

Generalization of Object Manipulation Skills Learned without Limb Motion

Christopher D. Mah^{1,3} and Ferdinando A. Mussa-Ivaldi^{1,2,3}

Departments of ¹Physical Medicine and Rehabilitation and ²Physiology, Northwestern University Medical School, and ³Sensory Motor Performance Program, Rehabilitation Institute of Chicago, Chicago, Illinois 60611

Recent work suggests that human subjects may learn mappings between object motion and exerted torque during manipulation of freely pivoting or unstable objects. In the present work, we studied an object manipulation task involving no arm movement to determine how subjects internally represent the force–motion relationship of an object during a skilled manipulation task. Human subjects learned to balance a simulated inverted pendulum. The simulation was controlled by pressing on a fixed force sensor, and applied forces resulted in motion of the simulated pendulum on a computer screen according to its equation of motion. Each subject initially learned the task in one arm posture and was tested 1 d later in a new arm posture. In one test condition, the effects of arm torque were matched to the original task, and in the other test condition, the simulation was unchanged. The pattern of skill transfer to different arm postures suggested that subjects had learned joint torque responses rather than a general model of the object interface forces. A second experiment showed that the advantage of training with matched arm torques was object specific, because torque-matched training on a tracking task involving similar forces was not a substitute for training in the balancing task.

Key words: object manipulation; isometric force; internal model; motor skills; balancing; human

Introduction

Object manipulation is central to many of the activities of everyday life. Often, during tool usage and other common examples of object manipulation, objects are firmly grasped, so that the object moves with the arm. Thus, when humans move objects, they typically make desired limb movements while producing the systematic force (Shadmehr and Mussa-Ivaldi, 1994) necessary to move the object. During motor adaptation, which involves a moving limb, generalization experiments show that knowledge of the applied force is represented as a pattern of arm torque (Gandolfo et al., 1996; Malfait et al., 2002). In this work, we investigated how a motor skill that involved a relationship between force and object motion would generalize to new arm configurations when the object did not move with the arm.

During control of unstable or lightly grasped objects, an input force may have a different effect depending on what state the object is in. Because object properties may vary greatly from object to object, efficient object control seems likely to depend on knowledge of the specific dynamic properties of each manipulated object (Dingwell et al., 2002; Mah and Mussa-Ivaldi, 2003). Thus, effective object control may require an internal model of how the object is expected to move in response to input forces or exerted arm torques (Wolpert and Miall, 1996). In the present

work, isometric forces were used as input to a simulation of the object. Thus, we studied an object manipulation task for a freely pivoting unstable virtual object (the inverted pendulum) that was constructed so that the task did not involve arm movement. The results confirmed the hypothesis that humans can form an internal model of unstable, pivoting, ungrasped objects without arm motion. In addition, we found that subjects learned patterns of joint torques and not interface forces.

We found that, after learning to manipulate the virtual object in one position, the subjects could not effectively manipulate the same object with the arm in a different configuration. This suggests that the subjects did not construct an object representation on the basis of the relationship between object visual motion and interface forces. However, subjects were able to retain much of the learned skill when the virtual object in the new arm configuration was programmed to respond to the same pattern of joint torques—but different end-point forces—generated by the subject during the training in the initial position. This pattern of generalization suggests that subjects learned to match the observed object motion with a motor command to the arm muscles rather than with a force to be impressed by the hand at the interface with the object. These results extend the finding that learning of predictable forces applied to the hand is represented in intrinsic limb coordinates (Gandolfo et al., 1996; Shadmehr and Mousavi, 2000) to the representation of motor skills involving object motion.

Materials and Methods

Nine subjects (five males, mean \pm SD age, 31.6 \pm 6.19 years; four females, mean \pm SD age, 32.5 \pm 7.85 years) participated in experiment 1 after giving informed consent according to the standards of the Institu-

Received Aug. 19, 2002; revised April 15, 2003; accepted April 21, 2003.

This work was supported by National Science Foundation Grant BES 9900684 and National Institute of Neurological Disorders and Stroke Grant NS35673. We thank R. Haner, M. Knikou, S. Li, J. L. Patton, and N. Suresh for comments on this manuscript.

