

# Manipulating Objects With Internal Degrees of Freedom: Evidence for Model-Based Control

JONATHAN B. DINGWELL, CHRISTOPHER D. MAH, AND FERDINANDO A. MUSSA-IVALDI  
*Sensory Motor Performance Program, Rehabilitation Institute of Chicago and Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, Illinois 60611*

Received 1 June 2001; accepted in final form 7 March 2002

**Dingwell, Jonathan B., Christopher D. Mah, and Ferdinando A. Mussa-Ivaldi.** Manipulating objects with internal degrees of freedom: evidence for model-based control. *J Neurophysiol* 88: 222–235, 2002; 10.1152/jn00454.2001. There is substantial evidence that humans possess an accurate and adaptable internal model of the dynamics of their arm that is utilized by the nervous system for controlling arm movements. However, it is not known if such model-based strategies are used for controlling dynamical systems outside the body. The need to predict events in the external world is not restricted to the execution of reaching movements or to the handling of rigid tools. Model-based control may also be critical for performing functional tasks with non-rigid objects such as stabilizing a cup of coffee. The present study investigated the strategies used by humans to control simple mass-spring objects. Subjects made straight line reaching movements to a target while interacting with a robotic manipulandum that simulated the dynamics of a one-dimensional mass on a spring. After learning, neither hand nor object kinematics returned to those of free reaching, suggesting that this task was not learned as a perturbation of free reaching. Although there are control strategies (such as slowing the movement of the hand) that would require little or no knowledge of object dynamics, subjects did not adopt these strategies. Instead, they tailored their motor commands to the particular object being manipulated. When object parameters were unexpectedly altered in a way that required no changes in kinematics to successfully complete the task, subjects nonetheless exhibited substantial kinematic deviations. These deviations were consistent with those predicted by a model of the arm-plus-object system driven by a low-impedance controller that incorporated an explicit inverse model of arm-plus-object dynamics. The observed behavior could not be reproduced by a controller that relied on modulating hand impedance alone with no inverse model. These results were therefore consistent with the hypothesis that subjects learn to control the kinematics of manipulated objects by forming an internal model that specified the forces to be exerted by the hand on the object to induce the desired motion of that object.

## INTRODUCTION

Humans often interact with and manipulate objects that contain hinges, pivots, flexible attachments, or other non-rigid elements where the movements of the object are different from the movements of the hand manipulating that object. For example, when carrying a bucket of water or a briefcase, the bucket moves relative to the handle in a way that cannot be directly controlled by the hand. Similar examples include balancing a pole on the tip of one's finger, carrying a cup of

coffee, or even using a lasso or bullwhip. The goal in these object-manipulation tasks is usually to affect some particular motion of the object being manipulated rather than of the hand itself. When a rigid object is held firmly in the hand, controlling the movement of the object is equivalent to controlling the movement of the hand. For non-rigid objects, however, the motion of the object is governed only indirectly through the interaction of the motions of the hand with the internal dynamics (i.e., degrees of freedom) of the object. Anecdotal evidence suggests that humans and other primates can and do learn to manipulate such dynamically complex objects efficiently but does not answer the question of how this process occurs. It is not yet known if humans develop and use internal representations of the dynamical behavior of non-rigid objects during manipulation. These questions were addressed in the present work.

When planning and executing reaching movements, the human motor control system must account for the mechanical properties of the arm. The results of a number of experiments suggest that this process is accomplished using a detailed "internal model" of arm dynamics. In general, an internal model can be defined as a neural representation of a dynamical system that predicts the consequences of the neural commands acting on that system (Imamizu et al. 2000; Wolpert et al. 1995). In this context, a *forward* model is one that predicts the next state of the system (i.e., its positions and velocities) based on the current state and a particular motor command, whereas an *inverse* model is one that estimates the value of the motor command required to move the system into a particular desired state (Imamizu et al. 2000; Miall and Wolpert 1996; Wolpert and Kawato 1998; Wolpert et al. 1995). In the present paper, the term "model-based controller" is used to describe control schemes that may include forward or inverse models or both. It must be emphasized that the CNS could execute movements without making use of any internal model by relying on feedback error-correction alone. Such a strategy could work for executing movements at low speed and with limited interaction forces between the limb and its environment. However, the inherent noise and time delays in sensory feedback signals makes accurate control of rapid or complex movements by feedback alone very difficult (Humphry and Reed 1983; Hogan 1988; Hogan et al. 1987).

Address for reprint requests: J. B. Dingwell, Dept. of Kinesiology and Health Education, The University of Texas, Austin, Texas 78712-1204 (E-mail: jdingwell@mail.utexas.edu).

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

Mechanical robots can be programmed to control un-actuated joints (Lynch and Mason 1999; Lynch et al. 2000; Schaal and Atkeson 1994). This demonstrates that such tasks are theoretically possible. When humans learn to manipulate similar objects with un-actuated degrees of freedom and unknown dynamics, they must either guess the dynamics of the object (i.e., form an internal model of those dynamics) based on their experiences or employ some strategy that does not depend on knowledge of object dynamics. The simplest model-independent strategy that subjects might adopt would be to slow their movements to a speed sufficient to allow visual feedback to become a viable source of information for on-line control. This strategy would be applicable to a broad class of statically stable objects (such as the bucket, briefcase, or coffee cup). However, this strategy would also be slow and inefficient and would fail outright when attempting to control statically unstable objects such as an inverted pendulum.

A second possible model-insensitive strategy that subjects might adopt would be to globally increase arm impedance to enforce a specific kinematic trajectory for the hand (Shadmehr and Mussa-Ivaldi 1994) regardless of the forces imposed on the hand by the object. This strategy corresponds to a nearly ideal position control of hand movements and still depends on acquiring a feasible hand trajectory that satisfies the task constraints. Such a trajectory could be found by trial and error without invoking any explicit model of the object's dynamics. Once a feasible hand trajectory was found, increasing arm impedance about this nominal trajectory would allow the motor controller to enforce the same hand trajectory in the presence of a wide range of interface forces encountered at the hand. For certain tasks, such as manipulating a mass on a spring, where subjects must control the mass as well as the hand, each enforced hand trajectory would produce an equivalent object trajectory only for a limited class of "kinematically equivalent" mass-spring objects; i.e., objects with different inertia and stiffness parameters but with the same resonant frequency (see METHODS). The present experiments exploited this feature of mass-spring dynamics to determine if subjects used such a position control strategy when manipulating mass-spring objects. Specifically, the effects of unexpectedly changing object mechanical parameters were examined and experimental results were compared with predictions made by different hypothesized control laws.

Numerous studies have shown that when making point-to-point reaching movements in various external force fields, such as viscous fields (Conditt and Mussa-Ivaldi 1999; Conditt et al. 1997; Gandolfo et al. 1996; Shadmehr and Brashers-Krug 1997; Shadmehr and Moussavi 2000; Shadmehr and Mussa-Ivaldi 1994; Thoroughman and Shadmehr 1999) or Coriolis fields (Cohn et al. 2000; Dizio and Lackner 1995; Lackner and Dizio 1994), humans adapt their motor commands to precisely cancel out the effects of these fields. This adaptation process can take up to several hundred movements to complete, depending on the nature of the field. Similar adaptive responses are found when humans manipulate rigid objects that impose inertial force fields on the arm, although the adaptation typically takes place much faster. Humans can adjust grip force to accommodate the simple inertial loads imposed by typical rigid objects within as little as 135 ms during a single movement (Bock 1990, 1993). When subjects lift objects of unusually high densities, they adapt their responses within fewer than five

movements (Gordon et al. 1993). Even for more complex modifications of the arm's inertial properties, adaptation is typically completed within 40–50 movements (Sainburg et al. 1999). Furthermore, precise predictive grip force modulation requires that the feedback forces produced by a manipulated rigid object mimic those of a real rigid object (Blakemore et al. 1998; Witney et al. 2000).

These studies have demonstrated that humans maintain an internal model of the dynamics of their arm that adapts when the physical (i.e., inertial, viscous, and/or Coriolis) properties of their arm are altered. However, in each of these contexts, the kinematics of the system being controlled by the CNS (i.e., the arm) remained unaltered. This mechanical system possessed the same number of degrees of freedom whether these force fields were present or not. It is not known if the adaptive mechanisms discussed in the preceding text can be extended to situations involving manipulation of *non-rigid* objects where the CNS must learn to control not only the behavior of the arm but also of those additional degrees of freedom introduced by the object. Indeed, a control strategy that is successful for controlling one mechanical system may fail when that system is coupled to a second mechanical system and the resulting combined system being controlled is therefore changed (Hogan et al. 1987). Furthermore, the degrees of freedom of the object are not directly actuated by muscles the way that limb segments are and therefore cannot be controlled directly. Instead, the controller must learn to act through the physics imposed by the object. When an object introduces degrees of freedom external to the body, observation of the states (i.e., positions and velocities) associated only with subject's limb is not generally sufficient to predict the hand-object interaction forces. In such a context, these interaction forces can only be represented as time-dependent forces. However, in experiments that applied strictly time-dependent force fields to the arm, subjects did not learn the time-dependent field but instead attempted to compensate for the forces imposed as if they were dependent on arm state variables (Conditt and Mussa-Ivaldi 1999). This observation that subjects do not construct appropriate representations of time-dependent forces raises the possibility that despite having an accurate and adaptable model of their own arm, humans may not be able to use similar models when dealing with environments or objects that introduce additional state variables external to the body.

The purpose of the present study was to determine if the strategies humans use to successfully manipulate non-rigid objects are consistent with a controller that attempts to predict the dynamical behavior of the object (i.e., one that constructs an internal model of object dynamics). This objective was addressed by having subjects perform a goal-directed reaching task where they learned to manipulate an object with one internal degree of freedom: a mass on a spring. It was hypothesized that subjects would learn to perform this task by implementing a model-dependent control strategy that relied on an internal representation of the dynamics of the mass-spring object being manipulated. The present experiments tested for evidence of the alternative model-independent hypotheses that subjects would slow down to minimize perturbations imposed by the hand-object interaction forces or might globally increase arm impedance to resist those perturbations and enforce a prespecified kinematic plan. The results indicate that subjects did not adopt either of these two alternative model-independent

strategies. Although subjects moved more slowly in the object-manipulation task than when reaching with the hand alone, they did exhibit systematic reductions in total movement time, which indicated that they were acquiring the capacity to predictively control the behavior of the object. Subjects also exhibited statistically significant and systematic deviations in hand trajectories when unexpectedly exposed to kinematically equivalent objects (see METHODS). Analysis of subjects' responses to these unexpected changes in object parameters refuted the possibility that they were controlling the object simply by increasing the impedance of the hand about some nominal trajectory. Instead, the present findings were consistent with the hypothesis that humans control non-rigid objects using low-impedance force control laws similar to those used when the arm is subjected to force perturbations (e.g., Shadmehr and Mussa-Ivaldi 1994). These findings expand on previous work on the use of internal model strategies for controlling arm movements in external force fields by demonstrating that similar control strategies might also apply to the control of dynamical systems that extend beyond our own bodies.

