

Learning of Visuomotor Transformations for Vectorial Planning of Reaching Trajectories

John W. Krakauer,¹ Zachary M. Pine,² Maria-Felice Ghilardi,^{3,4} and Claude Ghez³

Departments of ¹Neurology and ³Neurobiology and Behavior, Columbia University, College of Physicians and Surgeons, New York, New York 10032, ²Center on Aging, University of California, San Francisco, and Geriatrics Division, San Francisco Veterans Administration Medical Center, San Francisco, California and ⁴INB-CNR, Milan, Italy

The planning of visually guided reaches is accomplished by independent specification of extent and direction. We investigated whether this separation of extent and direction planning for well practiced movements could be explained by differences in the adaptation to extent and directional errors during motor learning. We compared the time course and generalization of adaptation with two types of screen cursor transformation that altered the relationship between hand space and screen space. The first was a gain change that induced extent errors and required subjects to learn a new scaling factor. The second was a screen cursor rotation that induced directional errors and required subjects to learn new reference axes. Subjects learned a new scaling factor at the same rate when training with one or multiple target distances, whereas learning new reference axes took longer and was less complete when training with multiple

compared with one target direction. After training to a single target, subjects were able to transfer learning of a new scaling factor to previously unvisited distances and directions. In contrast, generalization of rotation adaptation was incomplete; there was transfer across distances and arm configurations but not across directions. Learning a rotated reference frame only occurred after multiple target directions were sampled during training. These results suggest the separate processing of extent and directional errors by the brain and support the idea that reaching movements are planned as a hand-centered vector whose extent and direction are established via learning a scaling factor and reference axes.

Key words: vectorial planning; motor learning; visuomotor transformations; reaching movements; psychophysics; generalization

In planning reaches to visual targets the nervous system transforms information about target location into time-varying sets of muscle activations and joint torques that bring the hand to the desired position. Converging psychophysical and neurophysiological evidence suggests that it accomplishes this via a series of sensorimotor transformations in which the target and the movement are recoded in a series of successive representations of extrinsic and intrinsic space (Soechting and Flanders, 1989). At early stages of planning, the spatial location of the target is remapped from retinotopic into egocentric (eye-, head-, or shoulder-centered) coordinates (McIntyre et al., 1997; Carrozzo et al., 1999). Vectorial planning hypotheses posit that target information is combined with hand position information (Ghilardi et al., 1995; Vindras et al., 1998) to form a simplified hand-centered plan of the intended movement trajectory as an extent and direction in extrinsic space (Gordon et al., 1994a; Vindras and Viviani, 1998). Movement extent is determined by linearly scaling a stereotyped bell-shaped velocity profile, whereas movement duration is set by task context (Ghez and Krakauer, 2000). Importantly, planning an extent and a direction from the hand requires establishing a scaling factor relating target distance to a peak velocity and hand-centered reference axes relative to an egocentric reference frame. For movements to be accurate in a variety of tasks with different spatial characteristics, both operations must be under adaptive control. For example, when using a computer, if the screen and pad are displaced from in front of the body and the distance of the head from the screen changes, both

the reference frame and scaling factor must change to remain accurate with the computer mouse.

This paper examines whether, as would be predicted in a vectorial framework, errors in extent and direction are processed differently for adaptive learning. To address this question we compared the time course of adaptation and the degree of generalization across work space parameters for two conditions that separately perturbed the scaling and reference axes of the visuomotor transformation. In one condition, systematic extent errors were introduced by changing the gain of the hand path display; in the other, directional biases were introduced by rotating the displayed hand path around the initial position (Cunningham, 1989; Roby-Brami and Burnod, 1995; Pine et al., 1996). The change in display gain required subjects to rescale vector amplitude, whereas the rotation required reorienting vector direction. We focused our analysis on early trajectory variables rather than end points to reduce the contributions of feedback (Gordon et al., 1994a; Messier and Kalaska, 1999) and to separate vectorial parameters from position parameters (Paillard, 1996). The emphasis in all experiments was to determine whether learning achieved for a given target remained local or generalized to other locations in the work space, because patterns of generalization provide insight about the representation of internal models in the nervous system (Imamizu et al., 1995; Ghahramani and Wolpert, 1997).

Parts of this paper have been published previously (Krakauer et al., 1996, 1997, 1999b; for review, see Ghez et al., 2000).

MATERIALS AND METHODS

Subjects

Fifty-nine right-handed subjects (48 men and 11 women; aged 18–40 years) were in the study. All were naive to the purpose of the experiments, signed an institutionally approved consent form, and were paid to participate. To avoid unwanted crossover effects, we examined separate groups of subjects for adaptation to gain changes and rotation. Separate groups of subjects were also used when comparing the effects of target number on the time course and generalization of learning. In the single-target training experiments (for both gain and rotation), the same six subjects were trained on four separate single targets but only on one of these targets on any given

Received June 28, 2000; revised Sept. 7, 2000; accepted Sept. 18, 2000.

This work was supported by National Institutes of Health Grants NS 22713, NS 02138, and NS 01961. We thank Jaiek Oh and Hao Huang for computer programming. We thank Cathleen Song, Thomas Frontera, and Milana Veytsman for technical assistance with the experiments and the data analysis. We thank Ning Qian, James Gordon, Michael Kim, and David Krakauer for helpful comments on this manuscript.

Correspondence should be addressed to Dr. Claude Ghez, Center for Neurobiology and Behavior, New York State Psychiatric Institute, Columbia University, College of Physicians and Surgeons, 1051 Riverside Drive, New York, NY 10032. E-mail: cpg1@columbia.edu.

Copyright © 2000 Society for Neuroscience 0270-6474/00/208916-09\$15.00/0

day. Several days intervened between each training day. This was done to minimize crossover effects for single targets.

Apparatus

Subjects sat facing a computer monitor (17 × 12 cm) at eye level (distance, 72 cm) and controlled a screen cursor by moving a hand-held indicator across the surface of a horizontal digitizing tablet (sampling rate, 200 Hz) with their right arm. The elbow was kept at shoulder level by an airslid supporting the forearm, the shoulder was restrained to prevent translation in the A-P plane, and the wrist was splinted in the neutral position. In all experiments, except one, which examined generalization of learning across the work space, the starting configuration was with the shoulder at 45° and the elbow at 90°. An opaque shield prevented subjects from seeing their hand or arm at all times.

General experimental procedure

Subjects were required to reposition the screen cursor from a common central origin to a series of peripheral circular targets displayed on the screen (see below and figures for details of particular target sets; all distances refer to screen distances). Subjects were instructed to make straight and uncorrected out-and-back movements with sharp reversals in the target and to pause briefly in the starting position before moving to the next target. The signal to move was a tone synchronous with a particular target changing from white to black. The tones occurred at an interval of 1.25 sec. Subjects could be tested with or without screen cursor feedback of hand position. An air jet positioned above the starting position directed a stream of air onto the knuckle of the forefinger, allowing subjects to recenter their hand when visual feedback was absent.