Correspondence should be addressed to Dr. Christopher D. Mah, Sensory Motor Performance Program, Rehabilitation Institute of Chicago, 345 East Superior Street, #1406, Chicago, IL 60611. E-mail: c-mah@northwestern.edu.
Copyright © 2003 Society for Neuroscience 0270-6474/03/234821-05\$15.00/0

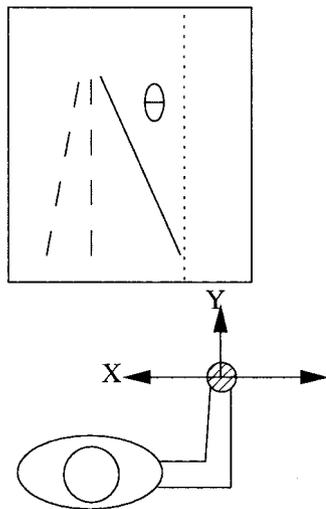


Figure 1. Experimental setup. Subjects controlled a virtual cart–pole system by pressing on a force sensor in either the X or Y directions. Arm segments were at an angle of ~ 0 or 90° with respect to the body. A leftward inclination of virtual pole from the vertical (dotted) corresponded to a positive angle θ . The force drove the entire cart–pole system to the left or right across the screen and caused corresponding changes in the angle of inclination of the pole as the pole pivoted about its base (dashed, solid transition).

tional Review Board of Northwestern University. Seven of the nine subjects were right-handed, and two were left-handed by self-report. With their dominant hands, subjects held the handle of a force sensor clamped to a fixed base. Input from the forces exerted on the force sensor controlled the position of the cart in a one-dimensional virtual cart–pole software simulation (Barto et al., 1983; Mehta and Schaal, 2002). The setup is illustrated schematically in Figure 1.

Subjects were told to maintain the pole as upright as possible near the center of the display. Trials lasted a maximum of 18 sec. The mass of the 3-m-long pole was 5 kg, and the mass of the cart was 0.15 kg. Virtual forces always pushed the base of the pole toward the right- or left-hand side of the display. However, the simulation responded only to the component of force exerted on the force sensor along one fixed axis called the control axis. The control axis was either left/right relative to the subject (X) or toward/away (Y) from the subject. The simulation was different from a physical object in the following ways. Exerted force vectors $>20^\circ$ from the control axis were set to zero and ignored by the simulation. The pole fell (rotated) one-half as fast as it would have fallen in a physically accurate simulation, and the displayed angle of inclination was 20 times the physical angle. If the cart–pole object left the computer screen on the right ($X = 22.5$ cm), it reappeared on the left ($X = -22.5$ cm) and vice versa. These modifications were introduced to adjust the difficulty of the task and to ensure that the vector of exerted forces was approximately aligned with the control axis. The simulation was not modified across or within subjects aside from rotating the control axis.

The pole simulation was displayed as three closely spaced vertical red lines that appeared initially at the center of the computer screen. A small amount of noise was added to the input force to ensure that the pole would fall if the subject did not act. Motion caused by the added noise was imperceptible to the subjects. Subjects practiced for two blocks of 100 trials each on day 1, using one of the control axes, and returned the next day for a second session. In the second session, the arm was placed in a new configuration, and subjects performed the task for 100 trials with X as the control axis and 100 trials with Y as the control axis. Figure 2 depicts the approximate arm configurations and control axes chosen for each of the nine subjects. On the second day, the combination of arm configuration and control axis was chosen to produce two possible conditions. In the matched force (F) condition, the control axis was the same as the one used the previous day, but with the arm in a new posture. In the matched torque (T) condition, both the arm posture and control axes were altered so that if the elbow angle was 90° on days 1 and 2, and exerted