## METHODS

### Data collection

Six young healthy subjects (3 male and 3 female; age =  $30.7 \pm 4.6$  yr; height =  $1.75 \pm 0.11$  m; weight =  $77.6 \pm 23.4$  kg) with no known neuromuscular disorders affecting their upper extremities participated in the experiment after providing written informed consent. Experiments were performed using a 2-degree-of-freedom robotic manipulandum (Fig. 1A), described in detail elsewhere (Conditt and Mussa-Ivaldi 1999; Conditt et al. 1997; Scheidt et al. 2000). Visual feedback of the relevant task variables was presented on a video monitor mounted above the robot's motors (Fig. 1A). Subjects made one-dimensional 20- to 25-cm reaching movements (Fig. 1B) with their dominant arm while firmly holding the handle of the manipulandum. Reaching movements were directed away from the subject's body along a line passing through the center of rotation of the shoulder and parallel to the sagittal plane (the positive  $y$  axis). Each subject's arm was supported by a sling attached to the 8-ft. ceiling and adjusted so that movements were made approximately in the horizontal plane. The manipulandum was programmed to produce forces that simulated a one-dimensional undamped mass on a spring (Fig. 1C) that oscillated in the  $y$  direction only. During movements made with the mass-spring object, subjects received on-line real-time visual feedback of both their hand position and the position of the virtual mass (Fig. 1B).

The second-order equation of motion for this mass-spring object was

$$M_O \ddot{y}_O + K_O (y_O - y_H) = 0 \quad (1)$$

where  $M_O$  was the object mass  $K_O$  was the object spring stiffness, and  $y_O$  and  $y_H$  were the  $y$  positions of the object and the hand, respectively. The first-order state space form of this equation was obtained by defining the object state variables as  $q_1 = y_O$  and  $q_2 = \dot{y}_O$

$$\begin{bmatrix} \dot{q}_1 \\ \dot{q}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -K_O/M_O & 0 \end{bmatrix} \cdot \begin{bmatrix} q_1 \\ q_2 \end{bmatrix} + \begin{bmatrix} 0 \\ K_O/M_O \end{bmatrix} \cdot y_H \quad (2)$$

This equation was integrated in real-time to compute the instantaneous position of the object ( $q_1$ ) for each corresponding position of the hand ( $y_H$ ). The forces exerted on the hand by the motors were then computed from

$$F_y = K_O (q_1 - y_H) \quad (3)$$

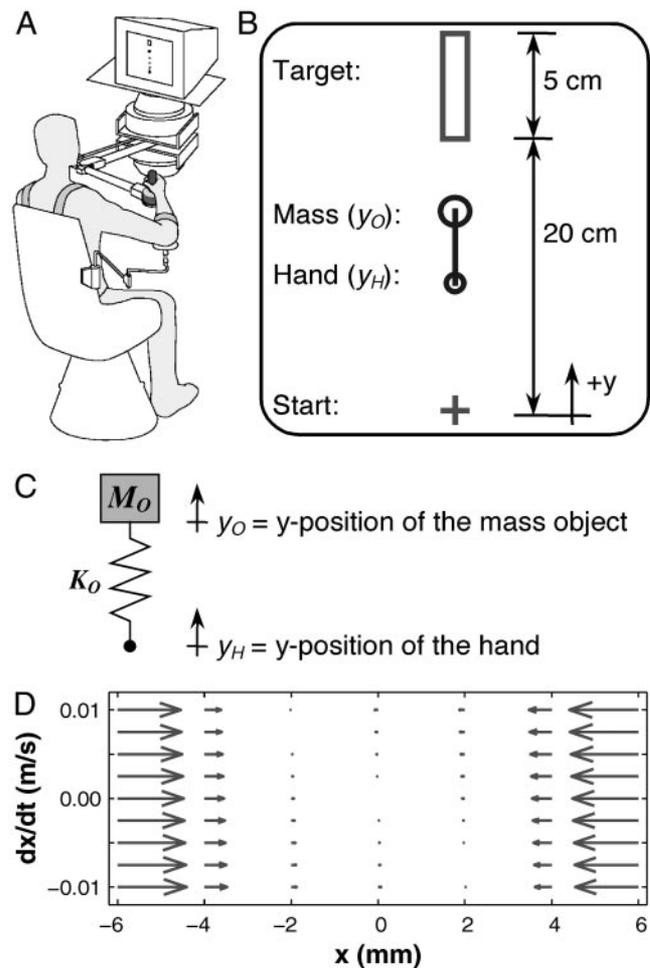


FIG. 1. A: representation of a subject using the 2 degree-of-freedom robotic manipulandum. Visual feedback of relevant task variables was displayed on a video monitor mounted above the robot's motors. B: representation of the screen display used in the experiment. Subjects saw the start and goal targets and a moving cursor that represented the current positions of both the hand and the mass (the text shown in this figure was not displayed on the screen). C: schematic diagram of the simulated mass-spring object attached (virtually) to the manipulandum handle (see Eqs. 1 and 2). D: quiver plot of the forces imposed by the mediolateral force field used to ensure that subjects made reasonably straight-line movements along the positive  $y$  axis.

The control input to the system, as defined in Eq. 2, was the  $y$  position of the handle;  $y_H(t)$ . These equations assumed a zero rest length for the spring, such that subjects experienced zero anterior-posterior forces when the handle and object were in the same  $y$  position (i.e.,  $F_y = 0$  when  $y_H = y_O$ ). Reaching movements were always initiated with the hand and object both at rest so that on any given trial, subjects could not tell prior to the onset of movement whether the object was present or not. To ensure that reasonably straight-line movements were made along the positive  $y$  direction, a mediolateral force field (Fig. 1D) was imposed along the  $y$  axis by generating forces acting on the handle in the  $x$  direction that mimicked a damped cubic spring (i.e.,  $F_x = -A \cdot x^3 - B\dot{x}$ ; where  $A = 50 \times 10^6 \text{ Nm}^{-1}$  and  $B = 50 \text{ Nsm}^{-1}$ ). This cubic spring produced restoring forces of  $F_x \leq 0.4 \text{ N}$  at lateral displacements of  $x \leq \pm 2 \text{ mm}$ , which rose sharply to  $F_x = 50 \text{ N}$  at  $x = \pm 10 \text{ mm}$ , thus providing a narrow "dead zone" in which subjects could move freely. The forces imposed by this force field were completely orthogonal to the forces imposed by the mass-spring object and were computed completely independently of those forces. After the first few trials, subjects made straight

enough movements that this mediolateral force field went effectively unnoticed.

Each subject completed five blocks of 150 movements each. Subjects were instructed to reach from the start location to the target (Fig. 1B) in  $0.8 \pm 0.2$  s. Subjects were given visual feedback on the computer monitor about their movement duration after each trial (too fast vs. just right vs. too slow). The first block of 150 movements was made primarily with the virtual object turned *off* (i.e.,  $F_y = 0$  and  $q_1 = y_H$ ). This was done to familiarize the subject with the reaching task and the desired timing of the movement. During 10 randomly selected trials of the last 100 trials of the first block, the object was unexpectedly turned *on* (i.e.,  $q_1$  defined from Eq. 2 and  $F_y$  from Eq. 3) to assess each subject's initial response to the task. Trial blocks two through five were all completed with the object turned *on* during all trials as subjects learned the new task. During these trials, subjects were instructed to bring both their hand and the mass to a complete stop within the target zone within the specified time frame ( $0.8 \pm 0.2$  s). Success in the task was defined as achieving a  $y$  velocity of  $\dot{y} < 0.02$   $\text{ms}^{-1}$  for both the hand and mass within the target zone for  $>0.15$  s. Subjects were given a maximum time limit of 2.5 s per trial to complete the task successfully, after which trials were terminated automatically. Subjects learned to perform the object-manipulation task with an object of  $M_O = 3$  kg and  $K_O = 120$   $\text{Nm}^{-1}$ , which had

a resonant frequency of  $f_0 = \frac{1}{2\pi} \cdot \sqrt{K_O/M_O} \approx 1$  Hz.

Equation 2 shows that the dynamics of the mass-spring object were defined exclusively by the ratio of the object's mass to its spring stiffness:  $K_O/M_O$ . Thus any "kinematically equivalent" object (i.e., any object with different  $K_O$  and  $M_O$  values, but the same  $K_O/M_O$  ratio and thereby the same resonant frequency) would exhibit the same output kinematics (i.e.,  $[q_1, q_2]^T$ ) when subjected to the same kinematic input,  $y_H(t)$ . In other words, the same hand trajectory,  $y_H(t)$ , would produce the same object trajectory,  $y_O(t)$ . To determine if the subjects in the present experiment were merely "stiffening up" to enforce a prespecified hand trajectory,  $y_H(t)$ , the last 120 trials of the fifth and final block of 150 movements included 12 randomly selected "catch trials" in which the  $K_O/M_O = (120 \text{ Nm}^{-1}/3 \text{ kg})$  object was unexpectedly replaced with either a  $K_O/M_O = (40 \text{ Nm}^{-1}/1 \text{ kg})$  object (6 randomly selected trials) or with a  $K_O/M_O = (200 \text{ Nm}^{-1}/5 \text{ kg})$  object (6 randomly selected trials).

During all trials, the position of the handle was recorded from position encoders mounted in the manipulandum. Handle velocities were obtained from handle positions using a finite difference formula. These data were used in real time to compute the position and velocity of the virtual mass of the simulated object and to compute the reaction forces that were applied to the hand. The positions and velocities of both the hand and virtual object and the interface forces applied to the hand were saved to disk for further analysis. All kinematic data were sampled at 100 Hz and then low-pass filtered using a zero-lag 6th order Butterworth filter with a cutoff frequency of 10 Hz.

### Data analysis

To quantify the relative rates of learning in each subject, plots of total movement time exhibited on each trial versus trial number were constructed. These data were fit with a single exponential function to quantify the trends in these learning curves

$$T_i = Ae^{(-i/\lambda)} + T_f \quad (4)$$

where  $T_i$  was the total movement time on trial  $i$  (where  $i = 1, 2, \dots, 600$  trials),  $\lambda$  defined the time constant of adaptation (in number of trials),  $T_f$  defined the final adapted movement time (i.e., the movement time that each subject asymptotically approached across the duration of the experiment), and  $A$  defined the gain of the adaptation process. Only those trials performed with the 3-kg mass-spring object were included in these calculations. The free parameters of Eq. 4 ( $A$ ,  $\lambda$ , and

$T_f$ ) were fit using a nonlinear least squares procedure (function "nlinfit" in Matlab). One model-independent strategy that subjects might have adopted in the present experiment would be to move slowly enough to minimize oscillations of the mass-spring object, thus allowing the use of on-line feedback control. If subjects adopted such a strategy, then one would not expect to see significant changes in overall movement times beyond the first few trials.

