J Neurophysiol 100: 1455–1464, 2008.
First published July 2, 2008; doi:10.1152/jn.90334.2008.

Reach Adaptation: What Determines Whether We Learn an Internal Model of the Tool or Adapt the Model of Our Arm?

JoAnn Kluzik,[1,3,4] Jörn Diedrichsen,[5] Reza Shadmehr,[2] and Amy J. Bastian[1,3]

[1]*Department of Neurology and* [2]*Department of Biomedical Engineering, Johns Hopkins School of Medicine;* [3]*Kennedy Krieger Institute, Baltimore, Maryland;* [4]*Sargent College of Health and Rehabilitation Sciences, Boston University, Boston, Massachusetts; and* [5]*Wolfson Centre for Clinical and Cognitive Neuroscience, School of Psychology, Bangor University, Bangor, Gwynedd, United Kingdom*

Submitted 5 March 2008; accepted in final form 28 June 2008

Kluzik J, Diedrichsen J, Shadmehr R, Bastian AJ. Reach adaptation: what determines whether we learn an internal model of the tool or adapt the model of our arm? *J Neurophysiol* 100: 1455–1464, 2008. First published July 2, 2008; doi:10.1152/jn.90334.2008. We make errors when learning to use a new tool. However, the cause of error may be ambiguous: is it because we misestimated properties of the tool or of our own arm? We considered a well-studied adaptation task in which people made goal-directed reaching movements while holding the handle of a robotic arm. The robot produced viscous forces that perturbed reach trajectories. As reaching improved with practice, did people recalibrate an internal model of their arm, or did they build an internal model of the novel tool (robot), or both? What factors influenced how the brain solved this credit assignment problem? To investigate these questions, we compared transfer of adaptation between three conditions: catch trials in which robot forces were turned off unannounced, robot-null trials in which subjects were told that forces were turned off, and free-space trials in which subjects still held the handle but watched as it was detached from the robot. Transfer to free space was 40% of that observed in unannounced catch trials. We next hypothesized that transfer to free space might increase if the training field changed gradually, rather than abruptly. Indeed, this method increased transfer to free space from 40 to 60%. Therefore although practice with a novel tool resulted in formation of an internal model of the tool, it also appeared to produce a transient change in the internal model of the subject's arm. Gradual changes in the tool's dynamics increased the extent to which the nervous system recalibrated the model of the subject's own arm.

INTRODUCTION

Human reaching is highly adaptable to novel contexts and dynamic environments. Adaptation of reaching movements to different mechanical environments has been studied by exposing subjects to novel force fields applied through the handle of a robotic arm (Shadmehr and Mussa-Ivaldi 1994). When subjects make reaching movements with a robot in the presence of novel dynamics, trajectories are initially perturbed and then gradually straighten with adaptation. Trajectories show mirror-image aftereffects when the field is unexpectedly removed. Similar patterns of adaptation and negative aftereffects have also been demonstrated for novel inertial loads (Sainburg et al. 1999) and novel Coriolis forces (Lackner and DiZio 1994). The aftereffects of force adaptation have been interpreted as evidence that the nervous system learns to anticipate and counteract novel forces by building a central representation, or internal model, of the limb and force-field dynamics. However, it is unclear to what degree this adapted internal model is a representation of the dynamics of the subject's own arm versus that of the novel tool (Imamizu et al. 2000, 2004; Kurtzer et al. 2005; Wolpert and Kawato 1998).

One way to test the degree to which force adaptation is tool specific is to examine how adaptation with a tool generalizes to reaches made in free space, i.e., after subjects let go of the robot. Limited generalization would suggest that the nervous system attributes the novel dynamics to the new tool, whereas broad generalization would suggest that the nervous system attributes the novel dynamics to the subject's own arm. A recent study showed that after adapting robot-held reaching movements to a force field, there was small but significant generalization to reaches made in free space (Cothros et al. 2006). In the present study, we explored whether transfer to the free condition could be enhanced by tightly controlling the contextual features of the reaching task during force adaptation.

To encourage transfer of force adaptation to free space, in all of our experiments we kept contextual cues related to hand grasp similar by asking subjects to hold the same handle in both the robot and the free-reach conditions. In *experiment* 1, the handle was rigidly attached to the robot. Participants did not have to control the orientation or height of the handle and they could partially rest their arm on the robot when it was present, whereas they had to actively maintain arm posture in the free-reaching task. In *experiment* 2, we removed this difference by requiring subjects to control the height of the hand as they reached, both with robot and without robot conditions. We predicted that this would make reaching with the robot more like reaching in free space and might therefore increase generalization to free space.

In the final experiment we attempted to increase the transfer to free space by introducing the force field so that it was less perceptible than the abrupt (i.e., full-strength) presentation. Gradual, versus abrupt, introduction of visuomotor perturbations has been shown to result in larger-magnitude aftereffects (Kagerer et al. 1997; Michel et al. 2007). Gradual presentation can also alter the coordinate system in which adaptation occurs (Malfait and Ostry 2004). Abrupt presentation of a force field results in adaptation in an extrinsic coordinate system that can transfer to other effectors, suggesting that subjects learn something about the robot (Criscamanga-Hemminger et al. 2003; Malfait and Ostry 2004). Gradual and implicit presentation of force fields results in adap-

Address for reprint requests and other correspondence: J. Kluzik, Boston University, Sargent College of Health and Rehabilitation Sciences, 635 Commonwealth Avenue, Boston, MA 02215 (E-mail: kluzikj@bu.edu).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked "*advertisement*" in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

tation in intrinsic arm coordinates that does not transfer to the other arm, suggesting that an internal model of the arm is adapted (Malfait and Ostry 2004). It may be that smaller errors induced by the robot are more easily attributed to one's own movement command versus an external influence (Diedrichsen et al. 2005). Therefore we conjectured that gradual imposition of the field might produce greater transfer to the free-reach condition.

Determining the degree to which force adaptation acquired with a robot generalizes to other reaching tasks has important implications for rehabilitation. Reaching with a robot is currently being explored as a training tool for rehabilitation in individuals who have had a stroke (Patton et al. 2001, 2006; Raasch et al. 1997; Reinkensmeyer et al. 2004). The degree of neural specificity of reaching adaptation will influence how robotic training will alter reaching movements when subjects no longer hold the robot. If the goal of training with a robotic device is to alter the control of reaching in a general context, then the degree of generalization from robot to free reaching is an important variable to determine the length and success of treatment. If, for example, the reach adaptation is completely tool specific, one may have to alter therapeutic intervention to increase the amount of generalization.

METHODS

Subjects

Thirty-eight healthy, right-handed subjects participated in this study. Sixteen subjects (11 females and 5 males; mean age 29.3 $\pm$ 4.8 yr) participated in *experiment* 1, 8 subjects (5 females and 3 males; mean age 33.8 $\pm$ 7.2 yr) participated in *experiment* 2, and 14 subjects (6 females and 8 males; mean age 27.6 $\pm$ 7.9 yr) participated in *experiment* 3. The Institutional Review Board at Johns Hopkins School of Medicine approved the experimental protocol. All subjects gave their informed consent prior to participation in experiments.