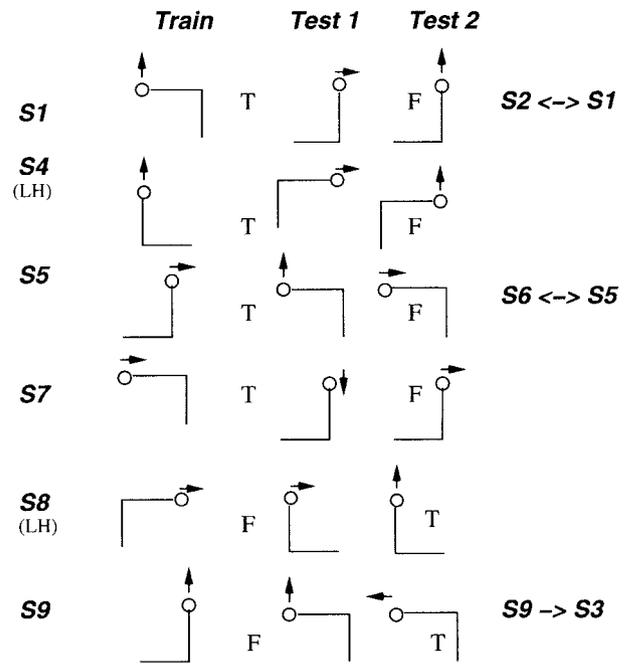


Figure 2. Experimental protocol. Arm configurations of experimental subjects (S1–S9) during training on day 1 (first column) and during testing on day 2 (second and third columns). The position of the hand is indicated by a circle. Arrows indicate the direction of exerted force that would move the cart in the cart–pole simulation to the right. T, Matched torque between training and test; F, matched force between training and test; LH, left-handed subjects. Subjects S1 and S2 received analogous test sequences with reversed test orders, as did subjects S5 and S6. Subject S9 had an inverted relationship between exerted force and screen movement at test 2. Subject S3 received the same sequence as S9 but without the inverted relationship seen by subject S9.

forces were aligned with the control axes, the same arm torques exerted on days 1 and 2 would result in the same motions of the simulated cart–pole object. This equivalence follows, because the changes in the arm posture we chose were expected to rotate the end-point forces resulting from fixed shoulder and elbow torques by $\sim 90^\circ$. There were a total of eight possible training and test orders, and eight subjects were randomly assigned to these test orders. In the matched torque condition of one of the protocols, the subject (S9) had to push to the right to make the object move to the left. To assess the effect of this reversal, the protocol of another subject (S3) replicated this condition with a force-to-screen relationship in which force to the right gave object motion to the right. Both subjects showed similar learning effects, and therefore, the force-to-screen relationship did not seem to be important. Because arm torques used by S3 during the test condition were opposite to those used in training, however, this subject's data were not included in the T condition and were not analyzed further. Four of these eight subjects were tested in the matched torque condition first on the second day, and four subjects were tested first in the F condition.

The subject's elbows were supported in the plane of the handle by a small platform. The subjects were directed to keep their wrists straight and to maintain contact with the handle during the task. It was not possible to constrain exerted force directions more accurately, because more narrow angular windows ($<20^\circ$) interfered with task performance. In addition, elbow angles varied between 90 and 120° because of limitations in placing subjects comfortably relative to the apparatus in three dimensions. Thus, the matched torque condition was approximate. However, it is important to note that if subjects had learned to exert Cartesian forces relative to the apparatus, arm posture would not have mattered at all.

To measure performance, we computed the SD of the pole angle within each trial over time (S). Perfect balancing performance would have resulted in a constant, unchanging pole angle ($S = 0$), whereas because of the definition of the SD, larger SDs occurred when the pole

angle varied from instant to instant during the trial. Such variation (large *S*) was typical of unstable balancing or unskilled performance. The SDs (*S*) reported here reflect the inclination angle of the pole displayed on the screen, in degrees. We also analyzed the time to failure (TTF) on each trial and the SD of the position (SDX) of the cart–pole object on the screen.

In experiment 1, balancing performance in a test condition with arm torques matched to training was compared with performance in a test condition for which arm torques were not matched to those used in training. It appeared possible, however, that training in the balancing task might enhance the subject’s ability to produce any desired force in the control axis direction. Therefore, experiment 2 examined whether arm-torque specificity of the balancing skill was just an enhancement of the ability to produce forces in a given direction. Because controlling the virtual object required the subject to produce complicated time-varying forces, it seemed unlikely that calibration with static forces would replace training with the virtual object. In experiment 2, subjects practiced a tracking task (which involved similar time-varying forces but different visual cues and task requirements), and we examined whether this would produce the same effect as balancing practice. We refer to the hypothesis that subjects only acquired an enhanced ability to produce time-varying forces in one direction as the dynamic calibration hypothesis.