The task requirements were to reach to the target and bring both the hand and object to a stop. Subjects were allowed to choose any specific trajectory they wished to satisfy these requirements. However, the additional un-actuated degree of freedom introduced by the mass-spring object restricted the choice of hand trajectories that subjects could adopt while remaining successful at the task. Because accurate control of the object could only be achieved through the choice of an appropriate hand trajectory, subjects were not expected to return to their preexposure hand kinematics. To examine the nature of the specific hand and object trajectories chosen by each subject, average  $y$  direction velocity profiles for the hand,  $\dot{y}_H(t)$ , in the preadaptation condition and for both the hand ( $\dot{y}_H$ ) and object ( $\dot{y}_O$ ) in the postadaptation condition were plotted and compared. For each average velocity profile, 95% confidence intervals (CI) were computed from ensembles of 20 trials. These velocity profiles were used for these comparisons because differences between test conditions were more readily apparent in these curves. Similar comparisons were made between movements of both the hand and object.

As discussed in the INTRODUCTION, another model-independent control strategy that subjects could use would be to globally increase arm impedance to enforce a specific kinematic trajectory for the hand (Shadmehr and Mussa-Ivaldi 1994). The present experiment directly tested for evidence of this possibility. If subjects had merely "stiffened up" to enforce a prespecified hand trajectory,  $y_H(t)$ , then unexpectedly replacing the 3-kg object with either of the kinematically equivalent 1- or 5-kg objects would not be expected to substantially alter task performance.  $y$ -direction hand-velocity profiles,  $\dot{y}_H(t)$ , obtained from the individual catch trials with the 1- and 5-kg objects were compared with 95% CI of the average hand-velocity profile obtained for the last 20 postadaptation trials for the 3-kg object. Deviations in the catch-trial hand-velocity profiles beyond the 95% CI for the postadaptation velocity profiles constituted statistically significant differences from the acquired postadaptation behavior. Accordingly, any such differences would provide evidence that subjects were not simply globally increasing hand impedance to enforce a specific hand trajectory.

### Simulating model-dependent control

The goal of these simulations was to determine if the exhibited catch-trial trajectories were consistent with a controller that relied on a prespecified feed-forward command to explicitly compensate for the forces imposed on the hand by the mass-spring object. This hypothesis was compared with the alternative possibility that the exhibited catch-trial trajectories could have been produced by a controller that relied on globally increasing hand impedance. The existence of such a predictive feed-forward control command would be consistent with a control strategy in which subjects were computing that command based on some prediction of the object's subsequent behavior (i.e., an internal model). The simplified one-dimensional model was determined to be adequate for obtaining a first-order approximation of this effect while maintaining computational efficiency.

Equation 1 takes as its control input the position of the hand as a function of time,  $y_H(t)$ . If the arm was an *ideal* position controller (i.e., one with infinite impedance), subjects could execute any desired  $y_H(t)$  with infinite precision regardless of the interface forces experienced at the hand. Because humans are *not* infinite-impedance position controllers, some degree of kinematic deviation was expected even if subjects were globally increasing their arm impedance. In previous experiments where subjects adapted their reaching movements to viscous force fields (Shadmehr and Mussa-Ivaldi 1994), subjects did

not use such a strategy but instead exhibited behavior consistent with a low-impedance controller that incorporated an explicit internal model of arm dynamics. To determine if the behavior exhibited by the subjects in the present experiment was similarly consistent with this form of controller, catch-trial movements were simulated for each subject using a simplified one-dimensional model of the arm plus object (Fig. 2A). The arm was modeled as a single point mass attached to the hand ( $M_H$ ) that was controlled by a single time-varying control force,  $C(t)$ . The second-order differential equations of motion for this double mass-spring system were

$$\begin{aligned} M_O \ddot{y}_O + K_O (y_O - y_H) &= 0 \\ M_H \ddot{y}_H - K_O (y_O - y_H) &= C(t) \end{aligned} \quad (5)$$

Defining a new set of state variables for this arm-plus-object model as  $[q_1 \ q_2 \ q_3 \ q_4]^T = [y_O \ \dot{y}_O \ y_H \ \dot{y}_H]^T$ , the equivalent first-order state space equations for this system were

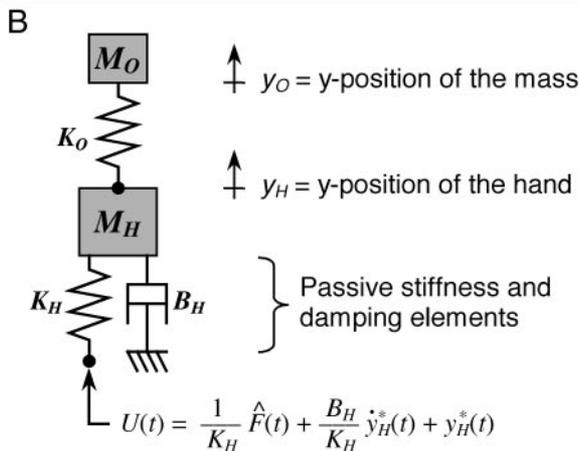
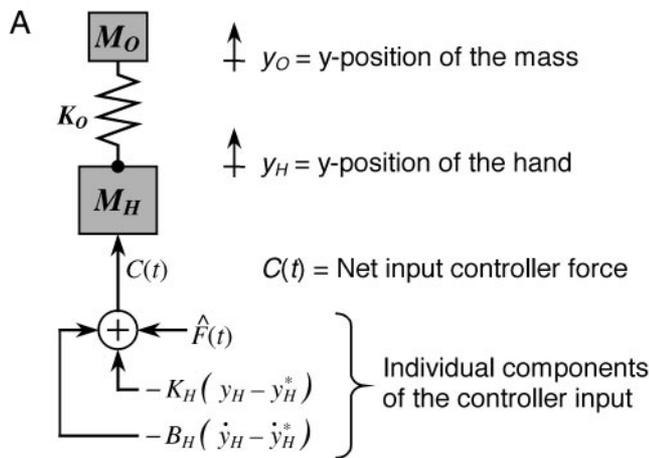


FIG. 2. A: schematic representation of the simulated 2-mass arm-plus-object model (Eqs. 5 and 6). The arm was modeled as a single point mass located at the hand ( $M_H$ ). The forces applied to the arm mass were produced by a controller consisting of an approximation of the inverse model of the arm-plus-object dynamics,  $\hat{F}(t)$  (Eq. 8); plus feedback terms necessary to provide stability (Shadmehr and Mussa-Ivaldi 1994). B: schematic diagram of a mathematically equivalent model in which the passive stiffness and damping of arm are included in the plant, rather than the feedback control loops (Eqs. 9 and 10). In this model, the control input is a single feed-forward command. Both models are mathematically identical and both include an estimate of the arm-plus-object inverse dynamics (i.e., a predictive internal model) in their feed-forward control inputs.

$$\begin{bmatrix} \dot{q}_1 \\ \dot{q}_2 \\ \dot{q}_3 \\ \dot{q}_4 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ -K_O/M_O & 0 & +K_O/M_O & 0 \\ 0 & 0 & 0 & 1 \\ +K_O/M_H & 0 & -K_O/M_H & 0 \end{bmatrix} \cdot \begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1/M_H \end{bmatrix} \cdot C(t) \quad (6)$$

Following the same form as the controller postulated by Shadmehr and Mussa-Ivaldi (1994; their Eq. 9), the control input for this model,  $C(t)$ , was defined to include three terms. The first of these terms was a feed-forward term,  $\hat{F}(t)$ , defined to be an estimate of the inverse dynamics of the arm-plus-object system. The remaining two terms were error-feedback terms that accounted for the passive stiffness and viscosity of antagonist muscles and their associated segmental reflexes (Shadmehr and Mussa-Ivaldi 1994). Thus  $C(t)$  was defined as

$$C(t) = \hat{F}(t) - K_H (y_H - y_H^*) - B_H (\dot{y}_H - \dot{y}_H^*) \quad (7)$$

where

$$\hat{F}(t) = \hat{M}_H \dot{y}_H^* - \hat{K}_O (y_O^* - y_H^*) \quad (8)$$

defined the estimate of the inverse of the arm-plus-object dynamics (i.e., the internal model) for movements made along the desired trajectory,  $[y_H^*(t) \ y_O^*(t)]$ . It should be noted that this system could be equivalently modeled by including the passive stiffness and damping in the plant (Fig. 2B) rather than in the feedback terms of the controller. The resulting equations for this equivalent system were

$$\begin{bmatrix} \dot{q}_1 \\ \dot{q}_2 \\ \dot{q}_3 \\ \dot{q}_4 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ -K_O/M_O & 0 & +K_O/M_O & 0 \\ 0 & 0 & 0 & 1 \\ +K_O/M_H & 0 & -(K_O + K_H)/M_H & -B_H/M_H \end{bmatrix} \cdot \begin{bmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ K_H/M_H \end{bmatrix} \cdot U(t) \quad (9)$$

where the control input then became

$$U(t) = \frac{1}{K_H} \hat{F}(t) + \frac{B_H}{K_H} \dot{y}_H^*(t) + y_H^*(t) \quad (10)$$

A straightforward substitution of terms demonstrates that the system defined by Eq. 6 with  $C(t)$  defined as in Eq. 7 is identically equivalent to the system defined by Eq. 9 with  $U(t)$  defined as in Eq. 10. The resulting system (Eqs. 9 and 10) was a linear time-invariant (LTI) system subject to a single time-varying input,  $U(t)$ . Note that both formulations of the model control system contain an estimate of the inverse dynamics of the arm-plus-object system,  $\hat{F}(t)$ , in their control input.

For these simulations, arm masses ( $M_H$ ) were estimated for each subject as a fixed percentage (5%) of their total body mass based on standard anthropometric data (Winter 1990; Table 3.1). Arm mass values ranged from 2.8 to 5.8 kg across the six subjects. Gomi and Kawato (1996, 1997) estimated endpoint (i.e., measured at the hand) stiffness ellipses for subjects making normal (i.e., no objects or perturbing force fields) anterior-posterior reaching movements. The anterior-posterior components of these stiffness ellipses shown in their results (their Fig. 4) had magnitudes ranging from approximately  $K_H = 100 \text{ Nm}^{-1}$  at proximal arm configurations to more than  $K_H = 1,000 \text{ Nm}^{-1}$  at the distal positions (Gomi and Kawato 1996, 1997). Studies examining arm stiffness properties during similar movement paradigms have reported either similar values (Bennett 1994; Bennett et al. 1992; Burdet et al. 2000), somewhat larger values (Xu and Hollerbach 1999), or somewhat smaller values (Mah 2001). When subjects actively co-contract to resist perturbations, postural endpoint stiffness can increase by as much as five to six times these values (Gomi and Osu 1998; Perrault et al. 2001). Estimates of hand viscosity ellipses have also been reported for postural tasks by Tsuji et al. (1995). The anterior-posterior components of the viscosity ellipses

shown in their results (their Fig. 11) had magnitudes ranging from approximately  $B_H = 5 \text{ N}\cdot\text{s}\cdot\text{m}^{-1}$  for proximal arm configurations to upwards of  $B_H = 25 \text{ N}\cdot\text{s}\cdot\text{m}^{-1}$  for distal arm configurations (Tsuji et al. 1995).

