Adaptive Representation of Dynamics During Learning of a Motor Task

Reza Shadmehr and Ferdinando A. Mussa-Ivaldi

Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, MA 02139 USA

Running Head: Representation of a learned motor task
Contents: 46 pages, including 1 appendix, 1 table, and 16 figures.

Key Words: Motor learning, reaching movements, internal models, force fields, virtual environments, generalization, motor control.

Published in: *Journal of Neuroscience*, May 1994, 14:3208–3224.

Send correspondence to: Dr. Reza Shadmehr, Room E25–201, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139. E-mail: reza@ai.mit.edu, Phone: (617) 253-0546, Fax: (617) 253-2964.

Acknowledgments: This work has been greatly enriched because of our interactions with Drs. Emilio Bizzi, Tomaso Poggio, Simon Giszter, Richard Held, Neville Hogan, Mike Jordan, and Eric Loeb. We are particularly grateful for the time and attention given to this project by Prof. Bizzi. Financial support was provided in part by Grants from the NIH (NS09343 and AR26710), and the ONR (N00014/90/J/1946). R. S. was supported by the McDonnell–Pew Center for Cognitive Neurosciences and the Center for Biological and Computational Learning at MIT.

June 14, 1994

Abstract

We investigated how the CNS learns to control movements in different dynamical conditions, and how this learned behavior is represented. In particular, we considered the task of making reaching movements in the presence of externally imposed forces from a mechanical environment. This environment was a force field produced by a robot manipulandum, and the subjects made reaching movements while holding the end–effector of this manipulandum. Since the force field significantly changed the dynamics of the task, subjects' initial movements in the force field were grossly distorted compared to their movements in free space. However, with practice, hand trajectories in the force field converged to a path very similar to that observed in free space. This indicated that for reaching movements, there was a kinematic plan independent of dynamical conditions.

The recovery of performance within the changed mechanical environment is motor adaptation. In order to investigate the mechanism underlying this adaptation, we considered the response to the sudden removal of the field after a training phase. The resulting trajectories, named *after–effects*, were approximately mirror images of those which were observed when the subjects were initially exposed to the field. This suggested that the motor controller was gradually composing a model of the force field, a model which the nervous system used to predict and compensate for the forces imposed by the environment. In order to explore the structure of the model, we investigated whether adaptation to a force field, as presented in a small region, led to after–effects in other regions of the workspace. We found that indeed there were after–effects in workspace regions where no exposure to the field had taken place, i.e., there was transfer beyond the boundary of the training data. This observation rules out the hypothesis that the subject's model of the force field was constructed as a narrow association between visited states and experienced forces, i.e. adaptation was not via composition of a look–up table. In contrast, subjects modeled the force field by a combination of computational elements whose output was broadly tuned across the motor state space. These elements formed a model which extrapolated to outside the training region in a coordinate system similar to that of the joints and muscles rather than endpoint forces. This geometric property suggests that the elements of the adaptive process represent dynamics of a motor task in terms of the intrinsic coordinate system of the sensors and actuators.

1 Introduction

Children start to reach for objects that interest them at about the age of 3 months. These goal-directed movements often accompany a "flailing" action of the arm. From a systems point of view, flailing can be seen as an attempt to excite the dynamics of the arm: to successfully make a reaching movement, the motor controller needs to find the appropriate force so that the skeletal system makes the desired motion. Effectively, this operation corresponds to inverting a dynamical transformation that relates an input force to an output motion. A controller may implement this "inverse transformation" via a combination of feedback and feedforward mechanisms: usually, the feedforward component provides some estimate of the inverse transformation— called the "inverse model" or simply the "internal model"—while the feedback component compensates for the errors of this estimation and stabilizes the system about the desired behavior (cf. Slotine 1985). Therefore, the internal model refers to an approximation of the inverse dynamics of the system being controlled. In the case of the infant, the action of flailing may be an attempt to explore this dynamics and build an internal model.

During development, bones grow and muscle mass increases, changing the dynamics of the arm significantly. In addition to such gradual variations, the arm dynamics change in a shorter time scale when we grasp objects and perform manipulation. The changing dynamics of the arm make it so that the same muscle forces produce a variety of motor behaviors. It follows that to maintain a desired performance, the controller needs to be "robust" to changes in the dynamics of the arm. This robustness may be achieved through an updating, or adaptation, of the internal model. Indeed, humans excel in the ability to rapidly adapt to the variable dynamics of their arm as the hand interacts with the environment. Therefore a task where the hand interacts with a novel mechanical environment might be a good candidate for studying how the CNS updates its internal model and learns dynamics.

The particular task which we have considered is one where a subject makes a reaching movement while the hand interacts with a field of forces. In a reaching movement, the problem of control can be seen as one of transforming information regarding a target position, as presented in the visual domain, into a torque command on the skeletal system to move the hand. This initially involves a set of coordinate transformations (so called "visuo-motor map", cf. Arbib 1976): work of Andersen et al. (1985) and Soechting and Flanders (1991) suggests that the target is transformed sequentially from a retino–centric vector into a head–centered and finally a shoulder–centered coordinate system. According to Gordon et al. (1993), the target is finally represented as a vector pointing from the current hand position (or end–effector position, for example, in the case that the hand is holding a long rod, Lacquaniti et al. 1982) to the target. At this point a *plan* is specified, describing a desired trajectory for the end–effector to follow: for unconstrained planar arm movements, there is strong evidence that this plan is a smooth hand trajectory essentially along a straight line to the target (Morasso 1981, Flash and Hogan 1985). The controller, acting on antagonistic spring-like actuators (cf. Bizzi et al. 1984, Hogan

1985, Shadmehr and Arbib 1992), then attempts to move the arm along the planned trajectory.

It is worth noting that for this task, adaptation may either occur in response to a change in the visual environment in which the target is presented (cf. von Helmholtz 1925, Cunningham 1989, Wolpert et al. 1993), or in response to a change in the mechanical environment with which the hand is interacting (cf. Lacquaniti et al. 1982, Ruitenbeek 1984, Flash and Gurevich 1992). Therefore, the problem of adaptation may be experimentally approached from two directions:

1. we may change the visual environment so that subjects have to modify the perceived *kinematics* of movement by changing the mapping of the target from ego–centric to a task based (e.g., hand-centered) coordinates, or

2. we may change the mechanical environment with which the hand interacts so that the subject's internal model of the arm has to adapt to the new *dynamics* of the system.

The first approach, i.e., changing the visually perceived kinematics, has received much attention because of the observations made by Held and colleagues (Held and Schlank 1959, Held 1962, Held and Freedman 1963) regarding adaptation of the visuomotor system to distortions produced by prism glasses. It had been noted that by wearing prism glasses, the visual scene could be shifted, for example, by x degrees laterally. This caused a change in the kinematic map relating target position to the arm's configuration. With the glasses on, initially a subject would reach to a target and miss it by x degrees, but after some practice, the subject would learn the appropriate kinematics and hit the target accurately. Predictably, when the glasses were removed, the subject would reach to a target and miss it by $-x$ degrees, displaying the persistence of the altered kinematic map (cf. Jeannerod 1988, pp. 52–57). This behavior has been termed an *after–effect* of adaptation.

Our work is along the second approach. We investigate how the motor control system responds when the arm's dynamics are changed. We address this issue by developing a paradigm where subjects make reaching movements while interacting with a virtual mechanical environment. From Lackner and Dizio (1992) it is known that after–effects exist when one performs arm movements in an environment where Coriolis forces are artificially increased. Here we show that as a subject practices arm movements in a force field, the controller builds an internal model of that field and uses this model to compensate for the expected forces during the movement. Our goal is to understand how the nervous system constructs this internal model and to reveal some of the properties of the motor adaptive process.

2 Materials and Methods

The purpose of our experiment was to observe how a subject adapted to the changed dynamics of a reaching task. A robot manipulandum whose handle was grasped by the subject produced these variable dynamics. A mathematical model was developed to provide a framework for

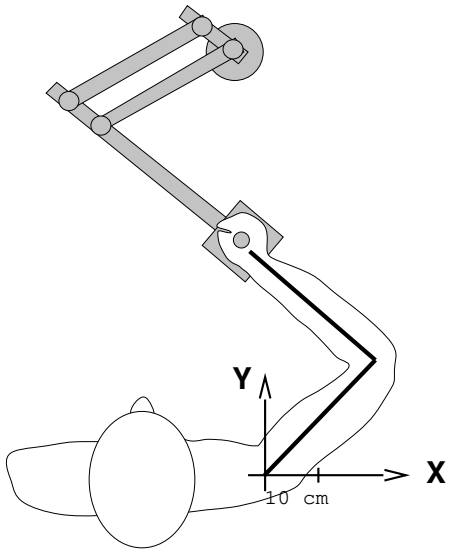

Figure 1: Sketch of the manipulandum and the experimental setup. Planar arm movements were made by the subject while grasping the handle of the manipulandum. A monitor, placed directly in front of the subject and above the manipulandum (not shown), displayed the location of the handle as well as targets of reaching movements. The manipulandum had two torque motors at its base which allowed for production of a desired force field.

describing the process of adaptive motor control. Both the experiments and the modeling procedures are described in this section.

2.1 Experimental setup

Eight right handed subjects with no known neurological history, ranging in age from 24 to 39, participated in this study. A schematic of the measurement apparatus is shown in Fig. 1: Subjects were seated on a chair that was bolted onto an adjustable positioning mechanism and instructed to grip the handle of a robot manipulandum with their right hand. Their shoulder was restrained by a harness belt, their right upper-arm was supported in the horizontal plane by a rope attached to the ceiling.

The manipulandum is a two degree of freedom, lightweight, low friction robot (Faye 1986) with a six-axis force-torque transducer (Lord F/T sensor) mounted on its end–effector (the handle). Two low inertia, DC torque motors (PMI Corp., model JR16M4CH), mounted on the base of the robot, are connected independently to each joint via a parallelogram configuration. Position and velocity measurements are made using two optical encoders (Teledyne Gurley) and tachometers (PMI), respectively, mounted on the axes of the mechanical joints. The apparatus includes a video display monitor mounted directly above the base of the robot (approximately at eye level with the subject). This was used to display the position of the robot's handle and give targets for reaching movements.

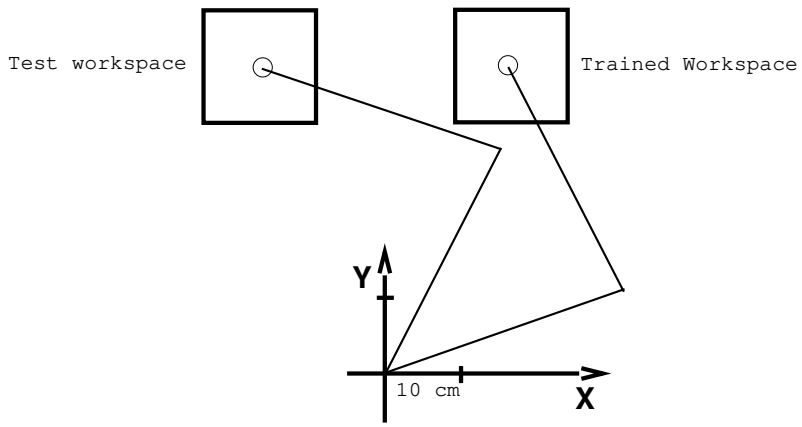

Figure 2: Configurations of a model two joint arm, representing typical kinematics of the human arm, at two workspace locations where reaching movements were performed. Typical shoulder and elbow angles at these two workspaces were 15 and 100 degrees at right and 60 and 145 degrees at left, using coordinates defined in Fig. 1.

2.2 Experimental procedures

Each subject participated in a preliminary training phase where the task was to move a cursor to a target. The cursor was a square of size 2×2 mm^2 on a computer monitor and indicated the position of the handle of the manipulandum. Targets were specified by a square of size 8×8 mm^2. The task was to move the manipulandum so as to bring the cursor within the target square.

Movements took place in two regions, each of the size 15×15 cm^2. The position of these regions is shown in Fig. 2, where they are labeled as the "left" and "right" workspaces. In order to avoid inertial artifacts associated with changing the operating configuration of the robot, workspaces were selected by moving the subject with respect to the robot.

Starting from the center of a workspace, a target at a direction randomly chosen from the set $\{0°, 45°, \ldots, 315°\}$ and at a distance of 10 cm was presented. After the subject had moved to the target, the next target, again chosen at a random direction and at 10 cm was presented. A *target set* consisted of 250 such sequential reaching movements. All targets were kept with in the confines of the 15×15 cm workspace. The targets represented a pseudo-random walk.