Nine subjects (three male, mean ± SD age, 28 ± 7 years; six female, mean ± SD age, 23.7 ± 3.4 years) participated in experiment 2. All of the subjects were right-handed by self-report. They were always tested in the arm posture shown for the training of S9 in Figure 2 on days 1 and 2, although force directions were varied. The pole-balancing task was as described in experiment 1. In the tracking task, isometric force on the handle drove a cart–pole object of the same mass as in experiment 1. However, the pole was locked upright, and no balancing was required. Subjects were asked to follow a target that undulated slowly across the screen according to a randomly chosen 15 term Fourier series [root mean square (RMS) speed, 0.069–0.072 m/sec] (Mah and Mussa-Ivaldi, 2003). Because application of a steady force would have caused the cart–pole object to accelerate indefinitely, this was not a simple force tracking task.

On day 1, subjects practiced the balancing skill for 100 trials with one control axis and practiced the tracking task along the other control axis for 100 trials. Training task orders were balanced. On day 2, subjects were asked to perform the balancing task for 100 trials along each control axis. There were eight possible combinations of force direction, training order, and test order. Subjects were randomly assigned to these combinations with the constraint that every test order and every training order should occur about equally often. Because *n* = 9 is odd, equality was not possible. The balancing task occurred first in the training of five subjects and second in the training of the other four subjects. On day 2, the pretrained balancing direction was tested first for four of the subjects and was tested second for the remaining five subjects. The analysis assessed any order effects.

Results

Figure 3 shows the time course of the displayed pole angle during selected balancing trials for subject S7 (experiment 1). This figure shows that a well trained subject could balance the pole effectively in one condition but exhibit growing oscillations followed by dramatic losses of control in other conditions. This subject was typical, and the pattern of results shown here was the same as the one we verified in the statistical analysis of all of the subjects. The top two panels show the first and the last three trials of practice on day 1, the bottom left panel shows the first three trials of the matched torque condition, and the bottom right panel shows the first three trials of the matched force condition. It is seen that the subject had difficulty balancing in the first three trials (top left; mean ± SD *S*, 60.6 ± 10.8) but showed stable balancing performance at the end of training on day 1 (top right; mean ± SD *S*, 16.6 ± 4.0). On day 2, initial trials in the matched torque condition were more stable (mean ± SD *S*, 8.0 ± 3.83) than in the matched force condition (mean ± SD *S*, 59.63 ± 12.12), indicating more effective balancing in the matched torque condition.

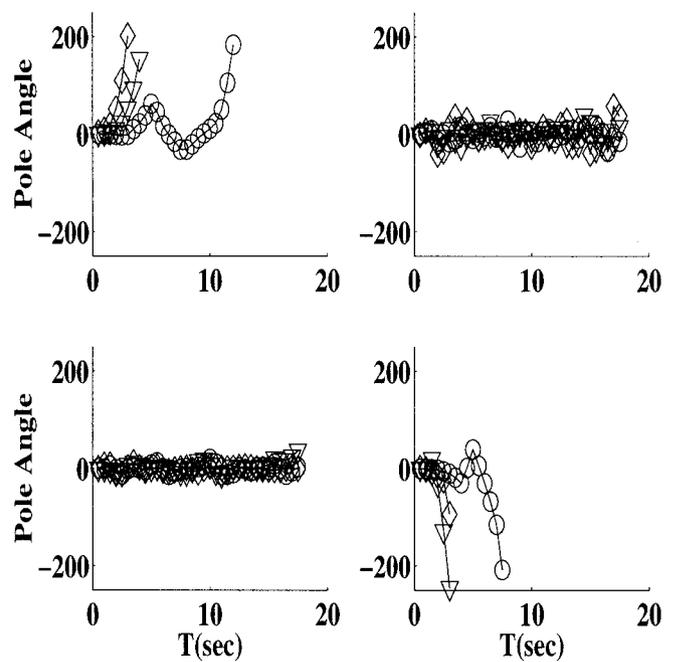


Figure 3. Balancing performance for a single subject. Time course of the displayed pole angle during balancing trials for subject S7. Each panel shows performance on three individual trials. ○, First trial; ◇, second trial; and ▽, third trial. The top two panels show the first (top left panel) and the last (top right panel) three trials on day 1. The first three trials in the matched torque condition on day 2 are shown in the bottom left panel, and the first three trials in the matched force condition on day 2 are shown in the bottom right panel. The angle is plotted every 0.5 sec over the trial. The displayed pole angle was 20 times the true angle. Trials ended when the angle exceeded ±360°, corresponding to a true angle of 18°. T, Time.