Reaching task

Subjects performed the reaching task using their right dominant arm while sitting in a firm-backed chair, with upright trunk and feet resting on the floor. Each subject was positioned with respect to the reaching targets such that when the hand was located at the center target, it was located at the body midline, the subject's shoulder was flexed about 60°, and the elbow was flexed about 120°.

Subjects grasped a hollow cylindrical handle and made 10-cm horizontal reaching movements, moving a 0.5-cm cursor from a 0.75 $\times$ 0.75-cm start box to a same-sized target box. The targets and the cursor showing the handle's location were projected onto a horizontal screen that was located directly in front of the seated subjects and about 2 cm above the top of the handle. The three-dimensional (3-D) position of the handle was recorded at 100 Hz with an Optotrak motion analysis system (Northern Digital [NDI], Waterloo, Ontario, Canada) and smoothed with a Gaussian filter before it was projected onto the screen.

For *experiments* 1 and 2, movements were made either to or from one of three targets located 10 cm directly (0°), 60° to the right, or 60° to the left of the straight-ahead position with respect to a center box (Fig. 1*A*). The directions of movement were chosen pseudorandomly. For *experiment* 3, only the straight-ahead and return-to-center directions of movement were included to simplify the experimental design. The target box was presented in red when it first appeared. Subjects were instructed to begin reaching when the target turned white and to then move the cursor straight into the center of the target box. Targets exploded when the cursor stopped inside of the target box within 400–500 ms of reach onset, acting as a reward for the subject. Color cues at the end of the reaching movement indicated when reaches were too slow or too fast. This feedback ensured that the movement velocity was similar across experimental conditions.

Force-field adaptation and transfer

Subjects performed the reaching task before, during, and after periods in which they adapted their reaching movements to viscous forces applied by the robot. The robotic device consisted of a custom-built 2 degree of freedom, low-friction, lightweight manipulandum with pneumatic force control (Diedrichsen et al. 2005). For the period of force adaptation, the robot was programmed to generate forces that were perpendicular to the direction of reach and that depended linearly on handle velocity.

Each subject performed the reaching task under three force conditions: *1*) the *robot-null* condition, in which subjects reached while the forces were turned off; *2*) the *robot-force* condition, in which subjects reached while the robot generated a viscous curl field; and *3*) the *free-space* condition, in which subjects held the same handle as in the robot conditions, but with the handle detached from the robot arm (Fig. 1*A*). To keep the kinematics of the reaching movement the same for the free-space condition, subjects were instructed to hold the handle at the same height as for the robot condition and were given audio feedback whenever the handle moved outside of an approximately 1.5-cm vertical window.

Experiment 1

Subjects were randomly assigned to one of two groups: Group 1 was tested for transfer of force adaptation in free-space condition first and the robot-null condition second. Group 2 was tested in the robot-null condition first (Fig. 1*B*). When tested for transfer of force adaptation to the robot-null condition, subjects were verbally informed that forces applied by the robot were now turned off.

Sessions began with a baseline, preadaptation period in which subjects reached under the robot-null condition and the free-space condition (Fig. 1*B*). Next, subjects underwent a period of adaptation, in which they reached with the robot arm in a force-field. The robot generated a clockwise (CW) curl field $F = Bv$, where $B = [0\ 9; -9\ 0]$ Ns/m and v is velocity. Catch trials, in which the force field was unexpectedly turned off, were interspersed randomly throughout the period of adaptation. Immediately after adaptation, subjects were tested for transfer. Subjects in Group 1 lifted the handle off the robot, observed as the robot was pushed out of the workspace, and then performed a block of trials in the free-space condition. For subjects in Group 2, the handle remained on the robot. The subjects were informed that no forces would be applied to the hand during the coming block of trials and then they performed a block of trials in the robot-null condition. After the first test for transfer, subjects readapted to the force field and then were tested for transfer in the remaining free or robot-null condition. The experimental session ended with reaches made under the robot-force condition. Periods of reaching in the force field that preceded and followed reaching in the free or robot-null conditions allowed us to quantify the degree to which the free and robot-null conditions washed out prior learning of the force field. To prevent fatigue during the experimental session, subjects were given brief rests between blocks of similar conditions.

Experiment 2: altered handle condition

The task and experimental procedure of *experiment* 2 were identical to that of *experiment* 1, except that participants were required to hold and control the height of the robot handle in both the free-space and robot conditions. In both experiments, subjects held a hollow, cylinder-shaped handle. In *experiment* 1, the cylinder rested directly on the robot arm, fixing the height of the handle in the vertical direction when subjects reached with the robot. During free-space reaching,

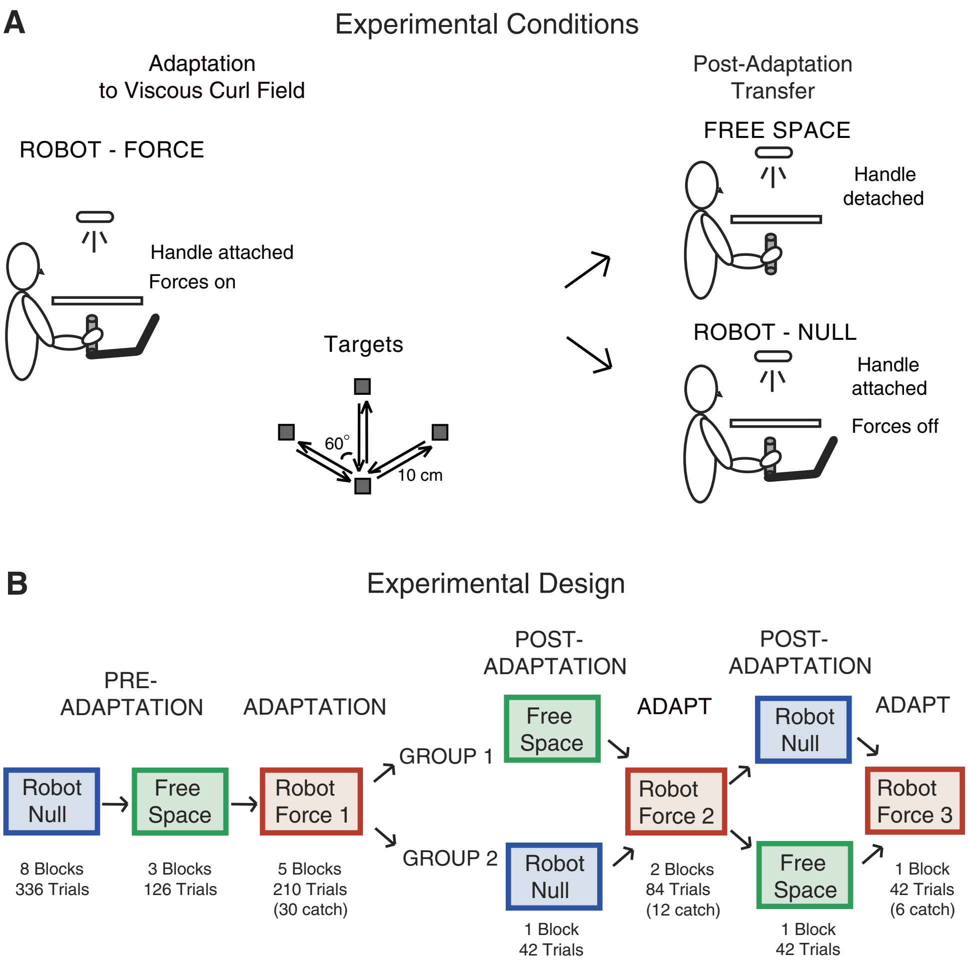

FIG. 1. Experimental protocols. *A*: in the *robot-force* condition, subjects adapted their reaching movements to a clockwise, viscous curl field. Targets were presented at a 10-cm distance in one of 6 directions from the start position. After a period of making reaches in the robot-force condition, subjects were tested for transfer of adaptation to either: *1*) the *free-space* condition, in which subjects continued to hold the handle, but the handle was completely detached from the robot arm and the robot arm was moved out of the work space; or *2*) the *robot-null* condition, in which subjects continued to make reaching movements with the robot arm, but with force field turned off. *B*: order and number of trials in which subjects were exposed to the robot-null (blue), free-space (green), and robot-force (red) conditions.