Each experiment was conducted with six subjects and consisted of four blocks of trials. The first was a “familiarization” block of 88 trials in which subjects moved to all targets in the relevant target set in the absence of any perturbations (cursor-feedback gain, 1/1; leftward hand movement caused leftward cursor movement) and with continuous cursor feedback. The second was a “control” block of trials that consisted of both feedback and no feedback targets. Feedback was provided to the target(s) that would be used for training in the subsequent training block; no feedback was provided to the other targets. Subjects made ~10 movements per target. The third was a “training” block in which subjects were trained on either one of two screen cursor perturbations over 144 trials. One perturbation altered the gain relationship between the distance moved on the screen and the distance moved on the tablet from 1/1 to 1.5/1. The other perturbation altered the direction of cursor movement relative to the hand movement on the tablet counterclockwise (CCW) by 30°. These gain and rotation values were chosen because they cause predicted linear errors of equal magnitude.

Training with the gain change and the rotation was with either one or more targets. When multiple targets were presented, they were at various distances and/or directions from the starting position. Data from the training blocks were used to determine the time course of learning. Finally, there was a “testing” block that was identical to the control block except that either the gain or rotation perturbation was imposed. Thus, in the testing block, subjects were provided with visual feedback (refresher feedback) to the targets they had trained to but not to any others. Differences between the control and testing blocks were used to generate the generalization data. For directional data, the training target direction was always 0°.

Time course protocol

Training blocks from six separate groups of six subjects were used to generate these data. Two groups were trained in the presence of a gain change of 1/1.5 with either one or eight targets (see Fig. 1A). The two directions were chosen so that movements were inertially equivalent, because an equal distance in the two directions required the same degree of shoulder and elbow rotation. The other four groups were trained in the presence of the 30° CCW rotation but with a different numbers of randomized targets: one, two, four, or eight (see Fig. 1B).

Gain generalization protocol

Two different groups of six subjects underwent the standard experimental blocks using to either one of two different target sets. One target set was of eight targets: four placed along a 45° line and four placed along a 135° line. The targets were spaced at 2.4, 4.8, 7.2, and 9.6 cm from the start position (see Fig. 4A). Subjects trained to the 2.4 and 9.6 cm targets in both directions on 4 separate days but to only one of them on a single day. The order of training across the 4 d was pseudorandomized. In the testing block, subjects were tested to the remaining targets in the absence of visual feedback.

The other target set consisted of targets arrayed in a circle of radius 4.2 cm. The training target was at either 45, 135, 225, or 315° from the start position (see Fig. 5). Subjects trained to all four of these targets but on separate days and in random order. Testing was to the remaining targets 0, ±22.5, ±45, ±90, and 180° relative to the training target and in the absence of visual feedback.

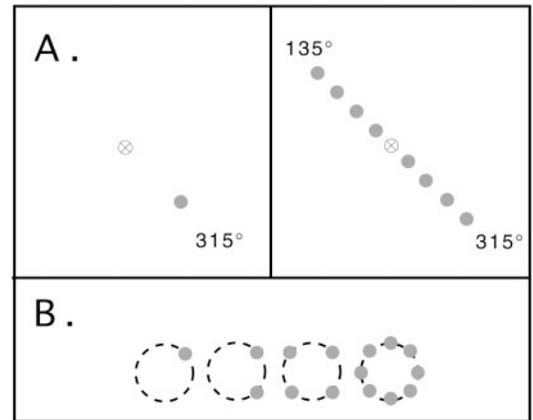


Figure 1. Target arrays for time course of learning experiment. *A*, One (left) and eight (right) training targets for gain learning. The crossed circle indicates the start position, and the targets are in gray. The targets were circular and were spaced at 2.4, 4.8, 7.2, and 9.6 cm from the starting position in both 135 and 315° directions. Single-target training was to the 7.2 cm target. *B*, One, two, four, and eight training targets for rotation learning. The targets were arrayed in a circle of radius 4.2 cm.

Rotation generalization protocol

Generalization across directions. The testing blocks for the four groups trained with different target numbers on the 30° CCW rotation used to study generalization of learning across directions. For the single-target group, the arrangement of the training and testing targets is shown (see Fig. 6A). For the two-target set, the testing targets were at 0, ±22.5, ±45, ±90, and 180° relative to the training target. For the four-target group, testing was at ±22.5 and ±45° relative to the training targets. For the eight-target group, testing was at ±22.5° relative to the training targets (see Fig. 6B). Subjects made 12 visits to each target. They were provided with refresher feedback to the training target every four movements in the one-, two-, and four-target testing and every other movement during the eight-target testing.

Generalization across distances. A separate group of subjects was used to examine how rotation learning generalizes across distances. The training block was to a single target at 45° and at a distance of 7.2 cm from the starting position. The testing block consisted of targets along the same directional axis but at distances of 2.4, 4.8, and 9.6 cm in the absence of visual feedback (see Fig. 7, inset). Subjects made 11 visits to the training target and 5 to the others. They visited the training target every four movements.

Generalization across arm configurations. A group of subjects was trained with a 60° CCW rotation with a single target at 45° in a circular array of radius 4.2 cm. Training continued until subjects could correct their directional error to less than the angle subtended by the target within two successive movements using cursor feedback. Subjects were then tested with the 60° rotation, in the absence of cursor feedback, to the original training direction and three others (135, 225, and 315°) both in the training configuration (shoulder = 45°; elbow = 90°) and in a new testing configuration (shoulder = 90°; elbow = 90°) (see Fig. 8A,B). Subjects were passively moved into the new configuration by displacing their chair laterally to the left. They maintained their hand position by use of the air jet.

Data analysis

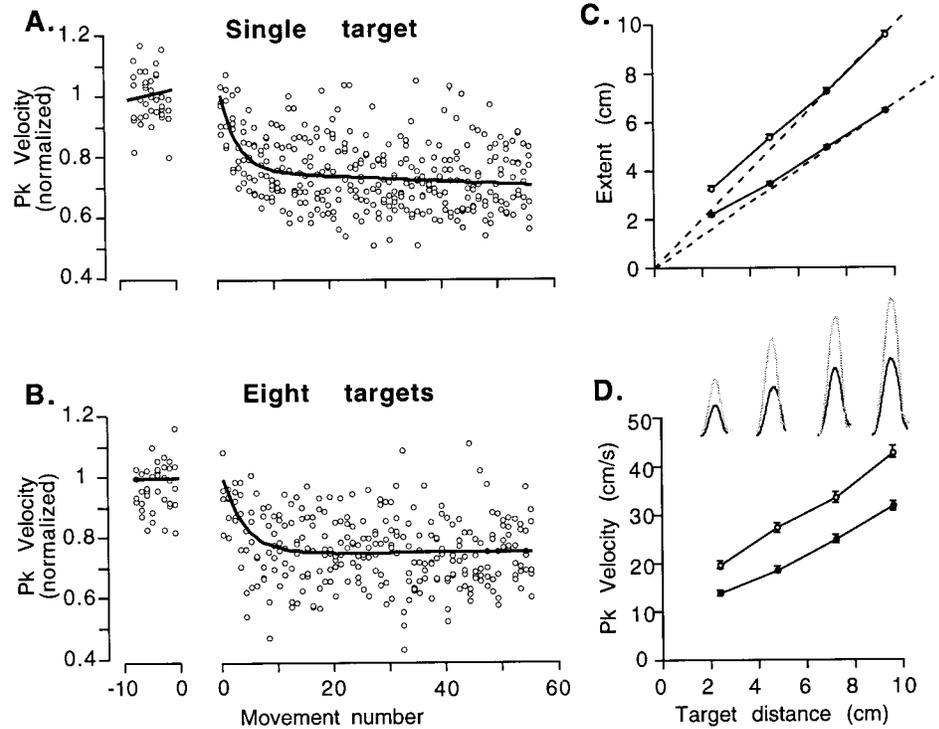
For each movement (from the onset of the change in hand velocity in the start circle to the velocity minimum when it returned near the starting position) we determined the hand locations and bin time at various critical points in the trajectory. These included the peak acceleration and velocity and the end point of the outward movement. The directional error for each movement was taken as the difference between the direction of the target from the initial hand position and the direction of the hand at the peak outward velocity from the same initial point. For group data, averages and SEMs were generated for each target.

