

Motor learning through the combination of primitives

F. A. Mussa-Ivaldi¹ and E. Bizzi^{2*}

¹*Department of Physiology, Northwestern University Medical School, Chicago, IL, USA*

²*Massachusetts Institute of Technology, Cambridge, MA, USA*

In this paper we discuss a new perspective on how the central nervous system (CNS) represents and solves some of the most fundamental computational problems of motor control. In particular, we consider the task of transforming a planned limb movement into an adequate set of motor commands. To carry out this task the CNS must solve a complex inverse dynamic problem. This problem involves the transformation from a desired motion to the forces that are needed to drive the limb. The inverse dynamic problem is a hard computational challenge because of the need to coordinate multiple limb segments and because of the continuous changes in the mechanical properties of the limbs and of the environment with which they come in contact. A number of studies of motor learning have provided support for the idea that the CNS creates, updates and exploits internal representations of limb dynamics in order to deal with the complexity of inverse dynamics. Here we discuss how such internal representations are likely to be built by combining the modular primitives in the spinal cord as well as other building blocks found in higher brain structures. Experimental studies on spinalized frogs and rats have led to the conclusion that the premotor circuits within the spinal cord are organized into a set of discrete modules. Each module, when activated, induces a specific force field and the simultaneous activation of multiple modules leads to the vectorial combination of the corresponding fields. We regard these force fields as computational primitives that are used by the CNS for generating a rich grammar of motor behaviours.

Keywords: force field; dynamics; module; spinal cord; cortex; internal model

1. INTRODUCTION

When we learn to move our limbs and to act upon the environment, our brain becomes to all effects an expert in physics. While we are still very far away from understanding how this feat is accomplished, great strides have been made in the last few decades through the combined efforts of biologists, computer scientists, physicians, physicists, psychologists and engineers. In this paper we review some of this progress. In particular we focus on one issue: What are the building blocks or, to borrow from linguistics, the ‘modules’ that the brain may use for generating the competence in physics that is necessary to act and move? And what do we know of how and where these modules are engraved into the circuits of the central nervous system (CNS)?

To illustrate the complexities of ordinary motor behaviours, let us consider the task that the CNS must solve every time a planned gesture is transformed into an action. If the goal is to move the hand from an initial position to another point in space, then clearly there are a number of possible hand trajectories that could achieve this goal: the solution of this elementary motor problem is not unique. Even after the CNS has chosen a particular path for the hand, its implementation can be achieved through multiple combinations of joint motions at the shoulder, elbow and wrist—again the solution is not

unique. Finally, because there are many muscles around each joint, the net force generated by their activation can be produced by a variety of combinations of muscles.

Perhaps what makes the issue of sensorimotor transduction such a complex problem is the fact that we have not found a satisfactory way to incorporate motor learning into our thinking about motor planning. While everybody agrees that throughout our life we learn a great variety of movements and that the memory of these movements is stored more or less permanently in the cortical areas of the frontal and parietal lobes and the cerebellum, we do not know whether we use fragments of what has been learned when we produce a motor response to a new contingency (Toni *et al.* 1998; Shadmehr & Holcomb 1997). In this paper we adopt the point of view that motor learning consists of tuning the activity of a relatively small group of neurons and that these neurons constitute a ‘module’. Combining modules may be a mechanism for producing a vast repertoire of motor behaviours in a simple manner.

2. THE PROBLEM OF INVERSE DYNAMICS

According to the laws of Newtonian physics, if we want to impress a motion upon a stone with mass m , we must apply a force, F , that is directly proportional to the desired acceleration, a . This is the essence of Newton’s equation $F = ma$. A desired motion may be expressed as a sequence of positions, x , that we wish the stone to assume

* Author for correspondence (emilio@ai.mit.edu).

at subsequent instants of time, t . This sequence is called a ‘trajectory’ and is mathematically represented as a function, $x = x(t)$. To use Newton’s equation for deriving the needed time sequence of forces, we must calculate the first temporal derivative of the trajectory, the velocity and then the second temporal derivative, the acceleration. Finally, we obtain the desired force from this acceleration. The above calculation is an example of what in robotics is called an ‘inverse dynamic problem’. The ‘direct’ dynamic problem is that of computing the trajectory resulting from the application of a force, $F(t)$. The solution of this problem requires a complex computational process, called integration, through which the motion of the stone, that is the function $x(t)$, is derived from the known acceleration, $a(t) = F(t)/m$.

