Modifiability of Generalization in Dynamics Learning

Andrew A. G. Mattar[1] and David J. Ostry[1,2]

[1]*Department of Psychology, McGill University, Montreal, Quebec, Canada; and* [2]*Haskins Laboratories, New Haven, Connecticut*

Mattar AA, Ostry DJ. Modifiability of generalization in dynamics learning. *J Neurophysiol* 98: 3321–3329, 2007. First published October 10, 2007; doi:10.1152/jn.00576.2007. Studies on plasticity in motor function have shown that motor learning generalizes, such that movements in novel situations are affected by previous training. It has been shown that the pattern of generalization for visuomotor rotation learning changes when training movements are made to a wide distribution of directions. Here we have found that for dynamics learning, the shape of the generalization gradient is not similarly modifiable by the extent of training within the workspace. Subjects learned to control a robotic device during training and we measured how subsequent movements in a reference direction were affected. Our results show that as the angular separation between training and test directions increased, the extent of generalization was reduced. When training involved multiple targets throughout the workspace, the extent of generalization was no greater than following training to the nearest target alone. Thus a wide range of experience compensating for a dynamics perturbation provided no greater benefit than localized training. Instead, generalization was complete when training involved targets that bounded the reference direction. This suggests that broad generalization of dynamics learning to movements in novel directions depends on interpolation between instances of localized learning.

INTRODUCTION

Generalization in motor function is reflected in the extent to which movements in novel situations are affected by previous experience. Generalization has been described between movements that differ in terms of speed (Goodbody and Wolpert 1998), amplitude (Goodbody and Wolpert 1998; Krakauer et al. 2000), direction (Bedford 1993; Donchin et al. 2003; Gandolfo et al. 1996; Ghilardi et al. 1995; Huang and Shadmehr 2007; Krakauer et al. 2000; Thoroughman and Shadmehr 2000; Thoroughman and Taylor 2005; Vetter et al. 1999), path (Conditt et al. 1997), workspace location (Hwang et al. 2003; Malfait et al. 2002; Shadmehr and Moussavi 2000), or the effector used (Criscimagna-Hemminger et al. 2003; Dizio and Lackner 1995; Krakauer et al. 2006; Malfait and Ostry 2004; Wang and Sainburg 2004a,b; Witney and Wolpert 2003). In the present paper we have focused on the extent to which the pattern of generalization for dynamics learning is modifiable.

Generalization of motor learning has been documented in studies involving the alteration of visual feedback during movement (Bedford 1993; Caithness et al. 2004; Ghahramani and Wolpert 1997; Ghahramani et al. 1996; Ghilardi et al. 1995; Krakauer et al. 1999, 2000, 2006; Tong et al. 2002; Vetter et al. 1999; Wang and Sainburg 2004a). Previous experiments have shown that this so-called visuomotor learning can generalize broadly across the workspace under certain conditions (Bedford 1993; Ghilardi et al. 1995; Krakauer et al. 2000; Vetter et al. 1999). For example, in a study in which the visually perceived extent of movements was scaled relative to their actual extent (i.e., a visuomotor gain perturbation), changes to movement amplitude generalized fully to movements in different directions and distances from a start location (Krakauer et al. 2000). Adaptations following other visuomotor perturbations show a more limited pattern of generalization. Changes in trajectory that compensate for discrepancies between actual and perceived movement direction (i.e., a visuomotor rotation) generalized to movements of different amplitudes in the training direction but showed less generalization to movements in other directions (Krakauer et al. 2000). Interestingly, the pattern of generalization for visuomotor rotation learning was modifiable. Specifically, the extent of generalization was sensitive to the distribution of directions in which the visuomotor rotation was encountered. As the training directions sampled larger amounts of the workspace, the extent of generalization increased (Krakauer et al. 2000). These findings suggest that visuomotor rotation learning is locally tuned to the training direction (also see Ghahramani and Wolpert 1997; Ghahramani et al. 1996) and that the breadth of this tuning is modifiable with experience.

Other studies have explored generalization of motor learning in response to the application of unexpected forces to the hand as subjects make reaching movements to targets (Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994). These novel dynamics cause errors in trajectory that are rapidly eliminated as patterns of muscle activity and their underlying control signals change (Gribble and Ostry 2000; Thoroughman and Shadmehr 1999). As is the case for visuomotor rotations, these newly learned dynamics generalize to movements in the training direction that differ in terms of speed or amplitude (Goodbody and Wolpert 1998), but generalization is less for movements in other directions. Instead, generalization of dynamics learning is tuned such that training affects movements in nearby directions more greatly than movements in distant directions (Donchin et al. 2003; Gandolfo et al. 1996; Huang and Shadmehr 2007; Thoroughman and Shadmehr 2000; Thoroughman and Taylor 2005). In the present study, we have tested the idea that as in the case of visuomotor rotations (Krakauer et al. 2000), the generalization of dynamics learning broadens as the distribution of training directions covers increasingly large amounts of the workspace.

Our subjects learned to compensate for forces applied to the hand during movements to one, two, or multiple targets. Generalization of dynamics learning was then tested in a reference direction. We found that the pattern of generaliza-

---

Address for reprint requests and other correspondence: D. J. Ostry, Department of Psychology, McGill University, 1205 Dr. Penfield Avenue, Montreal, Quebec, Canada H3A 1B1 (E-mail: ostry@motion.psych.mcgill.ca).

tion was not sensitive to the extent of the workspace explored during training. The magnitude of generalization was no different when subjects trained on a single target or on multiple targets throughout the workspace; however, the pattern of generalization was sensitive to the specific location of training targets. We found that the extent of generalization increased when training provided for the possibility for interpolation between instances of learning.