handle motion was unrestricted and subjects had to control the handle height (Fig. 1*A*). In *experiment* 2, we altered the handle such that it could slide up or down and could no longer rest on the robot arm (Fig. 3*A*). This equated the postural demands between the adaptation phase and reaching in free space.

Experiment 3: gradual versus abrupt exposure

In *experiment* 3, we investigated how transfer of adaptation to free space was affected by exposing subjects to the field gradually rather than abruptly. Two groups of subjects performed the experiment in a design that was similar to *experiment* 1 (Fig. 1*B*), with the following exceptions: *1*) targets were located in one of two locations with respect to the hand, 10 cm toward or away from the subject; *2*) a CCW curl field during the adaptation period was increased gradually for one group of subjects and abruptly for the other group; *3*) there were no catch trials for either group of subjects because catch trials would have introduced abrupt changes in force to the gradual condition; and *4*) the sequence of testing was the same, but the total number of trials per condition was fewer since there were two, not six, directions of reaching movements.

To rule out that any differences between transfer to the free and the robot-null conditions were specific to the particular set of forces in the first experiments, the robot generated a counterclockwise (CCW) curl field in *experiment* 3, where $B = [0\ -\alpha;\ \alpha\ 0]$. For the abrupt group, α was set equal to 9 throughout the adaptation period. For the gradual group, the magnitude of force was increased gradually using a method similar to that of Malfait and Ostry (2004). The value of α was gradually increased over 120 trials to a maximum of 9 according the equation $\alpha = n^k$, where n represents the nth trial and $k = \log(9)/\log(120)$. The maximal force-field strength was achieved in trial 120 and the adaptation period continued with an additional 60 trials at maximal field strength.

The trial sequence began with baseline robot-null and free periods (60 trials each), followed by a period of adaptation to the viscous field (180 trials). Immediately after adaptation, both groups of subjects were tested for transfer to the free condition (30 trials). Subjects then readapted to the viscous field with the forces at the full strength (60 trials), followed by a period of reaching in the robot-null condition (30 trials).

Data collection

To determine the position and velocity of the handle and of the subject's right arm and hand during reaching movements, we used an Optotrak motion analysis system (NDI) to record the 3-D position of infrared emitting diodes at 100 Hz. Markers were placed at the bottom of the handle and on the subject's right fifth metacarpophalangeal joint of the hand, the styloid process of the wrist, the elbow, and the acromion process of the shoulder. Kinematic data were used to characterize performance of reaching movements and to compare performance in the free-space and robot conditions.

Data analysis

MAGNITUDE OF TRANSFER. We quantified the magnitude of force adaptation and transfer based on the initial direction errors of the reach trajectories, which were defined as the magnitude of perpendicular displacement from a straight trajectory from the starting position to the target at 300 ms after the onset of reaching. The onset and the cessation of each reaching movement were determined based on a tangential velocity threshold of 2 cm/s that was maintained for a duration of ≥200 ms. An increase in the magnitude of adaptation was associated with a reduction in CW errors (CCW for *experiment* 3) during fielded trials and an increase in CCW errors during catch trials, in which forces were unexpectedly turned off.

For *experiments* 1 and 2, transfer of force adaptation to the free and robot-null conditions was quantified for each subject as the ratio of the postadaptation error (y_{post}) to the catch error of the late-adaptation period (y_{catch})

$$\text{Transfer Index} = y_{post}/y_{catch}$$

The term y_{post} is the average error that was made in the first six postadaptation trials. The term y_{catch} is the average error that was made in catch trials during the last two blocks of the field-adaptation period, a period in which subjects had reached an asymptote in their learning. Catch trials were matched for direction to the postadaptation trials. Preadaptation baseline errors were subtracted from the postadaptation and catch errors to correct for any individual biases. Note that the magnitude of the catch error indicates the amount of learning that occurred during adaptation. Therefore a transfer index of 1 would indicate 100% transfer. Normalizing the magnitude of error in the transfer period to the magnitude of late-adaptation catch trials accounted for individual differences in stiffness of the arm, velocity of the reaching movement, and learning of the force field.

For *experiment* 3, there were no catch trials during the period of adaptation. The absence of catch trials was essential for minimizing the subjects' awareness of the applied forces in the gradual condition.