In some cases, the manipulandum was programmed to produce forces on the hand of the subject as the subject performed reaching movements. These forces, indicated by the vector f, were computed as a function of the velocity of the hand:

$$f = B\,\dot{x} \tag{1}$$

where $\dot{x}$ was the hand velocity vector, and B was a constant matrix representing viscosity of the imposed environment in endpoint coordinates. In particular, we chose B to be:

$$B = \begin{bmatrix} -10.1 & -11.2 \\ -11.2 & 11.1 \end{bmatrix} \text{N.sec/m}$$

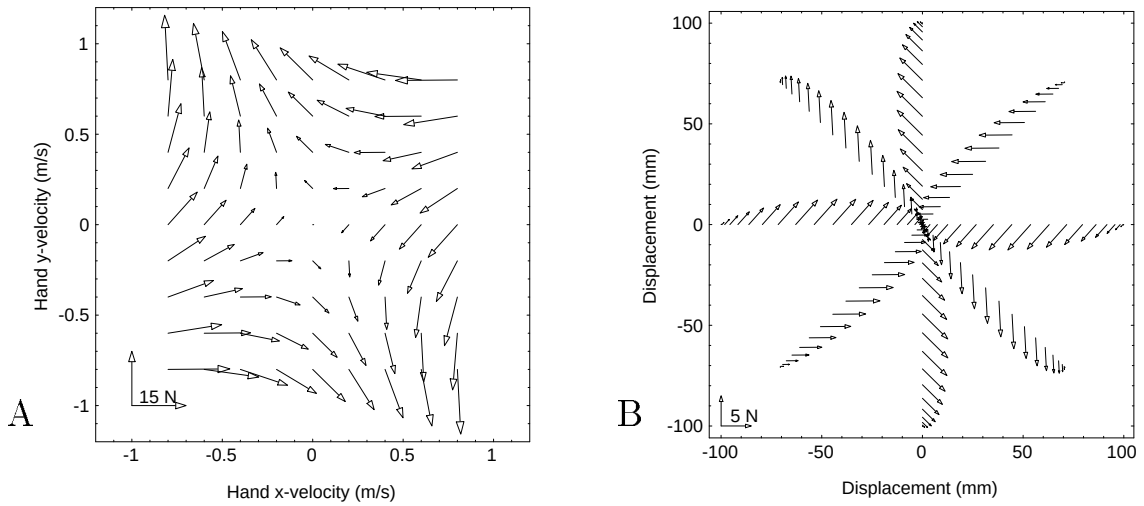

Figure 3: An environment as described by the force field in Eq. (1). **A:** The force field. **B:** Forces acting on the hand during simulated center–out reaching movements. Movements are simulated as being minimum–jerk with a period of 0.5 sec and amplitude of 10 cm.

Using this matrix, the forces defined by Eq. (1) may be shown as a field over the space of hand velocities (Fig. 3A). For example, as a subject made reaching movements in this field, the manipulandum produced forces shown in Fig. 3B (here we have assumed that the movements are minimum jerk, as specified by Flash and Hogan (1985), with a period of 0.5 seconds).

Note that in the field defined by Eq. (1), forces which act on the hand are invariant to the location of the workspace in which a movement is done, i.e., the forces are identical in the left and right workspaces of Fig. 2. Therefore, we say that the force field defined in Eq. (1) is translation invariant in endpoint coordinates.

In some cases, a different kind of a force field was produced by the manipulandum, one which was not translation invariant in endpoint coordinates. This field was represented as a function of the velocity of the subject's shoulder and elbow joints during the reaching movements:

$$\tau = W\ \dot{q} \tag{2}$$

where τ was the torque vector acting on the subject's shoulder and elbow joints, $\dot{q}$ was the subject's joint angular velocity, and W was a constant matrix representing viscosity of the imposed environment in joint-coordinates of the subject. We say that the field described by Eq. (2) is translation invariant in joint-coordinates. Indeed, note that the torque field in Eq. (2) is equivalent to the following force field (i.e., forces acting on the hand):

$$f = \left(J(q)^T\right)^{-1}\ W\ \dot{q} \tag{3}$$

where $J(q) = \partial x/\partial q$, is the configuration–dependent Jacobian of the configuration mapping from q to x, and the superscript T indicates the transpose operation. Because the Jacobian changes as a function of the angular position of the limb, f varies depending on the workspace where a reaching movement is performed. In particular, we chose W so that the force field

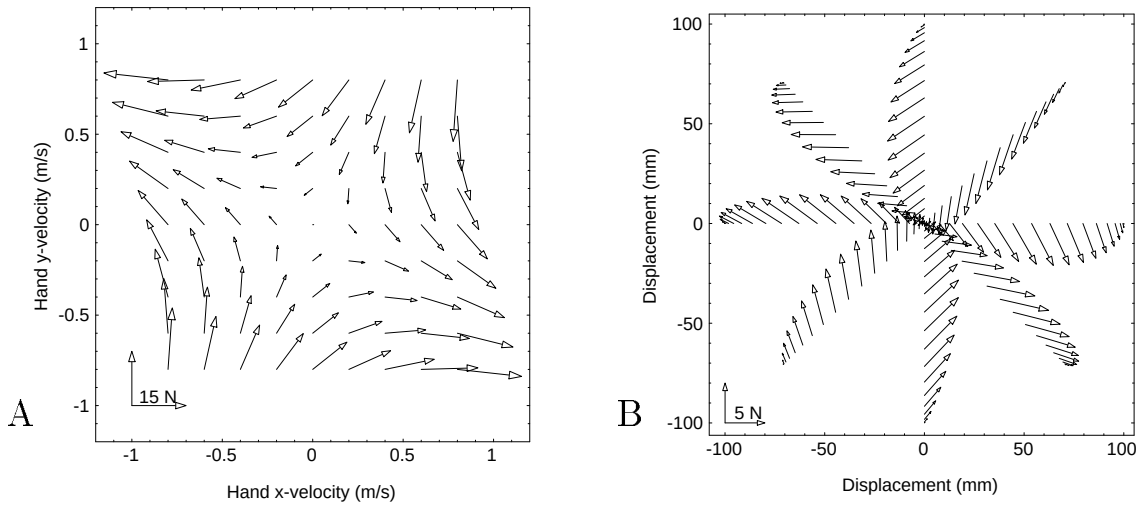

Figure 4: An environment described by the field in Eq. (3), plotted as it would appear in the left workspace of Fig. 2. **A:** The force field. **B:** Forces acting on the hand while making reaching movements in the left workspace of Fig. 2 from the center to targets about the circumference of a circle. Movements are simulated as being minimum-jerk with a period of 0.5 sec and amplitude of 10 cm.

which resulted from Eq. (3) at the right workspace was almost identical to the field produced by Eq. (1). To accomplish this, the matrix W was calculated for each subject as:

$$W = J_0^T \, B \, J_0$$

where J_0 is the Jacobian evaluated at the center of the right workspace. For a typical subject, we derived the following W matrix:

$$W = \begin{bmatrix} 1.66 & 0.64 \\ 0.64 & -1.54 \end{bmatrix} \text{ N.m.sec/rad}$$

When the above joint-viscosity matrix was used to define an environment, the resulting force field depended upon the position of the workspace where movements were being made. At the right workspace, this field (Eq. 3) was almost identical to that produced by Eq. (1) (a correlation coefficient of 0.99, see Appendix I). However, at the left workspace, the forces produced by Eq. (3) were substantially uncorrelated (nearly orthogonal) to that of Eq. (1). The force field produced by Eq. (3) is plotted for movements in the left workspace in Fig. 4A. Fig. 4B shows the forces acting on the hand for typical reaching movements.

We trained subjects with either the end–point or the joint translation invariant fields at the right workspace. Subsequently, we tested them in the field they had not been trained on at the left workspace. Hence, we defined two distinct groups of subjects. Those in **Group 1** were exposed to a field which was translation invariant with respect to the position of the hand (Eq. 1). Subjects of **Group 2** were exposed to a field which was translation invariant with respect to the angular position of the subject's joints (Eq. 3).

Our first objective was to compare movements during conditions of no-visual feedback before and during the initial exposure to a field. For 48 randomly chosen members of the target set,

hereafter referred to as the *no-vision target set*, the cursor position during the movement was blanked, removing visual feedback during the reaching period. For the remaining members of the target set, hand position was shown continuously to the subject. Initially, we quantified the performance in a null field, i.e., with the torque motors turned off, by presenting a target set in the right workspace. Upon completion, the hand was moved to the left workspace and another target set presented. These hand trajectories represented performance of the subjects in the null field.

Following this, the hand was returned to the right workspace and the target set was again presented, except that for 24 randomly chosen members of the no-vision target set, the manipulandum produced the force field assigned to the subject's group. For the remaining targets of this set a null field was present. These hand trajectories during the no-vision target set represented baseline performance in the force field.

The next objective was to observe performance of the subject in response to continuous exposure to the force field: With the hand at the right workspace and with the manipulandum producing the force field, a target set was presented. The force field was present for all targets except for 24 randomly chosen members of the no-vision target set, where the null field was present. The purpose of these 24 targets in the null field was to record any *after-effects* of adaptation to the force field. The target set was repeated 4 times (total of 1000 movements) while the manipulandum produced the field. This provided time for the subject to adapt.

Having completed the adaptation phase of the experiment, the subject's arm was moved to the left workspace with the objective of observing any *transferred after-effects*. 72 targets were presented sequentially and with no visual feedback. 24 randomly chosen members of this target set were in a null field. Another 24 randomly chosen members of this target set were in the force field on which the subject had been trained. The remaining members of this target set were in the force field which the subject had not been trained on. In Fig. 5 the experimental procedure is summarized.

2.3 Producing the force fields

In order for the manipulandum to produce a given force field, the microcomputer collected position and velocity information from the manipulandum's joints (represented by ϕ and $\dot{\phi}$) at a rate of 100 Hz. This information was needed in order to convert the desired endpoint force field into the torques to be applied by the motors. To produce the force field described by Eq. (1), we used the following expression:

$$\tau_R = J_R^T \, B \, J_R \, \dot{\phi}$$

where τ_R is the torque vector commanded to the motors, $J_R = \partial x / \partial \phi$, i.e., the Jacobian of the robot's kinematics, and $\dot{\phi}$ is the joint angular velocity vector of the manipulandum. Note that J_R is a function of robot joint angles ϕ, and from its definition it follows that $\dot{x} = J_R \, \dot{\phi}$. In

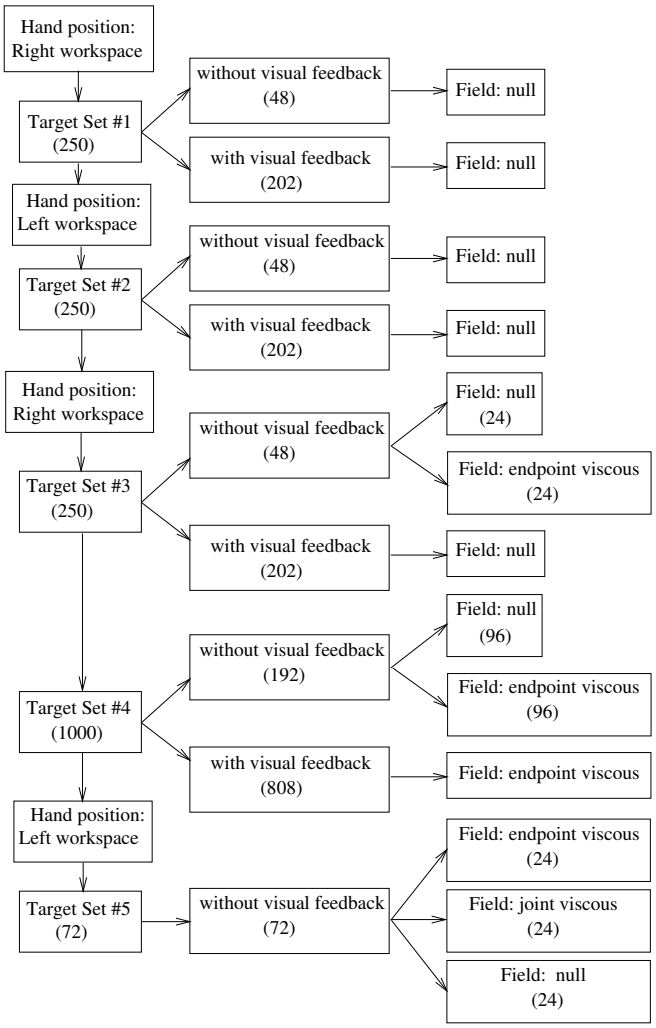

Figure 5: Summary of the experimental procedure for subjects in Group 1. The adaptation period was during target set number 4 where an end-point viscous field was present. Subjects in Group 2 underwent an identical procedure except that during the training period a joint-based viscous field was present.

order to produce the force field described by Eq. (3), the following control law was used:

$$\tau_R = J_R^T \, J^{-T} \, W \, J^{-1} \, J_R \, \dot{\phi}$$

where J is the subject's Jacobian matrix function. Calculation of J required knowledge of the subject's arm kinematics: At the beginning of each session, we measured the lengths of the subject's upperarm and forearm as well as the location of the shoulder with respect to a fixed point with respect to the workspace of the manipulandum. These data were sufficient to provide an estimate for J at each position of the hand.

