

Functional Stages in the Formation of Human Long-Term Motor Memory

Reza Shadmehr¹ and Thomas Brashers-Krug²

Departments of ¹Biomedical Engineering and ²Psychiatry and Behavioral Sciences, The Johns Hopkins University, Baltimore, Maryland 21205-2195

Previous research has demonstrated that the primate CNS has the ability to learn and store multiple and conflicting visuo-motor maps. Here we studied the ability of human subjects to learn to make reaching movements while interacting with one of two conflicting mechanical environments as produced by a robotic manipulandum. We demonstrate that two motor maps may be learned and retained, but only if the training sessions in the tasks are separated by an interval of ~5 hr. If the interval is shorter, learning of the second map begins with an internal model appropriate for the first task and performance in the second task is significantly impaired. Analysis of the after-effects suggests that with a short temporal distance, learning of the second task leads to an unlearning of the internal model for the first. With the longer temporal distance, learning of the second task starts with an unbiased internal model, and performance approaches that of naives. Furthermore, the memory

of the consolidated skill lasts for at least 5 months after training. These results argue for a distinct change in the state of resistance of motor memory (to disruption) within a few hours after acquisition. We suggest that motor practice results in memories that have at least two functional components: soon after completion of practice, one component fades while another is strengthened. A further experiment suggests that the hypothetical first stage is not merely a gateway to long-term memory, but also temporary storage for items of information, whether new or old, for use in the near-term. Our results raise the possibility that there are distinct neuronal mechanisms for representation of the two functional stages of motor memory.

Key words: motor learning; motor memory; consolidation; short-term memory; long-term memory; reaching movements; internal models; virtual environments; motor control

In novelty stores, one can find an object that appears to be a heavy brick but is actually constructed of light plastic. When one is asked to rapidly move the “brick,” the result is a flailing-like arm motion. This observation suggests that in programming the motor output to the muscles of the arm, the CNS uses an internal model (Wolpert et al., 1995b) to predict the mechanical dynamics of the task (Gottlieb, 1994). In theory, the internal model (IM) is an association from a desired trajectory for the hand (Wolpert et al., 1995a) to a pattern of muscle torques (Shadmehr and Mussa-Ivaldi, 1994). Because, in principle, this map is unique for the objects that we have learned to interact with, “motor memory” may be thought to contain, at least in part, a collection of IMs where visual information serves as an identifying cue that allows for binding of an appropriate association (Gordon et al., 1993), i.e., recall.

Because we routinely use our hands to interact with a remarkably diverse variety of mechanical systems, the ability to learn and recall IMs is likely a fundamental property of the motor system. Indeed, practice of arm movements with a novel mechanical system leads to formation of an IM for that system. The evidence for this comes from EMG studies (Milner and Cloutier, 1993; Gottlieb, 1994; Thoroughman and Shadmehr, 1996), and from studies that have quantified movement trajectories when the me-

chanical system's dynamics have been unexpectedly changed (Sanes, 1986; Lackner and Dizio, 1994; Shadmehr and Mussa-Ivaldi, 1994; Gandolfo et al., 1996). Once the IM is acquired, it becomes available for “recall”; performance is significantly improved when tested 24–48 hr later (Shadmehr et al., 1995). In this report, we show that the improvement in performance persists for at least 5 months, suggesting the formation of long-term motor memories.

Although we know little about the processes that culminate in long-term motor memory formation (Halsband and Freund, 1993; Salmon and Butters, 1995), a feature of memory across the animal kingdom is that it continues to develop after practice has stopped (Seeds et al., 1995). In general, memory appears to progress functionally from a short-lived fragile form to a long-lasting stable form (Bailey and Kandel, 1995; DeZazzo and Tully, 1995). Phases of memory are often distinguished with respect to their sensitivity to new experiences and susceptibility to interference and injury (Tully et al., 1994; Hammer and Menzel, 1995). The progression to long-term memory is referred to as consolidation, and the time during which information becomes consolidated has been used to functionally define short-term memory (Fuster, 1995).

Does formation of motor memory progress from a short-term, fragile form to a long-term, stable form? The distinction is not merely semantic. Differences in the functional properties of phases that culminate in long-term memory have suggested that training sets in motion events that develop in parallel biochemical pathways (Tully et al., 1994) in possibly distinct anatomical sites (Rose, 1991; de Belle and Heisenberg, 1994). For example, during storage of “declarative” information (Squire, 1992), distinct regions of the brain are believed to encode the memory during the

Received July 23, 1996; revised Oct. 3, 1996; accepted Oct. 8, 1996.

This work was funded in part by grants from the Office of Naval Research and the Whitaker Foundation to R.S. This work has been greatly enriched because of our interactions with Dr. Emilio Bizzi. Kurt Thoroughman, Maurice Smith, and Kasra Akhavan-Toyserkani provided technical support of the Hopkins manipulandum.

Correspondence should be addressed to Dr. Reza Shadmehr, Traylor 419, The Johns Hopkins University School of Medicine, Baltimore, MD 21205-2195.

Copyright © 1996 Society for Neuroscience 0270-6474/96/170409-11\$05.00/0

short-term and long-term stages (Squire et al., 1984a; Alvarez and Squire, 1994; Guigon and Burnod, 1995; McClelland et al., 1995). The hippocampus (Zola-Morgan and Squire, 1990; Kim et al., 1995) and its inputs (Rashidy-Pour et al., 1996) appear to play a time-limited role in consolidation of declarative memory. Learning of visuo-motor skills, however, does not appear to depend on the integrity of medial temporal lobe structures (Corkin, 1968; Gabrieli et al., 1993). Furthermore, interventions that interfere with formation of long-term declarative memory appear to spare retention of visuo-motor skills (Squire et al., 1984b). However, motor memories are also vulnerable (Lewis et al., 1951; Lewis and Miles, 1956; Heilman and Gonzalez-Rothi, 1985; Clark et al., 1994). Our results have suggested recently that retention of a newly acquired IM could be disrupted when a second IM, anti-correlated to the first, was learned (Brashers-Krug et al., 1995b). However, if IM₂ was learned beyond a critical time window (~4 hr) after acquisition of IM₁, it had little effect on recall of IM₁ (Brashers-Krug et al., 1996). In other words, within hours, the representation of IM₁ became gradually less vulnerable to the “intervention” caused by learning of IM₂. Furthermore, the ability to learn IM₂ became progressively better with temporal distance from IM₁ (Brashers-Krug et al., 1996). That is, subjects had an easier time learning IM₂ when 4 hr had passed since learning IM₁. This is surprising; if learning of IM₂ involves some form of unlearning of IM₁, then one would expect from the initial labile form of IM₁ that it should be easier to learn IM₂ when the temporal interval between the two training sessions is short. The opposite was observed. To understand this paradox, we report on further experiments that use the concept of after-effects to quantify the effect of time on representation of a recently acquired motor skill.

MATERIALS AND METHODS

The purpose of our experiment was to reveal functional properties of processes that lead to formation of long-term motor memory. The motor task we considered was one in which human subjects learned to make reaching movements while holding the handle of a robot manipulandum (Fig. 1). A mathematical model was developed to provide a framework for the human/robot force interaction.

Experimental setup. Sixty right-handed subjects with no known neurological history, ranging in age from 19 to 37 years, participated in this study. The procedures were approved by the Johns Hopkins University Joint Committee on Clinical Investigation, and all subjects signed an informed consent form.

Subjects learned to make reaching movements while interacting with a force producing manipulandum. A schematic and photo of the measurement apparatus are shown in Figure 1. Each subject was seated on a chair that was bolted onto an adjustable positioning mechanism and instructed to grip the handle of a robot manipulandum with the right hand. The right upper-arm was supported in the horizontal plane by a rope attached to the ceiling.

The Hopkins manipulandum is a two degree of freedom, portable, lightweight (0.8 kg for the shoulder link and 1.3 kg for the elbow link, including the force transducer), low-friction (0.02 and 0.06 N·m·sec viscous friction for shoulder and elbow joints) mechanism built based on the mechanical design principles of the MIT device (Faye, 1986; Charnnarong, 1991; Hogan et al., 1992) used in our previous works. Two low-inertia, DC brushless torque motors (Kollmorgen Corp., model RBEH-3003) driven by a pair of digital pulse-width-modulated servoamplifiers (Kollmorgen, model FAST Drive) were mounted on the base of the robot and independently delivered torque to the robot's shoulder and elbow joints via a parallelogram configuration. Robot's shoulder and elbow joint position measurements were made using absolute optical encoders (Gurley Precision Instruments) with a resolution of 0.0055°. Robot's shoulder and elbow joint velocity measurements were made using a system composed of incremental optical encoders, interpolators, and digital integrators that resulted in a resolution of better than 440,000

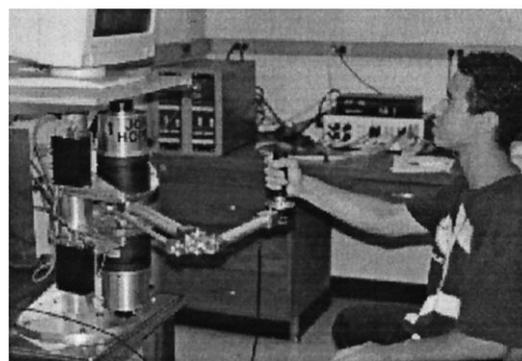
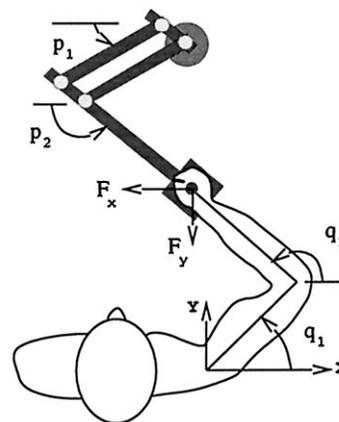


Figure 1. The robot manipulandum and the experimental setup. The manipulandum is a very low-friction, planar mechanism powered by two high-performance torque motors. The subject grips the handle of the robot. The handle houses a force transducer. The video monitor facing the subject displays a cursor corresponding to the position of the handle. A target position is displayed, and the subject makes a reaching movement. With practice, the subject learns to compensate for the forces produced by the robot.

counts per revolution. The handle of the robot housed a 6-axis force/torque transducer (Assurance Technologies, Inc).