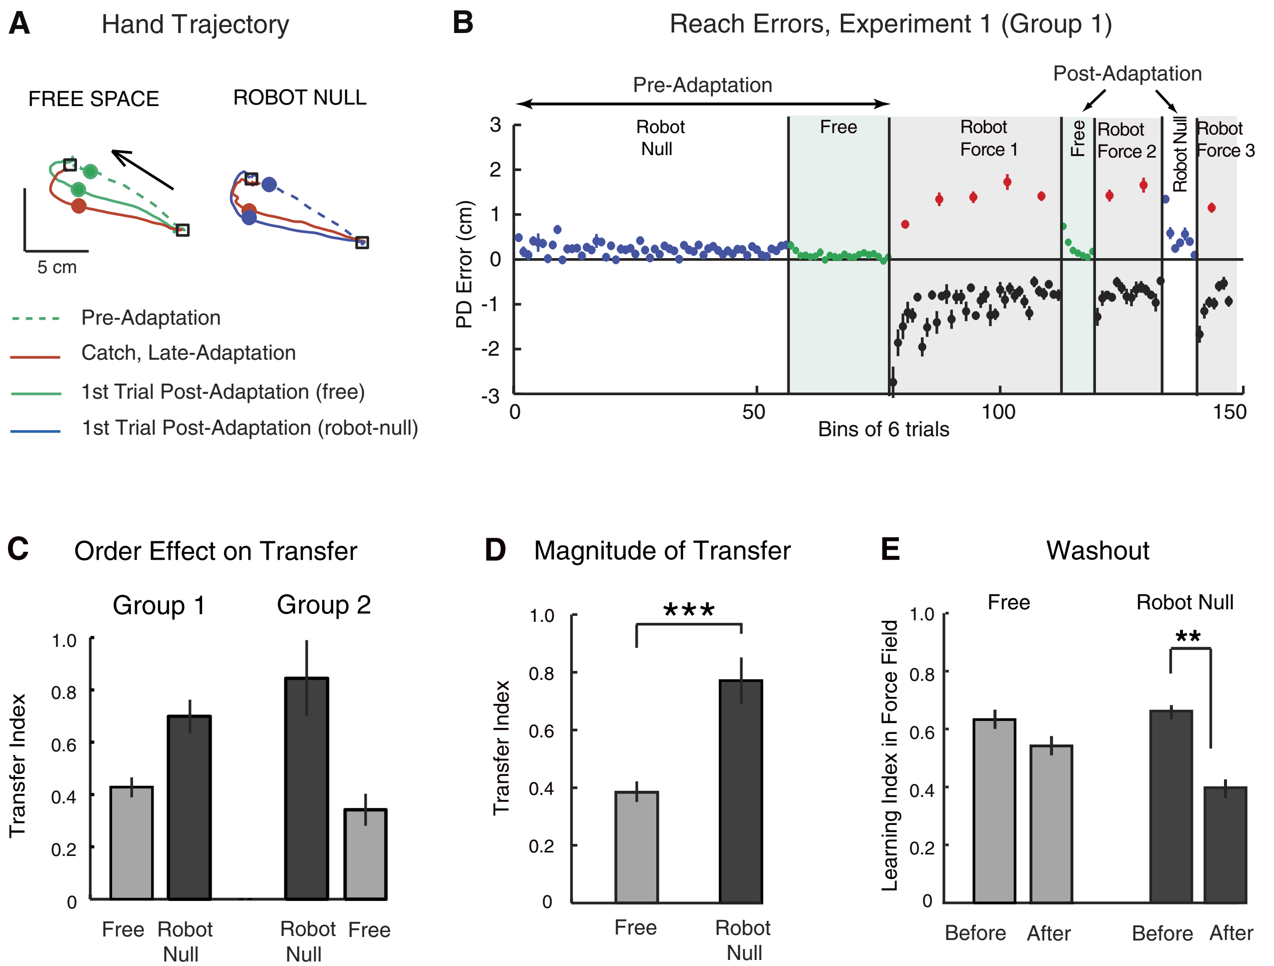

FIG. 2. Comparison of the magnitude of transfer of adaptation between the free-space and robot-null conditions. *A*: examples of a single subject's hand trajectories, plotted in the horizontal plane: before adaptation (dashed lines); late in the adaptation period in a trial in which the force field had been unexpectedly turned off (catch trial, red lines); and immediately after the adaptation period in the free condition (green line) or in the robot-null condition (blue line). The filled circles indicate the time of measurement of the perpendicular displacement (PD) error from a straight trajectory. *B*: mean (±SE) perpendicular displacement errors, binned in groups of 6 trials for subjects in Group 1 of *experiment* 1. The robot-null condition is shown in blue and the free-space condition is shown in green. During the adaptation period, trials in which the force field was turned on are shown in black; catch trials in which the force field was turned off are shown in red. Positive values indicate errors made in the counterclockwise direction. *C*: the test for transfer to free-space reaching took place after the first adaptation block for subjects in Group 1 and after the second adaptation block for Group 2. The transfer index (±SE) is the size of the aftereffect in the first 6 trials of generalization, divided by the size of the catch trials late in adaptation. The bar plots of the average (±SE) transfer index for the 2 groups show that the order of testing did not affect the relative magnitude of transfer of force adaptation to the free-space and robot-null conditions. *D*: the bar plots compare the amount of transfer of force adaptation between the free-space (gray) and robot-null (black) conditions for all subjects who participated in *experiment* 1. *E*: the bar plots compare the group average (±SE) learning index before and after subjects reached in either the free-space (gray) or robot-null (black) conditions. The learning index is the ratio of the magnitude of catch-trial errors to the magnitude of the difference between catch-trial errors and fielded-trial errors. Increasing values of the learning index indicate better compensation for and learning of the viscous force field. Washout of learning is indicated when the learning index is smaller after than before a block of reaching in the free-space or robot-null conditions.

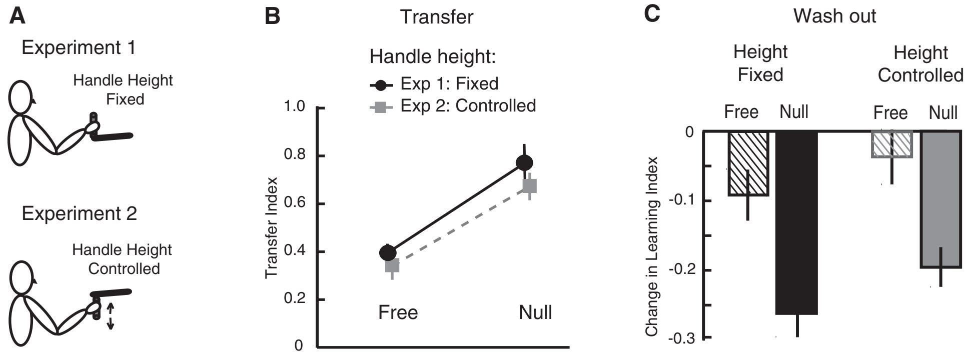

FIG. 3. *A*: in *experiment* 1, the handle height was fixed. In *experiment* 2, subjects controlled the handle height. Subjects were given audio feedback whenever the handle height moved outside of an approximately 1.5-cm window. *B*: the average (±SE) transfer of force adaptation to reaching in the free-space condition was a similar proportion of transfer to reaching in the robot-null condition, regardless of whether subjects were required to control the handle's height in *experiment* 2 (gray squares connected by dashed line) or whether the handle height was fixed in *experiment* 1 (black circles connected by solid line). *C*: the drop in the learning index (mean ± SE), indicating washout of learning, that occurred when subjects reached either in the free-space (hatched bars) or in the robot-null (solid bars) conditions was equivalent, irrespective of whether the handle's height and orientation had to be controlled. In both experiments, reaching in the robot-null condition produced greater washout of previous learning than reaching in the free condition.

Thus we compared the magnitude of postadaptation error, corrected for preadaptation baseline errors, to determine whether the gradual and abrupt groups differed in the amount of transfer of force adaptation.

LEARNING INDEX AND WASHOUT. We compared how much reaching in the free or robot-null conditions washed out retention of prior force-field adaptation. For *experiments* 1 and 2, we examined washout by calculating a learning index to quantify how well each subject compensated the force field during the periods of force adaptation (Smith and Shadmehr 2005)

$$\text{Learning Index} = y_{\text{catch}}/|y_{\text{catch}} - y_{\text{field}}|$$

The terms y_{catch} and y_{field} represent errors made during catch trials and fielded trials, respectively. When subjects have learned to perfectly compensate the force field, errors should be small in fielded trials and large in catch trials. Thus a learning index near 1 reflects nearly perfect compensation and smaller values reflect poorer compensation of the force field. The learning index was computed for sets of 12 reaches that were performed under the robot-force condition and that took place immediately before and immediately after the subject performed a single block of reaches under the free or robot-null conditions. The before and after sets of reaches each included nine fielded and three catch trials and were matched for reach direction.

Since gradual exposure to the viscous field precluded the use of catch trials during adaptation in *experiment* 3, we could not calculate learning indices to compare differences in washout of field learning between the gradual and abrupt groups. Instead, we analyzed washout by comparing the magnitude of reach error in the force field immediately before (mean of last 12 reaches) and immediately after (mean of first 6 reaches) the subjects had performed 30 trials in the free condition.

