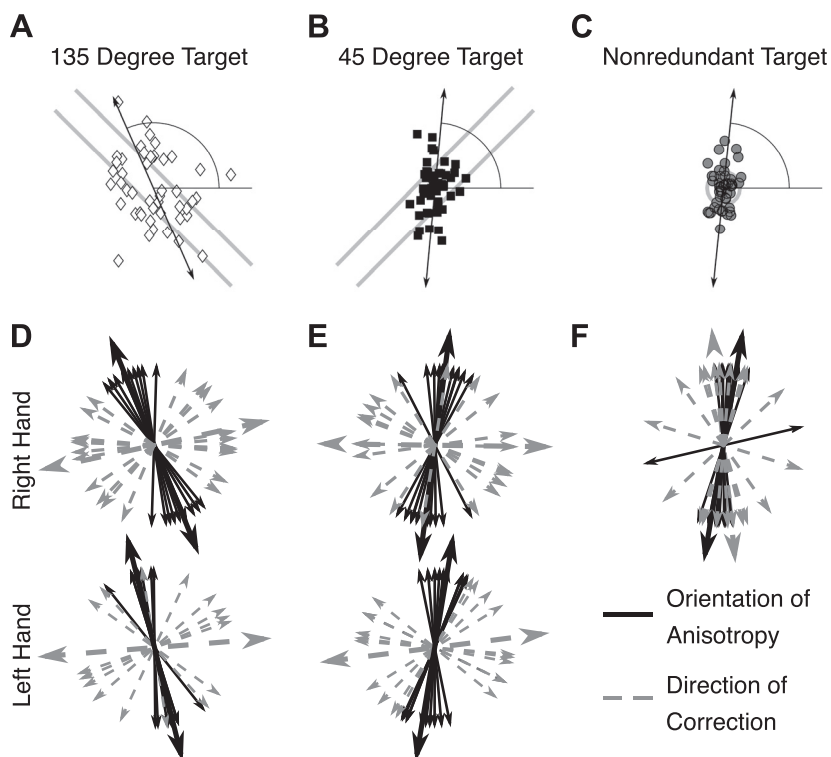

Fig. 2. Endpoint distributions during the baseline block. *A–C*: data for a representative subject reaching toward the 135-degree target (*A*), the 45-degree target (*B*), and the nonredundant control target (*C*). Light gray lines represent the target regions. Horizontal lines indicate the horizontal axis, and oblique lines indicate the principle component regression for the observed distributions. *D–F*: summary data for the anisotropy analysis. Black solid lines represent the principle components of the endpoint distributions. Gray dashed lines represent the direction along which the most trial-by-trial correction was observed based on a lag-1 autocorrelation. Short lines represent individual subjects, whereas the longer line represents the group average. Data from the right hand are shown at *top* and from the left hand at *bottom* for the 135-degree target for both hands (*D*), the 45-degree target for both hands (*E*), and the nonredundant control target for the right hand only (*F*).

along the vertical axis (Fig. 3*A*). Averaged over the first 5 trials, the endpoint was shifted by 47.4 (SD 28.5) mm. Participants learned to compensate for this perturbation (Fig. 3*A*), with the endpoint value dropping to 17.5 (SD 19.7) mm above the baseline value when averaged over the final 5 trials of the exposure block. Interestingly, the perturbation was not fully corrected along the vertical axis for the rectangular targets [$t(20) = 4.1$, $P = 0.001$]. In contrast, the perturbation was more completely corrected for the circular target in the control experiment, with the endpoints dropping from an initial perturbation of 33.1 (SD 13.8) mm above baseline values to 8.2 (SD 9.2) mm above baseline values over the final 5 trials.

The pattern for the horizontal component is more complex. Note that the force field perturbation, if not fully compensated along the vertical axis, should result in an endpoint that is to the left of the 45-degree target and to the right of the 135-degree target. With the introduction of the force field at the

A
Endpoints over Time (Vertical, Right Hand)

B
Endpoints over Time (Horizontal, Right Hand)

C
Endpoints over Time (Vertical, Left Hand)