Direct problems are the bread and butter of physicists, who may be concerned, for example, with predicting the motion of a comet from the known pattern of gravitational forces. Unlike physicists, the brain deals most often with inverse problems: we routinely recognize objects and people from their visual images—an ‘inverse optical problem’—and we find out effortlessly how to distribute the forces exerted by several muscles to move our limb in the desired way—an inverse dynamic problem.

In the biological context, the inverse dynamic problem assumes a somewhat different form from the case of the moving stone. One of the central questions in motor control is how the CNS may form the motor commands that guide our limbs. One proposal is that the CNS solves an inverse dynamic problem (Hollerbach & Flash 1982). A system of second-order nonlinear differential equations is generally considered to be an adequate description of the passive limb dynamics. A compact representation for such a system is

$$D(q, \dot{q}, \ddot{q}) = \tau(t), \quad (1a)$$

where q , \dot{q} and \ddot{q} represent the limb configuration vector—for example, the vector of joint angles—and its first and second time derivatives. The term $\tau(t)$ is a vector of generalized forces, for example, joint torques, at time t . Conceptually, this expression is nothing else than Newton’s $F = ma$ applied to a multi-articular rigid body. In practice, the expression for D may have a few terms for a two-joint planar arm (see figure 4*b*) or it may take several pages for more realistic models of the arm’s multi-joint geometry. The inverse dynamic approach to the control of multi-joint limbs consists in solving explicitly for a torque trajectory, $\tau(t)$, given a desired trajectory of the limb, $q_D(t)$. This is done by plugging $q_D(t)$ on the left side of equation (1):

$$\tau(t) = D(q_D(t), \dot{q}_D(t), \ddot{q}_D(t)). \quad (1b)$$

Another significant computational challenge comes from the need to perform changes of representation, or, more technically, coordinate transformations, between the description of a task and the specification of the body motions. Tasks, such as ‘bring the hand to the glass of water on the table’, are often described most efficiently and parsimoniously with respect to fixed reference points in the environment. For example, the glass may be 10 cm to the left of a corner of the table. The hand may be 20 cm to the right of the same corner. So, the hand will

need to be displaced 30 cm along a straight line in the left direction. This is a very simple description of the needed movement. However, this description cannot be used to derive the joint torques, as specified by equation (1*b*). To this end, one must represent the trajectory of the hand in terms of the corresponding angular motions at each joint. This is a complex transformation known in robotics as ‘inverse kinematics’ (Brady *et al.* 1982).

Does the brain carry out similar inverse dynamic calculations for moving the arm on a desired trajectory? A clear-cut answer is still to come but several alternatives have emerged from studies in robotics and computational neuroscience.

3. SOLUTIONS BASED ON FEEDBACK

Many of the problems that the brain must face to control movements are indeed similar to those that engineers must solve to control robots. In spite of the great differences between the multi-joint vertebrate system and current robotic arms, the field of neuroscience, unquestionably, has derived benefits from the theories and procedures that have guided the construction of man-made limbs. For instance, from early on, neuroscientists have been influenced by the notion of feedback. Feedback control is a way to circumvent the computation of inverse dynamics. At each point in time, some sensory signal provides the information about the actual position of the limb. This position is compared with a desired position and the difference between the two is a measure of the error at any given time. Then, a force may be produced with amplitude approximately proportional to the amplitude of the error in the direction of the desired position. This method of control is appealing because of its great simplicity.

Multiple feedback mechanisms have been found in both vertebrates and invertebrates. These mechanisms were discovered by Sherrington at the beginning of the last century (Sherrington 1910). They have been shown to control the muscles’ level of contraction, the production of force and the position of joints. Sherrington observed that when a muscle is stretched the stretch is countered by an increase in muscle activation. This ‘stretch reflex’ is caused by sensory activity that originates in the muscle spindles—receptors embedded within the muscle fibres. Sherrington put forward the daring hypothesis that complex movements may be obtained by combining stretch reflexes as well as other reflexes in a continuous sequence or ‘chain’. In this way, movement patterns as complex as the locomotion cycle could be generated by local reflexes, without central supervision. A similar idea was later proposed by Merton (1972), who suggested that central commands via the gamma system might initiate the execution of movement, not by directly activating the muscles, but by triggering a stretch reflex through the modulation of muscle spindle activities. Both Sherrington and Merton’s hypotheses are attempts at explaining movements as automatic responses to sensory feedback, thus limiting the role and the arbitrariness of voluntary commands.