METHODS

Subjects and apparatus

In all, 160 right-handed subjects (114 females, overall mean age 22.05 ± 3.80 yr) made horizontal reaching movements while holding the handle of a two-joint robotic device (InMotion2, Interactive Motion Technologies, Cambridge, MA). Sixteen-bit optical encoders (Gurley Precision Instruments, Troy, NY) sensed the position of the robot at 400 Hz. The position signal was low-pass Butterworth filtered at 20 Hz and numerically differentiated to compute hand velocity. The robot was programmed to deliver forces to the hand during movement through torque motors connected to the shoulder and elbow joints of the robot (see following text).

Procedure

Subjects made center-out reaching movements to targets (radius 1.5 cm) arranged around a circle (radius 15 cm) and separated by 45°. The center of this circle was defined by shoulder and elbow angles of 45 and 90° relative to the frontal plane and upper arm, respectively. On each trial, subjects were required to rest in the central start position for 1,200 ± 300 ms until a target was illuminated. Subjects were then required to move to the target within 350 ± 50 ms (indicated by auditory feedback) and stay within its boundaries for an additional 750 ms. The robot returned the hand to the start position before the next trial.

The logic of the experimental design was to test how performance in a reference direction was affected by training in directions (or combinations of directions) at various angular distances from the reference direction. The experimental session was divided into three consecutive phases. In the baseline phase, subjects made 25 movements in the reference direction. During the baseline phase the robot did not apply forces to the hand (a null force field). Next, during the training phase, subjects were assigned to a group depending on the target(s) to which they made training movements. They began the training phase by making 10 pretraining movements in a null field to each training target. They then made 150 movements to each training target in a clockwise force field. Target order was randomized when training involved more than one target. For four of the groups (the multitarget and full-interpolation groups; see Table 1 and Fig. 1) training was limited to 50 movements in each

direction to prevent fatigue. Note that our statistical analysis revealed that there were no differences in movement curvature at the end of the training phase [$F_{(6,153)} = 1.90$, $P > 0.05$ for movements in the force field; $F_{(6,151)} = 1.61$, $P > 0.05$ for normalized catch trials; see details in following text], which suggests that the extent of learning did not differ between short and long training conditions (50 vs. 150 movements). Moreover, in a control study involving 32 new subjects who made only 50 training movements per target, we found that subjects adapted to the force field as well as subjects who made 150 movements in these same directions. Specifically, movement curvature was no different after 50 or 150 movements per target for final movements in the force field [$F_{(1,56)} = 0.27$, $P > 0.05$] and for the final normalized catch trial [$F_{(1,56)} = 1.00$, $P > 0.05$].

The clockwise force field applied velocity-dependent loads according to

$$\begin{bmatrix} f_x \\ f_y \end{bmatrix} = \begin{bmatrix} 0 & 15 \\ -15 & 0 \end{bmatrix} \begin{bmatrix} v_x \\ v_y \end{bmatrix}$$

where x and y are the lateral and sagittal directions, v_x and v_y are movement velocities, and f_x and f_y are the forces (in Newtons) applied to the hand. Catch trials in the training direction, in which the clockwise force field was substituted with a null field, were randomly introduced within the training phase with a frequency of one catch trial in every 10 trials.

Immediately after the training phase, subjects made movements in the test phase of the experiment. Subjects made 25 movements to the reference target (the same target as in the baseline phase) in a null field. We measured the curvature of movements made in the test phase to determine the extent to which movements in the reference direction were affected by previous training in other directions.

Experimental conditions

Twenty groups of eight subjects each were tested in this experiment. The conditions to which subjects were assigned are given in Table 1 and presented graphically in Fig. 1. Performance was evaluated in one of two reference directions, located at 135° for half of the subjects and at 315° for the other half. To determine how movements in the reference direction were affected by training in other directions, we tested performance after training to single or multiple targets (Fig. 1). Subjects in the single-target condition made training movements to targets located at one of −90, −45, 0, +45, or +90° relative to the reference target. Subjects in the two-target condition made training movements to targets located at both −90 and −45° or both +45 and +90° relative to the reference target. The rationale was to test whether training to two targets resulted in greater generalization of learning than training with single targets alone. Subjects in the interpolation condition made training movements to targets located both −45 and +45° from the reference target. Here we tested the extent to which interpolation between training that bounded the reference target affected subsequent movements. In two further conditions, we tested

TABLE 1. *Experimental conditions*

|	Actual Directions		Relative Directions
Reference direction	135°	315°	0°
Single-target condition	45°	225°	−90°
	90°	270°	−45°
	135°	315°	0°
	180°	360°	+45°
	225°	45°	+90°
Two-target condition	45 and 90°	225 and 270°	−90 and −45°
	180 and 225°	360 and 45°	+45 and +90°
Multitarget condition	All but 90 and 135°	All but 270 and 315°	All but −45 and 0°
Interpolation condition	90 and 180°	270 and 360°	−45 and +45°
Full-interpolation condition	All but 135°	All but 315°	All but 0°

FIG. 1. Experimental conditions. Subjects in the single-target condition made movements to targets located −90, −45, 0, +45, or +90° relative to the reference target. To determine whether generalization of learning increased after more extensive training, subjects in the 2-target condition made training movements to targets located both −90 and −45° or both +45 and +90° relative to the reference target. In an extension of the 2-target condition, subjects in the multitarget condition made training movements to all targets but the reference target and the adjacent target in the clockwise direction. Dependence of generalization on interpolation between instances of previous training was assessed in the interpolation and full-interpolation conditions, in which subjects made training movements to targets that bounded the reference target or to all targets but the reference target, respectively. Although this figure depicts conditions relative to a reference target located at 315°, the data shown in Figs. 2, 3, and 4 and the findings presented in the RESULTS section are pooled over a second set of conditions in which the experiment was repeated relative to a reference target located at 135°.

variants of the two-target and interpolation conditions. In the extension of the two-target condition (the multitarget condition), subjects made training movements to all targets except the reference target and the adjacent target in the clockwise direction. This allowed us to determine the extent of generalization to movements in the reference direction following a near-complete exploration of the workspace. The extension of the interpolation condition (the full-interpolation condition) involved training in all directions except the reference direction. Here the goal was to determine whether interpolation between training movements benefited from a thorough exploration of the workspace. Table 1 gives the complete list of conditions, in terms of both the actual directions of training and the directions of training relative to the reference target.