D
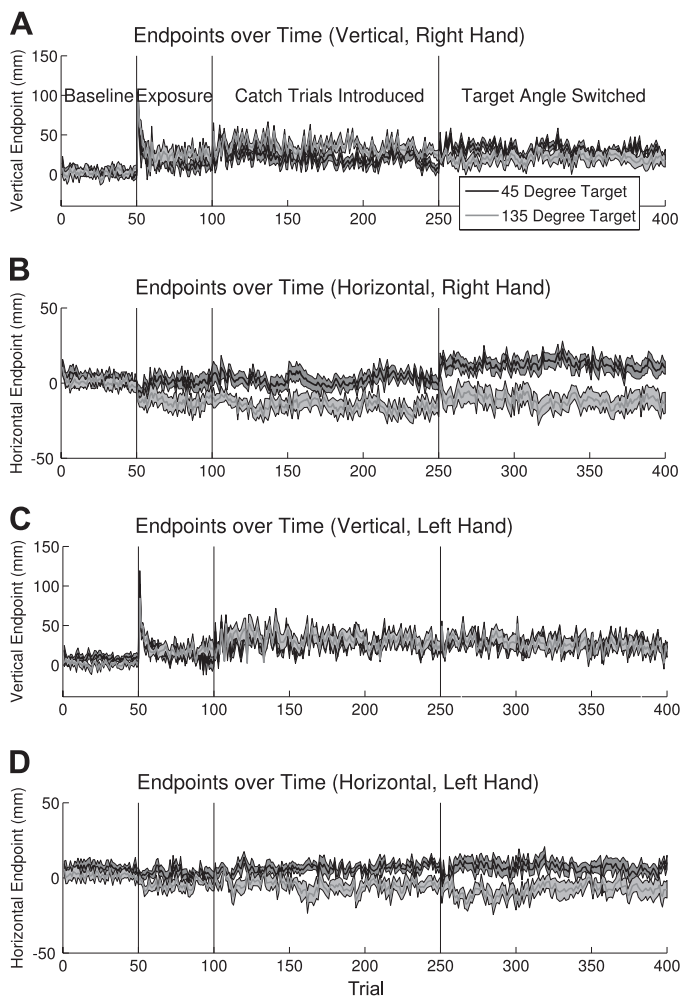Endpoints over Time (Horizontal, Left Hand)

Fig. 3. Changes in endpoint location over the course of practice. Black lines represent data for reaches toward the 45-degree target, and gray lines represent reaches toward the 135-degree target, regardless of the training target. *A*: the vertical coordinate immediately increases and then rapidly drops when the force field is applied during the exposure block. The pattern is similar for the 2 targets. *B*: the horizontal coordinate diverges for the 2 targets during the exposure block, and this separation is maintained over the test blocks. Negative numbers indicate endpoints to the left of center. *C*: for the left hand, the vertical coordinate shows the same immediate increase and rapid drop when the force perturbation is introduced. *D*: the horizontal coordinate for the left hand diverges across the test blocks.

start of the learning block, participants exhibited an immediate shift to the left. This deviation was likely a biomechanical consequence of an upward perturbation of the right arm. Importantly, a target-specific effect on the horizontal distribution of endpoint locations becomes evident over the course of learning, with the deviations resulting in endpoint locations that are brought closer to the target. For the 45-degree target, the endpoints shifted gradually in the rightward direction (black lines in Fig. 3*B*). In contrast, for the 135-degree target, the endpoints remained shifted to the left (gray lines in Fig. 3*B*). The divergence of the two functions along the horizontal axis increased the likelihood that the endpoint location would fall within the target surface. This profile is consistent with the hypothesis that the participants' response to the perturbations incorporated properties of the target orientation. Comparing the change in the horizontal endpoint between the first five trials in the exposure block and the last five trials in the exposure block failed to reveal a significant difference [$t(19) = 1.64$, $P = 0.12$]. However, the difference between baseline performance and the final five endpoints in the exposure block was influenced by the target angle [$t(19) = 2.17$, $P < 0.05$].

The average magnitude of the force pulse, across subjects, was 6.6 (SD 0.8) N. Participants reached the force maximum of 9 N on an average of 25 (SD 33) trials out of 320 reaches performed in the force field, with only 6 of 21 subjects saturating the forces on more than 10% of trials. The forces provided by the robot were only vertical, with no leftward torques presented to participants. We hypothesized that the initial leftward perturbation was the result of biomechanical interactions within the arm such that an upward force resulted in a slight adduction of the arm, which motivated the repetition of this experiment while asking participants to use the left arm.

After three test blocks, the target orientation was reversed such that participants initially trained with a 45-degree target were tested with the 135-degree target, and vice versa. Participants rapidly adjusted to the new target. This adjustment did not entail a change in the vertical coordinate of the endpoints. To statistically evaluate this, we compared the average final vertical coordinate in the final block with the first target (*block 3*) and the first block with the second target (*block 4*). A 2 × 2 repeated-measures ANOVA with factors of trained target (45 or 135 degrees) and current target (45 or 135 degrees) revealed a significant effect of trained target [$F(1,19) = 5.0$, $P < 0.05$] with participants trained on the 135-degree target reaching, on average, 13 mm higher than participants trained on the 45-degree target. There was no main effect of current target [$F(1,19) = 0.66$] and no interaction [$F(1,19) = 0.73$] of these two factors. Considering the horizontal coordinate, we observed only a significant main effect of current target [$F(1,19) = 14.3$, $P < 0.005$], with participants reaching 7 mm to the right when aiming for the 45-degree target and 11 mm to the left for the 135-degree target. Since neither coordinate showed a significant interaction between the training and performance, we elected to group these conditions to increase our power when considering catch trials.