In the present study, reaching movements were initiated from a posture in which the hand was located proximal to the body. Therefore values of  $K_H = 100 \text{ Nm}^{-1}$  and  $B_H = 5 \text{ N}\cdot\text{s}\cdot\text{m}^{-1}$  were used to simulate the catch-trial behavior that would be expected if subjects had *not* globally increased their arm impedance. Likewise, values of  $K_H = 500 \text{ Nm}^{-1}$  and  $B_H = 25 \text{ N}\cdot\text{s}\cdot\text{m}^{-1}$  were used to simulate the catch-trial behavior that would be expected if subjects had co-contracted to enforce a position control strategy. Although *in vivo* arm stiffness and damping in humans varies as a function of arm posture (Gomi and Kawato 1996, 1997; Tsuji et al. 1995), it was assumed that this range of variation would be adequately captured within the bounds established by these values. It should be emphasized that the present experiments were neither designed nor intended to estimate actual hand stiffness and viscosity values during the mass-spring-object-manipulation task. These simulations were only used to provide a first-order estimate of the likelihood that the kinematic trajectories exhibited during the catch trials were consistent with a predictive control strategy that compensated for the interface forces imposed on the hand by the mass-spring object rather than with a strategy that relied on globally increasing hand impedance to resist those perturbations.

Desired hand and object trajectories,  $[y_H^*(t) \ y_O^*(t)]$ , were defined for each subject from the average trajectories obtained from the last 20 postadaptation (i.e.,  $M_O = 3 \text{ kg}$ ;  $K_O = 120 \text{ Nm}^{-1}$ ) trials prior to the introduction of the catch trials (i.e., trials 11–30 in the 5th block of movements). To obtain sufficiently smooth hand trajectories from the inverse dynamics (Eq. 8), these average trajectories were first interpolated from 100 to 1,000 Hz and then zero-lag low-pass filtered at a cutoff frequency of 10 Hz. Controller trajectories,  $U(t)$ , were then computed directly from these data using Eq. 10 and the original learned object parameters ( $M_O = 3 \text{ kg}$  and  $K_O = 120 \text{ Nm}^{-1}$ ). These control trajectories were then used as the input to the forward dynamic simulation (Eq. 9), where the  $K_O/M_O = (120 \text{ Nm}^{-1}/3 \text{ kg})$  learned object parameters were replaced with either the  $K_O/M_O = (40 \text{ Nm}^{-1}/1 \text{ kg})$  or the  $K_O/M_O = (200 \text{ Nm}^{-1}/5 \text{ kg})$  catch-trial object parameters. Simulations were generated in Matlab using the “lsim” command. Simulated catch-trial hand-velocity profiles were compared directly to experimental catch-trial hand-velocity profiles to determine if the kinematics exhibited by the subjects in the present study were consistent with the hypothesis of a controller that incorporated a predictive model of the arm-plus-object dynamics.

To quantify the quality of the fits between each of the two model realizations and the data, an “index of relative fit” (IRF) was derived. First, the root-mean-square (rms) difference between the average catch-trial hand-velocity profiles,  $\langle \dot{y}_H^{\text{CT}}(t) \rangle$  (where  $\langle \cdot \rangle$  denotes the average over  $n = 6$  catch trials), and the model-predicted hand-velocity profile,  $\dot{y}_H^{\text{MP}}(t)$ , was computed. This difference measured how well each predicted catch-trial hand trajectory fit the actual catch-trial hand trajectories. If the model produced a perfect fit, this difference would be zero. The rms difference between the average catch-trial hand-velocity profiles,  $\langle \dot{y}_H^{\text{CT}}(t) \rangle$ , and the average postadaptation hand-velocity profiles,  $\langle \dot{y}_H^{\text{PA}}(t) \rangle$  (where  $\langle \cdot \rangle$  denotes the average over the last 20 postadaptation trials), was also computed. This difference quantified the average magnitude of the hand-velocity deviations exhibited by each subject during the catch trials and was used to control for the fact that because different subjects exhibited different postadaptation kinematics, they were also expected to exhibit somewhat different catch-trial behavior. The IRF was then defined as the ratio of these two rms differences

$$\text{IRF} = \frac{\text{rms}[\langle \dot{y}_H^{\text{CT}}(t) \rangle_{n=6} - \dot{y}_H^{\text{MP}}(t)]}{\text{rms}[\langle \dot{y}_H^{\text{CT}}(t) \rangle_{n=6} - \langle \dot{y}_H^{\text{PA}}(t) \rangle_{n=20}]} \quad (11)$$

Thus  $\text{IRF} \rightarrow 0$  if the predicted catch-trial hand movements closely matched the actual catch-trial hand movements, whereas  $\text{IRF} \rightarrow 1$  if the predicted catch-trial hand movements closely matched the post-adaptation hand movements. IRF values much more than 1 indicated poor model fit (i.e., either that the model predicted catch-trial deviations much bigger than those actually observed or in the opposite direction to those observed). For each subject and catch-trial condition (1 vs. 5 kg), the value of IRF was computed over the first 500 ms of the movements for both the low- and high-impedance realizations of the arm-plus-object model. Differences in IRF between low- and high-impedance model predictions were tested using a one-sided paired *t*-test for the IRF data computed for both 1- and 5-kg catch trials, where the data were paired within subjects.

### Simulating model-independent control

It is reasonable to ask whether or not a control strategy that did not include an explicit representation of the inverse dynamics of the arm-plus-object system,  $\hat{F}(t)$ , might also be capable of reproducing the catch-trial behavior observed in the present experiment. To test this possibility, simulations of the arm-plus-object system (Eq. 6) were also performed by assuming an alternative controller,  $C_A(t)$ , that did not include an internal model [i.e.,  $\hat{F}(t) = 0$ ], but instead relied solely on setting the overall gains ( $K_H$  and  $B_H$ ) for the net hand stiffness and viscosity feedback terms

$$C_A(t) = -K_H(y_H - y_H^*) - B_H(\dot{y}_H - \dot{y}_H^*) \quad (12)$$

This controller translated into an equivalent control input that could be applied to the equivalent LTI model with time-varying input (Eq. 9)

$$U_A(t) = + \frac{B_H}{K_H} \dot{y}_H^*(t) + y_H^*(t) \quad (13)$$

Simulations based on this alternative controller were generated in a manner similar to the process described in the preceding text. Control inputs,  $U_A(t)$ , were computed directly from the interpolated and filtered postadaptation data for each subject using Eq. 13. To determine what feedback gains ( $K_H$  and  $B_H$ ) were required to allow this controller to adequately reproduce the original learned behavior, these controller trajectories were first used to drive the forward dynamic simulation (Eq. 9) with the learned object parameters (i.e.,  $M_O = 3 \text{ kg}$  and  $K_O = 120 \text{ Nm}^{-1}$ ). A wide range of arm impedance gains (i.e.,  $100 \text{ Nm}^{-1} < K_H < 1,000 \text{ Nm}^{-1}$  and  $5 \text{ N}\cdot\text{s}\cdot\text{m}^{-1} < B_H < 50 \text{ N}\cdot\text{s}\cdot\text{m}^{-1}$ ) were tested. To simulate the corresponding catch-trial behaviors, the learned object parameters were replaced with either set of catch-trial object parameters, and the forward dynamic simulations were repeated. Simulated postadaptation and catch-trial hand-velocity profiles based on this alternative controller were compared directly to experimental hand-velocity profiles and to the catch-trial trajectories predicted by the original internal model controller (Eq. 7). These comparisons were made to determine if the kinematics exhibited by the subjects in the present study could have been generated by a controller that relied solely on adjusting hand stiffness and damping rather than trying to learn a predictive model of the arm-plus-object dynamics.

## RESULTS

### Subjects exhibited learning behavior

The trial-to-trial total movement times ( $T_i$ ) computed from the object-manipulation trials demonstrated that most subjects learned to perform the task with reasonable skill (Fig. 3). However, time constants of learning (Eq. 3) varied widely between subjects ( $18 < \lambda < 442$  trials). Because the mass-

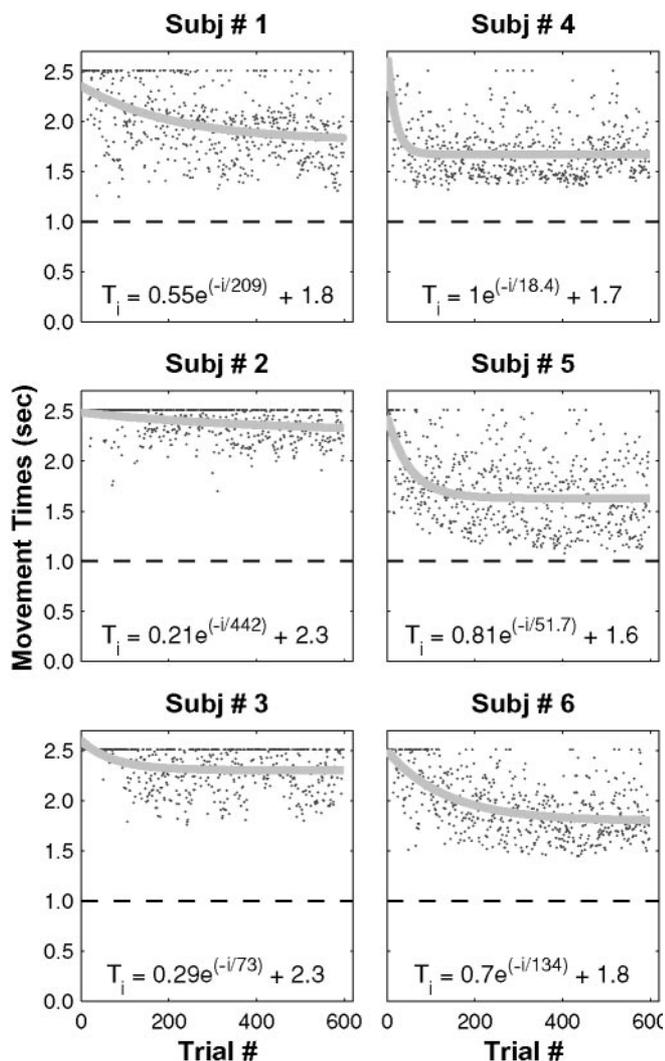


FIG. 3. Movement time results for each of the 6 subjects tested. Dashed horizontal lines indicate desired movement time ( $\leq 1.0$  s.). Thick gray lines represent exponential (Eq. 3) curve fits to the data. Equations for each curve fit are shown at the bottom of each subplot. Although learning rates varied, all subjects exhibited some capacity to improve their overall performance over the course of the experiment.

spring object itself contained no damping, completing the reaching task successfully in *any* amount of time required subjects to effectively dampen out the oscillations of the object. While no subject successfully completed the task within the desired time limit ( $0.8 \pm 0.2$  s) during the duration of the experiment, all subjects exhibited lower movement times with practice. This suggests that these reductions in movement time were an expression of learning.

These improvements in total movement time might indicate that subjects were learning the appropriate feedback gains to use in a control scheme based exclusively on feedback error correction. Alternatively, they might reflect the fact that subjects were first acquiring and then gradually improving their capacity to control the behavior of the object by learning an internal model of object behavior. In either case, these results demonstrate that these subjects did indeed learn (by some means) to improve their performance on the designated task.