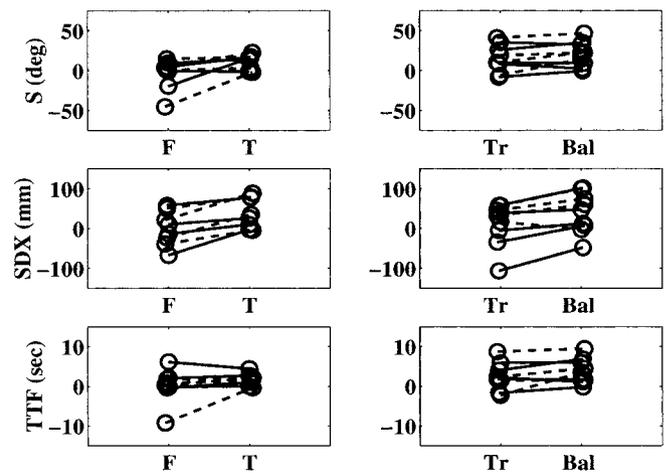


Figure 4. Individual balancing performance in the test conditions. Individual differences between performance measures (*S*, *SDX*, and *TTF*) in the test conditions and in the naive condition. The *S* and *SDX* panels show naive minus test, and the *TTF* panel shows test minus naive. Left panels show results for experiment 1 and right panels show results for experiment 2. The top panels show the results for *S*, the middle panels show *SDX*, and the bottom panels show *TTF*. F, Matched force condition (experiment 1); T, matched torque condition; Bal, matched balancing condition (experiment 2); and Tr, matched tracking condition. Subjects whose test order was the reverse of that shown are indicated by a dotted line, and the rest are indicated by solid lines.

The left panels of Figure 4 show individual changes in performance in experiment 1 between the first practice block on day 1 (naive) and the (day 2) F and T conditions with respect to the *S*, *SDX*, and *TTF* measures. The right panels of Figure 4 show the corresponding results for previous balancing (Bal) and tracking (Tr) practice in experiment 2. A matched group *t* test of *S* be-

tween the first and second training blocks on day 1 of experiment 1 ($t_{(7)} = 3.20$; $p < 0.016$) showed that performance improved on day 1.

To determine whether practice on day 1 transferred to the T and F conditions on day 2, we performed matched group t tests. We compared performance in the first block of day 1 (naive) with the test blocks. The subtraction $S_{\text{naive}} - S_{\text{F}}$ gave a nonsignificant mean difference of 4.11, indicating that we could not detect any effect of previous practice in the F condition. The subtraction $S_{\text{naive}} - S_{\text{T}}$ gave a mean difference of 10.75 ($p < 0.016$), indicating a strong effect of previous practice on the T condition on day 2. Finally, the *a priori* comparison $S_{\text{F}} - S_{\text{T}}$ gave a mean difference of 18.18 ($p < 0.04$), indicating that subjects performed better in the matched torque condition on day 2. These differences cannot be attributed to specific arm configurations or test order, because every possibility was used and ANOVA tests for the main effects of arm configuration and of test order were not significant. Thus, transfer of skill to the T condition was greater than transfer of skill to the F condition.

In experiment 2, subjects practiced manipulating an object (either balancing or tracking) in both of the control directions used in testing. An analysis of S in the five blocks of 20 trials in the first balancing practice session showed that subjects improved in performance over the session ($F_{(4,32)} = 8.29$; $p = 0.0001$). To assess whether subjects in experiment 2 had learned the tracking task, we computed the RMS distance between the cart–pole object and the moving target over each block of 20 trials during the training session. The RMS error decreased ($F_{(4,31)} = 7.82$; $p < 0.001$) over trial blocks decreasing from a mean (across subjects) of 16.7 cm (SD = 10.1 cm) in block 1 to 8.39 cm (SD = 8.3 cm) in block 5.