Catch trials. The six test blocks included catch trials, in which the perturbing vertical force was not presented. These trials provide a probe of the participants' underlying internal model of the task environment. For comparison, we used the trials that immediately preceded the catch trials (precatch trials). Note that the force field was presented on the precatch

trials. Assuming the planning process anticipated the upward perturbation, we expected that the average endpoints on catch trials would fall below the endpoints on precatch trials. By examining the horizontal coordinate on these trials, we could assess whether an online correction process incorporated target information. Specifically, will an online correction in the absence of an expected vertical perturbation include a horizontal component that increases the likelihood that the movement will end within the target surface?

Examination of the vertical dimension showed that the presence of the force field (precatch trials vs. catch trials) had a strong effect on the endpoint locations (Fig. 4A). Overall, the average endpoint location was 23.7 (SD 15.0) mm above the center of the target during precatch trials, and 10.9 (SD 15.6) mm below the center of the target during catch trials. The vertical endpoint was minimally affected by the target orientation.

Along the horizontal dimension (Fig. 4B), participants maintained the overall bias that they exhibited on the force field trials (to the right for the 45-degree target and to the left for the 135-degree target, consistent with aiming slightly above center). This effect adds further support to the hypothesis that the planned trajectory incorporates features of both the expected perturbation and target orientation. During catch trials, the trajectories for both targets moved in a rightward direction (relative to precatch trials). However, this effect was larger for the 135-degree target. Note that for the 135-degree target, a lower trajectory will cause the participant to be too far to the left. Thus the larger shift to the right increased the likelihood that the movement would terminate within the target.

To statistically analyze performance during the test blocks, we performed two 2×2 repeated-measures ANOVAs, with within-subject factors of target angle (45 or 135 degrees) and trial type (precatch or catch). For the vertical dimension, this analysis revealed a main effect of trial type [$F(1,20) = 126.4$, $P < 0.001$], confirming that the vertical endpoints were lower on catch trials compared with force field trials. There was no effect of target angle [$F(1,20) = 0.11$, $P = 0.77$]. Interestingly, the interaction approached significance [$F(1,20) = 3.8$, $P = 0.06$]. For the horizontal dimension, both main effects were reliable [target angle: $F(1,20) = 11.6$, $P < 0.003$; trial type: $F(1,21) = 47.3$, $P < 0.001$]. Moreover, the interaction term was highly significant [$F(1,20) = 34.2$, $P < 0.001$]. The main effect is consistent with the hypothesis that the participants

anticipated the force field and incorporated a horizontal deviation that was expected to move the hand closer to the target (rightward for 45-degree target, leftward for 135-degree target). The interaction, in which a rightward shift on catch trials was even more pronounced for the 135-degree target, indicates that the target information was also incorporated into online corrections.

During the control experiment, the average endpoint of the reaches to the circular target was 30.6 (SD 10.5) mm below the target during catch trials compared with an average endpoint 15.0 (SD 7.7) mm above the target on precatch trials. Thus, the catch trials produced a significant shift along the vertical (perturbed) direction [paired $t(9) = 13.0$, $P < 0.001$]. There was also a reliable shift of the horizontal endpoint on catch trials [paired $t(9) = 2.8$, $P < 0.05$]. The hand ended up 2.2 (SD 2.6) mm to the left of the target on catch trials compared with 0.5 (SD 2.4) mm to the right of the target during precatch trials. Note that this effect is quite small relative to that observed for the rectangular targets in the main experiment.

Postcatch trials. The internal model of the expected force perturbation was violated on the catch trials. This resulted in an error signal that should have affected behavior on the subsequent trial (Thoroughman and Shadmehr 2000), although the effects of these changes may have been small given that the catch trials are infrequent. Postcatch trials provided a second look at adaptation. We would expect adaptation to be evident in the movement heading before the force field is applied. To examine such trial-by-trial changes, we compared the trajectories on catch and postcatch trials. For this analysis, we reduced the effects of overall drift in performance by subtracting the trajectory during the immediately preceding precatch trial from each trajectory. Since catch trials were not predictable, we expected the average catch trial heading before the expected force onset to be identical to that of the precatch trial, equivalent to expecting a value of 0 after the adjustment. Since the expected vertical force was absent, participants ended up reaching to a location significantly lower than the target. Assuming this error influenced their planning for the subsequent trial, we expected the initial trajectory on the postcatch trial to be higher. Consistent with this prediction, catch trials were unaffected during the initial 30 mm of the movement (Fig. 5A). In contrast, an upward shift in the trajectory was evident on the postcatch trials (Fig. 5B).
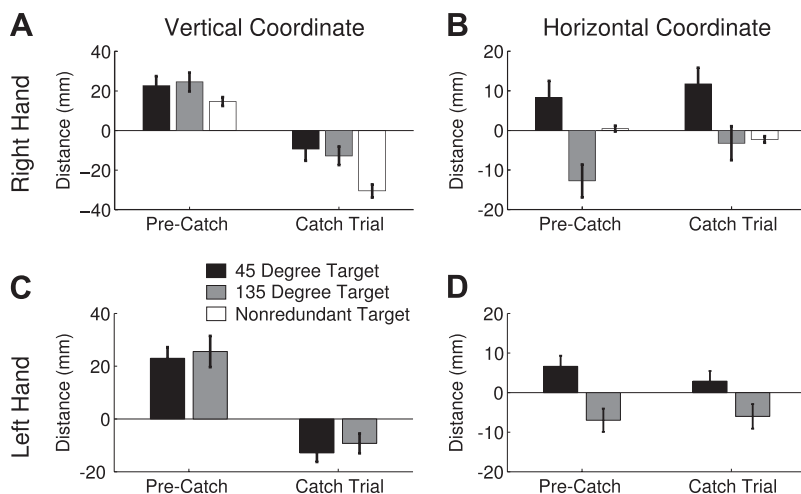