Experimental procedures. The experimental task was similar to that described in our previous reports (Shadmehr and Mussa-Ivaldi, 1994; Brashers-Krug et al., 1996). Subjects moved the cursor corresponding to the position of their hand to a target position that would appear at 10 cm in one of eight directions: four directions starting from the center of the monitor (0°, 45°, 90°, and 135°) and the four corresponding directions back to the center from each of those targets (180°, 225°, 270°, and 315°). Subjects were instructed that there was a timing goal for the task. The timer started as soon as subjects began their movement and stopped when they ended their movement at the target. If they reached the target in the allotted time (500 ± 50 msec), the target would make a distinctive sound. If they reached it too late, the target would turn blue, and if they reached it too soon, it would turn red. Generally, after 400 targets subjects were able to move at the required pace. There were no perturbing forces during these movements and no further analysis was made of this data. We refer to this situation when the robot motors were inactive as the *null field* condition. Visual feedback regarding hand trajectory was provided throughout the entire experiment.

Subjects returned on a subsequent day and were tested on the null field. The trajectories recorded during this condition are referred to as *baseline trajectories*; they are straight-line movements with “bell-shaped” linear velocity profiles and represent what we consider to be the desired trajectory of the biological adaptive control system (after introduction of perturbing forces, movement kinematics generally converge back to this trajectory). After measurement of baseline trajectories, a brief period of rest was provided (2–3 min), after which the subjects were told that the robot motors would now produce forces on their hand. Subjects were asked to move the handle (at their own pace and without any targets) and experience the forces for 10–15 sec, after which we began a target set. A

target set consisted of 192 targets, all in a force field, except for 33 randomly chosen targets during which no forces were present. The latter group allowed us to quantify subjects' after-effects.

To produce a forcefield, the motors were programmed as a function of hand velocity, \dot{x} , and a viscous matrix B . The effective force acting on the subject's hand was: $f = B\dot{x}$, where B was either equal to B_1 ,

$$B_1 = \begin{bmatrix} 0 & 13 \\ -13 & 0 \end{bmatrix} \text{ N} \cdot \text{m} \cdot \text{sec},$$

or B_2 , where $B_2 = -B_1$. The method used for producing such fields has been described elsewhere (Shadmehr and Mussa-Ivaldi, 1994).

Subjects were assigned into one of six groups. All groups except for group 1 practiced for three target sets in each of the two fields (group 1 practiced for 3 target sets only in B_1). The difference among the remaining five groups was the temporal distance between the practice sessions in B_1 and B_2 : this temporal distance was 5 min, 30 min, 2.5 hr, 5.5 hr, or 24 hr.

Group 1 subjects returned 5 months after their initial training and were again tested briefly in the null field (~20 targets) and then for a target set in B_1 . The remaining groups returned 1 week after completion of training in B_2 and were tested in B_1 . The group that had a temporal distance of 24 hr between B_1 and B_2 was tested 1–3 d later on B_2 .

Data analysis. We sampled the manipulandum's joint angles, joint velocities, and forces at the handle at a rate of 100 Hz and computed hand positions and velocities. Trajectories were aligned using a velocity threshold at the onset of movement. The performance measure was the similarity between the hand trajectory in the force field and a "typical" baseline trajectory (in the null field). This similarity was defined as a correlation between two time series of hand velocity vectors (Shadmehr and Mussa-Ivaldi, 1994). A typical baseline trajectory for a subject was found by correlating each trajectory with all the other trajectories for that target direction and finding the one with the highest average correlation.

An *after-effect* is the trajectory that results when a subject is expecting a force field but the robot is producing a null field. After-effects were analyzed using two indices. The first index was a distance measure that quantified how far the hand path had deviated from a straight line to the target. This distance was measured 300 msec into the movement. At this interval, the aftereffect is near its maximum deviation from a straight line. The second index was a force measure that quantified the difference between the force produced by the subject during an after-effect and the force recorded from the same subject for a typical movement in the baseline condition. This variable is a time series of force vectors. Because the force fields were always perpendicular to the direction of target, we computed the component of the force measure that was perpendicular to the target. The result was a time dependent scalar force variable.

Mathematical modeling. The purpose of the mathematical modeling was to predict force and position trajectories that result as an adaptive controller learns an internal model of the mechanical dynamics produced by the robot. The adaptive controller was modeled to reasonably estimate the biomechanical behavior of the human arm. We built on the ideas introduced in our previous work (Shadmehr and Mussa-Ivaldi, 1994). The current model takes into account the passive dynamics of the robot manipulandum as well as the passive dynamics of the subject's arm. This allows us to predict fairly accurately the patterns of motion and force generation in the case where we assume that the subject has learned a specific internal model of the task.

When we position a force transducer at the interaction point between the robot and the subject (i.e., the handle), we can write the dynamics of the four link system in Figure 1 in terms of the following coupled-vector differential equation:

$$I_r(p)\ddot{p} + G_r(p, \dot{p})\dot{p} = E(\dot{p}) + J_r^T F, \quad (1)$$

$$I_s(q)\ddot{q} + G_s(q, \dot{q})\dot{q} = C(q, \dot{q}, q^*(t)) - J_s^T F, \quad (2)$$

where I and G are inertial and coriolis/centripetal matrix functions, E is the torque field produced by the robot's motors, i.e., the environment, F is the force measured at the handle of the robot, C is the controller implemented by the motor system of the subject, $q^*(t)$ is the reference trajectory planned by the motor control system of the subject, J is the Jacobian matrix describing the differential transformation of coordinates from endpoint to joints, q and p are column vectors representing joint positions (e.g., q_1 and q_2) of the subject and the robot (Fig. 1), and the subscripts s and r denote subject or robot matrices of parameters, respectively.

In the null field, (i.e., $E = 0$) in Eq. 1, assume that a solution to this coupled system is $q = q^*(t)$, i.e., the arm follows the reference trajectory (typically a straight hand path with a Gaussian tangential velocity profile). Let us name the controller that accomplishes this task $C = C_0$ in Eq. 2. When the robot motors are producing a force field, i.e., $E \neq 0$, the arm's motion converges back to the reference trajectory if the new controller in Eq. 2 is $C = C_1 = C_0 - J_s^T J_r^{-T} \hat{E}$, where \hat{E} is an estimate of the force field environment as learned by the controller. The internal model composed by the subject is $C_1 - C_0$, i.e., the change in the controller after some training period.

We have suggested previously that a reasonable lumped model of the subject's biomechanical controller in the case of these targetted movements is (Shadmehr and Mussa-Ivaldi, 1994):

$$C_0 = \hat{I}_s(q^*)\ddot{q}^* + \hat{G}_s(q^*, \dot{q}^*)\dot{q}^* - K(q - q^*) - V(\dot{q} - \dot{q}^*), \quad (3)$$

where K and V are linear estimates of the subject's joint stiffness (at posture) and viscosity matrices (Mussa-Ivaldi et al., 1985). In this model, muscle forces produced by the arm are dependent on a feedforward model of the subject's passive dynamics (e.g., inertia of the arm) (Gomi and Kawato, 1996). The controller is stabilized around the desired trajectory $q^*(t)$ (presumably a smooth, straight-line motion to the target) by the stiffness and viscosity of the muscles and the spinal reflex pathways (Shadmehr et al., 1993).

We used the model of the controller in Eq. 3 coupled with the dynamics of the manipulandum and a typical subject's passive dynamics to simulate performance before and after adaptation. This allowed us to predict the forces that a subject's controller should produce if it had acquired an internal model of the forcefield. Parameter values for the model of the subject's arm were the same as that described in our previous report (Shadmehr and Mussa-Ivaldi, 1994). Parameter values for the robot were determined by using a derivation of the kinetic energy of the system in terms of the link lengths, masses, and center of masses of the four bars of the parallelogram. Despite the 12 unknowns, the mass parameters combine in the inertia matrix and reduce to 3 composite parameters (Slotine and Li, 1991). These parameters (along with friction and viscous parameters, which are comparatively small and were not used here), were estimated using a system identification technique. We estimated the inertia matrix of the robot to be:

$$I_r = \begin{bmatrix} a_1 & a_2 \cos(p_2 - p_1) \\ a_2 \cos(p_2 - p_1) & a_3 \end{bmatrix},$$

with p_1 and p_2 as Robot's joint angles (Fig. 1), $a_1 = 0.46$, $a_2 = 0.34$, and $a_3 = 0.094 \text{ kg/m}^2$, and link lengths of 0.460 and 0.344 m for the upper arm and forearm of the robot. We estimated the coriolis matrix of the robot to be:

$$G_r = \begin{bmatrix} 0 & -a_2 \sin(p_2 - p_1)\dot{p}_2 \\ a_2 \sin(p_2 - p_1)\dot{p}_1 & 0 \end{bmatrix}.$$

The desired trajectory in the simulations was assumed to be minimum jerk (Flash and Hogan, 1985) with a period of 0.5 sec.