### Postadaptation kinematics

Preexposure (no object) reaching movements performed by the subjects in the present study exhibited the bell-shaped velocity profiles (Fig. 4A) typically associated with such unperturbed reaching movements (Flash and Hogan 1985; Morasso 1981). As anticipated, unexpectedly exposing subjects to the mass-spring object caused severe disruption of these reaching kinematics (Fig. 4B). Figure 4B shows that the object oscillated at its resonant frequency of 1 Hz and that these oscillations were slowly damped out. Because the object was unanticipated in these trials, this damping behavior was likely the result of passive mechanisms related to mechanical properties of the arm and/or to the fact that subjects were not able to hold onto the handle with a perfectly rigid grasp.

By the end of the learning phase of the experiment, each subject had adopted a new pattern of movement (Fig. 5). Although different subjects exhibited different hand-velocity profiles, these movement patterns were quite consistent for each subject, as demonstrated by the low trial-to-trial variability. In contrast to previous studies where subjects adapted to force fields that depended only on arm state variables (e.g., Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994), no subject in the present study regained their original unperturbed kinematics. For every subject in the present experiment, final velocity profiles of the hand and object were both substantially different from their unperturbed hand kinematics. Although velocity profiles for the object were roughly bell-shaped for most subjects, they exhibited consistently lower peak velocities and were extended over a longer time frame compared with these subjects' preexposure hand movements. Hand-velocity profiles were distinctly non-bell-shaped in all subjects and exhibited different shapes for different subjects (Fig. 5). These changes in hand kinematics provide evidence of predictive control. They demonstrate that subjects were able to predict that if they moved their hand along one specific trajectory, then the object would follow another specific trajectory, the combination of these two trajectories leading to successful completion of the task goal. If subjects had been unable to learn these new appropriate hand trajectories, then they would not have been able to improve their overall task performance and thereby their overall movement times (Fig. 3).

### Subjects did not use high-impedance control

Although Fig. 5 demonstrates that each subject adopted a control strategy that relied on moving their hand along a specific trajectory to successfully complete the task, such movements could be executed in a number of ways. One way would be to globally increase arm impedance to neutralize the effects of the interface forces generated at the hand by the mass-spring object. An alternative approach would be for subjects to learn the temporal sequence of hand-object interface forces and adjust their feed-forward commands to compensate for them directly. The present study directly tested for evidence of these two possibilities by introducing kinematically equivalent objects (see METHODS) during randomly selected catch trials toward the end of the learning phase of the experiment. If subjects had globally increased their hand stiffness enough to enforce a specific hand trajectory regardless of

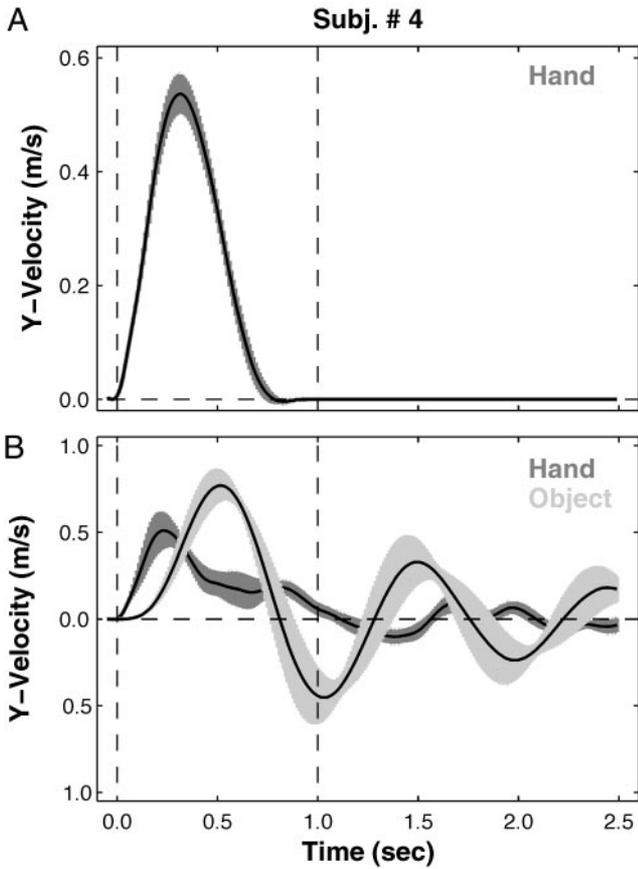


FIG. 4. Reaching kinematics for a typical subject (*subject 4*) prior to learning. *A*:  $y$ -direction velocity profiles for the hand ( $\dot{y}_H$ ) during preexposure trials exhibited the expected bell-shaped profile. *B*:  $y$ -direction velocity profiles for both the hand ( $\dot{y}_H$ ) and object ( $\dot{y}_O$ ) during initial exposure trials exhibited substantial disruption of the movement pattern. For each plot, dark lines represent averages and shaded areas represent  $\pm 95\%$  confidence intervals, each computed from ten trials. Dashed vertical lines indicate the start of movement ( $t = 0.0$  s) and the targeted movement time ( $t = 1.0$  s). All subjects exhibited similar patterns of behavior on initial exposure to the mass-spring object.

the perturbing forces encountered at the hand, then only minimal deviations would be expected during these catch trials. However, all subjects exhibited statistically significant and substantial deviations from their acquired postadaptation hand paths when making reaching movements with both the 1- and 5-kg “catch-trial” objects (Fig. 6). Catch trials with the 1-kg object always exhibited initially greater velocities during the first half of the movements ( $t \leq$  approximately 500–750 ms). With the exception of *subject 1*, catch trials with the 5-kg object always exhibited initially lower velocities during the first half of the movements.

*Model-based control predicts exhibited behavior*

To obtain first-order estimates of the kinematic deviations that would be expected during the catch trials if subjects had adopted a model-based control strategy, simulations were performed with a simplified one-dimensional model of the arm-plus-object (Fig. 2; *Eqs. 4–6*). The controller for this system was assumed to incorporate an explicit model of the arm-plus-object dynamics (*Eqs. 7 and 8*). By varying the magnitudes of arm stiffness ( $K_H$ ) and viscosity ( $B_H$ ) in this model, the relative

effects of these parameters on the resulting catch-trial kinematics were assessed. Simulated catch trials were generated for each subject (e.g., Fig. 7. *A* and *B*) based on values of arm stiffness and viscosity that were in the low to moderate range of those reported in the literature (i.e.,  $100 \text{ Nm}^{-1} < K_H < 500 \text{ Nm}^{-1}$  and  $5 \text{ Ns} \cdot \text{m}^{-1} < B_H < 25 \text{ Ns} \cdot \text{m}^{-1}$ ) (Burdet et al. 2000; Gomi and Kawato 1997; Tsuji et al. 1995). However, only when the arm model was assumed to incorporate low impedance (“LI model” in Fig. 7;  $K_H = 100 \text{ Nm}^{-1}$  and  $B_H = 5 \text{ Ns} \cdot \text{m}^{-1}$ ), did the model provide reasonably good predictions of the catch-trial behavior exhibited by these subjects. When the arm model was assumed to incorporate higher impedance (“HI model” in Fig. 7;  $K_H = 500 \text{ Nm}^{-1}$  and  $B_H = 25 \text{ Ns} \cdot \text{m}^{-1}$ ), the model predicted much smaller kinematic devi-

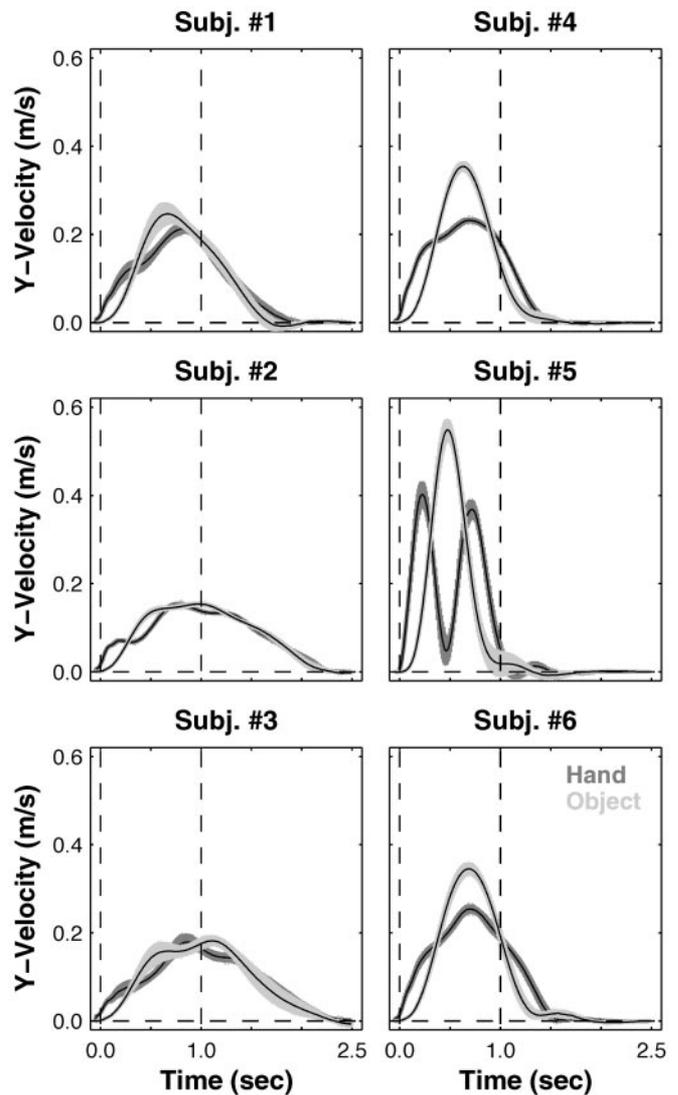


FIG. 5. Postadaptation reaching kinematics for all 6 subjects.  $y$ -direction velocity profiles for both the hand ( $\dot{y}_H$ ) and object ( $\dot{y}_O$ ) are shown. In each plot, dark lines represent averages and shaded areas represent  $\pm 95\%$  confidence intervals, each computed from the last 20 object-manipulation trials prior to the introduction of the catch trials. Dashed vertical lines indicate the start of movement ( $t = 0.0$  s) and the targeted movement time ( $t = 1.0$  s). All subjects exhibited postadaptation kinematics that were dramatically different from both their preexposure and initial-exposure kinematics (compare Fig. 4). These movement patterns were highly consistent for each subject, as indicated by the narrow 95% confidence intervals (CI).