2.4 Data analysis

We sampled hand positions and velocities at 10 msec intervals as the subject reached to a target. Trajectories were aligned using a velocity threshold at the onset of movement.

In order to compare hand trajectories, a technique was developed which quantified a measure of correlation between two sampled vector fields (see Appendix 1). We represented each trajectory as a time series of velocity vectors ($\dot{x}$ sampled at 10 msec intervals) and then compared the two resulting vector fields through a correlation measure. The same technique was also used to compare force fields. In particular, the endpoint viscosity matrix B in Eq. (1) was chosen such that when expressed in terms of a joint viscosity matrix W (through Eq. 3), the two resulting force fields were nearly identical at the right workspace (the correlation coefficient $\rho \approx 1$), while maximally different at the left workspace ($\rho \approx 0$). Specifically, the two fields had a correlation coefficient of 0.99 and 0.12 at the right and left workspaces, respectively.

In order to plot "typical" hand trajectories for a given target, we computed the expected value and standard deviation of the set of measured trajectories (each a time–series of velocity vectors) for that target. Our procedure consisted in deriving the expected value and standard deviation of the set of measured velocity vectors across the trajectories at each instant of time. The resulting velocity field was integrated from the start position of the movement to produce the average $\pm$ standard deviation of the hand trajectories for a given target.

2.5 Mathematical modeling

The purpose of the mathematical modeling was help describe the concept of an "internal model". We used this approach to simulate hand trajectories for reaching movements before the subject had adapted to the force field, as well as the after-effects when the subject had formed an internal model but the external field was suddenly removed.

Let us start by considering the arm's dynamics in generalized coordinates (cf. Spong and Vidyasagar 1989, p. 131): We indicate by q a point in configuration space (e.g., an array of joint angles) and by $\dot{q}$ and $\ddot{q}$ its first and second time derivatives. The dynamics of the motor–control system coupled (in parallel) with its environment can be described by the sum of the following terms: a time–invariant component, $D(q, \dot{q}, \ddot{q})$ and $E(q, \dot{q}, \ddot{q})$, representing the forces which depend on the "passive" or unmodulated system dynamics (bones, tendons, etc.) and forces which depend on dynamics of the environment, and a time–varying component, $C(q, \dot{q}, t)$, representing the forces which depend on the operation of the controller.

$$D(q, \dot{q}, \ddot{q}) + E(q, \dot{q}, \ddot{q}) = C(q, \dot{q}, t) \tag{4}$$

The force field represented by D is itself a sum of inertial, Coriolis, centripetal, and friction forces:

$$D(q, \dot{q}, \ddot{q}) = I(q)\, \ddot{q} + G(q, \dot{q}) \tag{5}$$

where I represents the system's mass in generalized coordinates (an inertia matrix, which may be a function of configuration), and G represents the rest of the position and velocity dependent forces (i.e., Coriolis, friction, etc.).

Let us consider a control system that is capable of guiding a limb along a desired trajectory $q^*(t)$ in the null environment $E = 0$. One way to obtain this tracking behavior is by picking the

right hand side of Eq. (4) to be an *ideal controller* specified by $I(q)\ddot{q}^*(t)+G(q,\dot{q})$. This simplifies Eq. (4) to $\ddot{q} = \ddot{q}^*(t)$, from which it follows that from some given initial position and velocity, the system will follow the desired trajectory. Note that this ideal control input describes a time varying force field: for a given desired acceleration, a force vector is assigned to each point in the state-space of the system. We name this ideal controller $\mathcal{D}$, i.e.,

$$\mathcal{D}(q,\dot{q},t) = I(q)\ddot{q}^*(t) + G(q,\dot{q}). \tag{6}$$

We call this controller "ideal" because it may well be that one cannot implement its field using the available actuators and local controllers. However, one may be able to approximate its force field, resulting in an internal model of the system dynamics. Let us call this internal model $\hat{\mathcal{D}}$, where for the system dynamics of Eq. (5), with a null environment, our internal model may be defined by the following field:

$$\hat{\mathcal{D}} = \hat{I}\,\ddot{q}^*(t) + \hat{G} \tag{7}$$

Note that the internal model is not a model of the dynamical system, but a model of the ideal controller for that dynamical system. Unfortunately, even with an exact model the system will be unstable about the desired trajectory: our controller will not be able to compensate for the slightest unexpected change in initial conditions or for any perturbation occurring during the movement. One way to overcome this is to define our controller C in Eq. (4) (assuming a null environment for now) so that it combines the internal model of Eq. (7) with an error–feedback system designed to provide stability about the desired trajectory:

$$C(q,\dot{q},t) = \hat{\mathcal{D}} - S(q-q^*(t), \dot{q}-\dot{q}^*(t)) \tag{8}$$

where S is a converging force field about the desired state of the system at time t, i.e., it has zero forces only when both of its arguments are zero (Slotine and Li 1991). This kind of representation for the controller is particularly well suited to the biomechanical system of the arm when we consider that the function S may be implemented via the stiffness and viscosity of antagonist muscles and their associated segmental reflexes:

$$C(q,\dot{q},t) = \hat{I}\,\ddot{q}^*(t) + \hat{G} - K\,(q-q^*(t)) - V\,(\dot{q}-\dot{q}^*(t))$$

where K and V are joint stiffness and viscosity matrices describing the behavior of the field S about the desired trajectory.

Now let us apply an environment $E \neq 0$ and consider the problem of finding a new controller such that $q^*(t)$ is still the solution for the coupled dynamics described by Eq. (4). The procedure is similar to the one just described: ideally, we would like to replace the right hand side of Eq. (4) by the field: $\mathcal{D}(q,\dot{q},t) + \mathcal{E}(q,\dot{q},t)$, where $\mathcal{E}$ is an ideal control input chosen such that the differential equation $E(q,\dot{q},\ddot{q}) = \mathcal{E}(q,\dot{q},t)$ has a solution $q^*(t)$ from a given initial position. We therefore express the new controller as:

$$C(q,\dot{q},t) = \hat{I}\,\ddot{q}^*(t) + \hat{G} + \hat{\mathcal{E}} - K\,(q-q^*(t)) - V\,(\dot{q}-\dot{q}^*(t)) \tag{9}$$

Table 1: Mechanical parameters of the simulated human arm

Upperarm	mass	1.93 kg
	center of mass	0.165 m
	inertia	0.0141 kg.m^2
	length	0.33 m
Forearm	mass	1.52 kg
	center of mass	0.19 m
	inertia	.0188 kg.m^2
	length	0.34 m
Stiffness	$\begin{bmatrix} -15 & -6 \\ -6 & -16 \end{bmatrix}$	N.m/rad
Viscosity	$\begin{bmatrix} -2.3 & -0.9 \\ -0.9 & -2.4 \end{bmatrix}$	N.m.sec/rad

where $\hat{\mathcal{E}}(q, \dot{q}, t)$ is our model of the environment, expressed as a first order time varying field:

$$\hat{\mathcal{E}} \sim \mathcal{E}(q, \dot{q}, t) \tag{10}$$

Assuming that the system was capable of producing the desired trajectory in the absence of an environment, then it is apparent that as $\hat{\mathcal{E}} \rightarrow \mathcal{E}$, the coupled dynamics is reduced back to the form of Eq. (4), of which the desired trajectory $q^*(t)$ was a particular solution. The idea then is to achieve a motor plan through a change in the dynamics of the system such that the new dynamics have an "attractor" at the to-be-learned trajectory. This formalism is very similar to the learning framework of Kelso, Saltzman and coworkers (Kelso and Schoner 1988, Saltzman and Kelso 1989, Schoner et al. 1992).

We used the controller in Eq. (9) coupled with the arm's dynamics to simulate performance before and after adaptation (e.g., the after-effects). The skeletal dynamics of Eq. (5) were simulated for each subject using an inertial matrix $I(q)$ as measured by Diffrient et al. (1978), and given for a typical subject in Table 1 (the Coriolis and centripetal forces which make up the G matrix can be derived from the inertia tensor, cf. Slotine and Li 1991, p. 400). For example, the differential equation describing the dynamics of the arm and the controller for movements in the force field of Eq. (1) were:

$$I(q)\,\ddot{q} + G(q, \dot{q}) + J(q)^T\, B\, J(q)\,\dot{q} = C(q, \dot{q}, t) \tag{11}$$

where C is defined in Eq. (9). Values for joint stiffness and visocity (K and V) were chosen based on measurements of Mussa-Ivaldi et al. (1985) and Tsuji and Goto (1993). The desired trajectory $q^*(t)$ was assumed to be minimum jerk in hand-based coordinates lasting 0.65 seconds. Values used for these variables are summarized in Table 1.

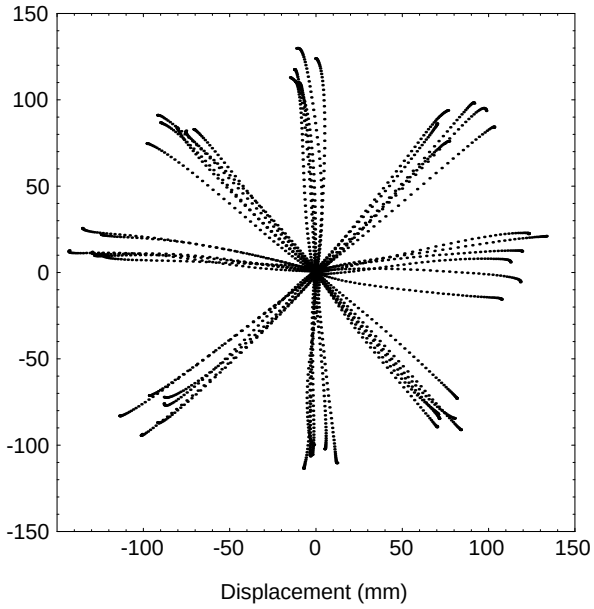

Figure 6: Typical hand trajectories at the right workspace in a null force field during no–visual feedback conditions. Dots are 10 msec apart.

3 Results

Reaching movements were made while the hand interacted with a mechanical environment. This environment was a programmable force field implemented by a light weight robot manipulandum whose end–effector the subject grasped while making reaching movements. When the manipulandum was producing a force field, there were forces which acted on the hand as it made a movement, changing the dynamics of the arm. When the manipulandum's motors were turned off, we say that the hand was moving in a "null field".

3.1 Hand trajectories before adaptation

Our first objective was to determine how an unanticipated velocity–dependent field affected the execution of reaching movements. The forces in the field (e.g., Eq. 1, as shown in Fig. 3A) vanished when the hand was at rest, that is, at the beginning and at the end of the movement. However, as shown in Figure 4B, a significant force was exerted midway, when the hand velocity was near maximum. How would this force influence the execution of a movement? Would subjects follow a pre–planned trajectory that was scarcely influenced by this perturbation or would they modify the movement and the final position in response to the perturbing force? To answer this question, we compared reaching movements in the null field with those in a force field. Trajectories in the null field are shown in Fig. 6. As observed in previous reports (Morasso 1981, Flash and Hogan, 1985), the hand path was essentially along a straight line to the target. The velocity profile (see Fig. 10A) had one peak, with approximately equal times spent to accelerate and decelerate the hand.

Once our subjects were familiar with the task of reaching within the null field, we began to

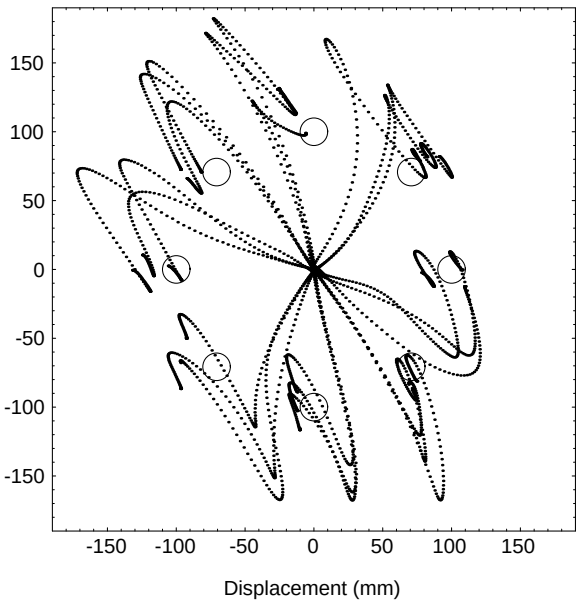

Figure 7: Performance during initial exposure to a force field: Shown are hand trajectories to targets at the right workspace while moving in the force field shown in Fig. 3. Movements originate at the center. All trajectories shown are under no–visual feedback condition. Dots are 10 msec apart.

introduce a force field in random trials. Note that subjects could not anticipate the presence of the field before the onset of the movement because the force field was not effective when the hand was at rest and no other clues were available. Furthermore, during the movement, the cursor indicating hand position was blanked, eliminating visual feedback. Figure 7 shows the hand trajectories of a typical subject when the movements were executed under the influence of the field shown in Fig. 3A. Figure 10B shows the tangential velocity of hand trajectories in this field. This field was designed to have opposing effects along two directions. At approximately 30 and 210 degrees the field produced resisting forces that opposed movement as a viscous fluid would do. At approximately 120 and 300 degrees the forces assisted the movement, thus producing a de-stabilizing effect.