Statistical comparisons for experiment 2 corresponded to those made in experiment 1. We tested differences in performance between the matched Bal condition and the dynamically calibrated or Tr condition as we had tested the difference between F and T conditions in experiment 1. Transfer of training to the Bal test block was assessed by comparing Bal performance with the naive condition via the subtraction $S_{\text{naive}} - S_{\text{Bal}}$. There was a highly significant effect (mean = 20.47; $p < 0.01$). There was also a smaller improvement between the naive balancing block and the Tr test condition: $S_{\text{naive}} - S_{\text{Tr}}$ (mean = 15.08; $p < 0.05$). Thus, in experiment 2 there was a strong previous practice effect in the matched balancing test condition, but it was also possible to detect a previous practice effect in the Tr condition. Because all of the subjects had practiced with the same virtual object and had also performed the tracking task with a similar object, this effect might have been due partially to dynamic calibration, to generic task knowledge, or to common elements between the Bal and Tr objects.

In contrast, any difference between Tr and Bal (described below) only reflected the specific effect of previous matched balancing practice. Unlike experiment 1, subjects showed a tendency to perform more poorly on the last test block regardless of test condition. After correcting for the main effect of test order, which showed a trend ($F_{(1,8)} = 3.49$; $p = 0.10$), the difference between the Tr and Bal conditions $S_{\text{Tr}} - S_{\text{Bal}}$ was significant ($F_{(1,8)} = 6.03$; $p < 0.05$). This difference could not be explained by the dynamic calibration hypothesis. The difference in performance between the Bal and Tr conditions and similar performance differences on SDX and TTF measures (described below) suggested that the dynamic calibration hypothesis did not completely account for the results of experiment 1.

Analyses of SDX and TTF measures bore out and strength-

ened the results for S. Analyses of SDX suggested that SDX was more sensitive than S to the effects of matched practice. In experiment 1, the performance comparison $SDX_{\text{F}} - SDX_{\text{T}}$ was highly significant ($t_{(7)} = 5.24$; $p = 0.0012$), indicating that arm-torque-matched practice benefited performance much more than non-matched practice. The comparison $SDX_{\text{naive}} - SDX_{\text{T}}$ was significant ($t_{(7)} = 2.86$; $p < 0.05$), whereas the comparison of naive with force-matched performance $SDX_{\text{naive}} - SDX_{\text{F}}$ was nonsignificant. These results corresponded closely to those from the analysis of S for experiment 1. In experiment 2, the comparison $SDX_{\text{Tr}} - SDX_{\text{Bal}}$ was highly significant ($t_{(8)} = 3.59$; $p < 0.01$), indicating that matched balancing practice benefited performance much more than dynamic calibration. The comparison of naive SDX with matched balancing practice $SDX_{\text{naive}} - SDX_{\text{Bal}}$ was significant ($t_{(8)} = 2.31$; $p < 0.05$), and the comparison $SDX_{\text{naive}} - SDX_{\text{Tr}}$ was nonsignificant. Thus, no previous practice effect was detected in the dynamically calibrated condition on the SDX measure, unlike the result for S. Although the display of SDX for experiment 2 in Figure 4 (middle right panel) suggested an interaction of test condition and test order, statistical tests of this interaction were not significant.

Analysis of the TTF measure (Fig. 4) gave results similar to the analysis of S, but this measure was less sensitive than S to the effects of previous practice, probably because many trials lasted for the 18 sec maximum time. Therefore, statistical analyses of TTF are not reported.

Discussion

Subjects were able to learn and retain the manipulation of an unstable object in a condition that involved little or no arm motion. The pattern of skill transfer to different postures and control axes suggested that the task was learned as joint torques or muscle commands and not as Cartesian interface forces. In experiment 2, subjects were allowed to practice the task in a nontest direction and given an opportunity to calibrate their arm torques in the test direction. These manipulations were designed to eliminate the arm torque specificity of the balancing skill, if this specificity were completely attributable to static or dynamic calibration. However, the balancing skill continued to show significant arm torque specificity.

Some performance benefits were observed in the Tr condition in experiment 2. Although this suggests that some of the effect in experiment 1 was attributable to dynamic calibration, it is not conclusive evidence for calibration in a strict sense. It is not surprising that performance improves when a relevant motor response is practiced. In a calibration, the subject would be gaining access to a hypothesized internal representation of force that already exists. If access to such a representation could be obtained through calibration, it seems that the results would have been dramatic. Thus, the persistence of any task specificity after 100 trials of dynamic calibration likely implies that skill transfer between tasks is not attributable to a muscle-independent internal representation of force. Together, these results show that some crucial elements of the object-balancing skill were arm torque specific and not attributable to calibration or generic task knowledge.