Gain adaptation time course data were compared across subjects by normalizing the peak velocities for each subject. For the one-target training, this was done by dividing each individual peak velocity in a subject's training block by the mean peak velocity to that particular target over the last half of the familiarization block. For training to eight targets, each individual peak velocity in the training block was divided by the mean peak velocity (V_{pk}), calculated over the last half of the familiarization block, to that same target distance.

Gain adaptation to a particular target distance was calculated as a percentage:

$$300 * \left(1 - \frac{V_{pk}(\text{testing})}{V_{pk}(\text{control})} \right)$$

Figure 2. Gain learning. The last 8 movements of the baseline block are shown followed by 56 consecutive movements at gain 1.5. Each plot shows group data. *A*, Learning curve for gain training to a single target. *B*, Learning curve for gain training to eight targets. *C*, The relationship between mean movement extent and target distance at a gain of 1 (open circles) and a gain of 1.5 (filled circles). The dashed lines represent accurate performance at the two gains. The movement extents closely matched the target distances except for the smallest movements, which were somewhat hypermetric. *D*, The relationship between mean peak velocity and target distance. The outward trajectories had stereotypical single-peaked velocity profiles that scaled with target distance (inset). It may be noted that the lines fitting the peak velocities at the four target distances do not intercept the y-axis at zero. This was not investigated specifically but may represent either a range effect or an intrinsic nonlinearity in programming of small but rapidly rising force impulses (Gordon and Ghez, 1987).



Rotation adaptation was calculated as a percentage:

$$100 * \left(1 - \frac{DirErrV_{pk}}{30^\circ} \right),$$

where $DirErrV_{pk}$ is the directional error at the peak velocity.

Generalization of gain and rotation adaptation to no-feedback testing targets (FB-) was measured by calculating their percentage adaptation relative to the adaptation to the training target in the testing block (refresher trials):

$$100 * \frac{\text{mean \% adaptation FB-}}{\text{mean \% adaptation FB+}}.$$

In the experiment examining the effect of arm configuration, movement-by-movement joint angle changes were computed from individual hand extents and directions and from subject limb segment lengths using trigonometry.

The effect of the gain and rotation changes on variability was examined by comparing the control and training blocks for the eight-target group. The variability in the rotation group was obtained by calculating the mean SD of the directional error at the peak velocity over the last 24 movements in the control and training blocks. The variability in the gain group was obtained by calculating the coefficient of variation for the peak velocity over the last 24 movements.

Path curvature was quantified by subtracting the directional error at the end point from the directional error at the peak velocity.

The time course data for individual subjects was, in the majority of cases, fitted better by double exponentials than by single exponentials. This was ascertained after analysis revealed consistently higher residuals when fitting individual subject data with single versus double exponentials. Thus, we chose to fit all our group data with double exponentials.

Differences across conditions were assessed using single- or two-factor ANOVAs with Bonferroni-Dunn *post hoc* tests significant at $p < 0.001$. Directional errors were computed relative to the target in question with clockwise (CW) errors being made negative and CCW errors being made positive.

RESULTS

Time course of learning gain changes and rotations

At the end of the familiarization block all subjects moved their hand out and back with straight paths, reversing direction in the target centers. Velocity profiles during the outward movements increased smoothly to a single peak before declining more rapidly to a minimum at movement reversal.

The first movements made after the display gain was increased were hypermetric on the screen by $46 \pm 4\%$ ($n = 12$), close to the

predicted value of 50%. Thereafter, movements became progressively smaller, and this was paralleled by a reduction in the normalized peak velocities whether targets were at one or at eight distances (Fig. 2*A,B*). As can be observed in Figure 2 from the fitted curves, the time course of adaptation across subjects was similar whether the gain change was learned with one or eight targets. This was confirmed statistically by comparing the mean peak velocities over the first and last 24 movements in the two conditions by a two-factor ANOVA and showing no significant interaction [$F_{(1,524)} = 0.49$; $p = 0.48$].

Peak velocities remained scaled to the target distance before and after adaptation (Fig. 2*C,D*), and velocity profiles were similar (Fig. 2*D*, inset). Subjects adapted to the new scaling factor by a change in the slope of the relationships of both movement extent and peak velocity to target distance. These results are similar to those in the monkey (Ojakangas and Ebner, 1991). The fact that adaptation was at least as rapid and complete with multiple distances as with a single-target distance suggests that errors produced in movements of one extent can be used to readjust the feedforward control of movements to another.

When the cursor display was rotated, the first movements showed the expected 30° CCW error. As shown in Figure 3, *A* and *B*, this bias was reduced over the ensuing movements, with both single- and multiple-target directions. Movement paths remained essentially straight and did not change significantly during the course of adaptation in any of the subjects (mean curvature in the eight-target group was $0.46 \pm 0.24^\circ$ for the familiarization block and $2.96 \pm 0.23^\circ$ for the training block). Thus, they did not attempt to correct the imposed directional errors via feedback during the movements themselves. Instead, they used visual feedback primarily to change the direction of subsequent movements adaptively.

In contrast to the effect of gain change, an increase in target number reduced the rate of adaptation to rotations. In Figure 3*B* it can be seen that learning a rotation to eight targets takes longer than learning to a single target (Fig. 3*A*) and is less complete at the end of 56 movements. We have plotted only the first 56 movements for the sake of clarity. The difference in the mean directional error over the first or last 24 movements as a function of the number of directions trained to (one, four, or eight) showed a highly signifi-