RESULTS

Experiment 1

Figure 2*A* shows a single subject example of reach paths before adaptation, during late catch trials, and immediately after adaptation. Figure 2*B* shows the averaged reach errors for Group 1. In early adaptation, reach trajectories in field trials deviated in a CW direction (negative numbers in Fig. 2*B*). With increased training, reaches gradually returned toward a straight trajectory. Toward the end of the adaptation period, when the force field was unexpectedly turned off in random catch trials, the trajectory showed an aftereffect, with deviation in the CCW direction (red dots in Fig. 2*B*).

For Group 1, the initial adaptation period was followed by a transfer test to reaching in free air. We observed significant aftereffects when we compared the pre- and postadaptation errors in the free condition ($P < 0.001$, Scheffé's post hoc repeated-measures ANOVA, time × condition). After the free condition, subjects held the robot again and readapted to the viscous forces. Following readaptation, subjects were told that the robot forces would now be turned off and were then tested for transfer to the robot-null condition. Despite the verbal instruction, the subjects still exhibited large aftereffects that were similar in magnitude to surprise catch trials during adaptation when forces were unexpectedly turned off. The transfer index for the first two postadaptation trials was near 1 (1.02 ± 0.13). Since the transfer index is the error of the initial, instructed, postadaptation trials divided by the error of late-adaptation, no-instruction catch trials, the magnitude of near 1 suggests complete transfer with minimal effects of the verbal instruction.

Groups 1 and 2 differed in the order of presentation of the postadaptation blocks of free-space and robot-null trials, although the order did not affect the results and the two groups showed similar transfer to the free and null conditions [Fig. 2*C*; $F_{(1,14)} = 0.098$, $P = 0.76$, main effect for group in a two-way ANOVA, condition × group]. Therefore the results that follow are for both groups combined.

There was significant generalization of adaptation to both the free and the robot-null conditions, as measured by an increase in CCW error from before to after adaptation. For the free condition, errors increased from 0.031 ± 0.031 to 0.635 ± 0.051 cm ($P < 0.0001$) and for the robot-null condition from 0.284 ± 0.027 to 1.505 ± 0.100 cm ($P < 0.0001$), when averaging error for the first six postadaptation trials. When the magnitude of postadaptation error was normalized to the magnitude of late-adaptation catch trials, as shown in Fig. 2*D*, transfer to the free-space condition (0.38 ± 0.04) was significantly smaller than transfer to the robot-null condition (0.77 ± 0.08) [$F_{(1,14)} = 25.35$, $P < 0.001$, main effect for condition in a two-way ANOVA, condition × group]. Therefore robot training had large effects on reaching when subjects moved the robot and were aware that the robot no longer produced forces (robot-null). The transfer was smaller, but still significant, when subjects reached in free space after having watched and felt the handle being physically disconnected from the robot.

It is possible that our transfer index, which was normalized to the amplitude of catch trial errors, could artificially over- or underinflate the extent of transfer in the case of very small or large extent of learning. However, similar results were found when the transfer index was calculated in a way that takes learning into account, normalizing postadaptation error to the absolute difference between errors in the catch and fielded trials. Transfer to free space (0.212 ± 0.020) was significantly smaller than transfer to the robot-null condition (0.419 ± 0.029; $t = -6.015$, $P < 0.0001$).

One possible explanation for the poorer transfer to free space is that reach kinematics may have differed between the free and robot conditions. The larger the change in arm posture, the smaller the transfer of adaptation (Shadmehr and Moussavi 2000). We anticipated this issue and kept many kinematic features of the task constant between the robot and the free-space conditions. We actively monitored arm posture and hand height to minimize differences between conditions. For example, when subjects reached in the free condition, there was only a small within-trial variation in handle height (0.82 ± 0.18 SD cm across subjects). Subjects also kept the handle height the same when reaching in the free condition as when reaching in the robot condition. The difference in subjects' wrist heights between the robot and the free conditions averaged 0.62 ± 1.4 cm SD and ranged from 0.78 cm lower to 3.14 cm higher in the free condition across all subjects. The peak velocity of reaching movements averaged between 45 and 55 cm/s for each of the force conditions, with mean ± SE peak velocities of 51.5 ± 2.0, 45.1 ± 1.6, and 49.8 ± 2.1 cm/s for the last block of reaches in the robot-force, the free, and the robot-null conditions, respectively. Furthermore, the magnitude of transfer to the free-space condition for each subject was poorly predicted by how much the kinematics changed when the subject transitioned from the robot to the free condition ($R^2 = 0.014$ for differences in hand height; $R^2 = 0.189$ for differences in peak velocity). Therefore, because arm posture and reach kinematics were kept similar between the robot-null and free conditions, these factors were probably not responsible for the reduced transfer in the free condition.

Reaching in the robot-null condition, but not the free-space condition, washed out a significant amount of prior force-field learning. Figure 2*B* shows comparable catch trial errors in the periods of force adaptation that immediately preceded (Robot Force 1) and followed (Robot Force 2) the block of reaches that were made in free space. In contrast, catch trial errors were reduced in the period of force adaptation that followed the block of trials made in the robot-null condition. Furthermore, errors made in field trials increased to a larger extent after subjects reached in the robot-null condition than after they reached in free space. We calculated a learning index of how well each subject compensated the force field during the periods immediately before and after the free and the null blocks of washout trials. The learning index takes into account the magnitude of errors in both catch and field trials. Figure 2*E* shows that reaching in the robot-null condition, but not in the free-space condition, caused a significant decrease in the learning index (Scheffé's post hoc analysis following repeated-measures ANOVA, $P < 0.01$).

Experiment 2

Despite our efforts to keep posture the same in the free-space and robot conditions of *experiment* 1, there was one significant difference between the two conditions: in the robot condition, the height of the hand was determined by the height of the handle attached to the robot and some of the weight of the arm could be supported by the robot. In the free-space condition, this task was left to the subject. We wondered whether this difference in contextual cues played a significant role in the amounts of transfer.

In *experiment* 2, subjects were required to control hand orientation and support the weight of their arm in all conditions, as shown in Fig. 3*A*. We found that this change in postural demands did not alter the results from *experiment* 1. Figure 3*B* shows that transfer of adaptation to reaching in free space was comparable in the two experiments (0.38 ± 0.04 for *experiment* 1 and 0.34 ± 0.06 for *experiment* 2). The replication of the results suggests that the transfer is a robust phenomenon. Regardless of handle condition, the transfer index for the free-space condition was significantly smaller than the transfer index for the robot-null condition (0.77 ± 0.08 for *experiment* 1 and 0.64 ± 0.06 for *experiment* 2), with an ANOVA main effect of $F_{(1,22)} = 31.97$ and $P < 0.0001$. The ANOVA showed no interaction effect between force conditions and handle conditions and no significant difference between the transfer indices of *experiments* 1 and 2.

Requiring subjects to control handle height in *experiment* 2 also had no effect on the pattern of washout of force-field adaptation. For both *experiments* 1 and 2, reaching in free space only minimally washed out prior learning of the field, but reaching with the robot when forces were off (robot-null) substantially and significantly washed out with prior learning. Figure 3*C* shows changes in the learning index from before to immediately after a block of trials of reaching under either the free-space or the robot-null condition. In both experiments, there was significantly greater washout after a period of reaching in the robot-null condition versus that in free space [repeated-measures ANOVA, $F_{(1,22)} = 18.60$, $P < 0.001$]. ANOVA of the change in learning index showed no significant interaction between force and handle conditions and no significant difference between the "height-fixed" or "height-controlled" handle conditions.