RESULTS

We report on experiments in which subjects learned to make reaching movements in two distinct dynamic environments. We find that the ability of subjects to learn movements in a second environment, and the ability to recall the skill acquired by practicing in the first environment, are influenced by the temporal distance between learning the first and second environments.

Learning control of a novel dynamic system

A typical subject's hand trajectory in the null field is shown in Figure 2A. Without the disturbing forces, subjects could readily make rapid and accurate movements to the targets. As previously noted (Flash, 1987), these movements were approximately in a straight line with a symmetric tangential velocity profile. However, once a field was introduced, movements became highly distorted. An example of the imposed force field and the resulting movements are shown in Figure 2, B and C. The force field (named B_1 in Materials and Methods) pushed the hand in a direction per-

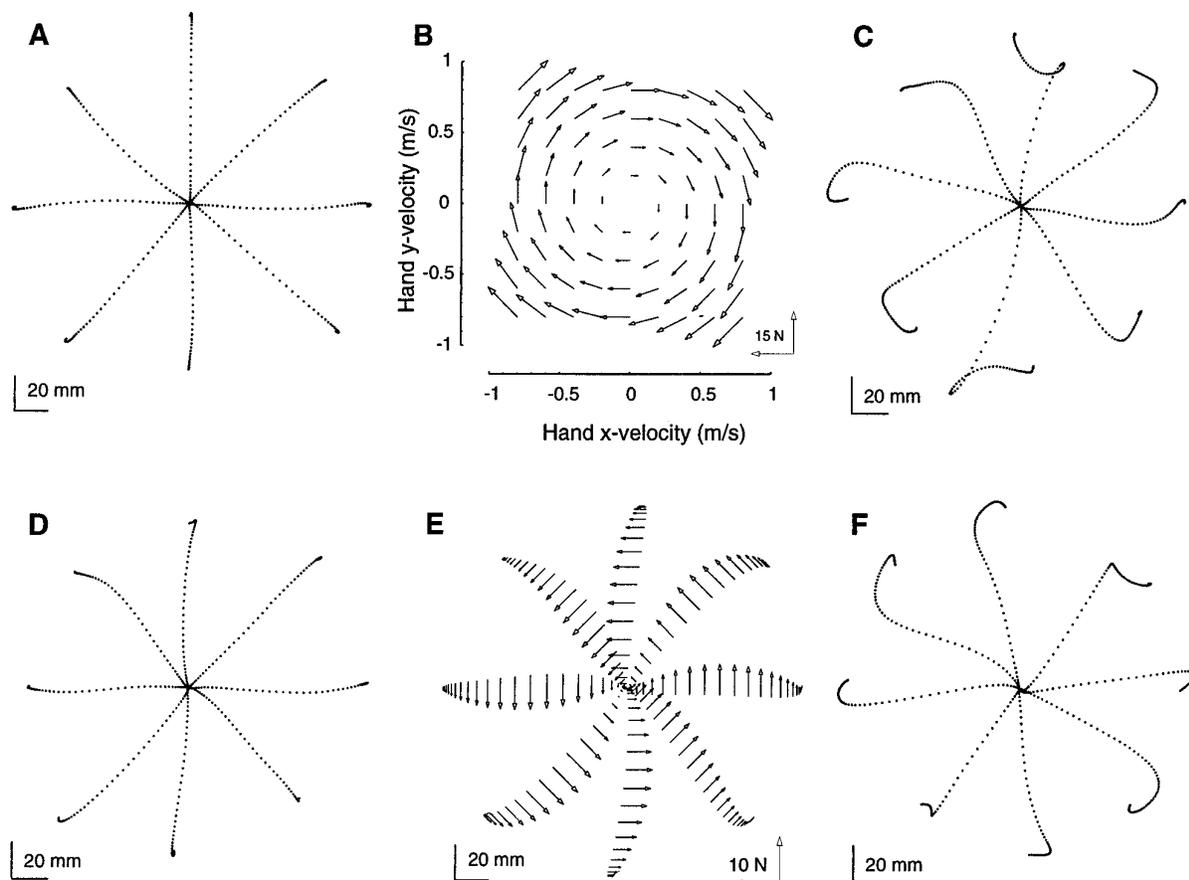


Figure 2. *A*, Hand path of a typical subject in the null field (the points in all hand paths are 10 msec apart). *B*, An example of a force field produced by the robot. The field is a linear function of hand velocity, and the *x*- and *y*-coordinates refer to that of Figure 1. *C*, Hand path of an untrained subject in the field. *D*, Hand path after 300 movements in the field. The trajectory in the field converges to the trajectory observed in the null field. *E*, Forces produced by a typical trained subject to counter the effect of the force field as a function of hand position for each movement. These forces are the projection of the forces measured at the interaction point between the subject and robot onto a direction perpendicular to the direction of target. *F*, While training in the field, random targets are presented with null field conditions. The results are after-effects.

pendicular to the direction of the target. The magnitude of the imposed force was a linear, increasing function of hand velocity. The resulting motion of the hand had a characteristic “hooking” pattern. In previous simulations, we observed similar patterns of motion when we assumed a biomechanical controller of the form detailed in Eq. 3 (Shadmehr and Mussa-Ivaldi, 1994). Note that our model controller produces a pattern of torques based on expectations of the dynamics of the task and is stabilized by the stiffness of the arm about the desired trajectory. This had led us to suggest that the hooks are not indicative of a second, corrective movement, but are attributable to the interaction between the stiffness and inertial characteristics of the subject’s arm and the imposed force field.

With practice, the hooks diminish and the hand trajectory in the field (Fig. 2*D*) becomes similar to that observed in the null field (Fig. 2*A*). Force measured at the interaction point of the robot and the subject suggest that, with practice, subjects learn to produce forces perpendicular to the direction of the target as the hand moves toward the target (Fig. 2*E*). These forces essentially cancel the imposed force field, allowing the hand to move along the desired trajectory. In principle, two biomechanical mechanisms may be responsible for this adaptation. By increasing stiffness of the arm, i.e., global muscular co-contraction, the subject can cancel most perturbing forces regardless of their direction.

Alternatively, the subject may learn to activate muscles so that in addition to the forces necessary to move the hand toward the target, perpendicular forces are generated to compensate for the expected dynamics of the force field. Only in the later scenario would we expect that a sudden removal of the force field should result in after-effects. Typical after-effects are shown in Figure 2*F*. This is an indication that the subject is learning to command a novel pattern of muscle forces in order to reach a target location. In the language of control theory, the subject is learning an IM that predicts a pattern of forces for a desired trajectory.

The after-effects give us a window through which we can examine the content of the IM. Normally, when moving in a null field, the amount of force that is produced by the subject perpendicular to the direction of target is rather small (Fig. 3*A*; this amount is nonzero because the inertia of the manipulandum and the arm is not isotropic). To make a straight-line movement in the field, the subject needs to produce significantly larger perpendicular forces. An example of forces produced by the trained subject is shown in Figure 3*B*. If this change in force production is achieved through learning of an internal model, then through simulation we can predict the pattern of forces that will result if we unexpectedly remove the force field. When the biomechanical controller has incorporated an IM of the field of Figure 2*B*, the change in the output of the system (force at the interaction point as compared

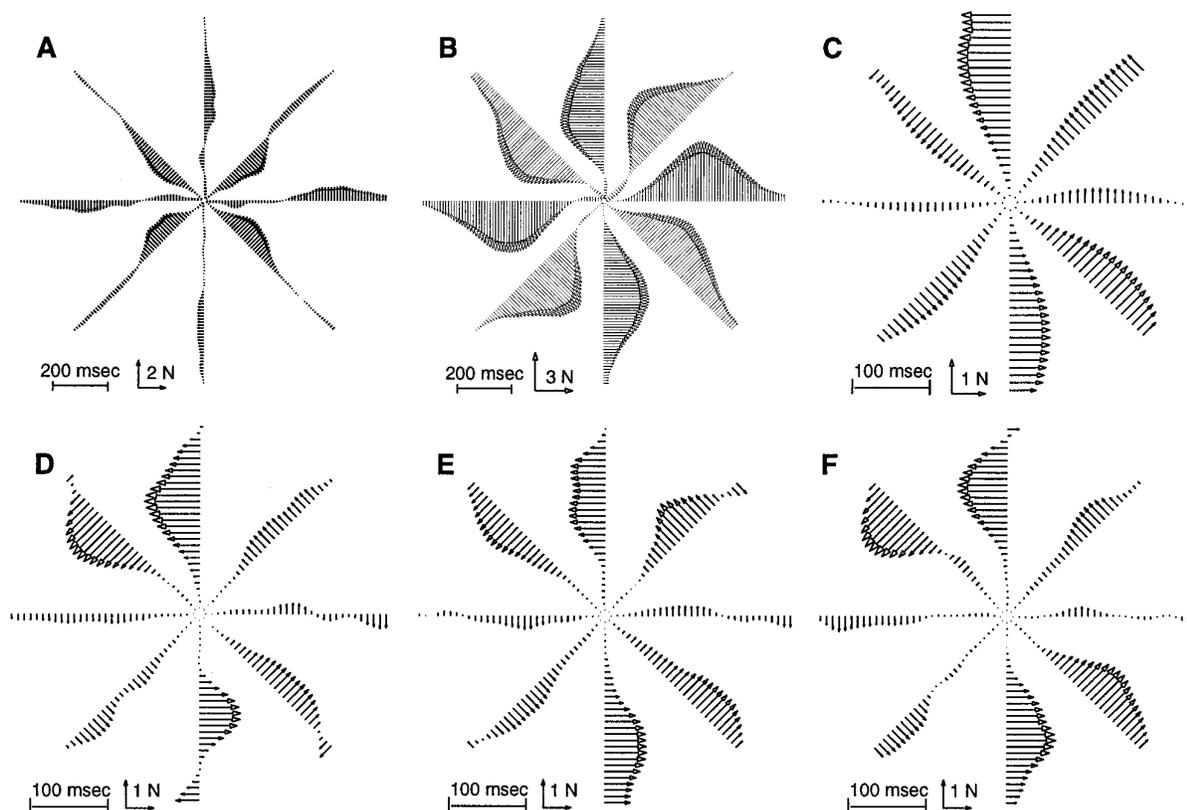


Figure 3. The component of the interaction force perpendicular to the direction of motion plotted as a function of time along straight-line paths to the targets. *A*, Interaction forces while moving in a null field. *B*, Interaction forces of a trained subject while moving in a field. *C*, Results for a simulation in which the controller of Eq. 3 had learned an IM of the force field and, unexpectedly, the field was removed, i.e., force predicted for the after-effects in the case that the controller had learned a perfect model of the field. Forces are plotted for the first 250 msec of movement. *D–F*, Forces recorded from three typical subjects during their after-effects (first 250 msec of movement). The motor output of subjects changed by roughly the same pattern and magnitude as the simulation had predicted. The vectors in all paths are 10 msec apart.