Note that the effect of the field on the hand trajectory was quite significant and may be divided into two parts. In the first part, the hand was driven off–course by the field and forced towards the unstable direction of the field. Movements to targets at $0, 225, 270$, and 315 degrees are pulled toward the unstable region at 300 degrees, while movements to the remaining targets are pulled toward the unstable region at 120 degrees. At the end of this first part, the field had caused the hand to veer off the direction of the target and the hand decelerated and stopped before making a second movement to the target. The pictorial effect of these two parts of the hand trajectory appeared as a "hook" that was oriented either clockwise or counterclockwise. The orientation and the overall appearance of this hook was found to depend upon the position of the target and the pattern of forces in the field, and was very similar among the 8 subjects.

One may interpret the hooks shown in Fig. 7 as "corrective movements" that are generated to compensate for the errors caused by the unexpected field. In light of the fact that no visual

feedback was available to the subjects during the movements shown in Fig. 7, this correction might imply some explicit reprogramming of the movement based on proprioceptive information detecting the error in the hand trajectory. Alternatively, this feature of the trajectory might be a byproduct of a "robust" control system implementing a *single program*: In this case, the program would be to simply move the hand along a desired trajectory to the target. The corrective movements might result because of the natural interaction between the mechanical properties of the arm, as imposed from the controller, and the force field produced by the manipulandum. To explore this scenario, we simulated the operation of a controller acting on the arm's skeletal system via antagonistic muscles within the force field. The controller, which is detailed in the Methods section (Eq. 9), was designed based on the assumption that the goal was to move the limb along a smooth, straight line trajectory to the target. We further assumed that the controller had, through years of practice, composed an accurate internal model of the skeletal dynamics. However, recognizing that there might be errors in this internal model, the controller used the viscoelastic properties of the muscles to make the system stable about this desired trajectory, i.e., the system resisted perturbations (whether external or due to model errors) as it moved along the planned trajectory. In our simulation, we initially assumed that the controller had no knowledge of the forces in the environment, i.e., $\hat{\mathcal{E}} = 0$. Then we calculated the desired joint–trajectories, $q^*(t), \dot{q}^*(t), \ddot{q}^*(t)$, corresponding to straight–line movements of the hand towards the 8 targets. Finally, given the parameters in Table 1, we integrated Eq. (4) for producing the motion of the hand in the force field.

The results of this simulation are shown in Fig. 8. We found that there was a striking resemblance between the result of the modeled control system (Fig. 8) and those measured in our subjects (Fig. 7). In particular, the presence of the "hooks" as well as their orientation is accurately accounted for by the modeled controller. The quantitative differences between model and data are likely a consequence of errors and simplifications in estimating mechanical parameters of the arm of each subject: for example, in Eq. 9, we assumed a constant stiffness K for the arm. This is true when the arm is near the desired position, i.e., when $q - q^*$ is small. However, it is known that K becomes progressively and significantly smaller as the distance between the actual and desired hand positions increases (Shadmehr et al. 1993). The simulations also suffer from the fact that our dynamical model neglects the small but non–zero forces due to the inertia of the manipulandum.

The observed corrective movements or hooks in Fig. 7 are consistent with the operation of a controller which is attempting to move the limb along a desired trajectory and bring it to a specified target position. Because this controller uses muscle viscoelastic properties to define an attractor region about the desired trajectory, the hand is eventually brought back to near the target position. The hooks result from the interaction of the viscoelastic properties of the muscles and the force field which perturbs the system from its desired trajectory. Indeed, the results of the model suggests that the subjects may be executing a single program, i.e., that of moving the hand along a specified plan.

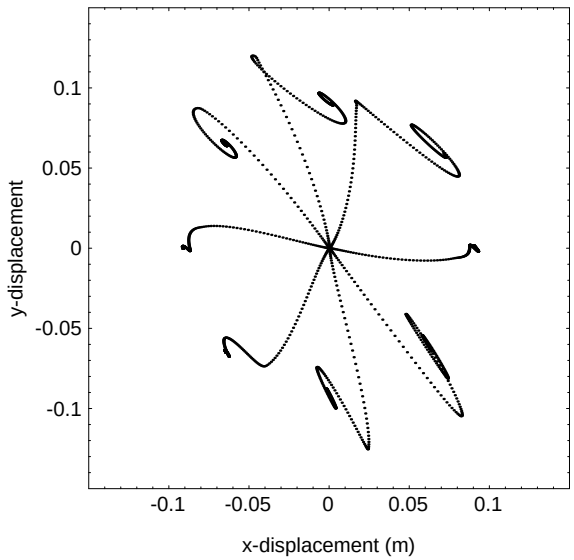

Figure 8: Simulation of hand trajectories in the force field of Fig. 3 before having formed an internal model, i.e., $\hat{\mathcal{E}} = 0$ in Eq. (8). Dots are 10 msec apart.

3.2 Adaptation to the force field

After measuring the movements of the arm in the null field as well as the initial responses to the unanticipated force field, we asked our subjects to keep executing reaching movements in the force field. We wish to stress that we did not give any instructions regarding the trajectory with which the targets should have been reached. Nevertheless, as the subjects practiced in the force field, the "hooks" shown in Fig. 7 eventually vanished and the hand trajectories became increasingly similar to those observed in the null field (Fig. 6). The progression of hand position traces as measured under conditions of no visual feedback and in the presence of the force field during the training period are shown in Fig. 9. Although the force field initially caused a significant divergence from the trajectory that was normally observed for a reaching movement, with practice, the subjects tended to converge upon this straight line trajectory. This recovery of the original unperturbed response constitutes a clear example of an adaptive behavior.

Further evidence of motor adaptation is offered by the significant change that occurred in the hand velocity profile at the onset of exposure to the force field, and after completion of the practice trials: Figure 10A shows the hand tangential velocity traces obtained when the hand was moving in a null field (corresponding to the hand position traces of Fig. 6). Consistent with previous studies (cf. Flash and Hogan 1985), these velocity traces are approximately along straight lines and symmetric in time. The hand velocity traces at the initial stage of practice in the force field (corresponding to the hand position traces of Fig. 7) are shown in Fig. 10B. In Fig. 10C we have the velocity traces near the end of the practice trials (corresponding to the hand position traces of Fig. 9D). Although the average velocity of the hand trajectory is now larger (as compared to Fig. 10A), the velocity trace for each target has essentially the same pattern as that observed for movements in a null field.

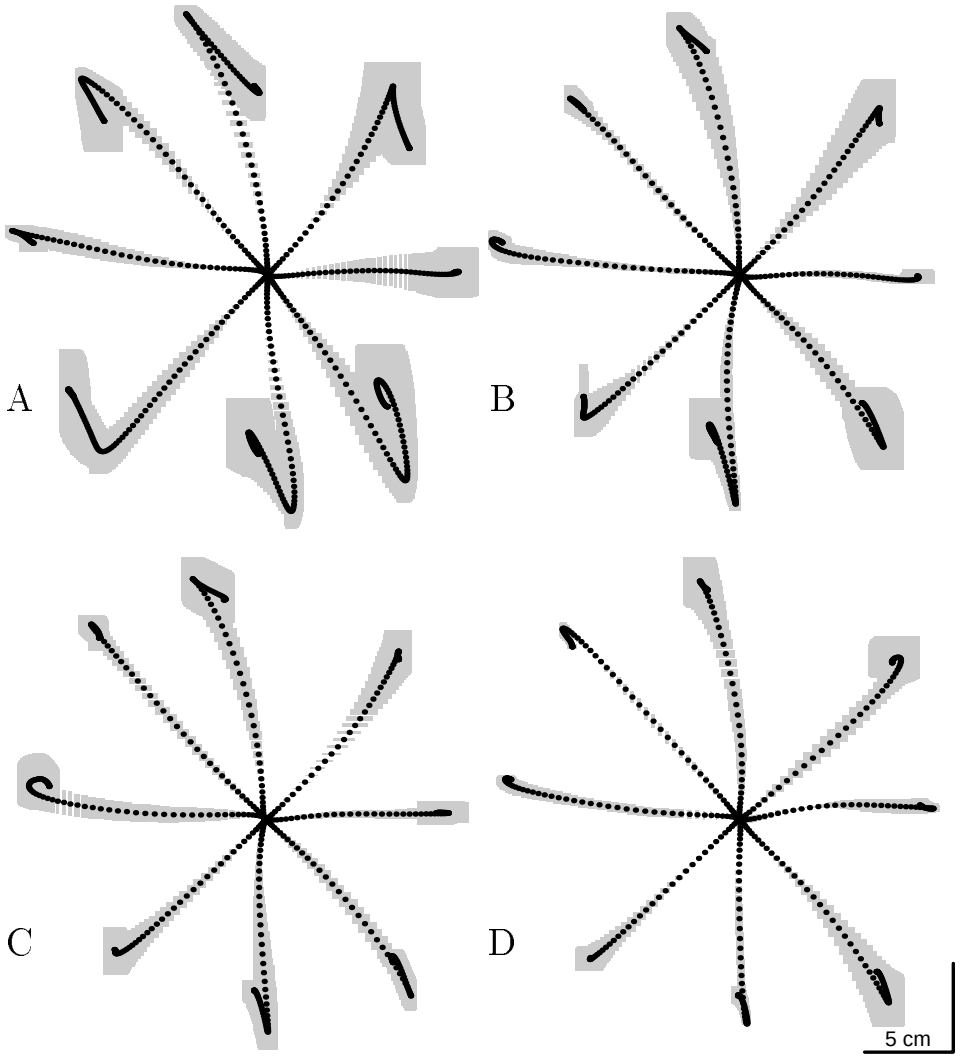

Figure 9: Average ± standard-deviation of hand trajectories during the training period in the force field of Fig. 3. Performance plotted during the first, second, third, and final 250 targets (A, B, C, and D, respectively). All trajectories shown are under no–visual feedback condition.

In order to quantify the time course of adaptation, we studied how the hand trajectories evolved as compared to those observed in the null field. For each subject, we compared the trajectories in the null field to those obtained as the subject practiced in the force field. This comparison was made through computation of a correlation coefficient between pairs of trajectories (Appendix I). We found that the average correlation between a trajectory in the null field and one in the force field increased with the amount of practice movements performed by the subject in the force field. The computed correlation coefficient for trajectories performed by all subjects are shown in Fig. 11. Remarkably, all the subjects displayed a strictly monotonic evolution of the correlation coefficient.