Measures and statistics

Throughout the experiment, we used movement curvature to track learning and transfer of learning. We quantified movement curvature using perpendicular error (PE), which is defined as the perpendicular deviation at peak tangential velocity from a straight-line linking movement start and movement stop (scored at 5% of peak tangential velocity). We assessed PE at peak tangential velocity to minimize the influence of feedback responses to movement error on each trial's measure of curvature. We examined other dependent measures of movement curvature (area bounded by the movement trajectory, initial angular deflection of the movement from a straight line, length of the movement path, PE 250 or 500 ms into movement, maximum PE) and found results consistent with those based on PE.

For statistical tests, we combined groups according to the absolute direction(s) of training movements relative to each subject's reference direction. This resulted in seven training conditions relative to the reference direction: 90° away, 45° away, 0° away (control condition), 45 and 90° away (two-target condition), +45 and −45° away (interpolation condition), from 45 to 270° away (multitarget condition), and all targets from 45 to 315° away (full-interpolation condition).

We used ANOVAs followed by Bonferroni-corrected pairwise comparisons to evaluate differences in performance between conditions. To rule out differences between subjects before training, we performed an ANOVA on movement curvature (PE) for the final five movements in the baseline phase. To determine whether groups differed in the extent to which they learned to compensate for the force field, we performed a pair of analyses. First, we performed an ANOVA on PE for the final five movements in the training phase to ensure that movement curvature immediately before the test phase was not different between groups. Next, because asymmetries exist in the extent to which the force field affects training and catch-trial movements to different targets (e.g., see Fig. 2 of Shadmehr and Brashers-Krug 1997; Fig. 4 of Malfait et al. 2002), we normalized catch-trial magnitude to the magnitude of initial movement curvature in the training phase. This normalized measure (ratio of PE on the final catch trial to PE on the initial training movement) quantifies the proportion of the initial load that was accounted for by learning. We performed an ANOVA on these normalized catch trials to ensure that the extent of learning did not differ between experimental conditions. Finally, we assessed transfer of learning from the training phase to the test phase by performing an ANOVA on PE for the initial movements made to the reference target in the test phase.

RESULTS

Here we tested the extent to which dynamics learning generalizes from a series of training movements to subsequent test movements in a reference direction. In the main experiment, we tested 160 subjects. Figure 2 shows how our measure of performance, perpendicular error (PE), changed over the course of the experiment for each experimental condition. Figure 2A shows performance for subjects who trained to single targets located −90, −45, 0, +45, or +90° from the reference target. Figure 2B shows performance for subjects
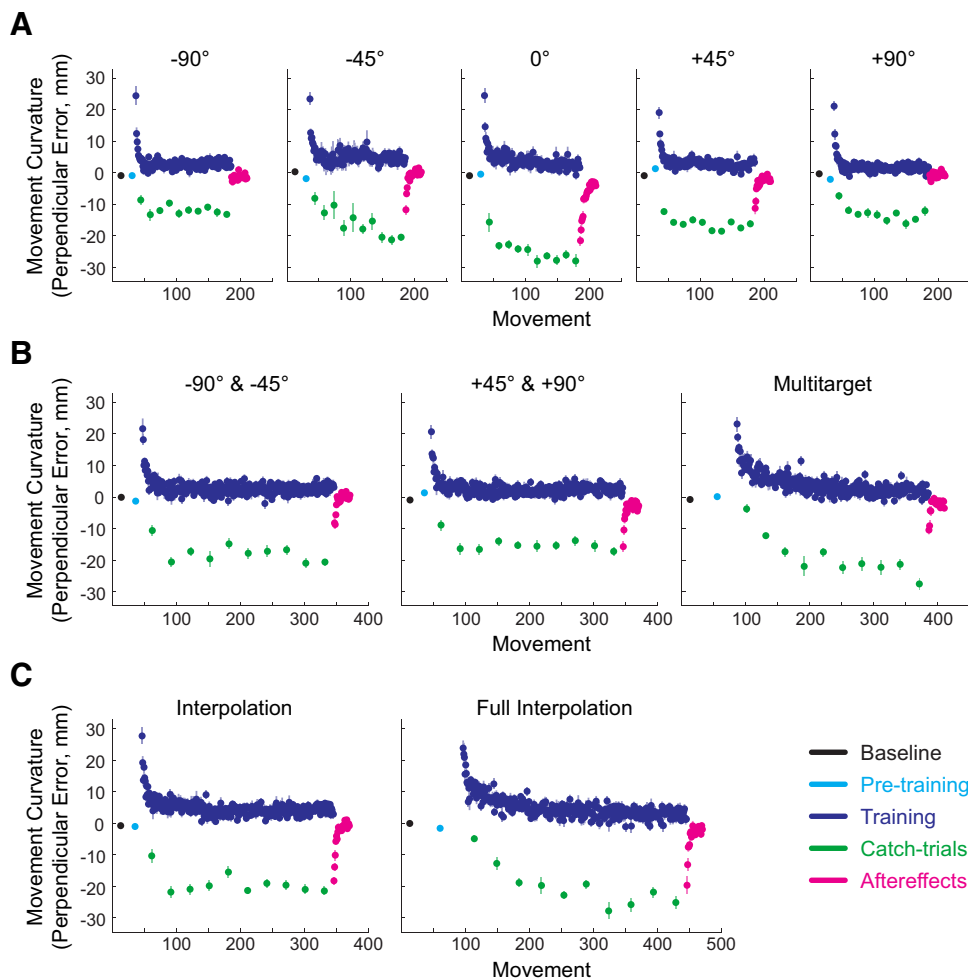
FIG. 2. Learning curves depicting changes in movement curvature over the course of training for each experimental condition. *A–C* show the following pattern of performance. In the baseline phase, subjects made 25 movements in the reference direction in a null field (movements depicted by the black dot). Subjects then made 10 pretraining movements per training target in a null field (movements depicted by the light blue dot). When the force field was turned on, subjects' movements were initially curved but eventually straightened (individual movements depicted by the dark blue dots). As they learned to compensate for the force field, curvature on catch trials increased. Catch trials were separated into 10 equally sized bins and are depicted by green dots. In the test phase of the experiment, subjects made movements to the reference target in a null field. These movements were curved opposite to the direction of the force field (movements depicted by the pink dots) indicating transfer of learning from the training phase to the test phase. Data points depict mean movement curvature ± SE. *A*: performance of subjects trained in the single target condition. *B*: performance of subjects trained in the 2-target or multitarget condition. *C*: performance of subjects trained in the interpolation condition or full-interpolation condition.