Fig. 4. Effect of catch trials on movement endpoints. Vertical (*A* and *C*) and horizontal (*B* and *D*) coordinates of the endpoints during precatch and catch trials are shown for the 45-degree target (black), the 135-degree target (gray), and nonredundant control target (white). The target orientation has a stronger effect on the horizontal component, with endpoints shifted toward the right for the 45-degree target and to the left for the 135-degree target. During catch trials, the endpoints for the oriented targets shifted to the left for reaching with the right hand (*B*), although the shift was stronger for the 135-degree target. For reaching with the left hand (*D*), the shift was more strongly rightward, and the interaction was again significant.

We also predicted a horizontal component of the adjustment in the postcatch trials. Since the endpoint analysis suggested that when participants reached to the 135-degree target the online correction was more strongly rightward compared with when they reached to the 45-degree target, we expected a similar effect on the postcatch trial. Looking first at the catch trial trajectories (Fig. 5C), the horizontal displacement was similar to that observed on the precatch trial (i.e., has a difference of 0) before the expected force onset. Immediately following this point, the trajectories deviated to the right for both targets. Near the end of the reach, there was an increase in the rightward deviation for the 135-degree target; for the 45-degree target, participants reduced the rightward deviation. On postcatch trials, there was a strong leftward deviation preceding force onset when participants reached toward the 45-degree target, with much less deviation, although still slightly leftward, when they reached toward the 135-degree target. Since participants missed to the right during catch trials, indeed more so for the 45-degree target, these adjustments are consistent with the notion that the motor system was using the error from the previous trial to adjust the planned trajectory, and doing so in a way that incorporated information concerning the orientation of the target.
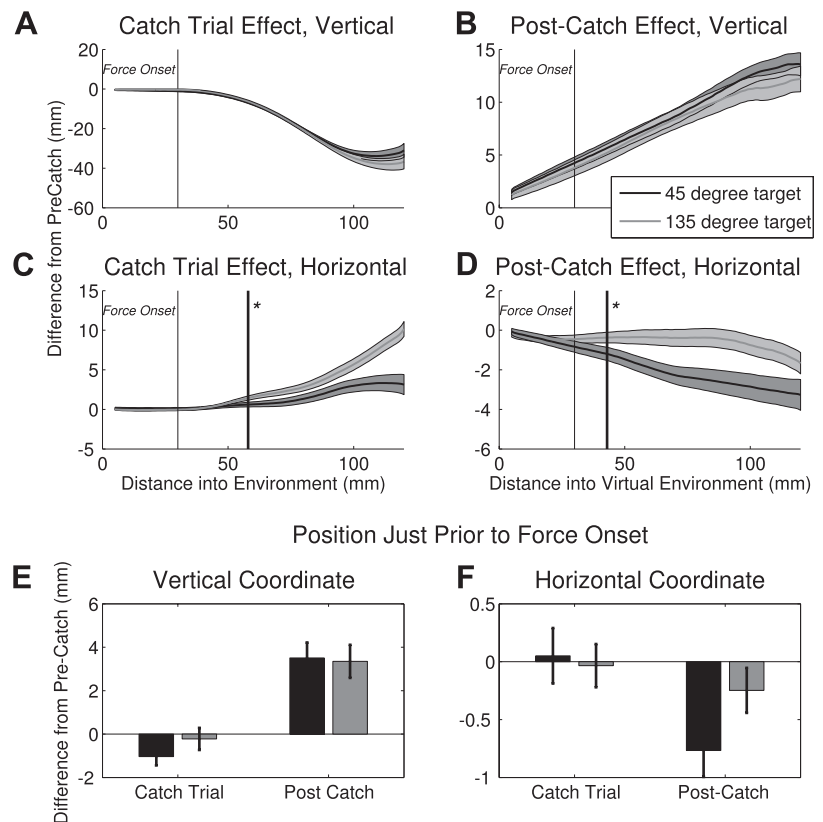
To quantify these effects statistically, we compared the normalized position at 30 mm into the reach, the first point at which the force field, if present, would be nonzero. For this analysis, we used two 2×2 repeated-measures ANOVAs, with within-subjects factors of target angle and trial type (catch or postcatch). For the vertical coordinate, a highly significant main effect of trial type was observed [$F(1,20) = 76.8, P < 0.001$]. The main effect of target angle was not significant [$F(1,20) = 0.55, P = 0.44$], nor was the interaction [$F(1,20) =$