to before adaptation conditions) is predicted to be a distinct pattern of counter-clockwise forces (Fig. 3*C*). These forces will be largest for targets at 90° and 270° and smallest for targets at 0° and 180°. The reason for this nonuniform pattern is the anisotropic behavior of the stiffness of the arm (Mussa-Ivaldi et al., 1985; Shadmehr et al., 1993). This stiffness has the largest influence on stabilizing the hand on movements to 0° and 180° and the smallest influence on movement to 90° and 270°. Indeed, we found that after practice (300 targets), the motor output of subjects (e.g., Fig. 3*D–F*) had changed by roughly the same pattern and magnitude as our simulation had predicted. This suggests that subjects were incorporating an IM of the novel dynamics in programming their motor output. We note, however, that the simulation was in agreement with the recorded forces only for the initial 200–250 msec into the movement. Beyond this, it is possible that long-loop reflexes (which are not modeled) or voluntary action begins to significantly influence the pattern of force generation.

Long-term motor memories

Learning of an IM allows the subject to move his/her hand along a desired trajectory. We assumed that the desired trajectory for each subject was their pattern of motion in the null field (baseline trajectories, as in Fig. 2*A*). Our performance index was a correlation between a subject's typical movement before imposition of the field with movements in the field (Shadmehr and Mussa-Ivaldi, 1994). Figure 4*A* shows the change in this index as a function of practice in all subjects. It is apparent that the majority

of the improvement is occurring in the first 150 movements (the first target set).

Does practice lead to long-term storage of the acquired internal model? We have shown previously that there is a significant improvement in the performance index when subjects are tested 24 hr after they are trained in a given field (Brashers-Krug et al., 1995b). Here we trained subjects on field B_1 ($n = 18$) and had them return at 24 hr to be tested on either the same field that they were trained in previously (control group, $n = 10$) or on a novel field B_2 ($n = 8$). Subjects in the control group also returned 5 months later and were tested on the same field in which they were trained. Results are shown in Figure 4, *B* and *C*: performance in the trained field was significantly higher when probed at 24 hr (Fig. 4*B*; $F_{(1,9)} = 17.99$, $p < 0.005$) and continued to be significantly higher at 5 months after the initial practice ($p < 0.005$; Fig. 4*C*). In comparison, performances of subjects that trained on B_1 and were tested 24 hr later on B_2 were not significantly different than the levels achieved by the naive subjects on B_2 ($F_{(1,7)} = 3.2$, $p > 0.1$). This suggests that the improvement in performance of the control group was not attributable to general familiarity with the experiment, but learning of an IM specific to the presented force field. This learning resulted in a long-term memory of the IM.

To determine whether subjects who practiced in two different fields (B_1 and B_2 , training sessions for the fields separated by 24 hr) formed long-term representations of both fields, we had sub-

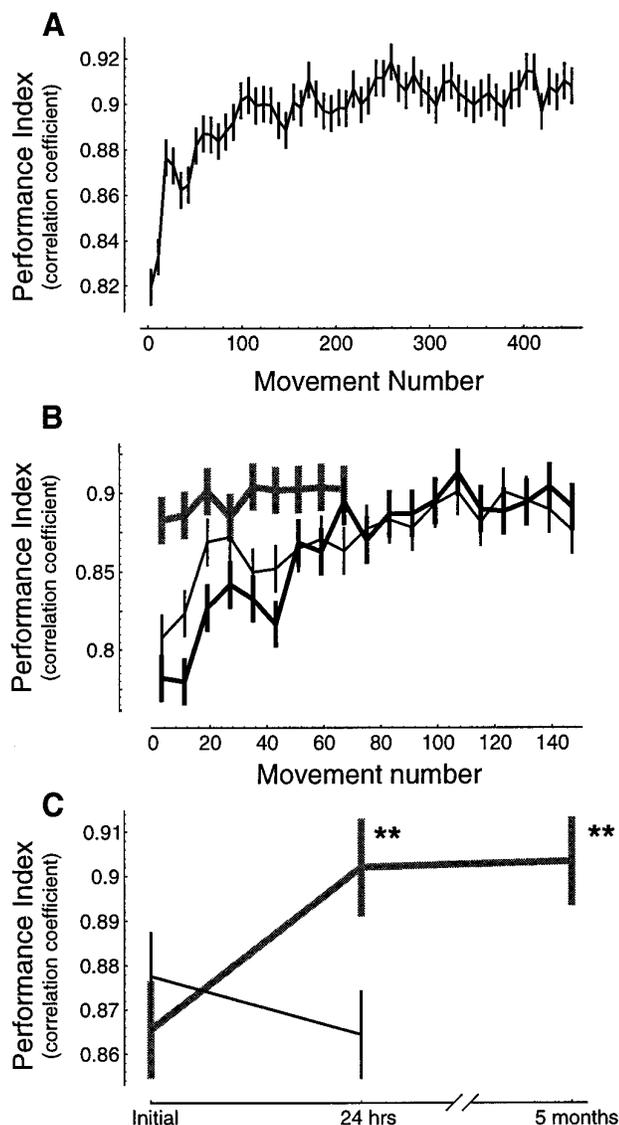


Figure 4. Performance during initial training in a force field and subsequent tests of recall. Performance index is a correlation between the hand trajectory in the force field and the hand trajectory in the null field (baseline trajectories, as in Fig. 2A). *A*, Mean performance \pm 95% confidence intervals for all subjects ($n = 60$). For each subject, groups of eight consecutive movements are binned together (there were 8 different directions of movement, and target directions were presented in random order). *B*, Initial performance in a field and performance 24 hr later. All lines are mean performance \pm 95% confidence intervals. *Thin line*, Performance of naive subjects ($n = 18$) in field B_1 . *Thick black line*, Performance of a subset of these subjects ($n = 8$) in a novel field, B_2 , measured 24 hr after training in B_1 . *Thick gray line*, Performance of the remaining subjects at 24 hr on field B_1 . *C*, Summary performance scores \pm 95% confidence intervals for the two groups of subjects. *Gray line* represents subjects that were tested on the field in which they were trained. *Black line* represents subjects that at 24 hr were tested on a novel field.

jects ($n = 6$) tested on field B_1 at an interval of 2 weeks and field B_2 at 3 weeks beyond completion of training. With respect to the performance during training, there was a significant improvement in performance during the recall sessions: mean performance index \pm 95% confidence interval = 0.89 ± 0.007 vs 0.92 ± 0.007 for field B_1 during training and recall ($F_{(1,5)} = 16.81$, $p < 0.01$), rejecting the null hypothesis that there is no improvement in performance during recall of B_1 as compared to initial training,

and 0.87 ± 0.008 versus 0.900 ± 0.008 for field B_2 during training and recall ($F_{(1,5)} = 40.63$, $p < 0.002$), rejecting the null hypothesis that there is no improvement in performance during recall of B_2 as compared to initial training. Therefore, when the training sessions were separated by 24 hr, subjects retained the IMs for both fields B_1 and B_2 .

Time course of consolidation

The idea that memories undergo a process of consolidation relies strongly on the observation that there are periods after acquisition of information during which the representation of the recently acquired material is fragile. With time, the representation becomes less susceptible to an intervention. For example, post-training treatments such as electric shocks (Squire et al., 1975), removal of key anatomical sites (Kim et al., 1995), or protein synthesis inhibition (Tully et al., 1994) retard this progression and often result in loss of the recently acquired information (Squire et al., 1981). These interventions, however, have little effect on recall once a window of time has passed since acquisition.

We tested for the stability of the acquired IM of field B_1 as a function of temporal distance to training in field B_2 . Subjects trained in field B_1 , and then trained in B_2 at 5 min ($n = 9$), 30 min ($n = 6$), 2.5 hr ($n = 7$), 5.5 hr ($n = 10$), or 24 hr ($n = 8$). They then returned 1 week later and were tested in field B_1 . Figure 5A shows the change in performance in B_1 during the recall session as compared to the initial learning for two groups: the group that learned B_2 5 min after B_1 , and the group that learned B_2 5.5 hr after B_1 . Whereas the 5 min group shows no recall of B_1 (mean performance not significantly different in recall vs initial learning, paired t test, $p > 0.4$), the 5.5 hr group shows significant recall (paired t test, $p < 0.02$). The data for all groups are summarized in Figure 5B. There is a significant effect of time on retention of B_1 ($F_{(49,44)} = 2.46$, $p < 0.05$). If B_2 is practiced 5–30 min after B_1 , we find no evidence for recall of B_1 . Recall becomes significant at 5.5 hr but approaches the level of recall observed in the control group only at 24 hr.