who trained in the two- or multitarget condition. Figure 2*C* shows performance for subjects who trained in the interpolation or full-interpolation condition. The following pattern of performance was observed in all conditions. In the baseline phase of the experiment, the robot applied a null field and subjects made movements to the reference target. In the pretraining phase, subjects made movements to the appropriate training target(s), also in a null field. When the clockwise force field was activated in the training phase, movements were initially curved consistent with the load. Over the course of training, movements straightened as subjects gradually learned to compensate for the externally applied loads. As movements in the force field straightened, curvature on catch trials (on which the load was unexpectedly removed) grew such that by the end of training, catch-trial curvature was equal in magnitude (but in the opposite direction) to initial movements in the force field. Immediately after the training phase, subjects made movements to the reference target in a null field (aftereffect trials). Curvature on initial test phase movements reflects the degree to which learning transferred from the training phase to the test phase of the experiment. A greater PE implies greater transfer and thus generalization of learning.

ANOVA revealed that there were no differences in movement curvature over the final five movements in the baseline phase of the experiment [$F_{(6,153)} = 0.50$, $P > 0.05$]. This suggests that before training, there were no preexisting differ-

ences between subjects in the various conditions. ANOVA likewise revealed that movement curvature did not differ between conditions over the final five movements in the training phase of the experiment [$F_{(6,153)} = 1.90$, $P > 0.05$]. Moreover, normalized catch trials (the ratio of PE on final catch trial to PE on the initial training movement) did not differ across conditions [$F_{(6,151)} = 1.61$, $P > 0.05$]. These results suggest that by the end of training, subjects did not differ in the extent to which they compensated for the force field. However, ANOVA indicated that subjects did differ in terms of movement curvature during the test phase of the experiment [$F_{(6,153)} = 28.98$, $P < 0.01$]. These differences are detailed in Figs. 3 and 4, which show differences in movement curvature over the first two trials in the test phase.

Figure 3 shows the effect of training in different directions on movements to the reference target in the test phase. The ordinate shows movement curvature. The abscissa gives the angular distance between the reference and training directions for the single-target condition. The pictographic labels indicate the direction(s) of training with filled circles. They also show three representative movements to the reference target in the test phase. Although the training directions and the representative movements are depicted here relative to a reference target at 315°, note that the results shown here (and also in Fig. 4) are pooled over two complete repetitions of the experiment, one relative to 315° and one relative to a reference target at 135°. The
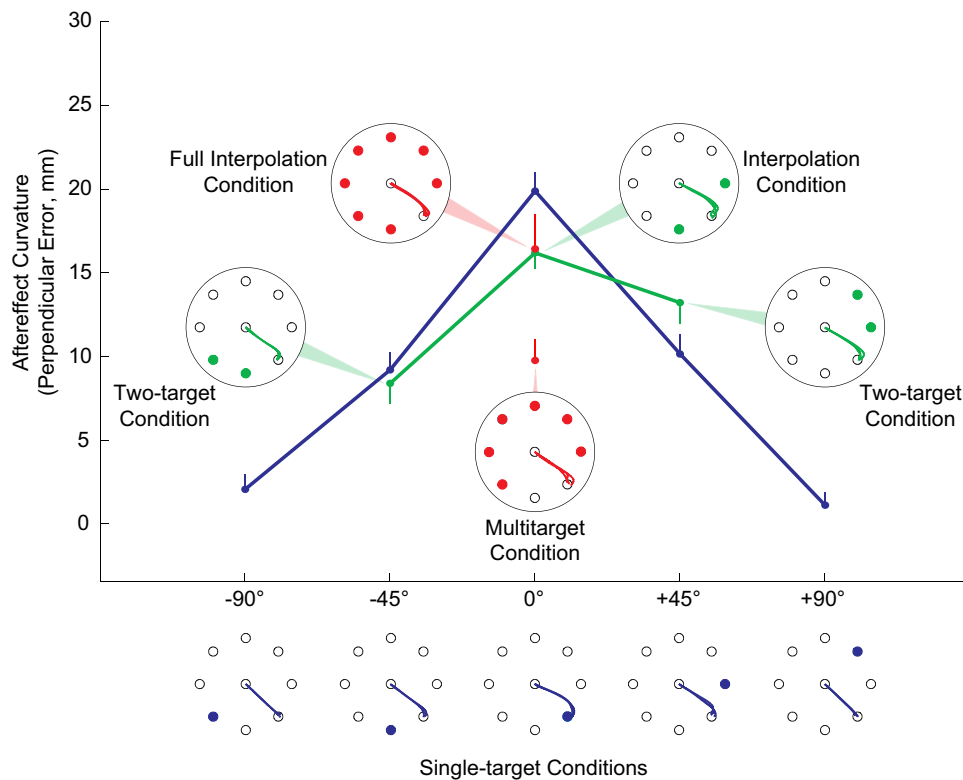
FIG. 3. Curvature in the reference direction as a function of the direction(s) in which subjects made training movements. Subjects made movements in the training phase before making movements in the reference direction. In the pictographic labels, the colored targets indicate the direction of training movements for each condition. Hand paths depict representative movements in the reference direction. In the single-target condition, relative to subjects who trained and were tested for transfer in the reference direction (0° data point), aftereffect magnitude was reduced as the separation between training and reference targets increased. Subjects in the 2-target condition, and indeed even subjects in the multitarget condition, showed aftereffects that were no larger than those demonstrated by subjects who trained only +45 or −45° from the reference target. In contrast, subjects in the interpolation condition and the full-interpolation condition, whose training included targets that bounded the reference direction (and thus provided a basis for interpolation), performed much like subjects who trained in the reference direction. Data points depict mean movement curvature ± SE for the initial 2 test movements.