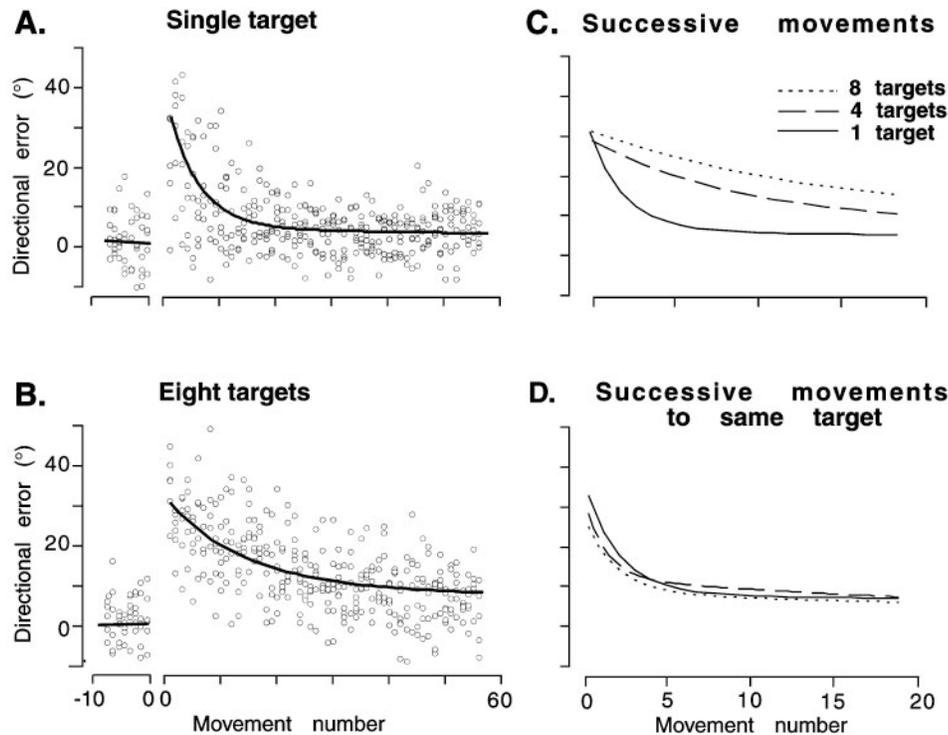


Figure 3. Rotation learning. Learning was measured by the progressive reduction in the directional error at the peak velocity. The last 8 movements of the baseline block are shown followed by 56 consecutive movements with the 30° CCW rotation. Each plot shows group data. *A*, Learning curve for rotation learning to a single target. *B*, Learning curve for rotation learning to eight targets. *C*, Learning curves for rotation learning with single, four, and eight targets, plotted for the first 18 moves of the training block. *D*, Learning curves for rotation learning plotted for consecutive moves to a single target for single-, four-, and eight-target training.

cant interaction by two-factor ANOVA [$F_{(1, 572)} = 16.53$; $p < 0.0001$].

One explanation for the slower learning of the rotation with multiple targets could be that directional errors made moving to a given target can only be used adaptively to correct movements in the same direction. If this were so, the rate of adaptation should decrease in proportion to the number of target directions but remain the same when plotted as a function of the number of visits to the same target. This was indeed the case (Fig. 3*C,D*). Adaptation with one, four, and eight targets showed no significant difference in the mean directional error over the first 18 visits to the same target by single-factor ANOVA [$F_{(2, 1201)} = 1.67$; $p = 0.19$]. This suggests that, unlike extent errors, directional errors are computed separately for each target direction and only used to adjust movements in the same direction.

In a previous study (Pine et al., 1996), we found that although extent variability did not change during adaptation to a gain change there was a marked increase in directional variability during the course of adaptation to rotations. This was not the case here. In the eight-target gain experiment, a two-factor ANOVA showed that the coefficient of variation decreased with target distance at both the control and new gains [$F_{(3,32)} = 4.82$; $p = 0.007$], but the interaction with condition was not significant [$F_{(3,32)} = 0.8$; $p = 0.5$]. In the eight-target rotation condition, the mean SDs in directional error for the last 24 movements of the control and training trials were not significantly different (unpaired t test, $p = 0.46$). We believe the difference between the two studies is attributable to the differences in feedback conditions. In the former study, feedback was in the form of knowledge of results; i.e., hand paths were displayed after the completion of each movement, rather than with continuous cursor feedback. In addition, the time between movements was more variable in the previous study, because the start was triggered manually (ranging from 5 to 6 sec instead of a fixed 1.25 sec). We have since observed that rotational adaptation tends to decay very rapidly in the few seconds between trials during the early phases of learning. We therefore attribute the increased variability of our previous study to the use of knowledge of results and to variable intertrial intervals.

Generalization of gain adaptation

As noted previously, the unchanged rate of adaptation to gain changes when errors are generated in movements of various extents

suggests that gain adaptation generalizes across target distances. Moreover, the hypothesis that such movements reflect mapping of the two-dimensional work space in a two-dimensional vector space raises the possibility that learning might generalize across directions as well. We first addressed this by training subjects with a single target and examining movements to targets at three other distances in two directions. In previous work we have shown that, despite systematic directional variations in limb inertia, subjects program the same impulse of force at the hand (Gordon et al., 1994b). This results in a lower peak velocity in the 135° direction, which has higher inertia because it requires motion of both the elbow and shoulder, compared with the 45° direction, which has lower inertia because it requires only elbow motion. However, end-point errors do not result because there is compensation in the movement time. We hypothesized that the CNS could only generalize across directions if these dynamic properties are taken into account.

Adaptation, although incomplete relative to the training target, was uniform across target distances (Fig. 4*A*). There was no significant effect of training target distance [$F_{(1,188)} = 0.04$; $p = 0.84$] or direction [$F_{(1,188)} = 2.22$; $p = 0.14$] on the percent adaptation (two-factor ANOVA), so we combined the data for all four training targets. A two-factor ANOVA showed no significant effect of test target distance [$F_{(3,160)} = 0.04$; $p = 0.99$] or direction [$F_{(1,160)} = 0.42$; $p = 0.51$] on the percentage adaptation. This was despite clear differences in the mean peak velocities and movement times at the new gain in the two testing directions as predicted from inertial anisotropy. The peak velocities were scaled but systematically lower in the high forearm inertia direction (135°) (Fig. 4*B*). However, the movement times in the two directions were inversely related to the peak velocities (Fig. 4*C*). Thus a gain relation learned with a single-target distance leads to the acquisition of a scaling rule that generalizes across distances and across two inertial configurations.

To verify that gain learning generalizes across all directions we examined learning of the gain change to a single-target distance and tested multiple directions arrayed in a circle. There was no significant difference in the percent adaptation by two-factor ANOVA for testing direction [$F_{(4,164)} = 2.4$; $p = 0.14$] or training direction [$F_{(2,164)} = 0.74$; $p = 0.47$] despite some falloff from the trained target, as illustrated in Figure 5. Thus, learning the gain in a single direction generalizes to all directions.

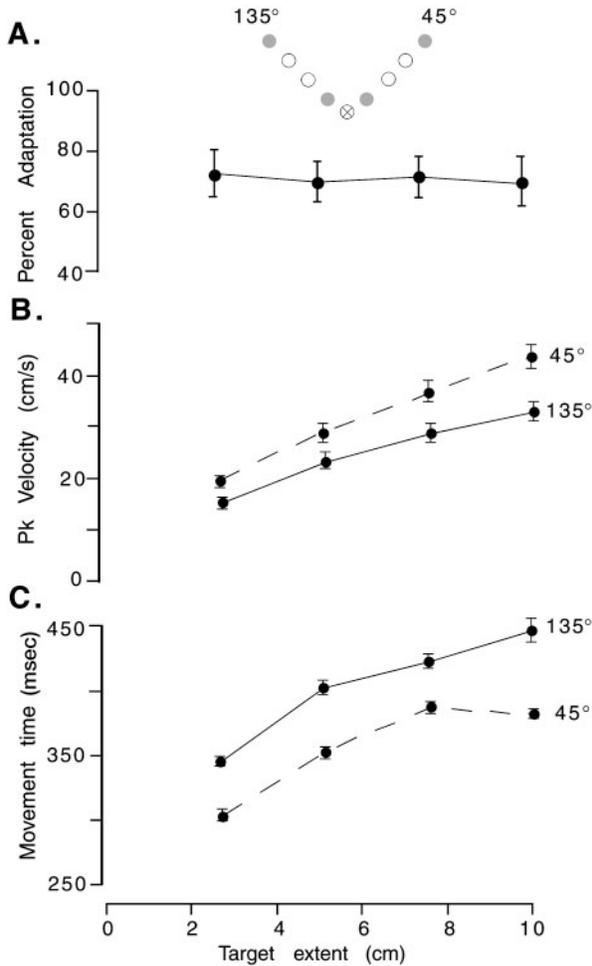


Figure 4. Gain generalization across multiple target distances. *A, Bottom.* The plot is of mean (\pm SEM) group data showing the percent adaptation to untrained target distances relative to adaptation to the training targets. The data are collapsed for the four different training targets. *Top.* The four different training targets (circles) are shown in gray, and the testing targets are in white. On any given training day only one of the gray targets was trained to, and the remaining seven targets were used for testing. *B.* Mean peak velocity for the untrained targets is plotted against target distance in the two testing directions. *C.* Mean movement time for the untrained targets is plotted against target distance in the two testing directions.