The time interval at which learning of field B_2 does not impair recall of B_1 is similar to what we had observed previously in a different group of 70 subjects (Brashers-Krug et al., 1996). In our previous work, this interval was estimated at 4 hr. Here, we find significant recall at 5.5 hr. There are two differences in the protocol of the current study and the previous work: (1) in the current setup, the subjects practiced 3 times longer on field B_1 before being exposed to B_2 , and (2) in the current study recall was measured 1 week after original training rather than at 24 hr (as in our previous work). The increased training on B_1 was chosen in the current protocol to ensure that the performance plateaued before B_2 was introduced (Fig. 4A). Recall was tested at 1 week rather than at 24 hr to ensure against any anterograde interfering effects that might be present after learning of B_2 . This testing of recall at 1 week is important because of a phenomenon called “release from inhibition”: it has been observed that in learning associations between pairs of words, learning to associate A with B followed by association of A with C leads to poor recall of A–B when tested at a short interval (hours after training) but leads to good recall at longer intervals (1 week; as compared to a group that only learned A–B) (Koppelaar, 1963). Therefore, it is possible that our previous observation regarding the poor recall of field B_1 (Brashers-Krug et al., 1996) might be attributable to a lingering anterograde interference from B_2 (see Fig. 8). The current study was designed with this concern in mind. The results of Figure 5

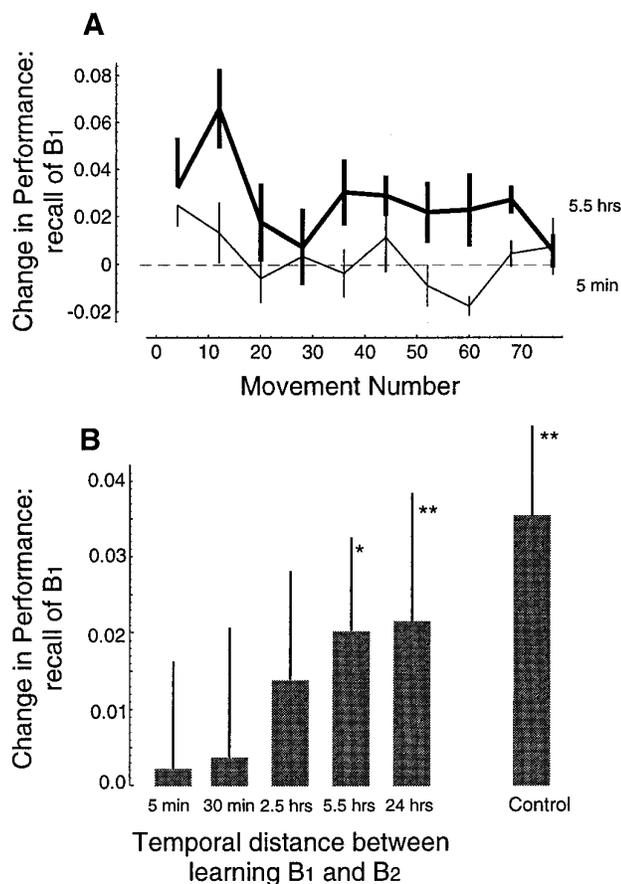


Figure 5. Performance during the test of recall for B_1 as a function of temporal distance between learning of B_1 and B_2 . B_1 was tested for recall 1 week after B_1 and B_2 were learned. *A*, Mean \pm SE improvement in performance for two groups of subjects. *Thin line* is for the group ($n = 9$) that practiced in B_2 at 5 min after completion of practice in B_1 . *Thick line* is for the group ($n = 10$) that practiced in B_2 at 5.5 hr after B_1 . *B*, The ability to recall B_1 is significantly dependent on temporal distance between B_1 and B_2 . Each bar is the mean \pm 95% confidence interval of change in performance as measured for a target set during the recall test versus during initial practice.

show that recall of B_1 , as measured a week after original learning, is significantly influenced by the time at which B_2 was learned.

Although the time of learning of B_2 influences the recall of B_1 , the link to consolidation would be strengthened if there was evidence that learning of B_2 within close temporal proximity of learning B_1 results in an unlearning of the IM for B_1 . The computational model for this kind of forgetting has termed the phenomenon “catastrophic interference” (Sutton, 1986). In this computational model of memory, forgetting occurs because the memories that represent the internal model of field B_1 (associating a desired trajectory to a specific pattern of muscle torques) are used for learning B_2 (Shadmehr et al., 1995). Because the two fields are anticorrelated, learning of B_2 would lead to a loss of memory for B_1 (e.g., massive changes in the weights of the network or patterns of activity). A prediction of this computational model is that at close temporal proximity, learning of B_2 takes place with an instantiated IM of B_1 , rather than a “tabula rasa.”

The after-effects give us a unique window into the contents of the IM being used to learn a field. For example, in a naive subject that is just beginning a target set in B_2 , there are no after-effects. It is with practice that after-effects develop. We quantified the size of an after-effect by measuring the distance that the trajectory

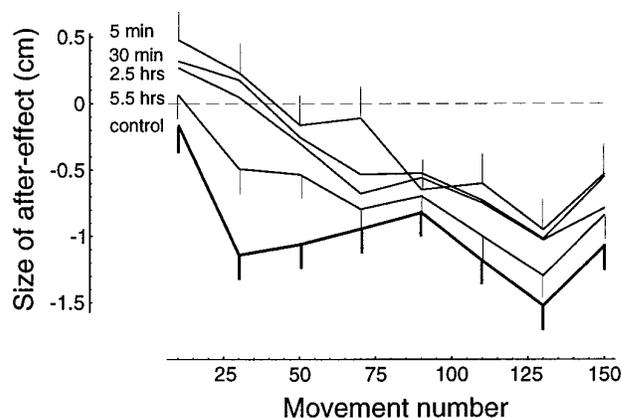


Figure 6. Size and direction of after-effects as subjects learn field B_2 at different time intervals after practicing in B_1 . Size is determined as the distance from a straight line (from the previous to the next target) at 300 msec into the movement. Direction is positive for an after-effect appropriate for field B_1 (i.e., counter-clockwise, as in Fig. 2) and negative for an after-effect appropriate for field B_2 . Plotted are the means and 95% confidence intervals. Each point represents the average after-effect for a group of subjects at a given movement number (bin size is 4). Because the size of after-effects depends on direction of target (e.g., Fig. 2*F*), the change for a given curve is not expected to be monotonic. However, the sequence of targets for all subjects is the same. Therefore, after-effects at a movement number may be directly compared among the different groups. The figure shows that at 5 min after learning B_1 , subjects begin learning B_2 with after-effects that are in the direction of B_1 . Control subjects who never learned B_1 begin learning B_2 with no after-effects, i.e., unbiased. With temporal distance between training sessions, initial after-effects begin to resemble those of the control group.

deviated from a straight-line path to the target. This distance was measured at 300 msec into the movement. The sign of this vector was positive if the after-effect was a counter-clockwise deviation from the straight line (appropriate for an IM of B_1 , as in Fig. 2*F*), and negative if it was a clockwise deviation. Figure 6 shows the progression of after-effect development for a group of naive subjects on field B_2 (the control group). This figure also shows the development of after-effects for the group of subjects that practiced in B_2 at 5 min after completion of practice in B_1 . In the 5 min group, subjects begin learning B_2 with an IM appropriate for B_1 . The rates of change in the after-effects are not different among the groups: the top four lines in Figure 6 are approximately parallel. The main difference between the four groups is the starting point. With temporal distance, the starting point gradually shifts toward that of the naives so, at 5.5 hr, there are no significant after-effects as learning of B_2 initiates. In other words, whereas in the 5 min group learning of B_2 starts with an IM of B_1 , in the 5.5 hr group, the IM is close to a “tabula rasa.”

Another way to quantify the contents of the IM that the subject is using to learn a given field is to measure the interaction forces between the subject and the robot. As Figure 3 demonstrates, in a given subject one can compare the interaction forces recorded during baseline movements (i.e., in a null field before introduction of the forces) with those during after-effects. The difference is an estimate of the change in the motor output of the subject, which is presumably attributable to adaptation. In our model of the biomechanical control system (Eq. 3), this variable is an estimate of the output of the subject’s IM. As we noted earlier, our measure of this change will very likely be an underestimation of the true learning because the stiffness properties of the arm will reduce the size of the after-effects. However, one can measure this

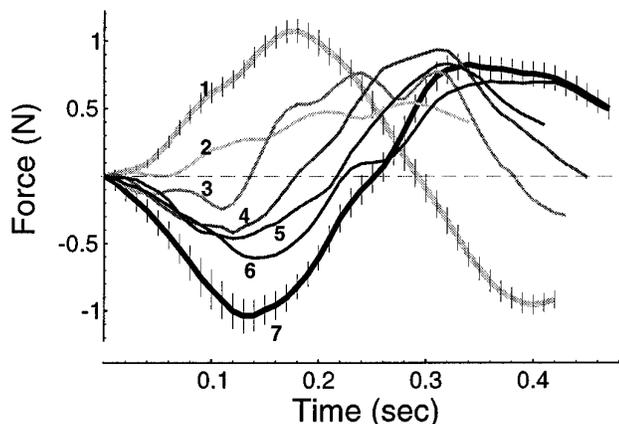


Figure 7. A measure of forces recorded at the interaction point between subjects and robot as a function of time into an after-effect. Each line is an average change in the interaction force during after-effects ($n = 5$) to targets at 270° and 225° for the first 80 movements in a given force field with respect to those recorded from the same group of subjects in the baseline conditions. This representation of the interaction force is the same as the measure shown in Figure 3D–F, with the difference that time is explicitly represented and the sign of the force is positive for a counter-clockwise vector and negative otherwise. *Line 1* represents mean \pm 95% confidence interval of the force produced by naive subjects that learned field B_1 . *Line 7* represents mean \pm 95% confidence interval for naive subjects that learned field B_2 . *Lines 2–6* are the forces produced by subjects that were learning B_2 at 5 min, 30 min, 2.5 hr, 5.5 hr, or 24 hr after learning B_1 , respectively.

variable and, assuming that arm stiffness is roughly equal among the subjects of all the different groups (Mussa-Ivaldi et al., 1985; Shadmehr et al., 1993), compare its time course.