However, both Sherrington’s ideas on compounding of reflexes and Merton’s hypothesis have taken a new form following subsequent experiments which clearly

demonstrated the generation of movements in the absence of sensory activities. For example, Taub & Berman (1968) found that monkeys can execute various limb movements after the surgical section of the pathways that convey all sensory information from the limb to the nervous system. Shortly thereafter, Vallbo (1970) was able to record muscle spindle discharges in human subjects and to compare these discharges with the activation of the muscles, as revealed by electromyography (EMG). Vallbo's study showed that, in a voluntary movement, muscle activation does not lag but leads the spindle discharges, contrary to the predictions of Merton's hypothesis.

In addition to the experimental findings described above, the idea that biological movements may be carried out by feedback mechanism has been challenged based on consideration about limb stability and reflex delays. It takes more than 40 ms before a signal generated by the muscle spindles may reach the supraspinal motor centres and it takes 40–60 ms more before a motor command may be transformed into a measurable contraction of the muscles. These transmission delays may cause instability (Hogan *et al.* 1987). The effects of delays are even greater when the limb interacts with the environment. For example, if a robotic arm were to contact a rigid surface, a delay of 30 ms would initiate a bouncing motion also known as 'chattering' instability. This instability is again due to the fact that the control system could detect the contact only after it has occurred. This would cause a back-up motion that would move the arm away from the surface. Then, the controller would move again towards the surface and so on in a repeated bouncing motion.

4. SOLUTIONS BASED ON FEED-FORWARD

An alternative to feedback control would be for the CNS to pre-programme the torques that the muscles must generate for moving the limbs along the desired trajectories. This method is often referred to as 'feed-forward control'. The torques needed to move the arm can only be computed after the angular motions of the shoulder, elbow and wrist have been derived from the desired movement of the hand—that is after an inverse kinematics problem has been solved. Investigations in robot control in the late 1970s and early 1980s showed that both the inverse kinematic and inverse dynamic problems may be efficiently implemented in a digital computer for many robot geometries (Brady *et al.* 1982). On the basis of these studies, Hollerbach & Flash (1982) put forward the hypothesis that the brain may be carrying out inverse kinematic and dynamic computations when moving the arm in a purposeful way. Their experimental investigation of arm-reaching movements, combined with inverse dynamics calculations, showed that all components of the joint torque played a critical role in the generation of the observed hand trajectories. In particular, Hollerbach & Flash found that while executing reaching movements the subjects were accurately compensating for the dynamic interactions between shoulder and elbow joints.

Evidence that the brain is carefully compensating for the interaction torques was further provided by more recent studies of Ghez and of Thach and their co-workers. Sainburg *et al.* (1993) studied the movements of subjects suffering from a rare peripheral neuropathy. A conse-

quence of this disease is the complete loss of proprioceptive information from the upper and lower limbs. These investigators found that the abnormal motions observed in these subjects could be accounted for by lack of compensation for the joint interaction torques. A similar conclusion was reached later by Bastian *et al.* (1996) about the movements produced by patients suffering from cerebellar lesions. In summary, a substantial body of evidence suggests that the CNS generates motor commands that effectively represent the complex dynamics of multijoint limbs. However, there are different ways for achieving this representation.

5. MEMORY-BASED COMPUTATIONS

A rather direct way for a robot to compute inverse dynamics is based on carrying out explicitly the algebraic operations after representing variables such as positions, velocity acceleration, torque and inertia. Something similar to this approach had been first proposed by Raibert (1978). He started from the observation that inverse dynamic can be represented as the operation of a memory that associates a vector of joint torques to each value of joint angles, angular velocities and angular accelerations. A brute-force approach to dynamics would simply be to store a value of torque for each possible value of position, velocity and acceleration—a computational device that computer scientists call a 'look-up table'. This approach is extremely simple and in fact look-up tables were implicit in early models of motor learning, such as those proposed by Albus (1971) and Marr (1969). However, a closer look at the demands for memory size in a reasonable biological context shows that the look-up table approach may be impracticable.