We examined generalization of gain adaptation across the work space in two subjects using the two arm configurations described in Materials and Methods for rotation generalization. There was no significant difference in the adapted peak velocities for the two configurations ($p = 0.71$). Therefore, gain generalizes across the work space as well as across distance and direction.

Generalization of rotation adaptation

Generalization across directions

As noted previously the lower rate of adaptation to rotation when the number of target directions increased suggested that learning was restricted to the direction of the target toward which movement had been made. We tested this hypothesis by examining directional errors to targets in untrained directions without visual feedback after training with one, two, four, or eight targets. See Figure 6*B, top*, for training target arrays. As predicted by the time course data, we found that training in a single direction led to only local learning of the rotation. It should be noted that the degree of adaptation at the end of the single-target rotation training block (81%) (Fig. 3*A*) was not significantly different from that achieved at the end of the single-target gain training block (89%) (Fig. 2*A*).

The percent adaptation falls off very steeply as the test target direction deviates from the training target direction. The same

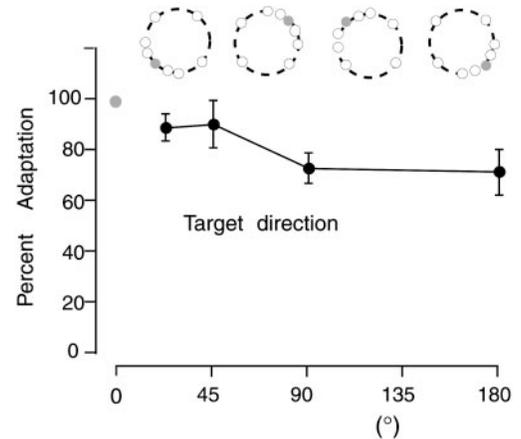


Figure 5. Gain generalization across multiple directions after training in a single direction. *Bottom.* The plot is of mean (\pm SEM) group data showing the percent adaptation to untrained directions relative to the training target. *Top.* The gray targets show the four different target directions for 4 different training days. The testing targets are in white.

falloff in adaptation was found with four different training directions and was symmetric around the training direction (Fig. 6*A*). This was confirmed by a two-factor ANOVA showing no significant effect for training target direction [$F_{(3,106)} = 1.2$; $p = 0.36$] and CW/CCW testing directions [$F_{(3,106)} = 1.2$; $p = 0.31$]. We thus combined data from the four training target directions and collapsed CW/CCW testing directions of equal magnitude in subsequent analyses.

As we increased the number of training directions, generalization increased with full generalization at eight targets (Fig. 6*B*). The effect of the number of training directions was highly significant by a single-factor ANOVA [$F_{(3,389)} = 59.9$; $p < 0.0001$]. This cannot be explained by differences in performance to the training target(s) during testing for the four groups [single-factor ANOVA, $F_{(3,19)} = 0.69$; $p = 0.57$]. Thus, as the number of training directions increases, there is increased generalization to untrained directions.

Generalization across distances

The finding that the learning of the rotation in a single direction does not generalize across directions raises the possibility that it is only a unique stimulus–response relationship that is being learned, i.e., a unique response to a particular target distance and direction. This, however, was not the case. We again trained subjects to a single target but tested them to three other target distances in the same direction. The performance was uniform across distances with no significant difference in directional error [single factor ANOVA, $F_{(3,400)} = 0.79$; $p = 0.50$] (Fig. 7).

Thus, in contrast to gain learning, the learning of a rotation with a single-target direction generalizes across distance but not direction. Generalization across directions requires the sampling of directional errors across multiple directions. It is possible that the improving performance in untrained directions as the number of training targets increases is caused by interpolation of local learning.

We performed one final experiment on a separate group of six subjects to address the issue of interpolation further. Subjects were trained with only two targets separated by 45°. They reached approximately the same level of performance to these targets as the one-, two-, four-, and eight-target groups did to their training target(s). We then tested them in the absence of feedback to a single target interposed between the training targets, i.e., separated by 22.5° as in the eight-target group. If subjects acquired the untrained direction by interpolation then we would have expected complete generalization to the interposed untrained target as in the

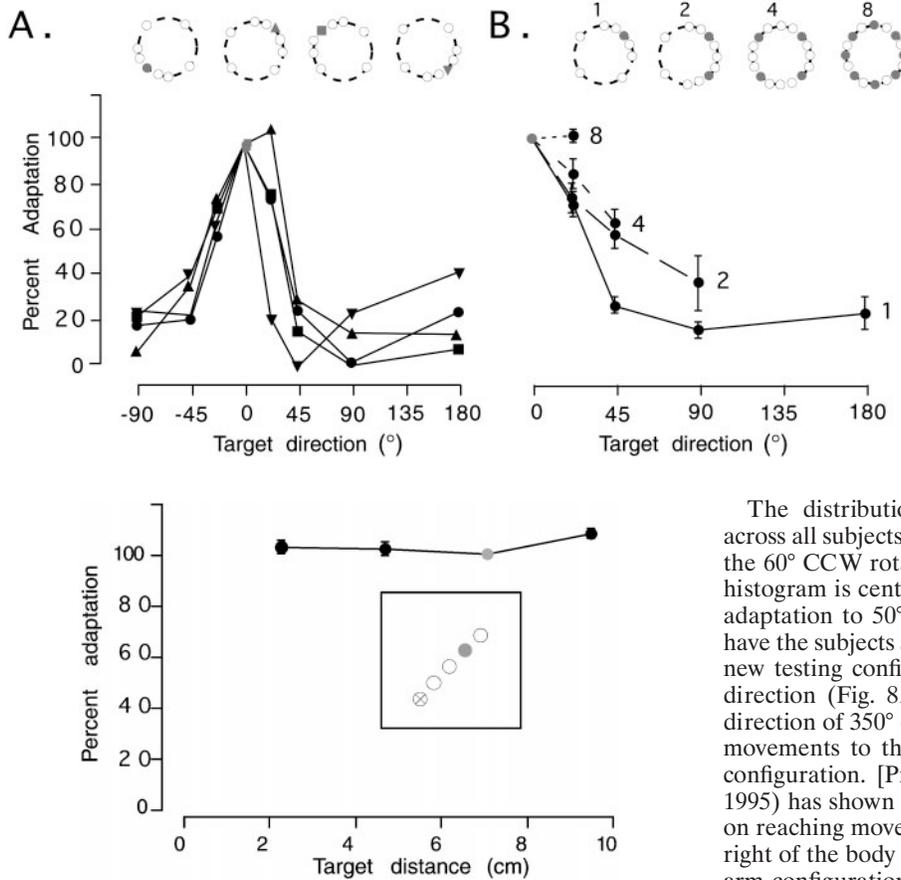


Figure 7. Rotation generalization across multiple target distances after training to a single distance of 7.2 cm. The plot shows mean (\pm SEM) group data of the percent adaptation to untrained distances 2.4, 4.8, and 9.6 cm relative to the training distance. *Inset*, The target array is shown with the training target in gray.

eight-target case. Instead, however, we saw a mean falloff in adaptation at the 22.5° direction, not significantly different from that seen in the four-target group (83.2 ± 8 vs $83.8 \pm 7\%$). This result suggests that a successful interpolation rule depends on a richer degree of information outside the interpolated area, as is afforded by the eight-target training.