Our subjects did not seem to be aware of the process of adaptation and of the change in their performance. The only subjective indication that some adaptive change had occurred

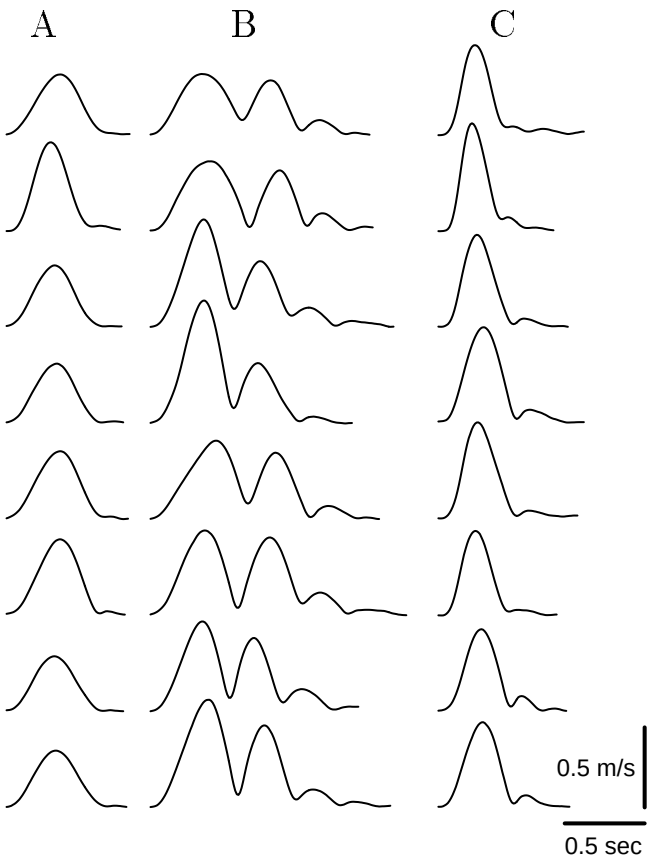

Figure 10: Tangential hand velocities before and after adaptation to the force field shown in Fig. 3. Traces are, from top to bottom, for targets at $0°, 45°, \ldots, 315°$. **A**: Hand velocities in a null field before exposure to the force field (corresponding to position traces in Fig. 6). **B**: Hand velocities upon initial exposure to the force field (corresponding to position traces in Fig. 7). Hand velocities after 1000 reaching movements in the field (corresponding to position traces in Fig. 9D).

was given by a reduction in the sense of effort associated with the task: during the first batch of 250 movements within the force field, some subjects reported an intense sense of effort. Paradoxically, this sense of effort diminished drastically after about 500 movements. At the end of the training period many commented that they were "not feeling" the field anymore.

3.3 After–effects

One way—although by no means the only way—for the subjects to recover the initial motor performance (what we have called the desired trajectory) after the exposure to the test field was by developing an *internal model* of this field. This internal model is the term $\hat{\mathcal{E}}$ in the expression of our model controller (Eq. 9). Indeed, if after the development of an internal model the test field is removed, then one expects to see a change in the resulting trajectory. This change is called an "after–effect" of the adaptation.

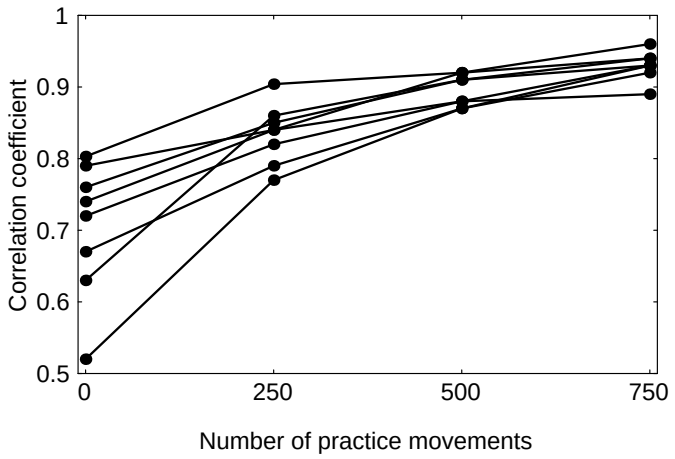

Figure 11: The average correlation coefficient for movements in a test force field as compared to movements in a null field, as a function of practice trials in the force field. Each line represents a subject.

We simulated the after–effect by setting $\hat{\mathcal{E}} = B\dot{x}^*(t)$ in our controller model (Eq. 9) and $E = 0$ in our dynamics model (Eq. 4). This simulation corresponds to the assumption that subjects developed an approximation of the force field and that this approximation led to after–effects as the null field was presented. Again, the commanded joint–trajectories corresponded to straight–line, minimum jerk movements of the hand towards the 8 targets. The results of this simulation are shown in Fig. 12. Qualitatively, one can see that the after effects are "opposite" to the initial perturbations induced by the field and shown in Fig. 8. In particular, (1) the hooks are oriented in opposite directions and (2) the metrics of the movements are reversed: long movements in Fig. 8 correspond to short movements in Fig. 12 and vice versa. These two features can be regarded as a strong property, almost like a "signature", of an internal model of the imposed force field.

Experimentally, we tested the hypothesis that adaptation in the subjects involved development of an internal model by removing the force field at the onset of movement and recording the after–effect. We found that the magnitude of the after–effects grew with the length of exposure to the force field: Figure 13 illustrates the temporal progression of after–effects, as measured under conditions of no visual feedback and in the null field, during the training period. The size of the after–effect, as indicated by the deviation of the hand trajectory from a straight line, grew with practice in the force field. By the final target set (Fig. 13D), the hand trajectory in the null field was significantly skewed. Remarkably, the observed after effects at the end of the adaptation period had the same qualitative features as those predicted by our simulation of an internal model within the null field (Fig. 12). In particular, by comparing Fig. 9 with Figure 13D, one can see that (1) all the hooks had reversed directions and (2) the metrics of movement has changed as in the simulation.

This finding is consistent with the hypothesis that subjects adapted to the force field by

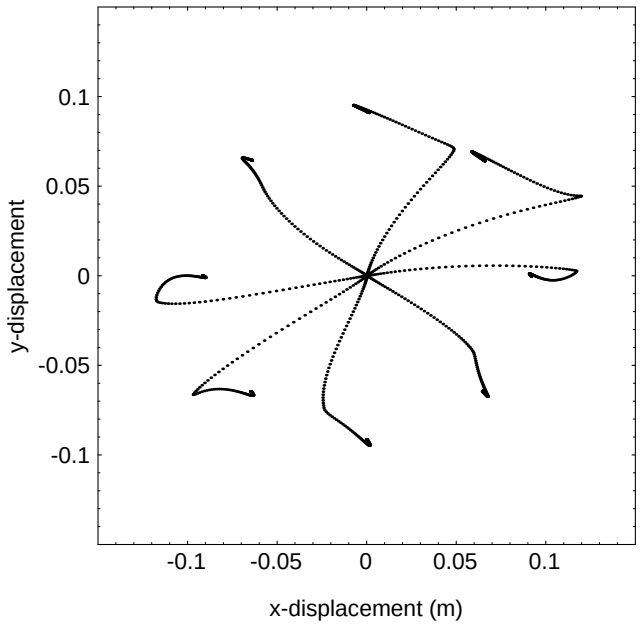

Figure 12: Simulated after–effect trajectories: Hand trajectories for the skeletal dynamics of Eq. (11) in a null force field with the controller of Eq. (12), assuming that the controller had formed an internal model of the force field shown in Fig. 3.

creating an internal model that approximated the dynamics of the environment. In addition, the data shown in Fig. 13 indicate that most of the development of this internal model took place early in the training period. From this observation one would expect that performances of the subjects in the force field should have shown most of its improvement rather early in the training. This is in agreement with the correlation curves shown in Fig. 11: in general, for all subjects the correlation coefficient increased most rapidly at the early stage of exposure to the field, indicating that the subjects had composed a fairly accurate internal model of the imposed force field by the midpoint of the training session.

3.4 Transferred after–effects

Our results indicate that adaptation occurred through development of an internal model of the applied field. What is the structure of this model and how is it represented in the nervous system? A priori, there are several hypotheses. This internal model can be regarded as a mapping between the state of the arm (position and velocity) and the corresponding force exerted by the environment. In an artificial system, one may implement such a mapping as a look–up table by storing away in memory the forces encountered at each state visited during the period of adaptation (cf. Raibert 1978, Miller 1978, Atkeson and Reinkensmeyer 1989). This type of local mapping has also been proposed in biological models, such as the one formulated by Albus (1975) for the cerebellum. In psychophysics, this kind of model is called a "specific exemplar model" and has been used to explain the process of motor learning (cf. Chamberlin and Magill 1992). Of course, if the internal model were a look–up table, adaptation would occur

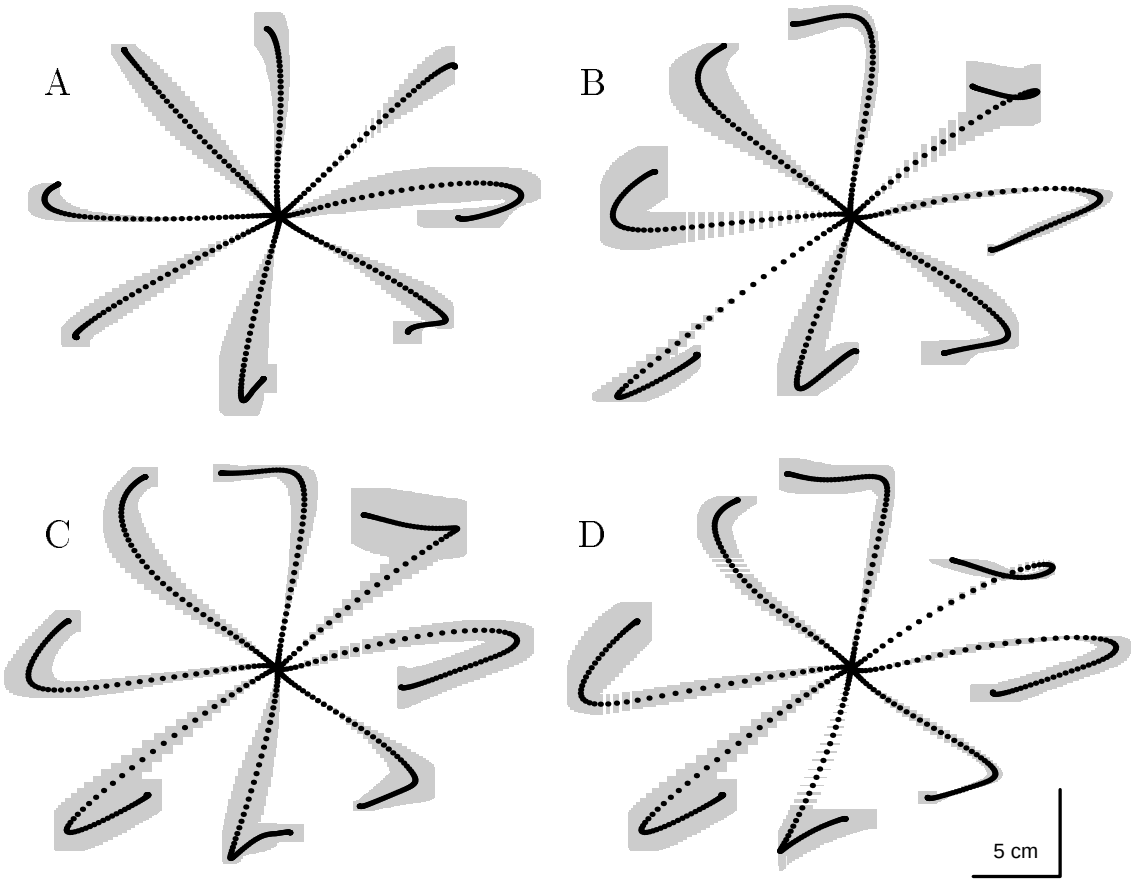

Figure 13: After effects of adaptation to the force field shown in Fig. 3 at the right workspace. Shown are average $\pm$ standard-deviation of the hand trajectories while moving in a null field during the training period for the first, second, third, and final 250 targets (A, B, C, and D, respectively). All trajectories shown are under no–visual feedback condition.

only at (or in the neighborhood of) the visited states. As a consequence, no after effect should be detectable if, after the adaptation, the null field was presented at some location outside the neighborhood visited during the training period.

To test this hypothesis regarding the representation of the internal model as a local association between states and forces, we asked our subjects to make reaching movements in the null field at the left workspace before and after having been exposed to the test field at the right workspace (workspaces are shown in Fig. 2). Fig. 14A shows a set of trajectories in the null field at the left workspace. These trajectories were obtained before the subject practiced movements in the force field at the right workspace. Figure 14B shows the average trajectories obtained from the same subject, in the same left workspace and with the same null field, but after the subject had adapted to the field in the right workspace. Clearly, there were substantial after–effects in the left workspace resulting from adaptation in the right workspace. This finding is not compatible with the hypothesis that subjects developed an internal model by building a look–up table, that is, a local association between visited states and experienced forces. On the contrary, the internal model appeared to extend and "generalize" quite broadly outside the

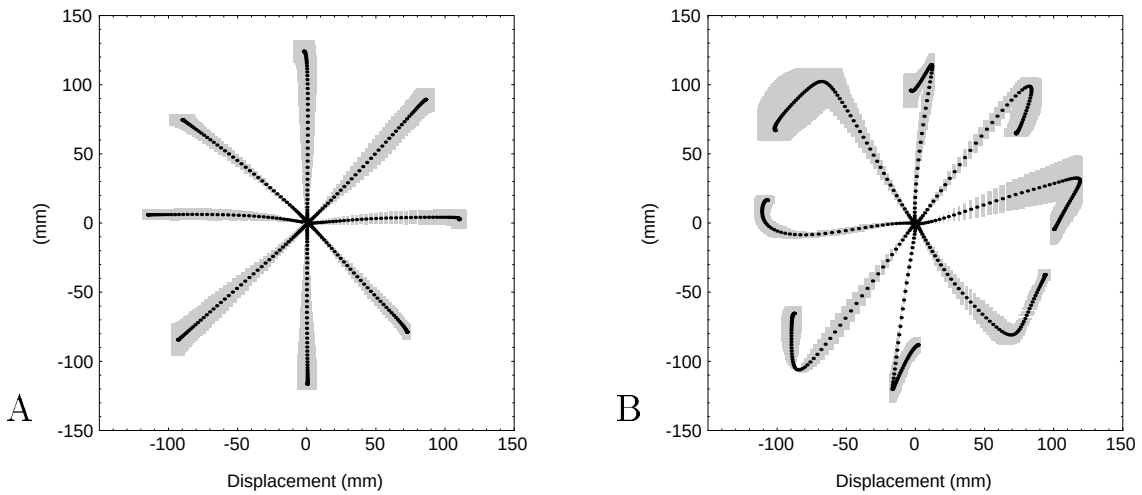

Figure 14: Transferred after–effects: Average ± standard-deviation of hand trajectories while moving in a null field at the left workspace. **A:** Before the subject practiced movements in the field of Fig. 3 at the right workspace. **B:** After the subject practiced movements in the field of Fig. 3 at the right workspace.

portion of workspace explored during the period of adaptation. This pattern of generalization, as evidenced by the transferred after–effects, was similar in all subjects, regardless of whether they had trained at the right workspace in an endpoint translation invariant field (Eq. 1) or a joint translation invariant field (Eq. 3).