2.27, $P = 0.11$]. For the horizontal coordinate, there was again a main effect of trial type [$F(1,20) = 18, P < 0.001$] but not target angle [$F(1,20) = 1.4, P = 0.24$]. The interaction between these two factors was reliable [$F(1,20) = 6.46, P < 0.02$]. This interaction indicates that the participants were adjusting their strategy to account for the previous error in task coordinates, rather than simply responding to the mismatch in the predicted force.

Latency of target effects. As noted above, statistical analysis of the horizontal component consistently revealed an effect of target orientation, supporting the hypothesis that responses to expected and unexpected perturbations incorporate information about behavioral redundancies. A final question of interest concerns the latency of these effects. To answer this question, we looked for the point at which the trajectories (specifically, the difference from the precatch trials; Fig. 5) diverged for the two targets. We performed this analysis separately for catch and postcatch trials using a permutation analysis (see *Data analysis*).

The trajectories along the vertical axis did not diverge for the two targets on catch trials. In contrast, a significant divergence of the horizontal component was observed 58 mm into the reach (Fig. 5C). On average, participants reached this point 179 (SD 31) ms after initiating the movement. When considered from the point of expected force onset, the latency is 75 (SD 10) ms. If only the first three test blocks for each subject are considered, the point of significant divergence occurs at 103 mm, or 200 (SD 25) ms after the time of the expected force onset. On postcatch trials, no divergence was observed for the vertical components of the trajectories. The horizontal components significantly diverged 43 mm into the reach (152 ms from start, 35 ms from force onset).

Fig. 5. Comparison of catch and postcatch trials. *A–D*: average trajectories along the vertical (*A* and *B*) and horizontal axes (*C* and *D*). To standardize the data, the trajectory from the precatch trial was subtracted from each catch and postcatch movement. Negative numbers for the horizontal component indicate a leftward change from the precatch trial. Arm position recorded 30 mm into the reach demonstrated strategic changes following a catch trial for the vertical coordinate (*E*) and a target dependency for the horizontal coordinate (*F*).

Reaching with the left hand. Certain features of the results obtained when participants reached with the right hand were at odds with what would be expected if the corrections were solely based on the target orientation. In particular, the upward force perturbation resulted in a slight rightward shift in the trajectory when participants reached with the right hand, independent of the target orientation. To test whether these effects might have a biomechanical origin, we repeated the experiment in a new group of participants who reached with the left arm. We did not observe an immediate leftward shift when the force perturbation was introduced, but there were indications of biomechanical effects. Importantly, we found that, when participants reached with the left arm, the main findings of interest were corroborated, with clear target-dependent effects visible.

When participants reached with the left hand, before any force perturbation was introduced, the pattern of results agreed with that of the right hand (gray lines in Fig. 2, *D* and *E*). For left hand reaches to the 45-degree target, the principle axis of the endpoints was oriented at 80.5 (SD 12.3) degrees. For left hand reaches to the 135-degree target, the orientation was 103.9 (SD 10.9) degrees. The angles in this case were influenced by the orientation of the target, confirmed by a Wilcoxon rank sum test (rank sum 111, $P < 0.001$). The direction of the most trial-by-trial correction when participants reached with the left hand was again mostly horizontal, oriented at 187.3 (SD 36.1) degrees for the 45-degree target and 184.6 (SD 48.8) degrees for the 135-degree target.

When the perturbation was introduced, learning along the vertical dimension proceeded as expected (Fig. 3*C*). There was no substantial initial deviation when we consider the horizontal endpoints, and furthermore, these endpoints did not diverge substantially during the first block containing a force perturbation (Fig. 3*D*). Considering the last five trials in the exposure block, no target-dependent effect was observed [$t(16) = 1.136$, $P = 0.28$]. However, over the course of the test blocks, the horizontal coordinates did diverge, similar to the divergence observed with the right hand.