Generalization across the work space

The demonstration that in adapting to rotation with targets in a single direction subjects learn a rule that can be applied across several target distances raises the question of whether this rule is learned in extrinsic or in intrinsic space. The rule could represent scaling of a learned joint profile with a linear synergy of shoulder and elbow muscle contractions (Bock, 1994; Gottlieb et al., 1997) or learning of a new directional axis in extrinsic space. The local learning of rotation across directions could be interpreted as supporting the possibility of learning in joint space. This is because a particular shoulder–elbow synergy would only be expected to apply over a narrow directional range. In a recent paper demonstrating limited generalization of a visuomotor transformation, the authors comment that the decay could have been in extrinsic or in joint-based coordinates (Ghahramani et al., 1996).

To answer this question we trained subjects on a 60° CCW rotation with a single target in one arm configuration and then tested them in another arm configuration but in the same target direction. If subjects learned a spatial axis in extrinsic space they would remain accurate in the new arm configuration. However, if they learned the rotation in joint space they would no longer be accurate to the target direction in the new arm configuration because the joint angle changes would no longer be appropriate.

Figure 6. Rotation generalization. *A*, Generalization across multiple directions after training in a single direction. The directional data are relative to the training target. *Bottom*, The plot is of mean (\pm SEM) group data showing the percent adaptation to untrained directions relative to the training target. *Top*, The four different training directions (45, 135, 225, and 315°) for 4 different days are shown by the gray symbols. The positioning of the testing targets (in white) is shown. *B*, Generalization across multiple directions after training in one, two, four, and eight directions. *Bottom*, The plot is of mean (\pm SEM) group data showing the relative percent adaptation in the untrained directions relative to the trained directions. When there was more than one training target, the mean performance to all the training targets was used to calculate the relative adaptation in untrained directions. Data were collapsed for clockwise and counterclockwise directions. *Top*, Training targets are shown in gray, and testing targets are in white.

The distribution of movement directions for all movements across all subjects shows that subjects adapted almost completely to the 60° CCW rotation in the training configuration (Fig. 8C). The histogram is centered on a movement direction of 355°, indicating adaptation to 50° of the imposed rotation (full adaptation would have the subjects at 345° on the tablet). When subjects shifted to the new testing configuration they remained accurate in the training direction (Fig. 8B). The histogram is centered on a movement direction of 350° (Fig. 8D). This represents a shift of 5° CW for all movements to the training target as compared with the training configuration. [Previous work in our laboratory (Ghilardi et al., 1995) has shown that there is a systematic clockwise bias imposed on reaching movements when the hand is displaced laterally to the right of the body midline. Indeed, in this previous study one of the arm configurations studied was the same as our testing configuration in the current experiment, and the mean bias was $-5.1 \pm 1.3^\circ$. When we subtract this anticipated bias from the movement directions in the training configuration and compare the resulting expected movement directions with actual directions in the testing configuration, there is no significant difference (unpaired *t* test, $p = 0.12$). Thus, the direction of movement in the testing configuration is the same as in the training configuration.] Because the elbow angle remained the same in the two configurations and the shoulder was rotated from 45 to 90°, then the anticipated new direction in joint space would be 45° CW of the training direction, i.e., 300° (see Fig. 8B). Thus, the average movement direction across all subjects was distributed around the anticipated movement direction if learning were of the training direction in extrinsic space and not distributed around the anticipated movement direction if learning had occurred in joint space.

Another way of analyzing the data was to calculate the joint angle changes at the shoulder and elbow for each individual movement from the known movement direction and extent and from the subjects' arm segment lengths. Figure 9 shows that the elbow and shoulder angle changes were significantly different for the two arm configurations (mean change in shoulder angle = 14.5°; mean change in elbow angle = -8.6°; both at $p < 0.0001$, unpaired *t* tests). The magnitude of the differences for these angle changes in the two configurations was such that it is not conceivable that the same patterns of joint torques could have generated them.

DISCUSSION

The experiments presented here demonstrate categorical differences in the time course and generalization of adaptation to induced errors in movement extent and direction. These differences suggest that the brain processes errors in extent and direction separately and in computationally distinct ways during learning. Our findings add to the idea that reaching movements are planned as a vector centered on the hand whose extent and direction are established via learning a scaling factor and reference axes.

The preservation of single peaked velocity profiles and straight

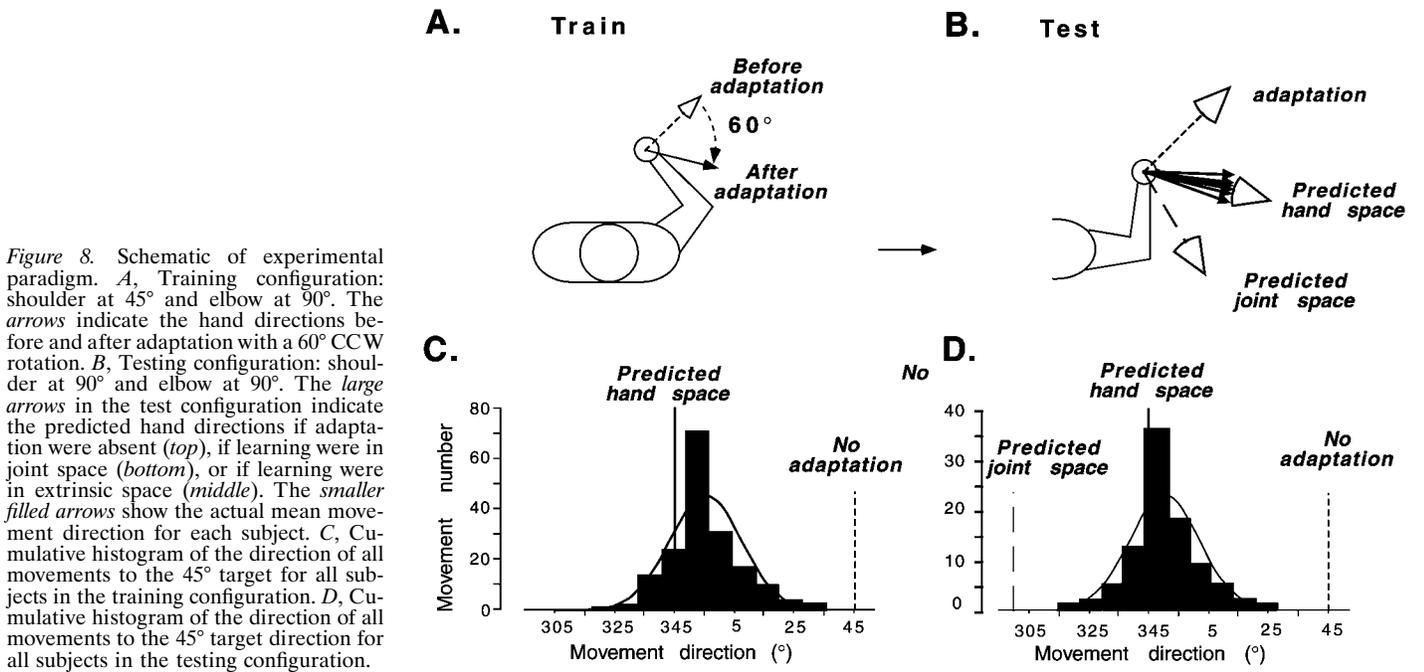


Figure 8. Schematic of experimental paradigm. *A*, Training configuration: shoulder at 45° and elbow at 90°. The arrows indicate the hand directions before and after adaptation with a 60° CCW rotation. *B*, Testing configuration: shoulder at 90° and elbow at 90°. The large arrows in the test configuration indicate the predicted hand directions if adaptation were absent (top), if learning were in joint space (bottom), or if learning were in extrinsic space (middle). The smaller filled arrows show the actual mean movement direction for each subject. *C*, Cumulative histogram of the direction of all movements to the 45° target for all subjects in the training configuration. *D*, Cumulative histogram of the direction of all movements to the 45° target direction for all subjects in the testing configuration.