Once we had established that the internal model was not merely a *local* association between states and forces, a question that remained was how the internal model extrapolated outside the region where the subject had trained. We consider two broad classes of generalizations. In one class, the generalization is the outcome of an inference about the mechanical properties of the environment. For example, if we are stirring a can of paint, from physics we know that we should experience the same forces on our hand (for a given hand trajectory) regardless of the location of the paint can in the workspace of our arm. In this sense, we would expect the viscous field representing the mechanical properties of the paint to be translation invariant in endpoint coordinates. This expectation would be reflected in the geometric structure of our internal model: the internal model would be a map between motion and forces in extrinsic coordinates. Consistent with the properties of the environment, it would predict identical forces acting on the hand when movements are done in the novel region of the workspace (as compared to movements in the region where we trained). As a consequence, the adaptation to a velocity–dependent field in the right workspace would also imply the adaptation to the same force field in the left workspace. In order to achieve this type of generalization, it is necessary to postulate existence of computations that transform predicted end–point forces (output of the internal model) into muscle torques.

Alternatively, adaptation may be through composition of an internal model that does not require further coordinate transformations; it simply represents the environment in terms of a

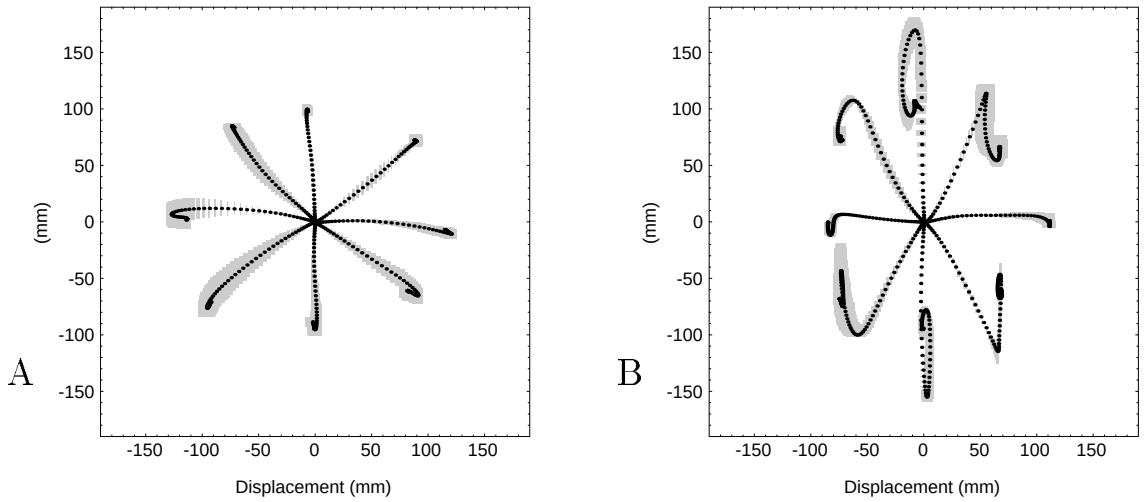

Figure 15: Average ± standard-deviation of hand trajectories during initial exposure to a field at the left workspace immediately after the subject practiced movements in the field shown in Fig. 3 at the right workspace. **A:** Performance at the left workspace in the field of Fig. 4. **B:** Performance at the left workspace in the field of Fig. 3.

map between motion and forces in the coordinate system of its sensors and actuators. This model would be implemented by a controller that, during execution of the task, effectively changes the dynamical behavior of the muscles (in this case, their apparent viscosity) to approximate and compensate for the force field during adaptation. Indeed, these changes in the apparent muscle behavior are bound to have a geometrically distinct effect beyond the region in which the subject was trained. According to this scenario, the internal model is translation invariant in an intrinsic coordinate system, and generalization is a side–effect of biomechanics.

Our experimental results clearly favor this second scenario where the forces in the environment are generalized in terms of an intrinsic coordinate system, i.e., in terms of torques on joints. The after–effects observed at the left workspace (Fig. 14B) were significantly different than those observed at the right (Fig. 13D). For example, compare movements to targets at 45°, 135°, 225° and 315° in each figure. These differences suggested that based on the internal model formed after practice in the right workspace, the subjects expected to interact with very different forces at the left workspace. We tested this hypotheses directly by having subjects which practiced in the field shown in Fig. 3A at the right workspace, make movements without visual feedback in the field shown in Fig. 4A at the left workspace. The results are shown for a typical subject in Fig. 15A: This subject belonged to Group 1, i.e., trained at the right workspace on the endpoint translation invariant field described by Eq. (1). Although forces in Figs. 3A and 4A are nearly orthogonal, the subject performed near perfectly ($\rho = 0.91$) at the left workspace in the field of Fig. 4A. The same subject's performance in the left workspace was poor ($\rho = 0.62$) in the field of Fig. 3A (shown in Fig. 15B). This indicated that the subject generalized the force field in terms of an intrinsic coordinate system.

The performance of all subjects in the two force fields at the left workspace was quantified

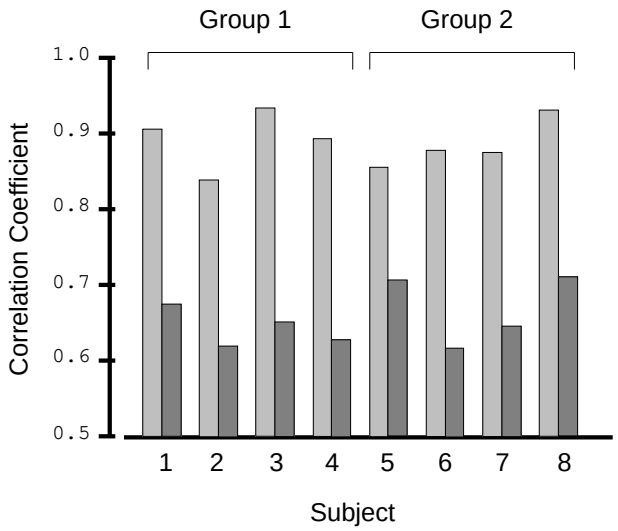

Figure 16: Summary of performance in the left workspace after training at the right workspace. Subjects in Group 1 trained on the field given by Eq. (1), while subjects in Group 2 trained in field given by Eq. (3). The two fields were essentially identical in the right workspace but orthogonal at the left. Shown are average correlation coefficients for movements in the left workspace in a force field as compared to movements in a null field for the same subject. Light gray bars are for movements in field given by Eq. (3) while dark gray bars are for movements in field given by Eq. (1). Performance was significantly better in both Groups when the force field was transferred to the left workspace in terms of joint torques rather than end–point forces.

by computing the correlation coefficient between the trajectories in each force field and the trajectory in the null field. These coefficients are shown in Fig. 16. The results consistently indicated that subjects retained the kinematic features of the adapted behavior when the environment was translated to the novel region of the workspace in joint coordinates, and not when this translation was in endpoint coordinates. This rejected the hypothesis that the internal model attributed a hand–based invariance to the environmental field.

4 Discussion

We used the paradigm of a programmable mechanical environment in order to study how the motor control system adapts to a change in the dynamics of a well rehearsed task. The task which we considered was a reaching movement where the hand interacted with a force field produced by a robot manipulandum. We chose a force field which significantly changed the dynamics of the task, resulting in a large change in the trajectory that the hand took in making a reaching movement (as compared to moving in a null field). The objective was to observe how the subjects responded to this change in the system dynamics.

We tested the hypothesis that in programming a reaching movement, the CNS initially specifies a desired trajectory of the hand and then uses an internal model of the limb's dynamics

to produce torques appropriate for moving the hand along this desired trajectory. When the limb's dynamics were changed (by imposing a force field on the hand), the internal model was no longer accurate, resulting in the hand moving along a trajectory that deviated from the desired behavior. This error led to gradual updating of the internal model so that it eventually approximated the new dynamics of the limb. We found evidence for the existence of a desired trajectory and that the motor controller achieved this desired performance via an explicit composition of an internal model.

4.1 Evidence for a desired trajectory

The task of moving the hand to a target position is ill-posed in the sense that the subject may choose from an infinite set of trajectories to achieve the goal. Yet, for two-dimensional movements with moderate accuracy requirements (such as our task), it has been demonstrated that subjects tend to move their hand smoothly and along a straight line (Morasso 1981, Soechting and Lacquaniti 1981, Flash and Hogan 1985). Reaching movements are characterized by fairly constant duration, whatever their direction or extent, and by a bell-shaped curve of the tangential hand velocity versus time (Morasso 1981). Here we confirmed this observation as subjects performed the task in a null field (Figs. 6 and 10A). In addition, we found that when the dynamics of the task were changed by imposing a force field onto the hand, the result was hand trajectories which deviated significantly from this smooth, straight line path, as is shown in the position traces of Fig. 7, and velocity traces of Fig. 10B. Nevertheless, through practice, the subjects' hand trajectories converged to the trajectory observed during null field conditions (Figs. 9 and 11). This convergence was gradual but monotonic in all subjects, consistent with an adaptive process whose goal was to compensate for the forces imposed by the field and return the hand's trajectory to that produced before the perturbation. This finding suggests that the kinematics observed in reaching movements is not merely a consequence of arm dynamics but reflects the presence of a plan, i.e., a desired trajectory.

4.2 Properties of the desired trajectory

The desired performance of a controlled system is usually established by a criterion, or optimization principle, expressed in a particular coordinate system (e.g., the coordinate system of the task, cf. Flash and Hogan 1985, Jordan and Rumelhart 1992, Jordan 1993). For skilled movements of the arm, this criterion appears to be one of smoothness. Specifically, in the context of reaching movements in the horizontal plane, Flash and Hogan (1985) have noted that the hand's trajectory is well described by a function which maximizes a measure of smoothness. In a similar work, Stein et al. (1988) have shown that in the single joint case, the optimal fit to joint velocity is a Gaussian function, which is also consistent with an optimization of smoothness (Poggio and Girosi 1990). Even in more complicated tasks such as reaching around obstacles, there is evidence that with practice, the trajectory of the hand becomes progressively

smoother (Abend et al. 1982, Schneider and Zernicke 1989). Therefore, this optimization of smoothness in terms of the trajectory of the hand serves as a possible computational principle that the CNS might be using to describe the desired performance during a reaching movement.

A characteristic of the above hypothesis is that the desired behavior of the arm is achieved via a purely kinematic principle, i.e., smoothness of the change in the position of the hand. This is appealing as it would imply a separation between the planning and the execution stages of the motor task: as long as the task is to move the hand to a target position, the desired trajectory remains a smooth, straight line path (in task coordinates), regardless of whether a force field is present. As Bernstein (1967) noted, this kind of separation of planning from execution is inherent to a hierarchical structure where a change in the dynamics of the controlled system does not affect the definition of the desired behavior.

Alternatively, one can postulate other computational principles which the CNS might be using to define a desired trajectory where the stage of planning is highly dependent on the stage of execution. For example, consider that the CNS could specify a desired trajectory for the hand such that the target is reached the most "effortlessly", where an effort is defined as a measure of energy, based on the physical cost of the movement (Nelson 1983), or based on changes in the forces or torques on the joints (Uno et al. 1989). In fact, it has been shown that the smoothness and straight line properties of the hand trajectory may be a by-product of a minimum torque-change criterion (Uno et al. 1989). However, in contrast to the previous approach, based on this scenario the desired trajectory would change as a function of the dynamics of the task, closely linking the process of planning to that of execution.