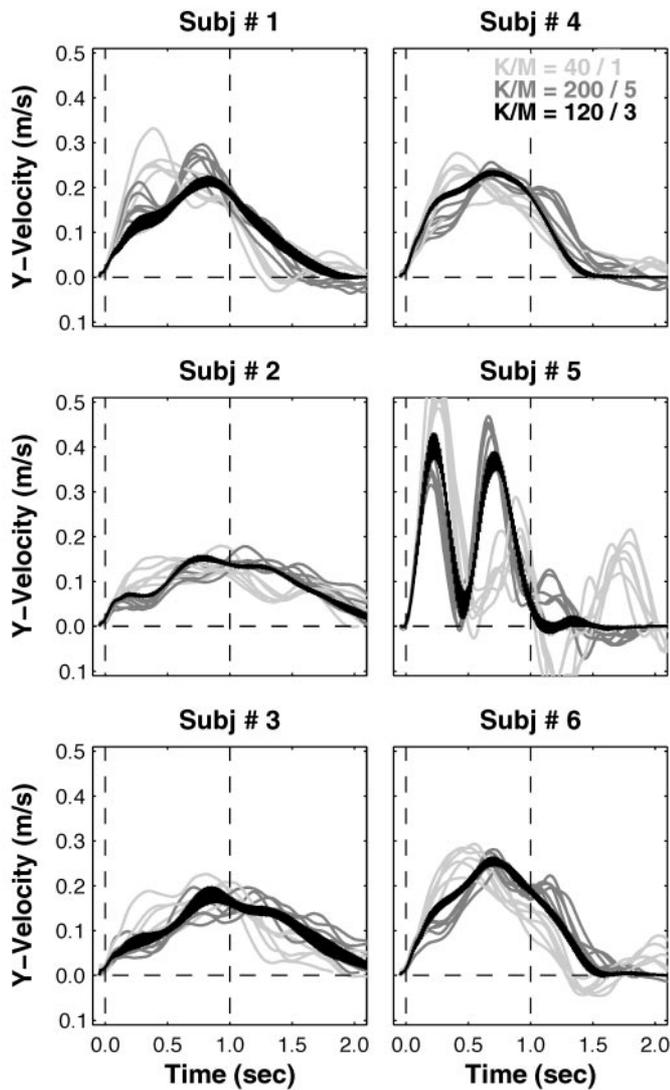


FIG. 6.  $y$ -direction velocity profiles for the hand ( $\dot{y}_H$ ) after learning the 3-kg object (95% CI shown in black) and for the individual catch trials with the 1 kg (light gray lines) and 5 kg (dark gray lines) objects for all subjects. Dashed vertical lines indicate the start of movement ( $t = 0.0$  s) and the targeted movement time ( $t = 1.0$  s). All subjects exhibited statistically significant deviations in their hand kinematics when exposed to either catch-trial object.

ations than observed experimentally. Still higher values of arm stiffness and damping lead to even smaller predicted deviations, as expected from the definition of these kinematically equivalent objects.

These findings were confirmed by the IRF scores computed over the first 500 ms of these movements for all subjects (Fig. 7C). In nearly all cases, the low-impedance (LI) model resulted in lower IRF scores (mean IRF = 0.460 excluding the 5-kg catch trials for *subject 1*), whereas IRF scores for the high-impedance (HI) model were almost always much closer to 1 (mean IRF = 0.896). These differences were highly statistically significant by the one-sided paired  $t$ -test ( $T = 5.43$ ;  $P = 1.03 \times 10^{-4}$ ). This finding indicates that the low-impedance predictions were a better fit to the actual catch trials, whereas the high-impedance model predicted catch-trial kinematics that would have more closely fit the postadaptation kinematics. The primary exception to these findings was in the 5-kg catch trials

for *subject 1* where, as shown in Fig. 6, these catch-trial trajectories exhibited deviations in the direction opposite to (i.e., higher rather than lower velocities) those of the other five subjects and of the model predictions. While it is not clear what exactly caused this subject to exhibit this particular behavior for these catch trials, it is clear that even these catch-trial

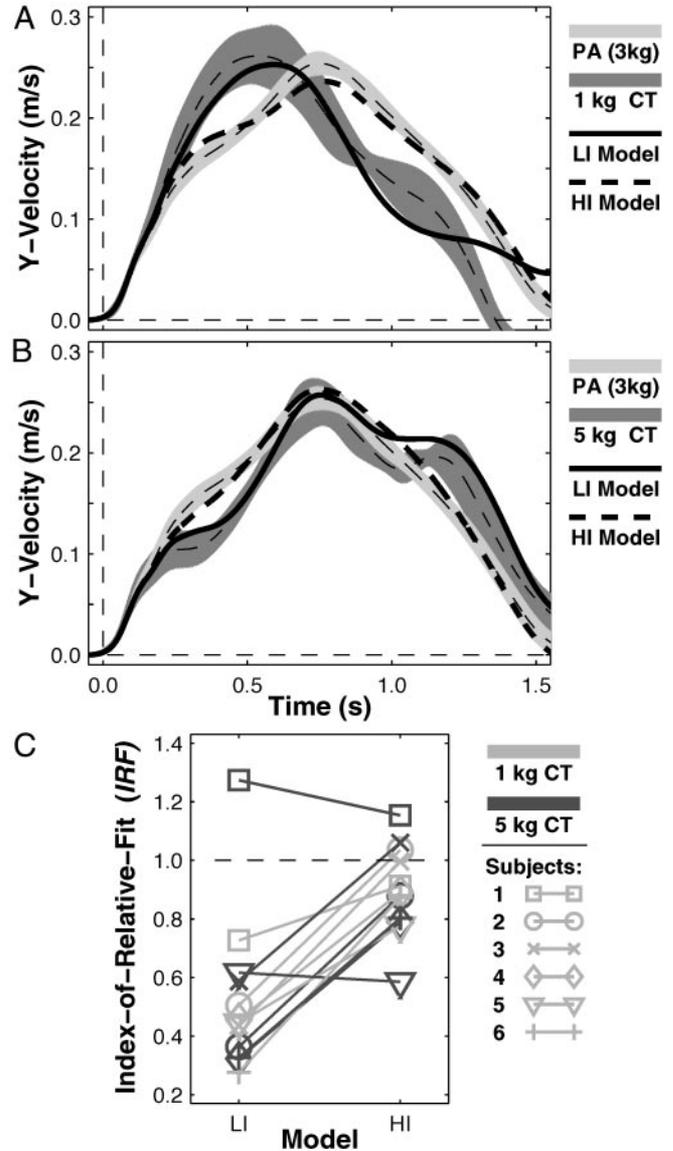


FIG. 7. *A*: example hand-velocity data (shaded areas) from a typical subject (*subject 6*) for 1-kg catch-trial data and predicted catch trials (black lines) based on the arm-plus-object model with either low impedance (LI) or high-impedance (HI) assumptions about the feedback gains in the internal model based controller (Eq. 7). *B*: example data (shaded areas) and associated model predictions (black lines) from the same subject for 5-kg catch trials. Light gray shaded areas in both subplots represent  $\pm 95\%$  CI for postadaptation trials ( $n = 20$ ), whereas dark gray shaded areas represent  $\pm 95\%$  CI for catch trials ( $n = 6$ ). Note that in both cases, the LI controller better predicts that actual catch-trial behavior, whereas the HI controller predicts much smaller deviations in hand kinematics. Similar results were obtained for all subjects. *C*: index of relative fit (IRF) scores for model predictions based on both LI and HI controllers for all subjects tested. In nearly all cases, the LI controller predictions more closely fit the actual catch-trial behavior (smaller IRF), whereas the HI controller predictions more closely fit the postadaptation behavior (IRF  $\rightarrow$  1).

kinematics could not have been generated by a control strategy that relied on globally increased arm impedance.

#### Feedback alone cannot predict exhibited behavior

The simulations shown in Fig. 7 indicate that the experimental catch-trial results were well predicted by a model system driven by a low-impedance controller that included an internal representation of the arm-plus-object dynamics. However, this finding does not preclude the possibility that the postadaptation and catch-trial trajectories exhibited by the subjects in the present study might be equally well predicted by a controller that did not include such an internal model. To address this question, simulations of the arm-plus-object system were performed using an alternative controller that relied solely on adjusting the overall gains for hand stiffness and viscosity for control (Eqs. 12 and 13). Two primary findings were evident from these simulations. In the first step of the modeling process, forward simulations of the postadaptation trajectories were generated by the arm-plus-object model using the original learned object properties ( $M_O = 3$  kg and  $K_O = 120$   $\text{Nm}^{-1}$ ) and the alternative controller. These simulations clearly demonstrated that the arm-plus-object model driven by the alternative controller was incapable of reproducing the original desired postadaptation behavior unless very large feedback gains ( $K_H \geq 1000$   $\text{Nm}^{-1}$  and  $B_H \geq 50$   $\text{Ns}\cdot\text{m}^{-1}$ ) were employed (Fig. 8A). However, the adoption of such large feedback gains leads to predictions of only minimal kinematic deviations during the catch trials (Fig. 8B). Thus if subjects had been using a control strategy based on setting the appropriate gains for a feedback controller that incorporated hand stiffness and viscosity alone, they could not have exhibited the catch-trial kinematics that they did in the present experiment.

Therefore the catch trial results exhibited in Fig. 6 were consistent with a controller that generated the appropriate sequence of forces required to achieve the desired kinematic result, with little or no increase in overall arm impedance beyond that exhibited during normal reaching (Fig. 7). These results could not be reproduced by assuming a controller that relied on setting hand stiffness and damping gains alone (Fig. 8). Thus subjects not only learned a new hand trajectory that allowed them to solve the task goals, but they also learned to directly compensate for the hand-object interface forces imposed by the mass-spring object along this trajectory. The catch-trial kinematics exhibited by these subjects demonstrated that they were employing a model-sensitive control law adapted to the specific object they had learned. This suggests that these subjects had formed a specific expectation of how the learned object would respond to the forces exerted on it. Thus in the most general sense, subjects had formed an internal representation of the mass-spring dynamics; i.e., they were able to predict the movements of the object in response to their executed motor commands (Imamizu et al. 2000; Wolpert et al. 1995). Subjects also had an expectation of the hand-object interface forces to be experienced and actively compensated for these forces. This suggests that subjects had also formed an internal representation of the sensory consequences of their actions; i.e., an efference-copy or re-efference component to the internal model (Wolpert and Kawato 1998; Wolpert et al. 1995).