Experiment 3: gradual versus abrupt exposure to the force field

In the adaptation period of *experiment* 3, the viscous force increased gradually over many trials for one group of subjects and increased abruptly to the maximum force for the other group (Fig. 4*A*). The errors in the gradual group during adaptation were small (0.3 to 0.5 cm CCW, in the direction of the field) and stayed at a nearly constant value throughout the course of adaptation. The gradual group demonstrated generally larger transfer to free space than did the abrupt group, with aftereffects that were 59.0% of the amplitude of the aftereffect in the robot-null condition, compared with 41.0% for that of the abrupt group (Fig. 4*A*). However, a closer examination of

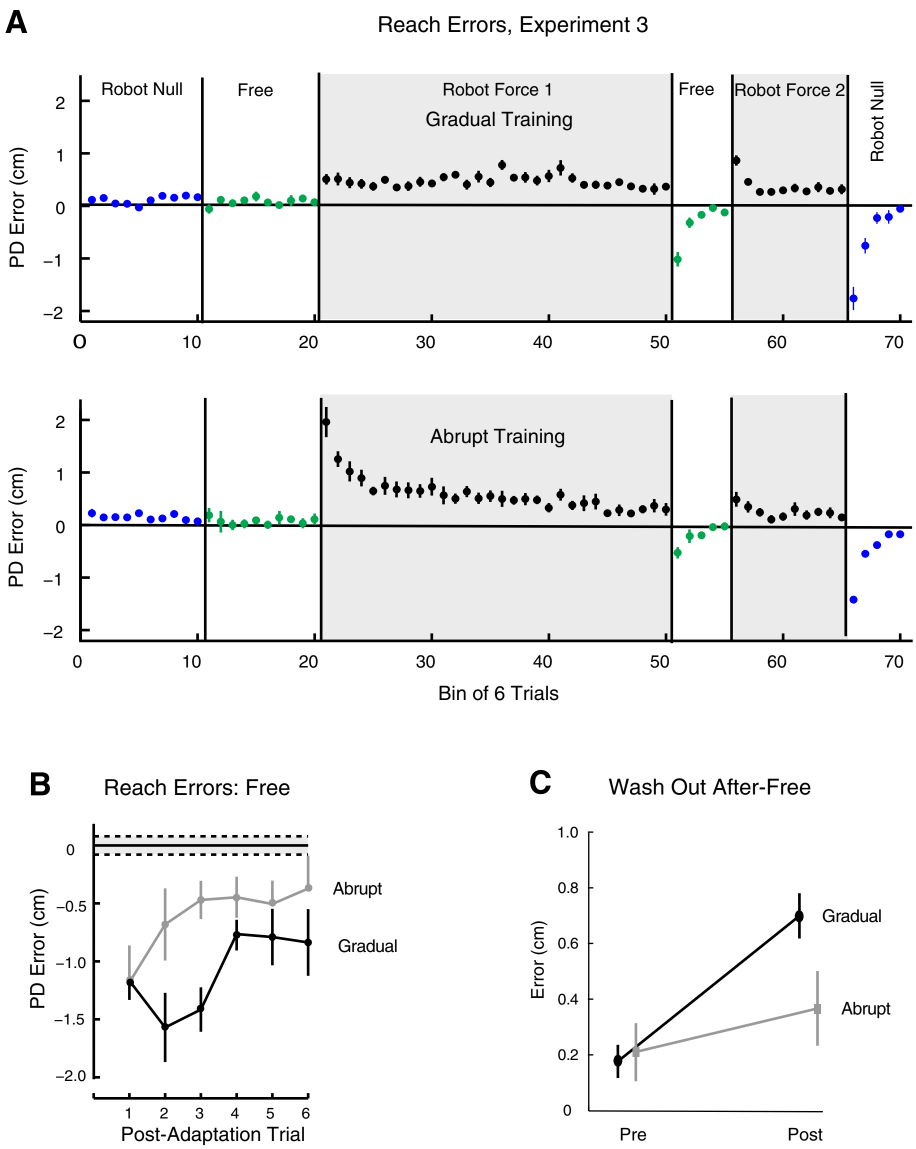

FIG. 4. Results of *experiment* 3. *A*: average error in the gradual group (*top plot*) in comparison to the abrupt group (*bottom plot*). *B*: comparison of transfer to the free condition between the gradual (black line) and abrupt (gray line) groups for the first 6 postadaptation reaches. The shaded area indicates reaching performance in free space prior to force adaptation (mean ± SD). *C*: average error immediately before and after the subjects performed a block of reaching in the free condition. The extent of washout is indicated by how much error increased from the before to the after periods. Reaching in free space washed out prior force adaptation to a greater extent for the gradual group (black) than for the abrupt group (gray).

the reach errors during the free condition suggests that the crucial difference between the two groups was the rate of washout of the aftereffects (Fig. 4*B*). During the free-reach block, the gradual and abrupt groups started out with nearly equal sized aftereffects ($t = -0.349$, $P = 0.733$), but these errors washed out more rapidly in the abrupt group than in the gradual group [$F_{(1,12)} = 7.290$, $P < 0.05$ group main effect for a two-way ANOVA, group × trial]. The larger transfer to the free condition by the gradual group could not be explained by differences in reach velocity. The mean peak velocity of the first six transfer trials was 48.58 ± 2.97 cm/s for the gradual group and 44.78 ± 1.13 cm/s for the abrupt group ($t = 1.06$, $P = 0.310$).

The smaller transfer to the free condition after abrupt learning suggests that an abrupt presentation of a perturbation encouraged adaptation of a model of the tool, rather than the arm. If this is the case, then washout in the free condition, a condition in which the tool is not present, should produce a smaller deadaptation on the prior learning in the abrupt case than that in the gradual case. Indeed, we found that although free reaching caused significant washout of the prior adaptation in the gradual group (post-free vs. pre-free performance with the robot, $P < 0.001$, Fig. 4*C*), the same washout period had no significant effect on the prior adaptation in the abrupt group [$P = 0.403$, Scheffé's post hoc analyses following two-way ANOVA, group × time period, with interaction $F_{(1,12)} = 9.997$, $P < 0.01$]. Thus gradual training more effectively adapted a general representation of the arm that was utilized in both free-space and robot reaching.

After relearning of the field with the robot (Robot Force 2, Fig. 4*A*), the field was turned off as the subjects continued to reach with the robot. The gradual and abrupt groups did not differ significantly from one another in the magnitude of transfer to the robot-null condition [$F_{(1,12)} = 2.47$, $P = 0.14$, group main effect for a two-way ANOVA, group × trial].

DISCUSSION

We have shown that when limb kinematics are tightly controlled, adaptation to novel dynamics associated with a handheld tool partially generalizes to free-space reaching, and this generalization can be increased if changes in the tool's dynamics are introduced gradually rather than abruptly (Fig. 4*B*). Our results suggest that the gradual condition produces greater adaptation in a general internal model of the arm than does the abrupt condition.