We computed an estimate of the output of the IM for the subjects in all groups during the after-effects of the first 80 targets in fields B_1 and B_2 . The sign of the output was set positive if the force vector was pointing counter-clockwise from a straight line to the target (appropriate for an IM of field B_1 , as in Fig. 3D–F) and negative if it was clockwise (appropriate for B_2). The result is shown in Figure 7. Lines 1 and 7 show the output of the IM for the naive subjects in fields B_1 and B_2 , respectively. These lines give us an estimate of what an unbiased IM will output after it has been trained with movements to 80 targets. The remaining lines are all measured while subjects were learning B_2 and are differentiated based on the temporal distance to learning B_1 . In the 5 min group, the IM used to learn B_2 is still mainly composed of B_1 . This evidence supports our contention that learning B_2 in close temporal proximity to B_1 takes place with an instantiated IM of B_1 . With temporal distance, subjects learn B_2 with an IM that can better estimate B_2 after a given number of movements. This predicts that performance should be better in B_2 as a function of temporal distance to B_1 . Results shown in Figure 8 demonstrate that this is indeed the case: performance in B_2 is significantly worse than in B_1 when temporal distance is 5 min (Fig. 8A). With time, the ability to learn the second field gradually improves (Fig. 8B; $F_{(39,35)} = 4.155$, $p < 0.01$).

The data on after-effect development (Fig. 6) suggest that the rate of learning an IM of B_2 was similar across the different groups; the difference was the initial bias from which the learning began: in computational terms, a first approximation would suggest that the IM used to learn B_2 had weights that were strongly initialized toward those appropriate for representation of B_1 . With temporal distance, learning of B_2 began with weights less

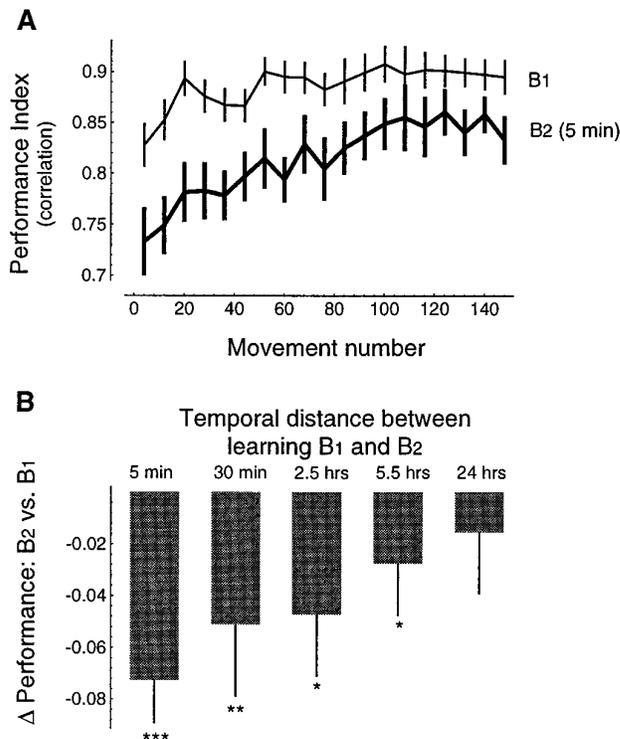


Figure 8. Performance of subjects in field B_2 as a function of time since learning B_1 . *A*, When B_2 is introduced 5 min after completion of training in B_1 , performance is worse than that recorded from B_1 . Plotted are mean \pm SE of correlations. *B*, Performance in B_2 is significantly dependent on temporal distance between B_1 and B_2 . Each bar is the mean \pm 95% confidence interval of change in performance as measured for the initial target set in B_1 and B_2 .

biased toward B_1 , approaching the “tabula rasa” of the naives in the control group. In other words, with time, there was a fading of an aspect of representation of B_1 . However, we know that long-term memory of B_1 did not fade with time (at least within a few months; Fig. 4C), suggesting that this fading component is not related to long-term memory of B_1 . In Figure 9 we combined the data on retrograde and anterograde effects of learning the two fields and show the two hypothesized aspects of memory for learning an IM for B_1 . The fading component has data points that are biases of the IM used to learn field B_2 . The rising component has data points that are the memories of B_1 that were retained after learning of B_2 . This represents a hypothesized time course for formation of long-term memory of B_1 .

The fading component of the memory of recently learned B_1 is presumably the reason why subjects at 5.5 hr can readily learn B_2 . To determine whether there is a relationship between this fading component and the consolidation process for formation of long-term memory of B_1 , we performed one last experiment. We recruited a new group of subjects ($n = 10$) and trained them in B_1 (3 target sets) on day 1 and had them return on day 2. On day 2, subjects were given a target set in B_1 , and then they practiced in three target sets in B_2 . When performance in B_2 was compared to performance in B_1 (as recorded 24 hr earlier), we found a significant reduction in performance ($F_{(1,9)} = 34.85$, $p < 0.001$, a comparison of performance in the first target set). Furthermore, the mean change in performance, -0.058 , was not significantly different than the change in performance observed when subjects’ only exposure to B_1 was 5 min before B_2 (the mean change for this

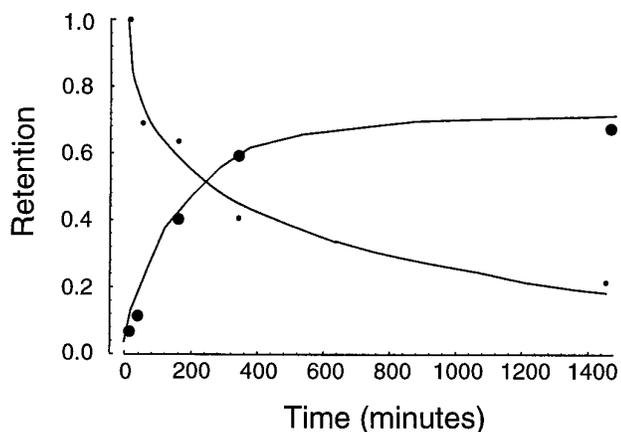


Figure 9. The data from Figures 5 and 8 have been combined to produce two hypothetical stages in formation of long-term motor memory. The data points that constitute the decaying stage are the amount of anterograde interference that was recorded (as a percentage of that recorded at 5 min) as subjects attempted learning of a second field at different times. The data points that constitute the rising curve are a measure of memory retained from the first field after learning of a second field at different times.

group was -0.0725 , as shown in Fig. 8). Therefore, the ability to learn a new field was not related to when B_1 was originally learned but, rather, when it was last practiced.

Although learning of B_2 is affected by the recently instantiated B_1 , if B_1 was originally learned 24 hr ago, then its long-term memory should not be affected by learning of B_2 . We tested for recall of B_1 on day 3. The subjects showed significantly improved performance compared to initial training ($F_{(1,9)} = 8.757$, $p < 0.02$). Mean improvement in this group ($+0.0316$) was not significantly different than the improvement that we had seen in our control subjects ($+0.034$, shown in Fig. 5). The long-term memory of B_1 was intact even though B_2 was learned immediately after B_1 was performed. This suggests a functional independence for the two hypothesized stages of motor memory.

DISCUSSION

The ability of the central nervous system to learn and store multiple and conflicting visuomotor maps has been demonstrated in both monkeys and man (Flook and McGonigle, 1977; McGonigle and Flook, 1978; Welch et al., 1993; Cunningham and Welch, 1994). For example, it has been shown that the CNS can learn and retain two conflicting visuomotor maps associated with left and right displacing prisms (McGonigle and Flook, 1978) and two different gains associated with the vestibulo-ocular reflex (Baker et al., 1987; Shelhamer et al., 1992; Tiliket et al., 1993). Here we demonstrated that two conflicting motor skills (what we have termed internal models) may also be learned and retained, but only if the training sessions in the two tasks are separated by a critical time interval of ~ 4 –5 hr. This time interval is in agreement with the data on the prism studies: in monkeys, adaptation was obtained only when alternate maps were presented far apart in time (24 hr) (Flook and McGonigle, 1977). In humans, after a single training session with a given prism, learning of a second visuomotor map (with a second prism) at close temporal proximity (10 sec) was significantly inhibited compared to naives, and a test of recall with the first prism at 3 d later showed no evidence of improvement (McGonigle and Flook, 1978).

In this study, we suggested that there is a critical time interval

required for learning and retention of two distinct IMs. Our results show that recall of IM_1 is affected by the temporal distance to learning of IM_2 . However, recall is the culmination of a chain of processes (e.g., perception of the task, integration of proprioceptive information, activation of motor memory, and action), and poor performance in a test of recall may not imply that the motor memory component has been affected (Bower et al., 1994); there is evidence that retrograde amnesia is sometimes not the result of consolidation failure (Miller and Marlin, 1984). This argument is based on two reports: (1) when reminder trials involving apparatus or other cues were presented during the retention interval after administration of an amnesic agent, retrograde amnesia was reduced (Lewis et al., 1968; Quartermain et al., 1970), and (2) performance improved when the experimenter provided cues regarding the correct response during test of recall (Postman and Stark, 1969; Bower and Mann, 1992).