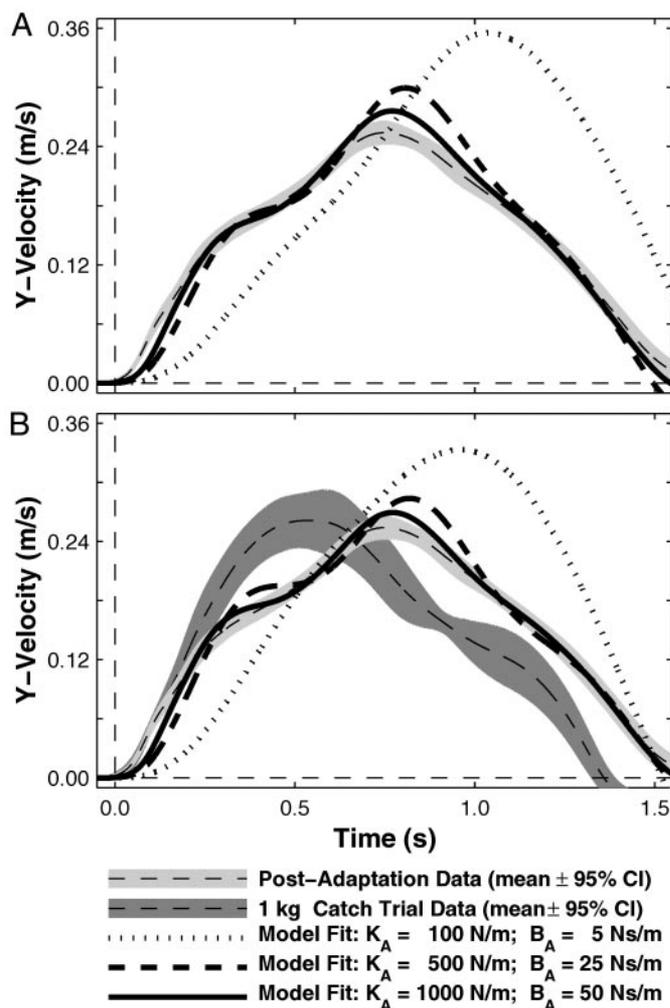


FIG. 8. A: example postadaptation hand-velocity data (shaded areas) from a typical subject (Subject #6) and predicted postadaptation catch trials (black lines) based on an arm-plus-object model driven by a model-independent controller (Eqs. 12 and 13). This alternative controller was capable of reproducing the original learned kinematics only when very large feedback gains ( $K_H \geq 1000$   $\text{Nm}^{-1}$  and  $B_H \geq 50$   $\text{Ns}\cdot\text{m}^{-1}$ ) were imposed. B: example data from the same subject for 1-kg catch trials (shaded areas) and associated model predictions (black lines) for these 1-kg catch trials. The model-independent controller did not reproduce the observed catch-trial behavior, regardless of the feedback gains employed. This controller was equally unable to predict the observed kinematics for the 5-kg catch trials. Similar results were obtained for all subjects.

#### DISCUSSION

Humans typically make point-to-point reaching movements in a stereotypical manner with the hand moving along a straight-line path with a bell-shaped velocity profile (Flash and Hogan 1985; Morasso 1981). This kinematic pattern is substantially disrupted when the mechanical properties of the arm are unexpectedly changed by the application of various external force fields, such as viscous fields (Conditt and Mussa-Ivaldi 1999; Conditt et al. 1997; Gandolfo et al. 1996; Shadmehr and Brashers-Krug 1997; Shadmehr and Moussavi 2000; Shadmehr and Mussa-Ivaldi 1994; Thoroughman and Shadmehr 1999) or Coriolis fields (Cohn et al. 2000; Dizio and Lackner 1995; Lackner and Dizio 1994). However, after sufficient exposure to these force fields, subjects typically recover their original unperturbed reaching kinematics, demonstrating

that they adapt their motor commands to precisely cancel out the effects of these fields. This adaptation is not accomplished by globally increasing arm impedance to resist perturbations (Shadmehr and Mussa-Ivaldi 1994) nor is it accomplished by memorizing the specific sequences of forces encountered (Condit et al. 1997). Similarly, when humans manipulate rigid objects, these objects impose inertial force fields on the arm. The perturbing effects of these inertial forces are adapted out in much the same way as those of viscous or Coriolis force fields (Blakemore et al. 1998; Bock 1990, 1993; Gordon et al. 1993; Sainburg et al. 1999; Witney et al. 2000)

While these studies provide substantial evidence that humans possess an accurate and adaptable internal model of the dynamics of their own arm, it is not known if humans use similar model-based control strategies to learn dynamical models of systems outside the body (i.e., those involving additional degrees of freedom not associated with the limbs). If the CNS constructs and uses internal models as a general strategy for controlling movement, then the use of such model-dependent strategies should extend to other types of motor-skill learning tasks as well. The present study was designed to determine if the strategy adopted by humans when learning to reach to a target while holding a mass-spring object in their hand is consistent with a controller that employs some form of internal representation of the object's dynamical behavior. It was hypothesized that, through practice, subjects would gradually develop and make use of an internal representation of the dynamics of the mass-spring object. The alternative hypotheses that subjects would minimize perturbations imposed by the object by slowing down or would globally increase hand impedance to resist those perturbations were also tested. Subjects did not adopt either of these two alternative control strategies but instead adopted a control strategy that was specifically adapted to the dynamics of the object they had learned, thus providing substantial evidence of internal model formation.

#### *Slower speeds and feedback control*

Each of the mass-spring objects used had a resonant frequency of approximately 1 Hz. Thus the chosen target time ( $T_i = 1.0$  s) corresponded to one natural oscillation of the object. Initial simulations had demonstrated that the task could be successfully completed with reasonably smooth force profiles well within physiological limits. However, even after a substantial amount of practice, subjects were still unable to successfully complete the required task within the desired time frame (Fig. 3). There are several possible reasons for this. First, it has been suggested that humans learn only *local* models, i.e., based only on the range of dynamical experiences to which they have been exposed (Gandolfo et al. 1996). This suggests that to construct a broad and accurate internal model of object dynamics, subjects would need to be exposed to those dynamics across a wide range of states. In the present experiment, subjects made only the single reaching movement in the positive  $y$  direction. Thus the restricted nature of the task itself limited the variety of subject's interactions with the mass-spring dynamics, and this may have in turn limited their capacity to form a complete and accurate internal model. Second, the reaching task itself was highly susceptible to small errors and/or perturbations. Thus even if subjects were able to construct an adequate internal model of mass-spring dynamics,

enacting such a strategy successfully would require a level of precision not typically available to biological systems. This possibility was substantiated by the large trial-to-trial variability exhibited in the movement time data (Fig. 3). Finally, it is possible and likely that subjects initially slowed their movement speeds to make use of visual and/or proprioceptive feedback for on-line control. The fact that their movement times continued to decrease across the duration of the experiment suggests that as learning progressed, subjects relied more on predictive feed-forward commands for control and less on on-line feedback control. Had subjects been given more practice over a longer period of time (perhaps several sessions over more days), their performance would likely have continued to improve.

The goal of the object-manipulation task used in the present experiment was for subjects to maintain control of the final state (i.e., position and velocity) of the object at the end of the reaching movement. The only direct sensory information available to the subject about these object states was through visual feedback. Thus to use on-line visual feedback control in the present experiment, subjects would have needed to move slowly enough for visual feedback to become a viable source of information for that control. For any feedback control system with delays, there is a maximum frequency that the controller can effectively respond to, which is approximately 1/20th of the inverse of the delay time (Hogan 1988; Hogan et al. 1987). For visual control, where the feedback delays are on the order of 200 ms, this translates into a maximum frequency of approximately 0.25 Hz. The object used in the present study oscillated at a resonant frequency of 1 Hz, well above this limit. Thus it is highly unlikely that subjects could have used a strategy based solely on visual feedback control alone. Two remaining alternative strategies that subjects might have employed would be to either actively cancel the disturbances imposed by the object by somehow internally estimating the object's dynamics or to co-contract their muscles to resist those forces. Indeed, when low-frequency (approximately 0.2 Hz) perturbations were applied to monkey forearms, these perturbations were resisted by reciprocal activation of agonist-antagonist muscle pairs (Humphry and Reed 1983). However, as the perturbation frequency was increased, muscle co-activation also increased, becoming significant for perturbation frequencies of 1 Hz and higher. In the present experiment, subjects experienced 1-Hz perturbations and thus in the absence of any internal model, would be expected to exhibit significant increases in muscle co-activation. The fact that no evidence of muscle co-activation was exhibited by these subjects (i.e., no increases in hand impedance above levels that would be anticipated during unperturbed reaching) strongly suggests that these subjects learned to control the movements of the object by predicting the object's dynamics.

#### *Adaptation vs. skill learning*

In previous experiments involving force fields that depended only on arm state variables, subjects' reaching kinematics always returned to their original unperturbed patterns after sufficient practice (Cohn et al. 2000; Condit and Mussa-Ivaldi 1999; Condit et al. 1997; Dizio and Lackner 1995; Lackner and Dizio 1994; Shadmehr and Brashers-Krug 1997; Shadmehr and Moussavi 2000; Shadmehr and Mussa-Ivaldi 1994; Thor-

oughman and Shadmehr 1999). This behavior has been termed “motor adaptation” because the effects of these force fields are gradually cancelled out by the controller so that the mechanical system being controlled (i.e., the arm) returns to its preadapted kinematics. In the present experiment, subjects’ hand kinematics did not return to their unperturbed trajectories but remained substantially altered even after several hundred trials (Fig. 5). Therefore these subjects did not “adapt-out” the effects of the mass-spring object but instead adopted a completely new kinematic pattern. This process was necessitated by the fact that subjects could not directly control the position of the mass but had to act through the dynamics imposed by the spring. This process is analogous to learning of a new motor skill where the goal of the controller is to find a coordination pattern appropriate for performing the specified task (Newell and Corcos 1993). Thus in the context of previous motor-adaptation experiments, the task used in the present study should be classified as a motor-skill learning task rather than a motor-adaptation task. The present experiment therefore extends previous findings that humans use low-impedance predictive control for executing arm movements beyond the context of motor adaptation and into the context of motor-skill learning, where control may also be applied to objects having additional degrees of freedom not associated with the limbs.

#### *Analogy of catch trials to aftereffects*

Subjects in the present experiment exhibited substantial kinematic deviations from their postadaptation behavior when they performed the reaching task with either of the two kinematically equivalent catch-trial objects (Fig. 6). These deviations were well predicted by a simplified one-dimensional model of arm-plus-object dynamics (Eqs. 5 and 6) driven by a controller that consisted of an inverse model of the arm-plus-object system (Eq. 8) working in parallel with error-feedback terms (Eq. 7) incorporating relatively low values of hand stiffness and damping (Fig. 7). The experimental results therefore demonstrated that these subjects were not globally increasing hand impedance to enforce a specific hand trajectory because the simulation results demonstrated that the perturbations introduced by the catch-trial objects could have been successfully resisted by even moderate increases in hand impedance above those levels typically exhibited during unperturbed reaching.

In previous experiments where subjects adapted their reaching movements in the presence of external force fields, when these force fields were unexpectedly removed, subjects exhibited kinematic deviations in the opposite direction to the deviations they exhibited when they were initially exposed to the force fields (Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994). These so-called “aftereffects” demonstrated that motor adaptation was accomplished by directly opposing the forces applied to the hand and not by adopting the more general strategy of globally increasing hand impedance to resist perturbations (Shadmehr and Mussa-Ivaldi 1994). The interpretation of the deviations exhibited in the catch trials from the present experiment is directly analogous to the interpretation of the aftereffects exhibited during the force-field learning studies. In both contexts, the trajectories exhibited by subjects during the catch trials were consistent with the hypothesis that subjects adjusted their feed-forward control commands to di-

rectly compensate for the sequence of hand-object interface forces encountered along the desired hand trajectory. This type of compensation necessarily implies that these subjects had an expectation of that sequence of interface forces, which implies that they had formed, in the most general sense, a predictive representation of the arm-plus-object system they were controlling (Imamizu et al. 2000; Miall and Wolpert 1996; Wolpert and Kawato 1998; Wolpert et al. 1995).