The field that we imposed on the hand during a reaching movement changed the dynamics of the arm drastically (cf. Fig. 7). Nevertheless, through practice, the subjects' hand trajectories converged upon the trajectory observed during unperturbed conditions. The only major difference was an increase in peak velocity (on average, an increase of 19% with respect to movements in a null field, cf. Fig. 10C), a phenomenon which has been linked to repetition of a motor task by other investigators (cf. Kerr 1992). This observed convergence to the unperturbed trajectory argues for an explicit description of a desired trajectory whose kinematics are essentially independent of the dynamics of the task, in line with the notion of a separation of the planning from an execution stage.

Recent results from Flash and Gurevich (1992) have provided evidence suggesting that there is an invariant kinematic plan for reaching when a static load is placed on the hand. Similarly, Lacquaniti et al. (1982) found that subjects who were asked to move a 2.5 kg weight did so, after some practice trials, along essentially the same trajectory as when moving without the weight. Our work has shown that even when the change in the dynamics of the limb is severe, the response is a convergence to the trajectory observed before the change, albeit this convergence may take place over a fairly long practice period (500 to 1000 movements, as shown in Fig. 11). This is similar to the conclusion reached for single degree of freedom movements by Ruitenbeek (1984), who found that when a subject interacted with a manipulandum with

variable dynamics, practice led to a trajectory that was invariant with respect to the dynamics of the manipulandum. These results are not compatible with the idea that the process of planning is mainly influenced by the dynamics of the task (Uno et al. 1989), as one would expect different planned trajectories for different environments since a change in the environment causes a change in the system's dynamics. Indeed, invariance of the plan with respect to the dynamics suggests that there may be specific elements in the motor control hierarchy which are concerned with the description of the task in terms of pure kinematics.

4.3 Adaptation through composition of an internal model

Convergence of the hand trajectories while interacting with the novel force field is an indication of the adaptation of the motor controller. We hypothesized that this adaptation was via composition of an internal model of the imposed force field. In this scenario, the internal model is a mechanism by which the nervous system predicts the forces that would be acting on the hand as it performs the task.

The force field which was imposed on the hand had the property of being dependent on the velocity of the hand, resulting in a situation where the subject did not know whether the field was "on" or "off" until the movement was actually initiated. However, during the training period, in 91% of the movements the field was on, presumably facilitating formation of a model of the force field which the CNS might use as a part of a control system to move the hand along the desired trajectory (for the remaining movements the field was off in order to measure any after–effects of adaptation). We suggested that this control system may be represented as the sum of three components: an internal model describing the dynamics of the skeletal system of the arm when moving in a null field, an internal model describing the dynamics of the force field imposed on the hand, and a viscoelastic or feedback system intended to stabilize the arm about the desired trajectory in case of errors in these models.

Initially, the subject had not formed a model of the force field, resulting in a discrepancy between the expected dynamics of the arm and the dynamics actually present. This "model error" led to trajectories (Fig. 7) which were significantly different than desired. Indeed, we found excellent correspondence between trajectories produced by the simulation (Fig. 8) and those observed in the movements of the subjects (Fig. 7). In particular, we observed that the responses to the sudden presentation of the field were characterized by a sharply curved trajectory that we described as a "hook". A possible interpretation for this hook would be that the hand starts the movement along a wrong direction and that the resulting error is corrected by a second movement. However, there is a simpler interpretation which does not make appeal to an explicit correction process. According to this, the corrective movement is a by-product of the interaction between the mechanical properties of the arm (stiffness and viscosity in Eq. 9) and the force field imposed on the hand. Presence of the hook as well as the initial error in movement direction are systematically predicted by our simulations which follow this later line

of reasoning. We favor this hypothesis only because of its computational simplicity as compared to the hypotheses which requires an explicit correction process.

If the adaptive process was via composition of an internal model of the imposed force field, then we argued that by removing the field, once again there should be a discrepancy between expected and actual dynamics of the system. Our simulations suggested that there would be *after–effects* of adaptation (Fig. 12). We found that when the field was unexpectedly removed, the subjects produced trajectories similar to those predicted by the simulation. The "magnitude" of the observed after–effects increased gradually with the practice period (Fig. 13). This progressive buildup of after–effects was further evidence that the CNS improved performance via an explicit composition of an internal model.

Of course, one may envision a system whose performance in response to a perturbation improves not because of an internal model, but because of an increase in the stiffness of the system about the desired trajectory. This alternative strategy may be achieved by an increase in the coactivation of the muscles. As a consequence, movements would become more insensitive to changes in the external forces. It is easy to show that modest increases in arm stiffness (about 3 folds with respect to the values measured in posture) leads to almost perfect performance in the force field. However, if this strategy is chosen as the mode of adaptation, then exposure to a force field would not cause an after–effect in a null field. The fact that practice does cause progressively larger after–effects (Fig. 13) is strong evidence against the hypothesis that the convergence of trajectories is due to a mechanism such as global coactivation of muscles. This is in agreement with measurements of van Emmerik (1991) and Milner and Cloutier (1993) where it has been shown that during learning of a novel movement the stiffness of the limb generally *decreases* with practice. In particular, Milner and Cloutier (1993) have shown that adaptation to an unstable viscous load is accompanied by a reduction in co-activation of antagonist muscles. This, along with the gradual increase in the after-effects favors the idea that improvement in performance was due to formation of an internal model of the imposed field rather than an increase in stiffness of the arm.

4.4 Transfer properties of the internal model

The description of a biological learning task can often be represented as approximation of a sensorimotor map. In the current experiment, the information contained in the internal model can be thought of as a map whose input is the state of the arm and whose output is a force. This output is the force, predicted by the internal model, which should be imposed by the environment as the arm passes through a given state. Therefore, the internal model is a sensorimotor map which approximates the force field imposed by the mechanical environment. The task for the subject is to learn to perform this approximation from a set of examples, where the examples are provided as the subject makes movements in the force field. How does the nervous system compose this sensorimotor map which represents the internal model?

From a computational point of view, a sensorimotor map may be implemented by a distributed technique inspired by the architecture of the nervous system: in this approach, the mapping is formed via interaction of a set of non–linear computational elements which represent neuron–like structures (cf. Barto 1989, Poggio 1990). For example, for the task of motor learning, combinations of non–linear basis functions have been used to implement an internal model which represents the inverse dynamics of a multi-joint limb (Raibert and Wimberly 1984, Kawato 1989, Jordan 1990, Shadmehr 1990, Kawato and Gomi 1992, Jordan 1993), mapping from states of the limb to an output force (e.g., Eq. 6). These results have provided an algorithm by which an internal model may be constructed. However, little has been learned regarding the properties of the computational elements with which the nervous system might be performing this adaptive process.

Consider that a property of the computational elements (e.g., basis functions or "neurons" in a neural network) used in learning such a sensorimotor map is their spatial bandwidth, i.e., the size of their support or "receptive field" in the input space.[1] This receptive field would indicate the region of the sensory space to which the element responds to. Because computation emerges from the superposition of the receptive fields of the activated elements, the size and location of the receptive fields greatly influence how the learning system interpolates between states which it has visited during training, and whether it can generalize to regions beyond the boundary of its training data (Poggio and Girosi 1990). Simply said, during the learning of the task, only the "weights" of those elements which are activated by the input are changed, and if these elements respond to only a narrow region of the sensory space, then the system can not generalize to a region outside the training data. In fact, research in visual perception has used the notion of generalization to make an inference regarding the receptive fields of the computational elements used by the visual system to learn a map: in a hyperacuity task, Poggio et al. (1992) have shown that if the computational elements have narrow receptive fields similar to those found in components of early vision, a subject should not be able to generalize to tasks which are slightly different than those on which the subject had been trained on—a prediction which agrees with results of experiments (Poggio et al. 1992). The implication is that for some visual recognition tasks, the nervous system learns a map by encoding information through the "low–level" elements which have fairly narrow receptive fields (akin to cells in a look–up table), and that this property of the computational elements leads to the inability of the composed map to generalize beyond the training region.

In our motor learning task, from the measured after–effects at the novel region we can state that the internal model generalized to well beyond the training region, leading to the suggestion that the elements with which the nervous system formed a model of the environmental forces had wide receptive fields. In other words, these elements produced a significant response for a region of the workspace that was outside the neighborhood where training data were provided.

[1] The support is that region of the function's domain where the output value is different from zero.

This property of the adaptive controller is inconsistent with the approach where motor learning takes place via construction of a look–up table in which local association is made between visited states (address of the memory cells in the table) and experienced forces (contents of the cells). On the contrary, adaptation is via computational elements which give the property of generalization to the internal model.

The after–effects at the left workspace suggest that the internal model generalized the environmental forces to a specific pattern. Interestingly, from the trajectory of after–effects (Fig. 14B), it was apparent that the expected force field at the novel region of the workspace was very different than the one that the subject had been trained on. We hypothesized that this difference could be accounted for if the field was generalized not in terms of forces on the hand, but in terms of torques on the joints. The idea was that perhaps the relative position of the computational elements in the motor control hierarchy dictated the coordinate system in which information about the environment was generalized: if these elements resided near the plan stage of the task, where a desired hand trajectory is specified, then they might encode the environmental dynamics as a mapping between the state of the arm and imposed forces in an extrinsic frame of reference. Assuming that these elements broadly encoded the input space, then local adaptation might produce an internal model which generalized to similar endpoint forces for similar endpoint trajectories. Alternatively, the computational elements might reside at a lower stage, perhaps near the effectors, where information is received in a coordinate system defined by the afferents and the muscles. Here the internal model would be a mapping between observed states of the arm and the imposed forces in an intrinsic frame of reference. As opposed to the high–level model, local adaptation here might produce a map which generalizes to similar joint torques for similar joint trajectories.

We tested the merits of these alternatives by a direct experiment. After practicing in a field at the right workspace, the subjects were asked to make movements at the left workspace. The field presented at the novel region (left workspace) was one of two kinds. In some trials, this field was a translation of the training field in endpoint coordinates, while in the other trials the field presented was a translation of the training field in joint coordinates. We found that the performance of the subjects was near optimum when the field was translated in joint coordinates (Figs. 15A and 16). This finding is in sharp contrast with the hypothesis that subjects adapted to the imposed field by building a model in endpoint coordinates. On the contrary, our finding suggests that the subjects represented the imposed force field as a map between motion and forces in the intrinsic coordinate system of the afferents and actuators.

Candidates for these low–level elements in the motor learning task are muscles and their associated spinal (Bizzi et al. 1991) and supra-spinal (Berthier et al. 1993) neural control pathways. For example, one of us (Mussa-Ivaldi 1992) has suggested that the behavior of spinal circuits may be categorized as computational elements in an approximation task. This idea is based on the observations of Giszter et al. (1993) where the input–output response of the neural circuits and the associated muscles in a frog's spinal cord have been, to some extent, quanti-

fied: each circuit is a collection of interneurons connected to a group of motor units. When a circuit is activated through microstimulation, the muscles generate a time–varying force. This force depends on the configuration of the limb and may be represented as a force field, e.g., an endpoint force as a function of the position of the tip of the limb. Therefore, computationally the behavior of the low–level elements in the motor control hierarchy is to produce an output force as a function of the input activation to the spinal neural circuitry and the position of the limb and time (Mussa-Ivaldi et al. 1990).

In a general framework, it seems more plausible to assume that the pattern of forces generated by such a spinal controller depends upon velocity of the limb as well as its position. The resulting time–varying force field is essentially a wave expressing the input–output behavior of a *motor computational element* within the central nervous system. In theory, a collection of these computational elements can be used in a motor learning task: A finding of the spinal microstimulation experiments (Bizzi et al. 1991, Giszter et al. 1993) has been that the output of the motor computational elements add when two are activated. Simultaneous stimulation of two separate sites resulted in the summation of the fields obtained from the separate stimulation of each site. Based on this property of superposition, a simple framework for motor learning in terms of these computational elements can be constructed (in relation to other theories in motor learning, each computational element can be thought of as a primitive movement, or motor schema, cf. Arbib 1985). Indeed, these low–level computational elements appear as reasonable candidates for the task of forming the sensorimotor map representing the internal model.