During catch trials, the endpoints when participants reached with the right hand shifted toward the right (Fig. 4*B*). When participants reached with the left hand, however, we observed an overall trend to deviate to the left on catch trials (Fig. 4*D*). This shift was consistent with the force perturbation causing a slight adduction of the shoulder (as the lack of force caused abduction). These catch trial data were analyzed statistically using a 2×2 repeated-measures ANOVA, with within-subject factors of target angle (45 or 135 degrees) and trial type (precatch or catch). For the vertical dimension, this analysis revealed a main effect of trial type [$F(1,17) = 106.4$, $P < 0.001$], confirming that the vertical endpoints were lower on catch trials compared with force field trials. There was no significant effect of target angle [$F(1,17) = 1.96$, $P = 0.18$] and no interaction [$F(1,17) = 0.11$, $P = 0.74$]. For the horizontal dimension, a significant main effect of target angle was observed [$F(1,17) = 8.64$, $P < 0.01$]. The main effect of trial type was not reliable [$F(1,17) = 2.51$, $P = 0.13$], but the interaction was highly significant [$F(1,17) = 48.7$, $P < 0.001$], providing a replication of the key finding of target-dependent responses during catch trials.

We evaluated the normalized position at 30 mm into the reach (the first point at which the force field, if present, would be nonzero) for catch and postcatch trials when participants

moved with the left hand. For this analysis, we used 2×2 repeated-measures ANOVAs, with within-subjects factors of target angle and trial type (catch or postcatch). For the vertical coordinate, a highly significant main effect of trial type was observed [$F(1,17) = 120.2$, $P < 0.001$]. The main effect of target angle was not significant [$F(1,17) = 1.11$, $P = 0.31$], nor was the interaction [$F(1,17) = 0.18$, $P = 0.68$]. For the horizontal coordinate, there was no significant main effect of trial type [$F(1,17) = 0.59$, $P = 0.45$] or target angle [$F(1,17) = 1.4$, $P = 0.26$]. However, the interaction was significant [$F(1,17) = 19.9$, $P < 0.001$], replicating the main result of a target-dependent effect on trial-by-trial adjustments.

In summary, when participants used the left hand, target-dependent effects were qualitatively similar to those observed with use of the right hand. To identify target-independent effects (to highlight the biomechanical differences), we averaged across the two targets for each hand (Fig. 6). During catch trials, the overall rightward shifts observed when participants reached with the right hand were reversed when they reached with the left hand. This is consistent with the hypothesis of a biomechanical interaction between the arm and manipulandum, an interaction that slightly masked the target-dependent effects.

DISCUSSION

Studies of sensorimotor adaptation have generally involved reaching movements that are restricted to a 2-D plane with the goal defined as a single point. This paradigm has been adopted in a variety of tasks including visuomotor adaptation (Fishbach and Mussa-Ivaldi 2008; Mazzoni and Krakauer 2006; Sober and Sabes 2003; Tseng et al. 2007) and force field learning (Izawa et al. 2008; Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994; Taylor and Thoroughman 2007; Thoroughman and Shadmehr 2000). Successful learning requires the formation of an accurate internal model of the disturbance. The current study replicated these results in a 3-D task. Participants in our task were quite sensitive to the vertical perturbation and rapidly adjusted their motor plan to anticipate this force over the initial exposure trials. The perturbing force was absent on catch trials. In these trials the endpoints of the
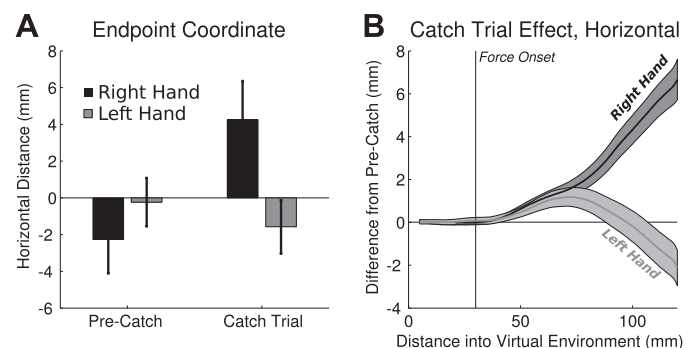

Fig. 6. Comparison of right and left hand performance to assess biomechanical effects of the force perturbation. *A*: independent of target orientation, the right hand shifted to the right when the force pulse was absent (catch trials), whereas the left hand shifted toward the left. This is consistent with the force pulse causing a slight adduction of the hand at the shoulder. *B*: when the entire trajectories during catch trials were considered, there was a slight rightward shift for both arms early in the reach, which became more pronounced for right arm movements and reversed to become an overall leftward shift when the left arm was used.

movements were much lower than on the standard trials, reflecting the fact that the participants were generating a downward correction in anticipation of the upward perturbing force. Moreover, the resulting error on the catch trials led to an adjustment on the next trial such that participants now appeared to be aiming for a higher location than on standard trials. This pattern of results is consistent with what has been observed in studies of planar force field adaptation (Fine and Thoroughman 2007; Taylor and Thoroughman 2007; Thoroughman and Shadmehr 2000).