individual data points show the dependence of movement curvature in the test phase on the direction(s) of movement in the training phase. The blue line is for the single-target condition. It can be seen that movement curvature in the test phase decreased monotonically as the distance increased between the reference and training targets. When the reference and training targets were separated by 90° there was effectively no curvature on the test trials. The green line shows performance of subjects in the two-target condition. It can be seen that the performance of subjects in the two-target condition (where subjects trained to targets located both 45 and 90° from the reference target) resembled the performance of subjects who trained to the 45° targets alone. Indeed, we saw no difference in performance between subjects who trained to the 45° target alone and those who trained in the multitarget condition in which targets spanned the workspace but did not bound the reference direction. In contrast, when subjects made training movements to targets that bounded the test direction (the interpolation condition), movement curvature increased and approached that observed in subjects who trained in the reference direction. The extent of generalization was similar for both the interpolation condition involving two targets that bounded the reference direction and the full-interpolation condition in which the entire workspace was explored.

Figure 4 summarizes the statistical analysis of the data presented in Fig. 3. Here we report how movement curvature in the test phase of the experiment depended on the separation between the training and reference targets. Conditions shown in different shades in Fig. 4 are reliably different from one another according to Bonferroni-corrected post hoc comparisons. Subjects who made training movements to a single target located 45° from the reference target were less curved on test movements ($P < 0.01$) than subjects who trained in the

reference direction. Subjects who trained to a single target 90° from the reference target showed almost no curvature in the test phase, suggesting that generalization of dynamics learning across 90° is extremely modest. Indeed curvature after training


FIG. 4. Adaptation to a force field in the training direction led to aftereffects in the reference direction. Degree of curvature in the test phase depended on the location of the training target(s) relative to the reference direction. Bars of different shades are reliably different from each other by post hoc tests. When 90° separated training and reference targets, movements in the test phase were straight. When 45° separated training and reference targets, aftereffect movements were curved but not to the same degree as when training and test movements were made to the same target. When subjects made training movements to targets located 45 and 90° from the reference target (2-target condition), aftereffect curvature was no different from that of subjects who trained only 45° away. Same was true for subjects in the multitarget condition. Subjects in the interpolation and full-interpolation conditions, whose training targets bounded the reference target, showed aftereffect curvature that was no different from that of subjects who made training movements to the reference target. Data points depict mean movement curvature ± SE for the initial 2 test movements.

at 90° from the reference direction was reliably less than curvature after training at 45° or in the reference direction ($P <$ 0.01 in both cases). Subjects in the two-target condition who trained to targets located at both 45 and 90° relative to the reference direction showed aftereffect curvature that was no different from that of subjects who trained only 45° away ($P >$ 0.05). Moreover, subjects in the multitarget condition showed exactly the same pattern; that is, their performance was no different from that of subjects who trained to a target located 45° from the reference direction ($P > 0.05$). This similarity in curvature in spite of a more thorough exploration of the workspace shows that generalization did not benefit from increased experience with the task. Instead the effects in the reference direction were limited to the localized effects of learning in the nearest training direction. In contrast, aftereffect curvature was greater ($P < 0.05$) when subjects' training movements bounded the test direction (in both the interpolation and the full-interpolation conditions). In both cases, the magnitude of movement curvature was no different from that of subjects who trained in the reference direction ($P > 0.05$ for both comparisons). The finding that both direct and interpolated training had similar effects on performance suggests that interpolation between instances of local learning may provide the basis for generalization of dynamics learning.

In a control study, we trained an additional 15 subjects to single targets located +135, −135, or 180° from the 315° reference target. In all cases, training beyond 90° resulted in minimal curvature on test phase movements that was no different from that of subjects who trained 90° from the reference target ($P > 0.05$).