This line of thinking suggests that poor recall may be attributable to inaccessibility of stored information, rather than its loss, and that with time or appropriate cues, information that once was inaccessible might become available (Koppelaar, 1963; Squire et al., 1981). Although we cannot rule out this possibility, there are four pieces of data from our study that argue for the idea that representation of a motor skill does undergo profound functional changes within a short window of time after acquisition.

- (1) If subjects are presented with B_2 shortly after learning B_1 , they learn B_2 with an IM appropriate for B_1 . The contents of the IM being used to learn B_2 (as inferred from the after-effects) suggests an unlearning of B_1 . With temporal distance, the learning of B_2 begins with an IM that approaches the tabula rasa of the naives.
- (2) Recall of IM_1 as measured 1 week after original learning shows a significant dependence on when IM_2 was acquired. This period of 7 d was chosen because it is significantly longer (~ 7 times) than the interval at which we detected an anterograde interference from IM_1 onto IM_2 .
- (3) Making movements in a force field provides continuous haptic, proprioceptive, and visual feedback to the subject regarding the nature of the forces present in the field. Yet when fields are learned in close temporal proximity, there is no evidence for recall as measured in a target set that included 192 movements.
- (4) Recall of IM_1 is not affected when it was learned 24 hr before learning IM_2 , even though subjects performed movements in B_1 moments before learning B_2 .

Taken as a whole, the above evidence, in our view, argues for a distinct change in the state of resistance of motor memory within a few hours after acquisition. Because the vulnerability to an intervention and the ability to learn a second task depend on time since acquisition, it is possible that the neuronal basis of motor memory changes after acquisition.

A number of mechanisms have been proposed to underlie memory formation in the central nervous system. These include long-lasting changes in synaptic efficacy (Bliss and Collingridge, 1993) and reverberation of activity in a collection of excitatory neurons (Hebb, 1949; Zipser et al., 1993). Hebb (1949) was the first to suggest a neural basis for the time dependent success of retrograde amnesic agents. In his framework, memories are stored for a period of time in a labile form of neuronal firing patterns generated through reverberating circuits. The firing pattern persists after completion of practice and leads to a more gradual development of synaptic plasticity, mediating long-term memory. A prominent example of synaptic plasticity is long-term potentiation (LTP). It has been shown that after inducing LTP, certain low-frequency stimuli can depotentiate the synapse (Fujii et al., 1991), effectively reducing the synapse's

efficacy to near baseline levels. These stimuli, however, are only effective if they are given within a small time window after potentiation of the synapse: 20 min after induction of LTP, the low-frequency stimuli depotentiate the synapse by 70%, whereas at 100 min, the depotentiation is only at 30%. There is a wealth of evidence for LTP (Asanuma and Keller, 1991; Kimura et al., 1994) and LTD (Castro-Alamancos et al., 1995) in the motor areas of the cortex and the cerebellum.

A first-approximation model of learning might begin with Hebb's ideas regarding the initial representation of memory as a labile form of neuronal firing patterns, and synaptic plasticity as the means for representing long-term memory. These two types of representations may form the neuronal basis of the two hypothesized stages of motor memory in Figure 9; according to this model, practice leads to recruitment of activity in neuronal circuits and establishes a reverberating pattern as the training comes to an end. This pattern gradually decays, but it serves as the teacher for a slower but more resistant form of memory storage (Alvarez and Squire, 1994), e.g., synaptic plasticity. We would expect that the initial stage to have a finite life and decay after completion of motor practice in task 1. If task 2 is attempted while the neuronal firing pattern is present, there will be interference; learning of task 2 will begin with a pattern appropriate for task 1, and performance will be impaired compared to naives. If time is allowed to pass after learning of task 1, changes in synaptic efficacy gain stability and serve as a more permanent representation of the motor memory for task 1. It is important to note that a model of memory that relies only on synaptic plasticity (e.g., LTP) would have trouble explaining our data: because the changes induced in synaptic efficacy are most fragile soon after they are established, it should be easy to learn IM_2 soon after learning IM_1 . However, we find that the opposite is true. The utility of a two-stage learning system has been elaborated recently in a formal computational model (McClelland et al., 1995).

Our last experiment shed some light on the role of the hypothesized initial stage. We noted that learning of IM_2 was impaired if movements were performed in B_1 just before B_2 , i.e., this impairment was just as severe for the case where IM_1 was just acquired versus the case where IM_1 was acquired 24 hr ago but was just recalled. It seems likely, therefore, that the hypothetical initial stage is not merely a gateway to long-term memory but also, at least in part, a temporary storage area for items of information, whether new or old, for use in the near-term. This is the description that has been used to define "working memory" (Fuster, 1995). A major function of this kind of memory is to hold information and update current information on a real-time basis (Goldman-Rakic, 1994).

It is possible that the neuronal basis of the hypothetical initial stage is mediated by regions distinct from that of the second stage, i.e., there may be a time-limited role associated with certain regions of the brain in maintaining motor memory (Mishkin et al., 1984). It has been argued that brain regions active during acquisition of motor memory are not necessarily the same as regions that will eventually store the memory (Pavlidis et al., 1993). For example, although the cortico-cortical projections from the somatosensory to the motor cortex play an important role in learning new motor skills, they may not be required for execution of existing motor skills (Aizawa et al., 1991). In humans, there is now mounting evidence from functional imaging studies of motor learning that indicate distinct motor areas are active during initial learning versus subsequent recall trials of a motor task (Grafton et al., 1994; Karni et al., 1995; Kawashima et al., 1995). It will be

important to ask whether changes in centers of neuronal activity correlate with functional changes in the stability of the recently acquired memory (Brashers-Krug et al., 1995a; Shadmehr and Holcomb, 1996). However, in all likelihood, classification of motor memory into only two discrete phases will turn out to be naive, because it has been argued that formation of stable memory is analogous to a developmental process in which extracellular signals initiate cascades of events, gradually modifying neuronal representation on a time scale of seconds to years (Dudai, 1989).

REFERENCES

- Aizawa H, Inase M, Mushiaki H, Shima K, Tanji J (1991) Reorganization of activity in the supplementary motor area associated with motor learning and functional recovery. *Exp Brain Res* 84:668–671.
- Alvarez P, Squire LR (1994) Memory consolidation and the medial temporal lobe: a simple network model. *Proc Natl Acad Sci USA* 91:7041–7045.
- Asanuma H, Keller A (1991) Neuronal mechanisms of motor learning in mammals. *NeuroReport* 2:217–224.
- Bailey CH, Kandel ER (1995) Molecular and structural mechanisms underlying long-term memory. In: *The cognitive neurosciences* (Gazzaniga M, ed), pp 19–36. Cambridge, MA: MIT.
- Baker J, Perlmutter SO, Peterson BW, Rude SA, Robinson FR (1987) Simultaneous opposing adaptive changes in cat vestibulo-ocular reflex direction for two body orientations. *Exp Brain Res* 69:220–224.
- Bliss TV, Collingridge GL (1993) A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361:31–39.
- Bower GH, Mann T (1992) Improving recall by recoding interfering material at the time of retrieval. *Learn Mem Cognit* 18:1310–1320.
- Bower GH, Thompson-Schill S, Tulving E (1994) Reducing retroactive interference: an interference analysis. *Learn Mem Cognit* 20:51–66.
- Brashers-Krug T, Krebs H, Rauch S, Savage C, Rubin R, Alpert N, Fischman A, Hogan N, Bizzi E (1995a) PET imaging of human motor learning. *Soc Neurosci Abstr* 21:517.
- Brashers-Krug T, Shadmehr R, Todorov E (1995b) Catastrophic interference in human motor learning. In: *Advances in neural information processing systems*, Vol 7 (Tesauro G, Touretzky DS, Leen TK, eds), pp 19–26. Cambridge, MA: MIT.
- Brashers-Krug T, Shadmehr R, Bizzi E (1996) Consolidation in human motor memory. *Nature* 382:252–255.
- Castro-Alamancos MA, Donoghue JP, Connors BW (1995) Different forms of synaptic plasticity in somatosensory and motor areas of the neocortex. *J Neurosci* 15:5324–5333.
- Charnnarong J (1991) The design of an intelligent machine for upper-limb physical therapy. Masters thesis, MIT.
- Clark MA, Merians AS, Kothari A, Poizner H, Macauley B, Gonzalez-Rothi LJ, Heilman KM (1994) Spatial planning deficits in limb apraxia. *Brain* 117:1093–1106.
- Corkin S (1968) Acquisition of motor skill after bilateral medial temporal-lobe excision. *Neuropsychologia* 6:255–265.
- Cunningham HA, Welch RB (1994) Multiple concurrent visual-motor mappings: implications for models of adaptation. *Hum Percept Perform* 20:987–999.
- de Belle JS, Heisenberg M (1994) Associative odor learning in *Drosophila* abolished by ablation of mushroom bodies. *Science* 263:692–695.
- DeZazzo J, Tully T (1995) Dissection of memory formation: from behavioral pharmacology to molecular genetics. *Trends Neurosci* 18:212–218.
- Dudai Y (1989) *The neurobiology of memory*. New York: Oxford UP.
- Faye IC (1986) An impedance controlled manipulandum for human movement studies. Masters thesis, MIT.
- Flash T (1987) The control of hand equilibrium trajectories in multi-joint arm movements. *Biol Cybern* 57:257–274.
- Flash T, Hogan N (1985) The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci* 5:1688–1703.
- Flook JP, McGonigle BO (1977) Serial adaptation to conflicting prismatic rearrangement effect in monkey and man. *Perception* 6:15–29.
- Fujii S, Saito K, Miyakawa H, Ito K, Kato H (1991) Reversal of long-term potentiation (depotentiation) induced by tetanus stimulation of the input to the CA1 neurons of guinea pig hippocampal slices. *Brain Res* 555:112–122.
- Fuster JM (1995) *Memory in the cerebral cortex: an empirical approach to neural networks in the human and nonhuman primates*. Cambridge, MA: MIT.