Within such adaptation tasks, a second strategy is also possible. Participants could adjust their movements from trial to trial with the goal of minimizing behavioral error. Importantly, learning an internal model or minimizing task error would lead to very similar behavioral changes when the movements are performed in a 2-D workspace toward a 1-D target. To investigate whether participants were able to correct for task error, we designed a task that allowed movement in a 3-D space and expanded the goal region. The 3-D workspace allowed us to create a redundant dimension that was task relevant and orthogonal to the perturbing force.

Assuming that the participants' goal was to terminate the movement within the target region (as opposed to an implicit singular point at the middle of the target region), behavioral error can be defined as the distance from the major axis defining the target. In this way, we asked whether corrective movements were sensitive to variation in the orientation of the target, even when the perturbing force was held constant. If learning was based on a force compensation strategy, we expected to see adjustments restricted to the vertical axis. In contrast, if corrective processes were sensitive to task redundancies, then we expected to see adjustments along the horizontal axis, and most importantly, these corrections should be in opposite directions for the two targets. The results presented, when assessed in terms of trial-by-trial adjustments as well as on-line corrections within a trial, indicate that participants incorporated goal-based information when making well-practiced movements in the presence of an external perturbation.

Between-trial changes. Over the course of the experiment, participants adjusted their movements such that the endpoints would result in better performance. This is evident in a comparison of the data from the redundant and nonredundant targets. When participants reached for the oblique rectangular targets in the baseline training block, the endpoints were distributed along the primary axis of the target, yet trial-by-trial adjustments occurred along the horizontal axis. When participants reached for the small circular target, the endpoints were slightly skewed upward, and the trial-by-trial corrections exhibited a similar endpoint anisotropy. We also observed target-dependent adjustments over the course of the initial exposure block, adjustments that resulted in improved performance. The unexpected upward perturbing force resulted in an endpoint that was higher than expected. Since the target was an oblique rectangle, the endpoint was to the left for the 45-degree target and to the right for the 135-degree target. Thus the behavioral changes observed over the course of the initial exposure block, i.e., moving lower and to the right for the 45-degree target and lower and to the left for the 135-degree target, were in a direction that attenuated both the force field and the behavioral error.

During the test trials, the force field was unexpectedly removed on selected trials. This led to an adjustment in performance on the subsequent trial. This effect is believed to reflect rapid changes in an internal model acquired during force field adaptation (Fine and Thoroughman 2007; Thoroughman and Shadmehr 2000). In the current study, a compatible effect was observed when the vertical component of the trajectory was considered. For the horizontal component, however, the target-based effects could be assessed. The results confirm that participants utilized the target information to improve performance. The orientation of the target significantly affected the observed adjustment on trials following catch trials.

When learning a new task, participants are able to exploit the tolerance afforded by the task space (Müller and Sternad 2004). The current investigation extends this idea by showing adjustments in behavior that occur when well-learned movements are produced in a novel context. Utilizing goal-based redundancies may be a general feature of motor control.

Within-trial changes. Within-trial changes are important for evaluating whether the exploitation of redundant task dimensions arises from rapid feedback mechanisms or from feedforward processes. Catch trials provided a window to address this question. The adjustment of the horizontal component of the trajectory indicated that participants were using task goal information to minimize endpoint error. The absence of the upward perturbing force resulted in a lower endpoint than had been predicted. Consistent with the hypothesis that online corrections incorporate goal-based information, there was an adaptive rightward horizontal shift when participants reached for the 135-degree target. The results were more ambiguous for the 45-degree target. In this case, the overall shift was also to the right, albeit significantly attenuated compared with the 135-degree condition. Inspection of the trajectories on catch trials (Fig. 4C) suggests an incomplete, late-onset leftward correction when participants reached to the 45-degree target.

Although the manipulandum applied a purely vertical force field, the consistent rightward shift of the endpoints on catch trials indicates that the force field induced a horizontal perturbation. This displacement was also apparent during the initial trials of the exposure block. We assume this reflects a biomechanical effect arising from the interaction of the manipulandum and arm. This conjecture is supported by the observation that there was an overall leftward shift in a separate group of participants tested with their left arm (Fig. 6).

The catch trials further provide a measure of the timing of corrective movements in the absence of any external, perturbing force. When all six test blocks were considered, the trajectories to the two targets diverged ~75 ms after the expected onset time of the missing perturbation. This delay is faster than what would be expected given the time required for supraspinal feedback corrections (Allum 1975), particularly if one considers the necessary delay between muscular activation and resultant limb movement. Interestingly, when we considered only the first three test blocks during which the movements were limited to a single target, the time of divergence was 200 ms, within the expected time for supraspinal feedback loops. The latency data suggest that the corrections may initially utilize feedback processing (Jacobs and Horak 2007). However, with practice, goal-based information can lead to an additional feedforward contribution. Indeed, the cost of recomputing trajectories during the movement to reduce an imminent

error may be greater than the costs associated with adjusting feedback gains in a feedforward manner.