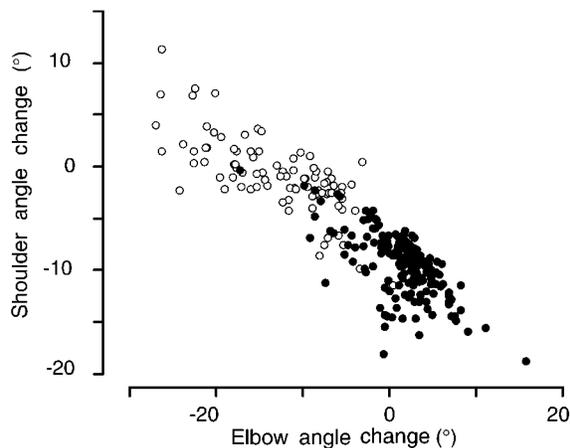


Figure 9. A scatter plot of the elbow versus shoulder angle change for the training and testing configurations for all six subjects. The filled circles represent the training configuration, and the open circles represent the testing configuration.

hand paths during adaptation to gain and rotations suggests that errors are mainly corrected using adaptive or feedforward mechanisms rather than via “on-line” feedback. Extent errors are used to adjust the pulsatile activation of segmental motor neurons accelerating the hand on successive trials and thereby rescaling the spatial mapping of target location to vector magnitude. Directional errors are used to adjust progressively the reference axis used to compute the directional error itself.

Although the time course of adaptation was sometimes fitted adequately with single-exponential functions, a double-exponential fit was typically needed to capture the initial rapid change and the later gradual reduction in mean error. This suggests two processes operating during the course of adaptation. It may be speculated that the initial rapid decline, when the errors are most salient, reflects a strategy intended to reduce visual errors rapidly with each subsequent movement. After errors come within the envelope of movement-to-movement variability, a second more gradual process appears to be implemented, in which successive changes depend on evaluation of errors made on several successive movements.

Consistent with the fact that the magnitudes of the gain and rotation perturbations were selected to produce equal linear errors,

the time course of adaptation for a single target was the same. However, the rate of adaptation was influenced differently when multiple targets were presented. Adaptation to the gain change occurred at the same rate for multiple target distances compared with one. With rotation, increasing target directions from one to eight produced proportional reductions in the rate of adaptation so that the time course of adaptation to any single target remained identical. This suggests that directional error information is being stored for the eight targets separately rather than being used to form a reference frame that might allow error information from one direction to be used to correct errors in another direction. This issue is discussed further below.

A prediction generated by the different effects of increased target numbers on the rate of adaptation is that the learning of a gain should generalize to targets at previously unvisited distances whereas learning rotation with a single target should generalize poorly to novel target directions. This was indeed what we found with complete gain generalization and only local learning of a rotation.

When two learning processes obey different rules for generalization, it suggests that the two processes are represented differently in the brain. Our results for gain generalization via scaling of the peak velocity across directions and distances are largely consistent with previous findings using movement end points (Bock, 1992; Bock and Burghoff, 1997) and a study of amplitude generalization in a velocity-dependent force field (Goodbody and Wolpert, 1998). The complete generalization seen for gain suggests that the scalar parameter, relating target distance to the amplitude of an activation profile, is explicitly estimated after adapting to a single target and applied to the whole work space. This may occur because the relationship between peak velocity and target distance is approximately linear throughout the work space under normal conditions. Studies of prism adaptation (Bedford, 1993) and vertical phoria adaptation (Schor et al., 1993) have also shown that generalized mappings are learned preferentially over isolated input–output relationships when linear interpolation or extrapolation is possible. We expect that subjects would find it difficult to learn two separate gains in two directions, distances, or configurations. However, specific experiments would have to be done to determine this.

In contrast, in the case of learning the rotation, the parameter, i.e., the angle of rotation, cannot be estimated from learning a single input–output pair. Computationally, the problem is “ill-posed” and requires function approximation, for example, radial

basis functions constrained for smoothness or, equivalently, multiple-layer perceptions (Poggio and Gorosi, 1990). These computational models yield intermediate degrees of generalization and have been found to correspond reasonably well with experimental results (Imamizu et al., 1995; Ghahramani et al., 1996). There is also psychophysical evidence to suggest that there are unique processing constraints for visuomotor and mental rotations as compared with other cognitive tasks (Pellizzer and Georgopoulos, 1993).

We found that performance to untrained directions was the same as that to trained directions only with eight equally spaced training directions. There was, however, a statistically significant improvement in the untrained directions as the number of training targets increased. This suggests that there is increasingly successful interpolation of local learning as more directions are sampled. Our time course data, showing no enhancement in the rotation learning rate with eight targets compared with one, suggest storage of errors and learning of eight local rules in separate working-memory buffers. This raises the possibility that the interpolation rule is not in effect early in learning (we could only compare training for the first 18 visits to a target) and may only begin to be established later. In addition, we have shown that even when adapting to a rotation with a single target, subjects are able to generalize across distances and arm configurations. This argues against tabular learning and more for learning of a directional vector. This vector can then be multiplied by a scalar for generalization across distances and also be translated across the work space.

Thus, we posit that whereas a new global gain parameter can be estimated from a single input–output pair, rotation parameters can only be estimated locally and that a full reference frame rotation is achieved by interpolation. These results strongly suggest that movements are planned as a vector with independent specification of extent and direction and that this is because of differences in the computational constraints for learning scaling factors and reference axes in extrinsic space. This clear separation of adaptation to extent and directional errors would not be expected if trajectory planning occurred in joint space. This separation between directional and scaling specification is consistent with what is known from single-unit studies demonstrating populations of neurons with preferred directions of movement (Georgopoulos et al., 1982) but with speed acting as a gain factor on the directional tuning curve (Moran and Schwartz, 1999). Rescaling movements in a given direction would involve up or down modulation of the activity of the same neuronal population, whereas a new directional axis would require either a new pattern of activity or a new population of neurons altogether.

We did not directly address the origin of the extrinsic reference frame, but previous work in our laboratory suggests that trajectory errors are computed relative to the hand rather than the shoulder or body midline (Gordon et al., 1994a). This conclusion is supported by our result showing that the learning of a rotation remained invariant around the hand despite a 45° rotation around the shoulder. However, a hand-centered reference frame for trajectory specification is compatible with concomitant specification of initial hand position and target position in an extrinsic egocentric reference frame. Indeed, we would argue that the hand-based reference frame is rotated relative to the egocentric reference frame.