- Gabrieli JDE, Corkin S, Mickel SF, Growdon JH (1993) Intact acquisition and long-term retention of mirror-tracing skill in Alzheimer's disease and in global amnesia. *Behav Neurosci* 107:899–910.
- Gandolfo F, Mussa-Ivaldi FA, Bizzi E (1996) Motor learning by field approximation. *Proc Natl Acad Sci USA* 93:3843–6.
- Goldman-Rakic PS (1994) Working memory dysfunction in schizophrenia. *J Neuropsychiatry Clin Neurosci* 6:348–357.
- Gomi H, Kawato M (1996) Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science* 272:117–120.
- Gordon AM, Westling G, Cole KJ, Johansson RS (1993) Memory representations underlying motor commands used during manipulation of common and novel objects. *J Neurophysiol* 69:1789–1796.
- Gottlieb GL (1994) The generation of the efferent command and the importance of joint compliance in fast elbow movements. *Exp Brain Res* 97:545–550.
- Grafton ST, Woods RP, Tszka M (1994) Functional imaging of procedural motor learning: relating cerebral blood flow with individual subject performance. *Hum Brain Mapp* 1:221–234.
- Guigon E, Burnod Y (1995) Short-term memory. In: *The handbook of brain theory and neural networks* (Arbib MA, ed), pp 867–871. Cambridge, MA: MIT.
- Halsband U, Freund H (1993) Motor learning. *Curr Opin Neurobiol* 3:940–949.
- Hammer M, Menzel R (1995) Learning and memory in the honeybee. *J Neurosci* 15:1617–1630.
- Hebb DO (1949) *The organization of behavior*. New York: Wiley.
- Heilman KM, Gonzalez-Rothi LJ (1985) Apraxia. In: *Clinical neuropsychology* (Heilman KM, Valenstein E, eds), pp 131–150. New York: Oxford UP.
- Hogan N, Krebs HI, Charnnarong J, Srikrishna P, Sharon A (1992) MIT-Manus: a workstation for manual therapy and training. *SPIE Telemanipulator Technol* 1833:28–34.
- Karni A, Meyer G, Jezzard P, Adams MM, Turner R, Ungerleider LG (1995) Functional MRI evidence for adult motor cortex plasticity during motor skill learning. *Nature* 377:155–158.
- Kawashima R, Roland PE, O'Sullivan BT (1995) Functional anatomy of reaching and visuomotor learning: a positron emission tomography study. *Cereb Cortex* 5:111–122.
- Kim JJ, Clark RE, Thompson RF (1995) Hippocampectomy impairs the memory of recently, but not remotely, acquired trace eyeblink conditioned responses. *Behav Neurosci* 109:195–203.
- Kimura A, Caria MA, Melis F, Asanuma H (1994) Long-term potentiation within the cat motor cortex. *NeuroReport* 5:2372–2376.
- Koppelaar RJ (1963) Time changes in the strengths of A-B A-C lists: spontaneous recovery? *J Verbal Learn Verbal Behav* 2:310–319.
- Lackner JR, Dizio P (1994) Rapid adaptation to coriolis force perturbations of arm trajectory. *J Neurophysiol* 72:299–313.
- Lewis D, Miles GH (1956) Retroactive interference in performance on the star discriminator as a function of amount of interpolated learning. *Percept Motor Skills* 6:295–298.
- Lewis D, McAllister D, Adams J (1951) Facilitation and interference in performance of the modified Mashburn apparatus. I. The effects of varying the amount of original learning. *J Exp Psychol* 41:247–260.
- Lewis DJ, Misanin JR, Miller RR (1968) Recovery of memory following amnesia. *Nature* 220:704–705.
- McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychol Rev* 102:419–457.
- McGonigle BO, Flook J (1978) Long-term retention of single and multistate prismatic adaptation by humans. *Nature* 272:364–366.
- Miller RR, Marlin NA (1984) The physiology and semantics of consolidation. In: *Memory consolidation: psychobiology of cognition* (Weingartner H, Parker ES, eds), pp 85–111. Hillsdale, NJ: Erlbaum.
- Milner TE, Cloutier C (1993) Compensation for mechanically unstable loading in voluntary wrist movement. *Exp Brain Res* 94:522–532.
- Mishkin M, Malamut B, Bachevalier J (1984) Memories and habits: two neural systems. In: *Neurobiology of learning and memory* (Lynch G, McGaugh JL, Weinberger NM, eds), pp 65–77. London: Guilford.
- Mussa-Ivaldi FA, Hogan N, Bizzi E (1985) Neural, mechanical and geometric factors subserving arm posture in humans. *J Neurosci* 5:2732–2743.
- Pavlides C, Miyashita E, Asanuma H (1993) Projection from the sensory to the motor cortex is important in learning motor skills in the monkey. *J Neurophysiol* 70:733–741.
- Postman L, Stark K (1969) The role of response availability in transfer and interference. *J Exp Psychol* 79:168–177.
- Quartermain D, McEwen BS, Azmitia EC (1970) Amnesia produced by electroconvulsive shock or cycloheximide: conditions for recovery. *Science* 169:683–686.
- Rashidy-Pour A, Motaghd-Larijani Z, Bures J (1996) Reversible inactivation of the medial septal area impairs consolidation but not retrieval of passive avoidance learning in rats. *Behav Brain Res* 72:185–188.
- Rose SP (1991) How chicks make memories: the cellular cascade from c-fos to dendritic remodeling. *Trends Neurosci* 14:390–397.
- Salmon DP, Butters N (1995) Neurobiology of skill and habit learning. *Curr Opin Neurobiol* 5:184–190.
- Sanes JN (1986) Kinematic and endpoint control of arm movements is modified by unexpected changes in viscous loading. *J Neurosci* 6:3120–3127.
- Seeds NW, Williams BL, Bickford PC (1995) Tissue plasminogen activator induction in purkinje neurons after cerebellar motor learning. *Science* 270:1992–1994.
- Shadmehr R, Holcomb HH (1996) Motor memory consolidation. In: *Neuroimage* (Toga AW, Frackowiak RSJ, Mazziotta JC, eds), p S590. San Diego: Academic.
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208–3224.
- Shadmehr R, Mussa-Ivaldi FA, Bizzi E (1993) Postural force fields of the human arm and their role in generating multi-joint movements. *J Neurosci* 13:45–62.
- Shadmehr R, Brashers-Krug T, Mussa-Ivaldi FA (1995) Interference in learning internal models of inverse dynamics in humans. In: *Advances in neural information processing systems*, Vol 7 (Tesauro G, Touretzky DS, Leen TK, eds), pp 1117–1124. Cambridge, MA: MIT.
- Shelhamer M, Robinson DA, Tan HS (1992) Context-specific adaptation of the gain of the vestibulo-ocular reflex in humans. *J Vestib Res* 2:89–96.
- Slotine J-JE, Li W (1991) *Applied nonlinear control*. Englewood Cliffs, NJ: Prentice Hall.
- Squire LR (1992) Declarative and nondeclarative memory: multiple brain systems supporting learning and memory. *J Cognit Neurosci* 99:195–231.
- Squire LR, Slater PC, Chace PM (1975) Retrograde amnesia: temporal gradient in very long-term memory following electroconvulsive therapy. *Science* 187:77–79.
- Squire LR, Slater PC, Miller P (1981) Retrograde amnesia following electroconvulsive therapy: long-term follow-up studies. *Arch Gen Psychiatry* 38:89–95.
- Squire LR, Cohen NJ, Nadel L (1984a) The medial temporal region and memory consolidation: a new hypothesis. In: *Memory consolidation: psychobiology of cognition* (Weingartner H, Parker ES, eds), pp 185–210. Hillsdale, NJ: Erlbaum.
- Squire LR, Cohen NJ, Zouzonis JA (1984b) Preserved memory in retrograde amnesia: sparing of a recently acquired skill. *Neuropsychologia* 22:145–152.
- Sutton RS (1986) Two problems with backpropagation and other steepest-descent learning procedures for networks. *Proc 8th Cognit Sci Soc* 8:823–831.
- Thoroughman KA, Shadmehr R (1996) Forward interference in motor learning. *Soc Neurosci Abstr* 22:897.
- Tiliket C, Shelhamer M, Tan HS, Zee DS (1993) Adaptation of the vestibulo-ocular reflex with the head in different orientations and positions relative to the axis of body rotation. *J Vestib Res* 3:181–196.
- Tully T, Preat T, Goynton SC, Del Vecchio M (1994) Genetic dissection of consolidated memory in *Drosophila*. *Cell* 79:35–47.
- Welch RB, Bridgeman B, Anand S, Browman KE (1993) Alternating prism exposure causes dual adaptation and generalization to a novel displacement. *Percept Psychophys* 54:195–204.
- Wolpert DM, Ghahramani Z, Jordan MI (1995a) Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study. *Exp Brain Res* 103:460–471.
- Wolpert DM, Ghahramani Z, Jordan MI (1995b) An internal model for sensorimotor integration. *Science* 269:1880–1882.
- Zipser D, Kehoe B, Littlewort G, Fuster J (1993) A spiking network model of short-term active memory. *J Neurosci* 13:3406–3420.
- Zola-Morgan S, Squire LR (1990) The primate hippocampal formation: evidence for a time-limited role in memory storage. *Science* 250:288–290.