In Fig. 2A, one can note that for subjects who made movements to single targets, the magnitude of curvature on test movements in the reference direction was correlated with the magnitude of curvature on final catch trials in the training direction. That is, subjects who showed large curvature during the test phase also showed large curvature on final catch trials. This could indicate that differences in aftereffect magnitude during the test phase were a consequence of directional differences in the extent to which subjects learned the force field. To test this possibility, we performed the following control study. We tested 20 new subjects in a variant of the single-target condition from the main experiment. Subjects were trained −90, −45, 0, +45, or +90° relative to a reference target located at 225° (i.e., halfway between the reference targets in the main experiment). After familiarization, subjects made 150 training movements followed immediately by 25 test movements in the reference direction. Learning curves depicted in Fig. 5A show that subjects learned to compensate for the force field in each direction. Figure 5B shows that, as in the main experiment, the magnitude of curvature in the reference direction varied depending on the direction of training. We combined groups according to the absolute direction of training, and ANOVA revealed that these differences in curvature were reliable [$F_{(2,17)} = 10.97$, $P < 0.01$]. Bonferroni-corrected post hoc tests showed that curvature on test movements after training in the reference direction was greater than that for subjects who trained 45 or 90° away ($P < 0.05$, $P < 0.01$, respectively). Subjects who trained 45° from the reference target showed a trend toward greater aftereffect curvature than subjects who trained 90° away (Fig. 5C). Importantly, we found that curvature on aftereffect movements in the test phase was largest for the condition that showed the smallest curvature on final catch trials (Fig. 5A). Moreover, whereas curvature on final catch trials increased as separation from the reference direction grew from 45 to 90°, curvature on initial test movements in the reference direction decreased. Thus this control study suggests that the magnitude of curvature on movements
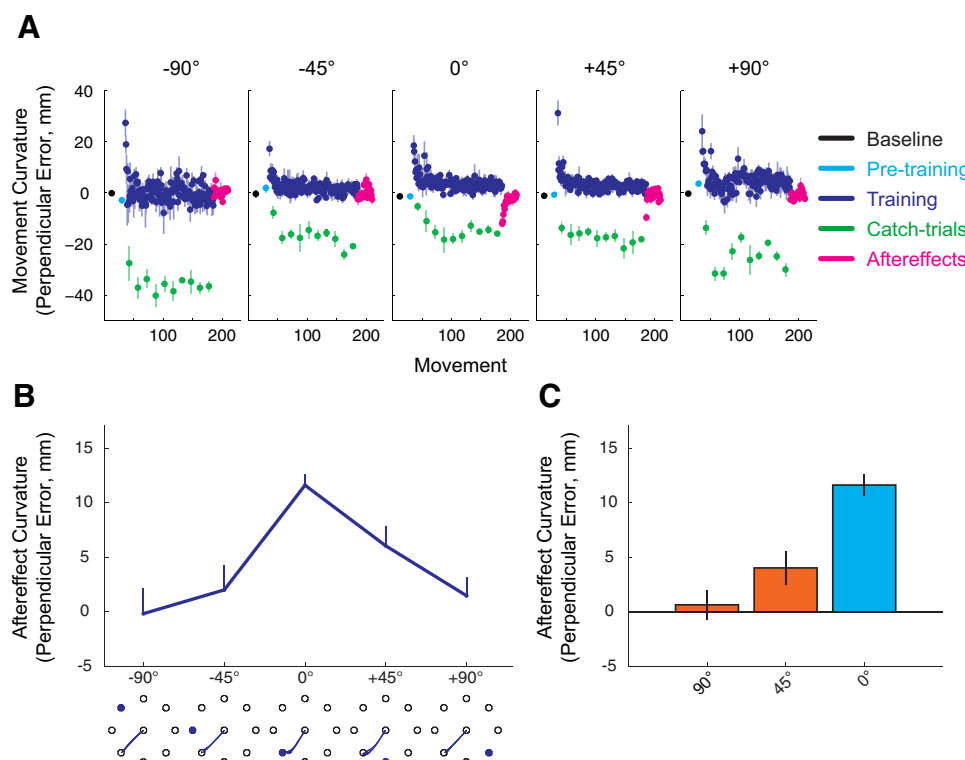

FIG. 5. A control experiment reveals that the magnitude of curvature on test movements does not depend on the magnitude of curvature on catch trials during training. Data in this figure are plotted as in Figs. 2, 3, and 4. A: learning curves depicting changes in movement curvature over the course of training. Subjects made training movements −90, −45, 0, +45, or +90° relative to a reference direction located at 225°. Subjects who made training movements −90 or +90° from the reference direction had the largest curvature on catch trials but the smallest curvature on aftereffects. B: as in the main experiment, we found that curvature on test movements decreased as the separation between the reference and training directions grew. C: curvature in the test phase depended on the direction of training. Bars of different colors are reliably different from each other according to post hoc tests. In A, data points depict mean movement curvature ± SE for each trial. In B and C, data points depict mean movement curvature ± SE for the initial 2 test movements.

in the reference direction is not tied to the magnitude of curvature on catch-trial movements during training.

Several features of the quantitative analysis merit comment. Whereas the data presented in Fig. 3 show the actual direction of training (+ or − relative to the reference direction), the analyses reported in Fig. 4 are based on the absolute separation between training and reference directions. Further, in Fig. 3 it looks as if there may be a directional asymmetry in the effects observed in the two-target conditions. In particular curvature in the reference direction appears to be affected to a greater extent by training to targets that were +45 and +90° away than targets at −45 and −90°. However, this asymmetry was not reliable in a statistical analysis. Specifically, ANOVA produced no evidence that the generalization gradient was asymmetric about the reference direction [$F_{(1,90)} = 2.76$, $P > 0.05$]. To explore further the possibility that transfer of learning was asymmetric, we carried out a control experiment in which we tested eight new subjects. Here we changed the direction of the load applied during the training phase to a counterclockwise force field and once again trained subjects at +45 and +90° or −45 and −90° relative to the 315° reference direction. We found that under these conditions, any evidence for an asymmetry in the extent of generalization disappeared. This suggests that any directional differences in aftereffect magnitude were due to the direction of loads applied in the training phase and not to a differential transfer of learning. Finally, we have presented results quantified over the first two trials in the test phase. We repeated our analysis throughout the initial seven movements in the test phase and found the same pattern of statistical differences presented in Fig. 4, in which the interpolation and full-interpolation conditions show complete generalization, the 45° two- and multitarget conditions show partial transfer and the 90° condition shows no transfer of dynamics learning.

DISCUSSION

We have assessed whether generalization of dynamics learning is affected by the distribution of directions in which subjects were trained. The single-target condition showed that the generalization gradient decreased steeply such that learning transferred minimally to movements ≥90° from the initial training direction. When subjects trained to an increasing number of targets that more fully sampled the workspace, generalization was no greater than would be expected after training to the nearest target alone. Generalization was complete for both interpolation groups, whose training involved targets that flanked later test movements by ±45°.