The generalization of gain learning across directions is of interest in the context of inertial anisotropy. Examination of mean peak velocities and movement times across two testing directions revealed that they differed in a manner anticipated from previous work on inertial anisotropy (Gordon et al., 1994a). In the high inertia direction (135°), the movement times were significantly longer, and the peak velocities were significantly lower when compared with the low inertia direction (45°; Fig. 4). However, despite differences in these kinematic planning variables, adaptation was the same in the two testing directions. A recent model has shown that direction-dependent variations in movement time, which compensate for inertial anisotropy, can be attributed to intrinsic plant

properties and segmental feedback (E. Todorov, personal communication). This means that for gain to generalize across directions the CNS must have a model of anisotropic effects. This conclusion is similar to that of Sabes et al. (1998), but although they argue that this is only true for certain cases such as obstacle avoidance, our result suggests that dynamics are taken into account in all point-to-point movements. Vectorial planning could not specify movement extent accurately without an internal dynamic model. In this sense, dynamics are taken into account so that only kinematic variables need to be specified in the planning process. This independence of acquisition of a rescaling rule from inertial anisotropy is consistent with our recent demonstration that learning a screen cursor rotation is independent of learning novel inertial dynamics (Krakauer et al., 1999a).

In conclusion, the data suggest that accuracy in reaching movements is achieved by using errors in extent and direction to update adaptively a vectorial representation of intended movement in extrinsic space.

REFERENCES

- Bedford F (1993) Perceptual and cognitive spatial learning. *J Exp Psychol Hum Percept Perform* 49:517–530.
- Bock O (1992) Adaptation of aimed arm movements to sensorimotor discordance: evidence for direction-independent gain control. *Behav Brain Res* 51:41–50.
- Bock O (1994) Scaling of joint torque during planar arm movements. *Exp Brain Res* 101:346–352.
- Bock O, Burghoff M (1997) Visuo-motor adaptation: evidence for a distributed amplitude control system. *Behav Brain Res* 89:267–273.
- Carrozzo M, McIntyre J, Zago M, Lacquaniti F (1999) Viewer-centered and body-centered frames of reference in direct visuomotor transformations. *Exp Brain Res* 129:201–210.
- Cunningham HA (1989) Aiming error under transformed spatial mappings suggests a structure for visual-motor maps. *J Exp Psychol Hum Percept Perform* 15:493–506.
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2:1527–1537.
- Ghahramani Z, Wolpert DM (1997) Modular decomposition in visuomotor learning. *Nature* 386:392–395.
- Ghahramani Z, Wolpert DM, Jordan MI (1996) Generalization to local remappings of the visuomotor coordinate transformation. *J Neurosci* 16:7085–7096.
- Ghez C, Krakauer J (2000) The organization of movement. In: *Principles of neural science* (Kandel ER, Schwartz JH, Jessell TM, eds), pp 653–673. New York: McGraw-Hill.
- Ghez C, Krakauer JW, Sainburg R, Ghilardi MF (2000) Spatial representations and internal models of limb dynamics in motor learning. In: *The new cognitive neurosciences* (Gazzaniga MS, ed), pp 501–514. Cambridge, MA: MIT.
- Ghilardi MF, Gordon J, Ghez C (1995) Learning a visuomotor transformation in a local area of work space produces directional biases in other areas. *J Neurophysiol* 73:2535–2539.
- Goodbody SJ, Wolpert DM (1998) Temporal and amplitude generalization in motor learning. *J Neurophysiol* 79:1825–1838.
- Gordon J, Ghez C (1987) Trajectory control in targeted force impulses. II. Pulse height control. *Exp Brain Res* 67:241–252.
- Gordon J, Ghilardi MF, Ghez C (1994a) Accuracy of planar reaching movements. I. Independence of direction and extent variability. *Exp Brain Res* 99:97–111.
- Gordon J, Ghilardi MF, Cooper SE, Ghez C (1994b) Accuracy of planar reaching movements. II. Systematic extent errors resulting from inertial anisotropy. *Exp Brain Res* 99:112–130.
- Gottlieb GL, Song Q, Almeida GL, Hong DA, Corcos D (1997) Directional control of planar human arm movement. *J Neurophysiol* 78:2985–2998.
- Imamizu H, Uno Y, Kawato M (1995) Internal representations of the motor apparatus: implications from generalization in visuomotor learning. *J Exp Psychol Hum Percept Perform* 21:1174–1198.
- Krakauer JW, Pine ZM, Ghez C (1996) Differences in generalization of adaptation to altered gains and display rotations in reaching movements. *Soc Neurosci Abstr* 22:899.
- Krakauer JW, Pine ZM, Huang H, Ghez C (1997) Learning a rotated reference frame requires moving in multiple directions and differs with joint system used. *Soc Neurosci Abstr* 23:202.
- Krakauer JW, Ghilardi MF, Ghez C (1999a) Independent learning of internal models for kinematic and dynamic control of reaching. *Nat Neurosci* 2:1026–1031.

- Krakauer JW, Pine ZM, Ghez C (1999b) Visuomotor transformations for reaching are learned in extrinsic coordinates. *Soc Neurosci Abstr* 25:2177.
- McIntyre J, Stratta F, Lacquaniti F (1997) Viewer-centered frame of reference for pointing to memorized targets in three-dimensional space. *J Neurophysiol* [Erratum (1998) 79:1135, 3301] 78:1601–1618.
- Messier J, Kalaska JF (1999) Comparison of variability of initial kinematics and endpoints of reaching movements. *Exp Brain Res* 125:139–152.
- Moran DW, Schwartz AB (1999) Motor cortical representation of speed and direction during reaching. *J Neurophysiol* 82:2676–2692.
- Ojakangas CL, Ebner TJ (1991) Scaling of the metrics of visually-guided arm movements during motor learning in primates. *Exp Brain Res* 85:314–323.
- Paillard J (1996) Fast and slow feedback loops for the visual correction of spatial errors in a pointing task: a reappraisal. *Can J Physiol Pharmacol* 74:401–417.
- Pellizzer G, Georgopoulos AP (1993) Common processing constraints for visuomotor and visual mental rotations. *Exp Brain Res* 93:165–172.
- Pine ZM, Krakauer J, Gordon J, Ghez C (1996) Learning of scaling factors and reference axes for reaching movements. *NeuroReport* 7:2357–2361.
- Poggio T, Gorosi F (1990) Regularization algorithms for learning that are equivalent to multilayer networks. *Science* 247:978–982.
- Roby-Brami A, Burnod Y (1995) Learning a new visuomotor transformation: error correction and generalization. *Brain Res Cogn Brain Res* 2:229–242.
- Sabes PN, Jordan MI, Wolpert DM (1998) The role of inertial sensitivity in motor planning. *J Neurosci* 18:5948–5957.
- Schor C, Gleason G, Maxwell J, Lunn R (1993) Spatial aspects of vertical phoria adaptation. *Vision Res* 33:73–84.
- Soechting JF, Flanders M (1989) Errors in pointing are due to approximations in sensorimotor transformations. *J Neurophysiol* 62:595–608.
- Vindras P, Viviani P (1998) Frames of reference and control parameters in visuomanual pointing. *J Exp Psychol Hum Percept Perform* 24:1–23.
- Vindras P, Desmurget M, Prablanc C, Viviani P (1998) Pointing errors reflect biases in the perception of the initial hand position. *J Neurophysiol* 79:3290–3294.