In conclusion, during adaptation to a force field which significantly changes the dynamics of a reaching movement, the CNS forms an internal model of the added dynamics. This internal model has the power to generalize well beyond the training region. The geometric property of this generalization is consistent with a representation of information in an intrinsic rather than extrinsic frame of reference. This choice of the coordinate system for the internal model suggests that the planning and control of a reaching movement are undertaken by fundamentally different computational elements in the nervous system: while the planned trajectory for the arm is in an extrinsic frame of reference, the model for the dynamics of the task (i.e., the internal model) is in an intrinsic frame. What results is a scenario where learning a motor task, say hitting a golf ball, entails both formation of an appropriate kinematic plan, i.e., golf club trajectory, and composition of a model of the task's dynamics so that the plan may be executed, i.e., forming an internal model of the club's dynamics. Here we have reported on some of the properties of the computational elements with which the nervous system forms the internal model for a task's dynamics. It remains to be seen whether computational elements which are involved in learning kinematics of a task produce a model which has a different geometric property than that which results when learning dynamics. Perhaps elements involved in learning kinematics and dynamics can eventually form a kind of alphabet for the language of movement.

Appendix I: Correlation of two trajectories

In order to compare hand trajectories, a technique was developed which measured the correlation between two sampled vector fields: we represented each trajectory as a time series of velocity vectors ($\dot{x}$ sampled at 10 msec intervals) and then compared the two resulting vector fields through a correlation measure. The same technique was also used to compare force fields. This technique was based on the notion of inner product of two sampled vector fields (Gandolfo and Mussa-Ivaldi 1993).

Empirically, a time series of vectors, as well as a vector field, may be regarded as a finite ordered set of vectors, sampled at subsequent instance of time, or in a given arrangement of spatial locations. A finite ordered set of vectors, U, is a mapping that assigns to each element, i, of the index set, $(1, \ldots, n) \in N$, a vector u_i. Then the expected value of U, denoted by $\varepsilon(U)$, is a mapping from the same index set to the set of vectors $\{v_i\}$, where

$$v_i = v = \frac{1}{n} \sum_{j=1}^{n} u_j$$

According to this definition, the expected value of U is a constant set ($v_i = v_j, \forall i, j$). It follows that: $\varepsilon(\varepsilon(U)) = \varepsilon(U)$. Now consider the task of comparing two sets U and Y, where $Y = (y_1, y_2, \ldots, y_n)$. Let us define the inner product of U and Y as the scalar:

$$< U, Y > = \sum_{i=1}^{n} u_i \cdot y_i$$

where the symbol $\cdot$ on the right side indicates the dot product operation between two vectors. We define the expected value of this inner product as:

$$\varepsilon(< U, Y >) = \frac{1}{n} < U, Y >$$

Then, we may use the above expressions for defining the co-variance of two vectorial sets:

$$\begin{aligned} \mathrm{Cov}(U, Y) &= \varepsilon(< U - \varepsilon(U), Y - \varepsilon(Y) >) \\ &= \varepsilon(< U, Y >) - < \varepsilon(U), \varepsilon(Y) > \end{aligned}$$

Furthermore, the correlation coefficient between two sets, $\rho(U, Y)$, is given by the ratio of the co-variance of the time series and the product of their standard deviations:

$$\rho(U, Y) = \frac{\mathrm{Cov}(U, Y)}{\sigma(U)\, \sigma(Y)}$$

where standard deviation of an ordered set of vectors is the scalar:

$$\sigma(U) = \varepsilon(\|U - \varepsilon(U)\|)^{1/2}$$

and $\|U\|$ is defined as: $\|U\| = (< U, U >)^{1/2}$. It follows that $-1 \leq \rho(U, Y) \leq +1$.

5 References

Abend W, Bizzi E, Morasso (1982) Human arm trajectory formation. Brain 105:331–48.

Albus JS (1975) A new approach to manipulator control: the cerebellar model articulation controller (CMAC) J Dynamic Syst Measure Control 97:228–33.

Andersen RA, Essick GK, Siegel RM (1985) Encoding of spatial location by posterior parietal neurons, Science 230:456–458.

Arbib MA (1976) Program synthesis and sensorimotor coordination, Brain Theory Newsletter, 2(2):31–33.

Arbib MA (1985) Schemas for the temporal organization of behavior. Human Neurobiology 4:63–72.

Atkeson CG, Reinkensmeyer DJ (1989) Using associative content–addressable memories to control robots. Proc. IEEE Conf Robotics & Automation, pp. 1859–1864.

Atkeson CG (1989) Learning arm kinematics and dynamics. Annual Review of Neuroscience 12:157–183.

Barto AG (1989) Connectionist learning for control: an overview. COINS Tech Rep 89–89, Univ Massachusetts, Amherst MA.

Bernstein N (1967) The coordination and regulation of movements. London: Pergamon Press.

Bizzi E, Accornero N, Chapple W, Hogan N (1984) Posture control and trajectory formation during arm movement. J Neurosci 4:2738–44.

Bizzi E, Mussa-Ivaldi FA, Giszter SF (1991) Computations underlying the execution of movement: a novel biological perspective. Science 253:287–291.

Brooks VB (1986) The neural basis of motor control. Oxford University Press, Clarendon UK.

Chamberlin CJ, Magill RA (1992) A note on schema and exemplar approaches to motor skill representation in memory. J Motor Behavior 24:221–224.

Cunningham HA (1989) Aiming error under transformed spatial mappings suggests a structure for visual-motor maps. J Exp Psychol 15:3:493–506.

Diffrient N, Tillery AR, Bardagjy JC (1978) Humanscale. Cambridge, MA:MIT Press.

van Emmerik REA (1992) Kinematic adaptations to perburbations as a function of practice in rhythmic drawing movements. J Motor Behavior 24:117–131.

Faye IC (1986) An impedance controlled manipulandum for human movement studies. S.M. Thesis, MIT Dept. of Mechanical Engineering, Cambridge, MA.

Flash T, Hogan N (1985) The coordination of arm movements: an experimentally confirmed mathematical model. J Neurosci 5:1688–1703.

Flash T, Gurevich I (1992) Arm movement and stiffness adaptation to external loads. Proc IEEE Engin Med Biol Conf, Orlando FL, 13:885–886.

Gandolfo F, Mussa-Ivaldi FA (1993) Vector summation of end–point impedance in kinematically redundant manipulators. Proc Int Robotics Systems, Yokahama, Japan, in press.

Gordon J, Ghilardi MF, Ghez C (1993) Accuracy of planar reaching movements: I. Independence of direction and extent variability. Exp Brain Res, in press.

Giszter SF, Mussa-Ivaldi FA, Bizzi E (1993) Convergene force fields organized in the frog's spinal cord. J Neurosci 13:467–491.

Grillner S (1981) Control of locomotion in bipeds, tetrapods and fish. In: Handbook of Physiology–the Nervous System II: Motor Control. (JM Brookhard and VB Mountcastle, eds.), American Physiological Society:Waverly.

Held R (1962) Adaptation to rearrangement and visual-spatial aftereffects. Psychologische Beitrage 6:439–450.

Held R, Freedman SJ (1963) Plasticity in human sensorimotor control. Science 142:455–462.

Held R, Schlank M (1959) Adaptation to disarranged eye-hand coordination in the distance-dimension. Am J Psychol 72:603–605.

von Helmholtz HV (1925) Treatise on physiological optics, Vol. 3, Optical Society of America, Rochester, New York.

Hogan N (1985) The mechanics of multi-joint posture and movement control. Biol Cybern 52:315-331.

Jacoby LJ, Brooks LR (1984) Nonanalytic cognition: Memory, perception, and concept learning. In: Psychology of Learning and Motivation (Bower GH, ed.), vol. 18, pp. 1–47, Academic Press, New York.

Jeannerod M (1988) The Neural and Behavioural Organization of Goal-Directed Movements, Clarendon, Oxford University Press, UK.

Jordan MI (1990) Motor learning and the degrees of freedom problem. In: Attention and Performance, XIII (Jeannerod M, ed), Erlbaum, Hillsdale, NJ, ch. 29.

Jordan MI (1993) Computational aspects of motor control and motor learning, In: Handbook of Perception and Action: Motor Skills (Heuer H, Keele S, eds), Academic: NY, pp. 87–146.

Jordan MI, Rumelhard DE (1992) Forward models: supervised learning with a distal teacher, Cognitive Science 16:307–354.

Kawato M (1989) Adaptation and learning in control of voluntary movement by the central nervous system. Advanced Robotics 3:229–249.

Kawato M, Gomi H (1992) A computational model of four regions of the cerebellum based on feedback–error learning. Biol Cybern 68:95–103.

Kelso JAS, Schoner GS (1988) Self–organization of coordinative movement patterns. Human Movement Science 7:27–46.

Kerr GK (1992) Visuomotor control in goal–directed movements. In: Approaches to the Study of Motor Control and Learning (JJ Summers, ed.), Elsevier, pp. 253–287.

Lackner JR, Dizio P (1992) Rapid adaptation of arm movement endpoint and trajectory to coriolis force perturbations, Soc Neurosci Abs, vol. 22.

Lacquaniti F, Soechting JF, Terzuolo CA (1982) Some factors pertinent to the organization and control of arm movements. Brain Res 252:393-397.

Milner TE, Cloutier C (1993) Compensation for mechanically unstable loading in voluntary writs movements. Exp Brain Res 94:522–532.

Morasso P (1981) Spatial control of arm movements. Exp Brain Res 42:223–227.

Mussa-Ivaldi FA (1992) From basis functions to basis fields: vector field approximation from sparse data. Biol Cybern 67:479–489.

Mussa-Ivaldi FA, Hogan N, Bizzi E (1985) Neural, mechanical, and geometric factors subserving arm posture in humans. J Neurosci 5:2732–2743.

Mussa-Ivaldi FA, Giszter SF, Bizzi E (1990) Motor-space coding in the central nervous system. Cold Spring Harbor Symp Quant Biol 55:827–835.

Mussa-Ivaldi FA, Giszter SF (1992) Vector field approximation: a computational paradigm for motor control and learning. Biol Cybern 67:491–500.

Nelson WL (1983) Physical principles for economies of skilled movements. Biol Cybern 46:135–147.

Poggio T (1990) A theory of how the brain might work. Cold Spring Harbor Symp Quant Biol 55:899–910.

Poggio T, Girosi F (1990) Networks for approximation and learning. Proc IEEE 78:1481-1497.

Poggio T, Fahle M, Edelman S (1992) Fast perceptual learning in visual hyperacuity. Science 256:1018–1021.

Raibert MH (1978) A model for sensorimotor control and learning. Biol Cybern 29:29–36.

Raibert MH, Wimberly FC (1984) Tabular control of balance in a dynamic legged system. IEEE Trans Systems, Man, Cybernetics SMC-14(2):334–339.

Ruitenbeek JC (1984) Invariants in loaded goal directed movements. Biol Cybern 51:11-20.

Saltzman E, Kelso JAS (1987) Skilled action: task dynamic approach. Psychol Rev 94:84–106.

Schneider K, Zernicke RF (1989) Jerk-cost modulations during the practice of rapid arm movements. Biol Cybern 60:221–230.

Schoner G, Zanone PG, Kelso JAS (1992) Learning as change of coordination dynamics: Theory and experiment. J Motor Behav 24:29–48.

Shadmehr R (1990) Learning virtual equilibrium point trajectories for control of a robot arm. Neural Computation 2:436–446.

Shadmehr R, Arbib MA (1992) A mathematical analysis of the force-stiffness characteristics of muscles in control of a single joint system. Biol Cybern 66:463–477.

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 13:43–62.

Slotine JJE (1985) The robust control of robot manipulators. Int J Robotics Res 4:49–64.

Slotine JJE, Li W (1991) Applied nonlinear control, Prentice Hall, Englewood Cliffs, New Jersey.

Soechting JF, Flanders M (1991) Deducing central algorithms of arm movement control from kinematics. In: Motor Control: concepts and issues (Humphrey DR, Freund HJ), Wiley, Chichester, pp. 293–306.

Soechting JF, Lacquaniti F (1981) Invariant characteristics of a pointing movement in man. J Neurosci 1:710–720.

Spong MW, Vidyasagar M (1989) Robot dynamics and control, Wiley:NY.

Stein RB, Cody FWJ, Capaday C (1988) The trajectory of human wrist movements. J Neurophysiol 59:1814–1830.

Thach WT, Goodkin HP, Keating JG (1992) The cerebellum and the adaptive coordination of movement. Ann Rev Neurosci 15:403–42.

Tsuji T, Goto K (1993) Estimation of human hand impedance during maintenance of posture, Trans Soc Instrument Control Engineers Japan, submitted (in Japanese).

Uno Y, Kawato M, Suzuki R (1989) Formation and control of optimal trajectory in human arm movement. Biol Cybern 61:89–101.

Wolpert DM, Ghahramani Z, Jordan MI (1993) On the role of extrinsic coordinates in arm trajectory planning: Evidence from an adaptation study. Computational Cognitive Science Tech Rep 9308, Massachusetts Institute of Technology.