Partial generalization of force adaptation to free reaches

When a new tool such as a robot imposes novel forces on the limb, adaptation generalizes to different movements of the limb with the tool (Conditt et al. 1997). The learner appears to assume that the forces depend on the state of the limb (i.e., its position, velocity, etc.) as represented in the intrinsic coordinates of joints and muscles, rather than state of the hand as represented in an extrinsic coordinate system (Malfait et al. 2002; Sainburg et al. 1999; Shadmehr and Moussavi 2000; Shadmehr and Mussa-Ivaldi 1994). Indeed, models that encode limb state with the sensitivities found in muscle afferent neurons reproduce these patterns of generalization (Hwang and Shadmehr 2005), even though in many cases such coding is not optimal for control of certain tools (Hwang et al. 2006a). This pattern of generalization suggests that in learning to reach with a new tool, the brain does not begin with a blank slate, but rather with a prior model that is highly structured. Because we observed that there was robust generalization from robot to free reaching in the different conditions of all three of our experiments, the neural representation of this tool-specific model probably partially overlaps with a general model of the arm.

The concept of multiple internal models, or representations, of limb and tool dynamics that is suggested by our results is not new. For example, the MOSAIC model (modular selection and identification for control) for learning movements proposes that our ability to flexibly adapt to changing and novel contexts depends on a modular organization that consists of multiple internal models that can be optimally blended to match the demands of a given task and environment (Haruno et al. 2001; Wolpert and Kawato 1998). Imaging data also suggest a modular organization of arm and tool representations that are learned and modified based on experience. Learning to make movements with different tools involves distinct, but partially overlapping, parts of the lateral cerebellum (Imamizu et al. 2003, 2007). Recent work suggests that the prefrontal and parietal areas may be involved in selecting and blending modules so that they are optimally matched for a given context (Imamizu et al. 2004, 2007). Our observations of partial transfer of force adaptation from the robot to the free condition are consistent with the concept of modular organization of internal models.

Subjects had aftereffects of force adaptation in both the robot-null and free conditions, even though they had explicit knowledge that no external force would be applied to the reaching hand. This suggests that force adaptation cannot easily be switched on or off by cognitive cues or knowledge. The aftereffects in the robot-null condition were large and decayed only gradually across many trials. Thus the prior sensorimotor experience with the context of the robot had a stronger influence on reaching behavior than the cognitive instructional cue that forces were turned off. This is in agreement with other studies showing that subjects do not easily learn to distinguish between two different force conditions for multijoint reaching movements if the force conditions are signaled only by purely cognitive, symbolic cues, such as color (Gandolfo et al. 1996; Osu et al. 2004; Shadmehr et al. 2005), a systematic alternating pattern of two different force fields (Karniel and Mussa-Ivaldi 2002; Wainscott et al. 2005), or a postural cue unrelated to task performance, such as thumb position (Gandolfo et al. 1996).

Which contextual cues limit generalization to free-space reaching in abrupt conditions?

In the abrupt condition, the degree of generalization to free space was incomplete, about 40%, and free space reaches did not washout prior learning with the robot. Thus subjects were able to rapidly alter their reaching strategy when the context changed. Which cues were responsible for this switching? Because we tightly controlled the kinematic and postural features of the reaches across the various conditions, these factors played no role in the differing amounts of generalization to robot-null and free-space conditions.

One potential cue was cognitive awareness that the handle was no longer attached to the robot. However, the awareness that all forces were turned off in the robot-null condition did not result in a similar immediate switch in motor behavior when transitioning from the force to the robot-null condition. Unfortunately, the implicit cues that allowed for this switching of internal models are still poorly understood. It is important that free-space reaching did not substantially wash out the prior force adaptation in any of our experimental conditions except when in the gradual condition. This suggests that in the abrupt condition, subjects largely acquired a tool-specific model that was linked to context of the robot (Cothros et al. 2006; Imamizu et al. 2000, 2004; Wolpert and Kawato 1998). This is consistent with the finding that reach adaptation with a new tool results in tool-specific aftereffects days and months after initial training (Shadmehr et al. 1997, 1998).

Gradual perturbations may change error assignment

Our data demonstrate that what was learned in the gradual force condition differed from the abrupt condition, since the former not only resulted in a larger transfer of force adaptation from the robot (tool) to the free condition, but reaching in the free condition also washed out the prior learning with the robot. This result could be related to how the nervous system assigns "blame" to the potential causes of error: an incorrect internal model of the arm or the tool. In the gradual condition, the nervous system may have attributed the errors to the representation of the arm's dynamics to a greater extent than when the context changed abruptly in one large step. In contrast, in the abrupt condition, a greater proportion of error may have been assigned to the representation of the dynamics of the tool (i.e., robot). This is consistent with work showing that when a force field is imposed gradually, training with the right arm generalizes to movements made with the right arm in different shoulder configurations, but does not generalize to the left arm (Malfait and Ostry 2004), i.e., subjects update the internal model of their right arm. However, abrupt presentation of a force field results in some transfer from the right to the left arm (Criscimagna-Hemminger et al. 2003), i.e., subjects learn an internal model of the tool.

Several differences between the gradual and the abrupt training conditions could account for why there was greater generalization from the robot to the free condition in the gradual case. When forces increased abruptly and in a single step, subjects made large errors in reaching. In contrast, in the gradual case, the force changes from one trial to the next were small and resulted in small performance errors. The brain may use the error's size to resolve how to assign "blame" for the error, assigning larger errors to an internal model of the tool. Such an error-dependent credit-assignment scheme was recently used to explain mechanisms of saccade adaptation (Chen-Harris et al. 2008). A second possibility for explaining different generalization patterns is that an abrupt step change in force produces cognitive awareness that is not present in the gradual condition. Previous work on reach adaptation suggests that the presence of cognitive awareness can improve performance (Hwang et al. 2006b) and result in a generalization pattern that has an extrinsic coordinate system (Malfait and Ostry 2004). This would be consistent with an internal model of the tool. Finally, a step change is more consistent with a change in the tool rather than a change in the body, and the nervous system may have the capacity to detect these time-varying properties to distinguish between body and world disturbances (Kording et al. 2007).

Implications for rehabilitation

The most important contribution of our study may be for rehabilitation, in that we found that the magnitude of errors experienced during training with a novel tool affects the generalization of that training to other conditions. Our results suggest that gradually changing training conditions that result in smaller trial-to-trial movement errors are more likely to lead to changes in neural representations of the body's dynamics, with broader generalization of the learning across conditions.

Although we found a significant extent of transfer of force adaptation from the robot to the free condition, we kept tight control of the task goal and of the movement kinematics. We expect that there would be less generalization to reaching tasks that involve handling of different objects and that are performed with different arm postures and different task goals (Ma et al. 1999). However, the fact that there are some aftereffects from the robot to the free condition in healthy controls is encouraging for rehabilitation of patients (Patton et al. 2006; Raasch et al. 1997) because, in some cases, short periods of adaptation in patients with cortical damage produce aftereffects that are both longer-lasting and much more general (Pisella et al. 2002; Rossetti et al. 1998).