In agreement with previous studies, our results have shown that dynamics learning generalizes such that the effects of training are greatest on nearby movements (Donchin et al. 2003; Gandolfo et al. 1996; Huang and Shadmehr 2007; Thoroughman and Shadmehr 2000; Thoroughman and Taylor 2005). Unlike generalization of visuomotor rotation learning (Krakauer et al. 2000), the shape of the generalization gradient for dynamics learning does not appear to be modifiable. Subjects' ability to compensate for the effects of a dynamics perturbation in the reference direction did not benefit from extensive exposure to the same perturbation over a large portion of the workspace. This is consistent with the idea that dynamics learning is highly localized and argues against the possibility that motor learning results in the development of a broadly generalizable dynamics representation.

We have used an experimental design in which the training phase was immediately followed by a test phase in which generalization of dynamics learning was assessed. This allowed us to examine how dynamics learning acquired in an uninterrupted phase lasting hundreds of movements affected subsequent movements in the reference direction. Our approach differs from previous studies in which generalization gradients were determined trial by trial, by modeling the sensitivity of the current movement to error on the previous movement in a different direction (Donchin et al. 2003; Huang and Shadmehr 2007; Thoroughman and Shadmehr 2000; Thoroughman and Taylor 2005). Unlike in these previous studies, here we did not find evidence for transfer of learning to movements further than 90° from the training direction. At a neuromuscular level, nearby movements involve similar patterns of muscle activation. Transfer of learning declines as the separation between training and test movements increases and the extent to which movements share underlying control signals (and thus patterns of muscle activation) is reduced. The idea that generalization of learning between movements depends on the similarity of their underlying motor commands is reflected in studies that have shown that adaptation is tied to the specific muscles involved in training (Gandolfo et al. 1996; Malfait et al. 2002; Shadmehr and Moussavi 2000; Shadmehr and Mussa-Ivaldi 1994).

Here, we have found narrow generalization of dynamics learning in a polar coordinate frame. As the angular separation between training and reference movements increased, the extent of generalization decreased to zero. This finding is in contrast to other studies that suggest dynamics learning can be encoded in cartesian coordinates (Thoroughman and Shadmehr 2000; Donchin et al. 2003; Thoroughman and Taylor 2005; Huang and Shadmehr 2007). In these studies, errors experienced to the right on movements outwards from the center led subjects to predict rightward forces on subsequent movements in all directions, including inward movements 180° away. That is, aftereffects were counterclockwise ≤90° from the training direction but clockwise for movements between 90° and 180° away. The results in the present study are not consistent with broad generalization of dynamics learning in Cartesian coordinates. In a control study we found no evidence for transfer of dynamics learning across separations of +135, −135, or 180°. Moreover, the aftereffects for subjects trained in the two- versus multitarget conditions were the same magnitude. The multitarget condition adds several targets to the two-target condition that, had generalization occurred broadly in Cartesian space, would produce aftereffects in the opposite direction because they lie >90° from the reference target. These more distant targets should thus mitigate the effects of the nearby targets to some extent. Instead, the lack of a difference between the two- and multitarget conditions argues against this possibility. The same argument can be applied to the full-interpolation condition, in which the addition of a number of distant (i.e., >90°) targets did not diminish curvature in the reference direction relative to the interpolation condition. Thus in the present data it appears that generalization did not occur in Cartesian coordinates. This difference between our findings and those of others is intriguing and, at present, its source is unknown. Our study differs from previous work in that gener-

alization was assessed in a block of null-field trials after training, rather than trial by trial by assessing the effects of catch trials throughout the training phase. Moreover, in the current study training movements were restricted to narrow parts of the workspace, whereas in previous studies movements were made throughout the workspace during training. Perhaps the presence or absence of generalization of dynamics learning in Cartesian coordinates depends on these aspects of the training phase.

We found that for both the interpolation and full-interpolation groups, dynamics learning generalized fully to movements in the reference direction. This is consistent with the idea that in the case of dynamics learning, the motor system is capable of combining control signals for movement (Atkeson 1989; Ghahramani and Wolpert 1997; Malfait et al. 2005; Mattar and Ostry 2007). In Malfait et al. (2005), performance could be accounted for by a linear interpolation of control signals for each muscle, weighted to account for the position of the arm relative to training locations. Similarly, in the present study transfer of learning to the reference direction may have resulted from a combination of the control signals for the two flanking directions. Importantly, the presence of training directions beyond those flanking the reference target may not be needed for generalization of dynamics learning because the flanking directions alone seem to be sufficient for successful interpolation.

In a previous study (Krakauer et al. 2000), subjects who trained to multiple targets in addition to those flanking the reference direction (analogous to our full-interpolation group) showed full generalization of visuomotor rotation learning. In contrast, subjects who trained only to flanking targets (analogous to our interpolation group) showed less generalization. Thus in contrast to the findings reported here, visuomotor rotation learning in directions that flanked the reference target was not sufficient for interpolation and full generalization. This difference may be due to the involvement of distinct neural processes in visuomotor and dynamics learning (Krakauer et al. 1999; although see Tong et al. 2002 for an alternate view). It may also be due to differences between the coordinate frame in which the visual targets were presented and the coordinate frame in which dynamics are learned. Unlike visuomotor learning, which occurs in eye-centered (Vetter et al. 1999) or Cartesian space (Ghahramani et al. 1996), dynamics learning proceeds in intrinsic, joint-based coordinates (Gandolfo et al. 1996; Malfait et al. 2002; Shadmehr and Moussavi 2000; Shadmehr and Mussa-Ivaldi 1994; although see Donchin et al. 2003; Huang and Shadmehr 2007; Thoroughman and Shadmehr 2000; Thoroughman and Taylor 2005 for evidence of dynamics learning in Cartesian space). As in previous studies that examined generalization in motor learning (Donchin et al. 2003; Huang and Shadmehr 2007; Krakauer et al. 2000; Thoroughman and Shadmehr 2000; Thoroughman and Taylor 2005), here we used visual targets that were evenly spaced in Cartesian coordinates. Because the transformation between Cartesian space and joint space is nonlinear, the separation between targets in joint-space was uneven. Thus training to targets separated evenly in the joint-based coordinate frame in which dynamics learning occurs may result in a pattern of generalization different from that described here.