#### *Other possible interpretations*

In the present study, subjects learned to complete the mass-spring object-manipulation task successfully by using a feed-forward control strategy that directly compensated for the interface forces imposed by the object on the hand along the desired trajectory. The presence of this feed-forward control command [ $U(t)$  in Eq. 10] required that subjects form a specific expectation about how the object would respond to the forces exerted on it. Having such an expectation satisfies the general definition of an internal model (Imamizu et al. 2000; Wolpert et al. 1995). Specifically, these findings suggest that this internal model was composed as a specific mapping between hand-object interface forces and state variables (i.e., positions and velocities) of both the hand and object. Furthermore, the modeling results demonstrated that the experimental results were consistent with a controller that incorporated an inverse model of the arm-plus-object dynamics. However, this does not imply that the controller employed by the human brain was necessarily of the exact same form.

In the forward simulations used to predict the postadaptation and catch-trial trajectories, the feed-forward command,  $U(t)$ , was an arbitrary time-varying control signal used to drive the plant (Eq. 9). In these simulations,  $U(t)$  was obtained by explicitly computing the inverse dynamics from the postadaptation trials (Eqs. 8 and 10). While the subjects in this experiment might have learned the correct  $U(t)$  in this same manner, they could also have learned the same  $U(t)$  by some alternate means that did not require the formation of an explicit inverse model of the arm-plus-object dynamics. For example, it is possible that subjects in the present experiment learned only a “local” model by memorizing the specific sequence of forces required to move this single object along the single desired trajectory that accomplished the task goal. Subjects could have done this in a manner consistent with reinforcement learning techniques by learning a “value function” that assigned an expected reward to each control action executed at each state (Sutton and Barto 1998). This direct-reinforcement learning approach is sometimes presented in contrast to learning a more explicit model of the controlled dynamics. While both alternatives represent (implicitly) a type of model for the controlled system, it was not possible to distinguish between these more specific alternatives within the context of the present experiment.

If subjects merely learned this type of local model of the mass-spring object’s behavior, then their ability to successfully manipulate this object would not generalize beyond the narrow boundaries of the learned task. When subjects learn to adapt their reaching movements in viscous force fields, they appear to learn these fields as a general mapping between arm states and the forces experienced at the hand. By learning such a mapping, subjects can generalize their learned reaching behav-

ior to new movements that reside within the same state-force space (Conditt et al. 1997; Shadmehr and Mussa-Ivaldi 1994). However, this adaptation is largely restricted to the range of states that are experienced during the learning process and decays smoothly and rapidly when subjects are asked to make movements involving arm states not visited during training (Gandolfo et al. 1996). If subjects learn to manipulate external objects by forming similar mappings between *object* state variables and the forces experienced at the hand, then this learned mapping might be expected to generalize to other movements that visit similar object states and forces but might not generalize to movements that visit regions of the state-to-force space that were not adequately explored during training. Further experiments will be necessary to explore the capacity of subjects to generalize their behavior in these types of object-manipulation tasks.

### Conclusions

The present study was designed to determine if humans adopt model-dependent control strategies when learning a novel motor-skill task in which the mechanical system being controlled (i.e., the arm) is altered by the addition of degrees-of-freedom external to the arm. Subjects from the present experiment were able to first acquire and then improve their capacity to control the behavior of the mass-spring object. Although different subjects exhibited different postadaptation hand kinematics, unexpectedly replacing the 3-kg training object with either of the 1- or 5-kg catch-trial objects resulted in substantial kinematic deviations from the postlearning movement trajectories. These deviations were consistent with those predicted by a model of the arm-plus-object system driven by a low-impedance controller that incorporated an explicit inverse model of arm-plus-object dynamics. The observed behavior could not be reproduced by a controller that relied on modulating hand impedance alone with no inverse model. These findings refuted the hypothesis that subjects had globally increased arm impedance to enforce a prespecified hand trajectory. Each of these findings is consistent with the hypothesis that learning was associated with the formation of an internal model of the dynamic behavior of the mass-spring object and inconsistent with either of the alternative hypotheses tested. Although there may still be other alternatives for acquiring successful control strategies for this task that do not rely on the implementation of an explicit inverse model as a part of the controller, further experiments will be required to test these alternatives. Thus the results of the present study extend previous work on force field adaptation by demonstrating that the learning of such internal models may also apply to the control of dynamical systems that extend beyond our own bodies.

Partial funding for this work was provided by National Institutes of Health Grants T32-HD-07418, F32-HD-08620-01, and NS-35673 and National Science Foundation Grant BES-9900684.

### REFERENCES

- BENNETT DJ. Stretch reflex responses in the human elbow joint during a voluntary movement. *J Physiol (Lond)* 474: 339–351, 1994.
- BENNETT DJ, HOLLERBACH JM, XU Y, AND HUNTER IW. Time-varying stiffness of human elbow joint during cyclic voluntary movement. *Exp Brain Res* 88: 433–442, 1992.
- BLAKEMORE SJ, GOODBODY SJ, AND WOLPERT DM. Predicting the consequences of our own actions: the role of sensorimotor context estimation. *J Neurosci* 18: 7511–7518, 1998.
- BOCK O. Load compensation in human goal-directed arm movements. *Behav Brain Res* 41: 167–177, 1990.
- BOCK O. Early stages of load compensation in human aimed arm movements. *Behav Brain Res* 55: 61–68, 1993.
- BURDET E, OSU R, FRANKLIN DW, YOSHIOKA T, MILNER TE, AND KAWATO M. A method for measuring endpoint stiffness during multi-joint arm movements. *J Biomech* 33: 1705–1709, 2000.
- COHN JV, DIZIO P, AND LACKNER JR. Reaching during virtual rotation: context specific compensations for expected coriolis forces. *J Neurophysiol* 83: 3230–3240, 2000.
- CONDITT MA, GANDOLFO F, AND MUSSA-IVALDI F. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J Neurophysiol* 78: 554–560, 1997.
- CONDITT MA AND MUSSA-IVALDI FA. Central representation of time during motor learning. *Proc Natl Acad Sci* 96: 11625–11630, 1999.
- DIZIO P AND 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.
- FLASH T AND HOGAN N. The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci* 5: 1688–1703, 1985.
- GANDOLFO F, MUSSA-IVALDI FA, AND BIZZI E. Motor learning by field approximation. *Proc Natl Acad Sci USA* 93: 3843–3846, 1996.
- GOMI H AND KAWATO M. Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science* 272: 117–120, 1996.
- GOMI H AND KAWATO M. Human arm stiffness and equilibrium-point trajectory during multi-joint movement. *Biol Cybern* 76: 163–71, 1997.
- GOMI H AND OSU R. Task-dependent viscoelasticity of human multijoint arm and its spatial characteristics for interaction with environments. *J Neurosci* 18: 8965–8978, 1998.
- GORDON AM, WESTLING G, COLE JK, AND JOHANSSON RS. Memory representations underlying motor commands used during manipulation of common and novel objects. *J Neurophysiol* 69: 1789–1796, 1993.
- HOGAN N. Planning and execution of multijoint movements. *Can J Physiol Pharmacol* 66: 508–517, 1988.
- HOGAN N, BIZZI E, MUSSA-IVALDI FA, AND FLASH T. Controlling multijoint motor behavior. *Exerc Sport Sci Rev* 15: 153–190, 1987.
- HUMPHRY DR AND REED DJ. Separate cortical systems for control of joint movement and joint stiffness: reciprocal activation and coactivation of antagonist muscles. *Adv Neurol* 39: 347–372, 1983.
- IMAMIZU H, MIYAUCHI S, TAMADA T, SASAKI Y, TAKINO R, PUTZ B, YOSHIOKA T, AND KAWATO M. Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature* 403: 192–195, 2000.
- LACKNER JR AND DIZIO P. Rapid adaptation to coriolis force perturbations of arm trajectory. *J Neurophysiol* 72: 299–313, 1994.
- LYNCH KM AND MASON MT. Dynamic nonprehensile manipulation: controllability, planning and experiments. *Intl J Robot Res* 18: 64–92, 1999.
- LYNCH KM, SHIROMA N, ARAI H, AND TANIE K. Collision-free trajectory planning for a 3-DOF robot with a passive joint. *Intl J Robot Res* 19: 1171–1184, 2000.
- MAH CD. Spatial and temporal modulation of joint stiffness during multijoint movement. *Exp Brain Res* 136: 492–506, 2001.
- MIALL RC AND WOLPERT DM. Forward models for physiological motor control. *Neural Networks* 9: 1265–1279, 1996.
- MORASSO P. Spatial control of arm movements. *Exp Brain Res* 42: 223–227, 1981.
- NEWELL KM AND CORCOS DM. Issues in variability and motor control. In: *Variability and Motor Control*, edited by Newell KM and Corcos DM. Champaign, IL: Human Kinetics Publishers, 1993, p. 1–12.
- PERRAULT EJ, KIRSCH RF, AND CRAGO PE. Effects of voluntary force generation on the elastic components of endpoint stiffness. *Exp Brain Res* 141: 312–323, 2001.
- SAINBURG RL, GHEZ C, AND KALAKANIS D. Intersegmental dynamics are controlled by sequential anticipatory, error correction, and postural mechanisms. *J Neurophysiol* 81: 1045–1056, 1999.
- SCHAAL S AND ATKESON CG. Robot juggling: an implementation of memory-based learning. *Control Syst Mag* 14: 57–71, 1994.

- SCHEIDT RA, REINKENSMAYER DJ, CONDITT MA, RYMER WZ, AND MUSSA-IVALDI FA. Persistence of motor adaptation during constrained, multi-joint, arm movements. *J Neurophysiol* 84: 853–862, 2000.
- SHADMEHR R AND BRASHERS-KRUG T. Functional stages in the formation of human long-term motor memory. *J Neurosci* 17: 409–419, 1997.
- SHADMEHR R AND MOUSSAVI ZMK. Spatial generalization from learning dynamics of reaching movements. *J Neurosci* 20: 7807–7815, 2000.
- SHADMEHR R AND MUSSA-IVALDI FA. Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14: 3208–3224, 1994.
- SUTTON RS AND BARTO AG. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 1998.
- THOROUGHMAN KA AND SHADMEHR R. Electromyographic correlates of learning an internal model of reaching movements. *J Neurosci* 19: 8573–8588, 1999.
- TSUJI T, MORASSO PG, GOTO K, AND ITO K. Human hand impedance characteristics during maintained posture. *Biol Cybern* 72: 475–485, 1995.
- WINTER DA. *Biomechanics and Motor Control of Human Movement* (2nd ed.). New York: Wiley, 1990.
- WITNEY AG, GOODBODY SJ, AND WOLPERT DM. Learning and decay of prediction in object manipulation. *J Neurophysiol* 84: 334–343, 2000.
- WOLPERT DM, GHARAMANI Z, AND JORDAN MI. An internal model for sensorimotor integration. *Science* 269: 1880–1882, 1995.
- WOLPERT DM AND KAWATO M. Multiple paired forward and inverse models for motor control. *Neural Networks* 11: 1317–1329, 1998.
- XU Y AND HOLLERBACH JM. A robust ensemble data method for identification of human joint mechanical properties during movement. *IEEE Trans Biomed Eng* 46: 409–419, 1999.