Implications for models of motor adaptation. Adaptation is commonly modeled as an error-based learning process. Models based on linear dynamical systems are capable of capturing much of the variability in performance (Cheng and Sabes 2006; Donchin et al. 2003; Thoroughman and Shadmehr 2000;). Although such models make simplifying assumptions about the nature of the error to be corrected, they can be used to ask which errors are most efficient for constraining learning. For example, in a 2-D visuomotor task, both visual feedback of endpoint error and the hand-to-cursor perturbation are acceptable error signals. However, Cheng and Sabes (2007) observed that a model in which these two inputs are combined did not sufficiently improve predictions of behavior over the two error signals alone. This suggests that for such simple tasks, the specification of the exact nature of the error signal is not critical for modeling behavior. However, the current results suggest that this may no longer be the case when task complexity is increased. Although the current data set is not well suited to detailed modeling, the results suggest that for reaching in a 3-D space to a redundant target, the signals underlying adaptation may need to incorporate information regarding both endpoint error and perturbation magnitude.

Recent experiments suggest that learning is not a singular process but may entail the parallel operation of multiple mechanisms. Models that incorporate multiple mechanisms are generally implemented to share an error signal while differing in the parameters representing retention and acquisition rates (Kording et al. 2007; Smith et al. 2006; but see Taylor and Ivry 2011). The current experiment suggests that models which include multiple mechanisms with nonidentical error signals would improve predictions about adaptation, particularly as the task becomes more complex. For example, a two-rate model in which the fast system learns from the perturbation while the slow system learns from the visual error would make very similar predictions in 2-D tasks while capturing the response to target redundancy observed in the current study. Future experiments with carefully constructed learning trajectories are necessary to explore the validity of this hypothesis.

Summary. The current study demonstrates that participants utilize a redundant target dimension to facilitate performance during learning. Furthermore, adjustments within a reach incorporate target-dependent information. These results underscore that the motor system is not merely reacting to the force perturbation during learning but optimizing behavior for the intended goal.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the author(s).

REFERENCES

Allum J. Responses to load disturbances in human shoulder muscles: the hypothesis that one component is a pulse test information signal. *Exp Brain Res* 22: 307–326, 1975.

Cheng S, Sabes P. Calibration of visually guided reaching is driven by error-corrective learning and internal dynamics. *J Neurophysiol* 97: 3057–3069, 2007.

Cheng S, Sabes P. Modeling sensorimotor learning with linear dynamical systems. *Neural Comput* 18: 760–793, 2006.

Cusumano J, Cesari P. Body-goal variability mapping in an aiming task. *Biol Cybern* 94: 367–379, 2006.

Diedrichsen J, Shadmehr R, Ivry R. The coordination of movement: optimal feedback control and beyond. *Trends Cogn Sci* 14: 31–39, 2010.

Donchin O, Francis J, Shadmehr R. 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, 2003.

Fine M, Thoroughman K. Trial-by-trial transformation of error into sensorimotor adaptation changes with environmental dynamics. *J Neurophysiol* 98: 1392–1404, 2007.

Fishbach A, Mussa-Ivaldi F. Seeing versus believing: conflicting immediate and predicted feedback lead to suboptimal motor performance. *J Neurosci* 28: 14140–14146, 2008.

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

Jacobs J, Horak F. Cortical control of postural responses. *J Neural Transm* 114: 1339–1348, 2007.

Kording K, Tenenbaum J, Shadmehr R. The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat Neurosci* 10: 779–786, 2007.

Lackner J, Dizio P. Rapid adaptation to Coriolis force perturbations of arm trajectory. *J Neurophysiol* 72: 299–313, 1994.

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

Müller H, Sternad D. Decomposition of variability in the execution of goal-oriented tasks: three components of skill improvement. *J Exp Psychol Hum Percept Perform* 30: 212–233, 2004.

Scholz J, Schöner G. The uncontrolled manifold concept: identifying control variables for a functional task. *Exp Brain Res* 126: 289–306, 1999.

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

Smith M, Ghazizadeh A, Shadmehr R. Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol* 4: e179, 2006.

Sober S, Sabes P. Multisensory integration during motor planning. *J Neurosci* 23: 6982–6992, 2003.

Taylor J, Ivry R. Flexible cognitive strategies during motor learning. *PLoS Comput Biol* 7: e1001096, 2011.

Taylor J, Thoroughman K. Divided attention impairs human motor adaptation but not feedback control. *J Neurophysiol* 98: 317–326, 2007.

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

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

Trommershäuser J, Gepshtein S, Maloney L, Landy M, Banks M. Optimal compensation for changes in task-relevant movement variability. *J Neurosci* 25: 7169–7178, 2005.

Trommershäuser J, Maloney L, Landy M. Statistical decision theory and the selection of rapid, goal-directed movements. *J Opt Soc Am A* 20: 1419–1433, 2003.

Tseng YW, Diedrichsen J, Krakauer JW, Shadmehr R, Bastian AJ. Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J Neurophysiol* 98: 54–62, 2007.