In the present study we have shown narrow generalization of dynamics learning. However, other studies have shown that dynamics generalizes more broadly under certain conditions.

In particular, dynamics learning can generalize to movements that differ in amplitude or velocity from the training movements (Goodbody and Wolpert 1998) and also between limbs (Criscimagna-Hemminger et al. 2003; Dizio and Lackner 1995; Malfait and Ostry 2004; Wang and Sainburg 2004b). In some studies, interlimb generalization has been observed in an extrinsic or world-based coordinate frame (Criscimagna-Hemminger et al. 2003; Malfait and Ostry 2004). As noted earlier, this is in contrast to the intrinsic, joint-based coordinate frame in which dynamics learning is encoded and generalizes (Gandolfo et al. 1996; Malfait et al. 2002; Shadmehr and Moussavi 2000; Shadmehr and Mussa-Ivaldi 1994). This suggests that the process by which dynamics learning generalizes from one arm to the other may be distinct from the process by which control signals within an arm are updated to compensate for forces. Indeed, recent evidence suggests that, unlike intralimb generalization, generalization between limbs may depend on cognition as transfer disappears when subjects are unaware of the dynamics learning process (Malfait and Ostry 2004).

Here, we have tested generalization of dynamics learning using an experimental design that fully separates the training and test conditions. We asked whether movements in the reference direction show an additional benefit from training in multiple directions within the workspace. We found that generalization of learning was the same whether subjects trained to one adjacent target, to one adjacent target plus a more distant target, or to one adjacent target plus an additional five targets that spanned the workspace. The motor system was unable to exploit its experience compensating for forces in multiple movements that of necessity had different patterns of muscle activation and thus different motor commands. Instead, the effects on movements in the reference direction were limited to those that propagated from nearby training alone. When training involved movements in directions that flanked the reference target, we saw full generalization of dynamics learning. Thus the extent to which newly learned dynamics generalizes is not modified by broad experience with those forces distributed throughout the workspace. Instead, broad generalization of dynamics depends on interpolation between instances of local learning.

REFERENCES

Atkeson CG. Learning arm kinematics and dynamics. *Annu Rev Neurosci* 12: 157–183, 1989.
Bedford FL. Perceptual and cognitive spatial learning. *J Exp Psychol Hum Percept Perform* 19: 517–530, 1993.
Caithness G, Osu R, Bays P, Chase H, Klassen J, Kawato M, Wolpert DM, Flanagan JR. Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. *J Neurosci* 24: 8662–8671, 2004.
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.

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.

Dizio P, Lackner JR. Motor adaptation to Coriolis force perturbations of reaching movements: endpoint but not trajectory adaptation transfers to the nonexposed arm. *J Neurophysiol* 74: 1787–1792, 1995.

Donchin O, Francis JT, 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.

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

Ghahramani Z, Wolpert DM. Modular decomposition in visuomotor learning. *Nature* 386: 392–395, 1997.

Ghahramani Z, Wolpert DM, Jordan MI. Generalization to local remappings of the visuomotor coordinate transformation. *J Neurosci* 16: 7085–7096, 1996.

Ghilardi M-F, Gordon J, Ghez C. Learning a visuomotor transformation in a local area of work space produces directional biases in other areas. *J Neurophysiol* 73: 2535–2539, 1995.

Goodbody SJ, Wolpert DM. Temporal and amplitude generalization in motor learning. *J Neurophysiol* 79: 1825–1838, 1998.

Gribble PL, Ostry DJ. Compensation for loads during arm movements using equilibrium-point control. *Exp Brain Res* 135: 474–482, 2000.

Huang VS, Shadmehr R. Evolution of motor memory during the seconds after observation of motor error. *J Neurophysiol* 97: 3976–3985, 2007.

Hwang EJ, Donchin O, Smith MA, Shadmehr R. A gain-field encoding of limb position and velocity in the internal model of arm dynamics. *PLoS Biol* 1: 209–220, 2003.

Krakauer JW, Ghilardi MF, Ghez C. Independent learning of internal models for kinematic and dynamic control of reaching. *Nat Neurosci* 2: 1026–1031, 1999.

Krakauer JW, Mazzoni P, Ghazizadeh A, Ravindran R, Shadmehr R. Generalization of motor learning depends on the history of prior action. *PLoS Biol* 4: e316, 2006.

Krakauer JW, Pine ZM, Ghilardi M-F, Ghez C. Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J Neurosci* 20: 8916–8924, 2000.

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

Malfait N, Gribble PL, Ostry DJ. Generalization of motor learning based on multiple field exposures and local adaptation. *J Neurophysiol* 93: 3327–3338, 2005.

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.

Mattar AAG, Ostry DJ. Neural averaging in motor learning. *J Neurophysiol* 97: 220–228, 2007.

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

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

Thoroughman KA, Shadmehr R. Electromyographic correlates of learning an internal model of reaching movements. *J Neurosci* 19: 8573–8588, 1999.

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

Thoroughman KA, Taylor JA. Rapid reshaping of human motor generalization. *J Neurosci* 25: 8948–8953, 2005.

Tong C, Wolpert DM, Flanagan JR. Kinematics and dynamics are not represented independently in motor working memory: evidence from an interference study. *J Neurosci* 22: 1108–1113, 2002.

Vetter P, Goodbody SJ, Wolpert DM. Evidence for an eye-centered spherical representation of the visuomotor map. *J Neurophysiol* 81: 935–939, 1999.

Wang J, Sainburg RL. Limitations in interlimb transfer of visuomotor rotations. *Exp Brain Res* 155: 1–8, 2004a.

Wang J, Sainburg RL. Interlimb transfer of novel inertial dynamics is asymmetrical. *J Neurophysiol* 92: 349–360, 2004b.

Witney AG, Wolpert DM. Spatial representation of predictive motor learning. *J Neurophysiol* 89: 1837–1843, 2003.