It is possible that greater amounts of practice over time might change the relative sizes of the observed effects. Because there was a delay of 1 d between initial practice and testing, it is not likely that the results of experiment 1 reflect unique features of short-term motor memory (Brashers-Krug et al., 1996). However, the observation of an order effect in experiment 2 suggests that arm-torque-specific skill components may have been more suscepti-

ble to learning interference. This is consistent with the hypothesis that such skill components may involve response retrieval or selection from a set of competing responses (Rowe et al., 2000; Mah and Mussa-Ivaldi, 2003). Difficulties in response selection might be more evident over a larger number of trials as the subject's experience begins to include a larger number of possible responses.

Several control mechanisms might be involved in learning and performing the balancing task. Subjects must learn which stimuli are relevant and which responses are useful, and they must learn to associate them (Barto et al., 1983). They may also learn an expectation of how an object will move when it is pushed or pulled (forward model), and they may learn how to produce a desired object motion by pushing or pulling the force sensor (inverse model). There is evidence that both types of object model (Mehta and Schaal, 2002; Mah and Mussa-Ivaldi, 2003) are involved in balancing tasks similar to the one we studied. Stimulus–response association models do not explain how the subject learns which stimuli are most relevant or how errors are corrected. The development of forward and inverse object models is likely to facilitate any association-learning mechanism by reducing dependence on exact sensory cues and by providing a means to find the correct response when an error is made.

Some of the results might be interpreted differently depending on which control mechanisms are believed to be most important. Subject S3 experienced a reversed relationship between stimuli and torque, but showed a large performance benefit in the reversed T condition ($S_F - S_T = 31.3$, compared with a group mean difference of 16.6). According to a stimulus–response model, this result suggests that the benefits of practice might be attributed to dynamic calibration or response practice. An account based on a forward or inverse object model, however, might allow compensation for distorted sensory cues, such as a reversed visual image, after a brief adaptation. Some evidence for this account was found in a previous study (Mah and Mussa-Ivaldi, 2003), which found that learning task elements with a reversed relationship between torque and object motion led to improved performance in a balancing task. Because all of the tasks in the cited study used

the same muscles, this result could not be explained by calibration.

Our main result concerning representation of an object model with joint torques is independent of the control mechanisms that subjects may have used to perform the task. Because subjects were not able to fully compensate for changes in posture relative to the object, our results suggest that an object model is represented using torques or muscle commands and the transformation necessary to apply the model in a new posture is not easily performed by the CNS. Although common experience suggests that an internal representation of absolute force exists, object manipulation in everyday life is performed with an abundance of visual and other cues that could mask the absence of a muscle-independent internal representation of force.

References

- Barto AG, Sutton RS, Anderson CW (1983) Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Trans Sys Man Cybern* 13:834–846.
- Brashers-Krug T, Shadmehr R, Bizzi E (1996) Consolidation in human motor memory. *Nature* 382:252–255.
- Dingwell JB, Mah CD, Mussa-Ivaldi FA (2002) Manipulating objects with internal degrees of freedom: evidence for model-based control. *J Neurophysiol* 88:222–235.
- Gandolfo F, Mussa-Ivaldi FA, Bizzi E (1996) Motor learning by field approximation. *Proc Natl Acad Sci USA* 93:3843–3846.
- Mah CD, Mussa-Ivaldi FA (2003) Evidence for a specific internal representation of the motion-force relationships during object manipulation. *Biol Cybern* 88:60–72.
- Malfait N, Shiffer DM, Ostry DJ (2002) Transfer of motor learning across arm configurations. *J Neurosci* 22:9656–9660.
- Mehta M, Schaal S (2002) Forward models in visuomotor control. *J Neurophysiol* 88:942–953.
- Rowe JB, Toni I, Josephs O, Frackowiak RSJ, Passingham RE (2000) The prefrontal cortex: response selection or maintenance within working memory? *Science* 288:1656–1660.
- Shadmehr R, Moussavi ZMK (2000) Spatial generalization from learning dynamics of reaching movements. *J Neurosci* 20:7807–7815.
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208–3224.
- Wolpert DM, Miall RC (1996) Forward models for physiological motor control. *Neural Netw* 9:1265–1279.