Conclusions

Adaptation of reaching movements to a viscous force field while holding a robotic arm generalized to affect reaches made without the robot in free air. When the forces were introduced abruptly, the extent of generalization to free space was small and reaches in the free condition only minimally washed out prior learning with the robot. When the forces were introduced gradually, generalization from the robot to the free condition increased. The larger generalization to the free condition in the gradual case suggests that the statistics of the environmental forces affected how the nervous system interpreted movement errors. In the gradual case, the nervous system increased the extent to which it attributed the error to a general model of the limb's own dynamics instead of exclusively to the tool's dynamics.

GRANTS

This project was supported by National Institutes of Health (NIH) Grants HD-49217 to J. Kluzik, HD-40289 to A. J. Bastian, and NS-037422 to R. Shadmehr. Development of the portable robot was funded by a Human Frontiers Science Program grant to R. Shadmehr. A. J. Bastian's facilities were supported in part by NIH National Center for Research Resources Extramural Research Facilities Construction Grant RR-15488.

REFERENCES

Chen-Harris H, Joiner WM, Ethier V, Zee DS, Shadmehr R. Adaptive control of saccades via internal feedback. *J Neurosci* 28: 2804–2813, 2008.

Conditt MA, Gandolfo F, Mussa-Ivaldi FA. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J Neurophysiol* 78: 554–560, 1997.

Cothros N, Wong JD, Gribble PL. Are there distinct neural representations of objects and limb dynamics? *Exp Brain Res* 173: 689–697, 2006.

Criscimagna-Hemminger SE, Donchin O, Gazzaniga MS, Shadmehr R. Learned dynamics of reaching movements generalize from dominant to nondominant arm. *J Neurophysiol* 89: 168–176, 2003.

Diedrichsen J, Hashambhoy Y, Rane T, Shadmehr R. Neural correlates of reach errors. *J Neurosci* 25: 9919–9931, 2005.

Gandolfo F, Mussa-Ivaldi FA, Bizzi E. Motor learning by field approximation. *Proc Natl Acad Sci USA* 93: 3843–3846, 1996.

Haruno M, Wolpert DM, Kawato M. MOSAIC model for sensorimotor learning and control. *Neural Comput* 13: 2201–2220, 2001.

Hwang EJ, Shadmehr R. Internal models of limb dynamics and the encoding of limb state. *J Neural Eng* 2: S266–S278, 2005.

Hwang EJ, Smith MA, Shadmehr R. Adaptation and generalization in acceleration-dependent force fields. *Exp Brain Res* 169: 496–506, 2006a.

Hwang EJ, Smith MA, Shadmehr R. Dissociable effects of the implicit and explicit memory systems on learning control of reaching. *Exp Brain Res* 173: 425–437, 2006b.

Imamizu H, Higuchi S, Toda A, Kawato M. Reorganization of brain activity for multiple internal models after short but intensive training. *Cortex* 43: 338–349, 2007.

Imamizu H, Kuroda T, Miyauchi S, Yoshioka T, Kawato M. Modular organization of internal models of tools in the human cerebellum. *Proc Natl Acad Sci USA* 100: 5461–5466, 2003.

Imamizu H, Kuroda T, Yoshioka T, Kawato M. Functional magnetic resonance imaging examination of two modular architectures for switching multiple internal models. *J Neurosci* 24: 1173–1181, 2004.

Imamizu H, Miyauchi S, Tamada T, Sasaki Y, Takino R, Pütz B, Yoshioka T, Kawato M. Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature* 403: 192–195, 2000.

Kagerer FA, Contreras-Vidal JL, Stelmach GE. Adaptation to gradual as compared with sudden visuo-motor distortions. *Exp Brain Res* 115: 557–561, 1997.

Karniel A, Mussa-Ivaldi FA. Does the motor control system use multiple models and context switching to cope with a variable environment? *Exp Brain Res* 143: 520–524, 2002.

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

Kurtzer I, DiZio PA, Lackner JR. Adaptation to a novel multi-force environment. *Exp Brain Res* 164: 120–132, 2005.

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

Ma H, Trombly CA, Robinson-Podolski C. The effect of context on skill acquisition and transfer. *Am J Occup Ther* 53: 138–144, 1999.

Malfait N, Ostry DJ. Is interlimb transfer of force-field adaptation a cognitive response to the sudden introduction of load? *J Neurosci* 24: 8084–8089, 2004.

Malfait N, Shiller DM, Ostry DJ. Transfer of motor learning across arm configurations. *J Neurosci* 22: 9656–9660, 2002.

Michel C, Pisella L, Prablanc C, Rode G, Rossetti Y. Enhancing visuomotor adaptation by reducing error signals: single-step (aware) versus multiple step (unaware) exposure to wedge prisms. *J Cogn Neurosci* 19: 341–350, 2007.

Osu R, Hirai S, Yoshioka T, Kawato M. Random presentation enables subjects to adapt to two opposing forces on the hand. *Nat Neurosci* 7: 111–112, 2004.

Patton JL, Mussa-Ivaldi FA, Rymer WZ. Robotic-induced improvement of movements in hemiparetics via an implicit learning technique. *Soc Neurosci Abstr* 832.16, 2001.

Patton JL, Stoykov ME, Kovic M, Mussa-Ivaldi FA. Evaluation of robotic training forces that either enhance or reduce error in chronic hemiparetic stroke survivors. *Exp Brain Res* 168: 368–383, 2006.

Pisella L, Rode G, Farnè A, Boisson D, Rossetti Y. Dissociated long lasting improvements of straight-ahead pointing and line bisection tasks in two hemineglect patients. *Neuropsychologia* 40: 327–334, 2002.

Raasch CC, Mussa-Ivaldi FA, Rymer WZ. Motor learning in reaching movements by hemiparetic subjects. *Soc Neurosci Abstr* 924.17, 1997.

Reinkensmeyer DJ, Emken JL, Cramer SC. Robotics, motor learning, and neurologic recovery. *Annu Rev Biomed Eng* 6: 497–525, 2004.

Rossetti Y, Rode G, Pisella L, Farnè A, Li L, Boisson D, Perenin MT. Prism adaptation to a rightward optical deviation rehabilitates left hemispatial neglect. *Nature* 395: 166–169, 1998.

Sainburg RL, Ghez C, Kalakanis D. Intersegmental dynamics are controlled by sequential anticipatory, error correction, and postural mechanisms. *J Neurophysiol* 81: 1045–1056, 1999.

Shadmehr R, Brandt J, Corkin S. Time-dependent motor memory processes in amnesic subjects. *J Neurophysiol* 80: 1590–1597, 1998.

Shadmehr R, Brashers-Krug T. Functional stages in the formation of human long-term motor memory. *J Neurosci* 17: 409–419, 1997.

Shadmehr R, Donchin O, Hwang EJ, Hemminger SE, Rao A. Learning dynamics of reaching. In: *Motor Cortex in Voluntary Movements: A Distributed System for Distributed Functions*, edited by Riehle A, Vaadia E. Boca Raton, FL: CRC Press, 2005, p. 297–328.

Shadmehr R, Moussavi ZMK. Spatial generalization from learning dynamics of reaching movements. *J Neurosci* 20: 7807–7815, 2000.

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

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

Wainscott SK, Donchin O, Shadmehr R. Internal models and contextual cues: Encoding serial order and direction of movement. *J Neurosci* 93: 786–800, 2005.

Wolpert DM, Kawato M. Multiple paired forward and inverse models for motor control. *Neural Networks* 11: 1317–1329, 1998.