The number of entries in a look-up table grows exponentially with the number of independent components that define each table entry. Being well aware of this problem, Raibert suggested splitting the arm dynamics computations in a combination of smaller subtables: one can obtain the net torque by adding (i) a term that depends on the joint angles and on the angular accelerations to (ii) a term that depends on the joint angles and on the angular velocities. These two terms may be stored in separate tables. Assuming a resolution of only ten values per variable, the control of a two-joint limb would require two tables with 10^4 entries each. For a more complete arm model, with seven-joint coordinates, each table would have 10^{14} entries. These are still exceedingly large numbers. A method for reducing the size of look-up tables was suggested by Raibert & Horn (1978), who represented the dynamic problem as a sum of three elements, each one requiring a table that depended only on the joint angles. Thus, the two-joint limb involved tables with 100 entries and the seven-joint limb tables with 10^7 entries.

6. THE EQUILIBRIUM-POINT HYPOTHESIS

The work of Raibert (1978) and Hollerbach (1980) showed that inverse dynamics of complex limbs may be computed with a reasonable number of operations and with reasonable memory requirements. However, this work did not provide any direct evidence that the brain is

ever engaged in such computations. Furthermore, on a purely theoretical basis, explanations based on computing inverse dynamics are unsatisfactory because there is no allowance for the inevitable mechanical vagaries associated with any interaction with the environment. For instance, if an external force perturbs the trajectory of the arm, dramatic consequences may follow. When we pick up a glass of water, we must update the pattern of torques that our muscles must apply to generate a movement of the arm. When we open a door, we must deal with a constraint, the hinge, whose location in space is only approximately known. One may say that most of our actions are executed upon a poorly predictable mechanical environment. It would then be erroneous to suggest that a stored pattern of neuromuscular activations corresponds to some particular movement. Instead, the movement that arises from that pattern is determined by the interaction of the muscle forces with the dynamics of the environment.

Hogan (1985*a*) developed this concept in a theory known as impedance control. Hogan's ideas relate to earlier experiments of Feldman (1966) and Bizzi and co-workers. In one of these experiments, Polit & Bizzi (1979) trained monkeys to execute movements of the forearm towards a visual target. The monkeys could not see their moving arm nor could they perceive it as their proprioceptive inflow had been surgically interrupted by the transection of cranial and thoracic dorsal roots—a procedure called 'deafferentation'. Surprisingly, Polit & Bizzi found that, despite such radical deprivation of sensory information, the monkeys could successfully reach the visual targets. What was more unexpected was that the monkeys could reach the intended target even when their arm had been displaced from the initial location just prior to the initiation of an arm movement. This result did not seem to be compatible either with the idea that goal-directed movements are executed by a pre-programmed sequence of joint torques or with the hypothesis that sensory feedback is essential to reach the desired limb position.

The performance of the deafferented monkey can be accounted for by the hypothesis that the centrally generated motor commands modulate the stiffness and rest-length of muscles that act as flexors and extensors about the elbow joint. As a consequence, the elastic behaviour of the muscles, like that of an opposing spring, defines a single equilibrium position of the forearm. A position that ultimately is reached in spite of externally applied perturbations, without need for feedback corrections. This result led to a question concerning the execution of target-directed movements. Are these movements executed just by setting the equilibrium point of a limb to the final target? Or does the descending motor command specify an entire trajectory as a smooth shift of the same equilibrium point? Bizzi *et al.* (1984) addressed this question in another experiment. If, as suggested by the first hypothesis, there is a sudden jump of the limb's equilibrium to the target location, an elastic force driving the hand towards the target would appear from the onset of the movement. This force would be directed all the time towards the target. The experiment of Bizzi and co-workers disproved this hypothesis. As in the work of Polit & Bizzi (1979), they instructed deafferented monkeys to

execute arm movements towards a visual target but with the vision of the arm blocked by an opaque screen. As soon as the EMG activity indicated the onset of a movement, a motor drove the arm right on the target. If this were the equilibrium position specified by the muscle commands at that time, the arm should have remained in place. On the contrary, the experimenters could observe an evident motion backward towards the starting position followed by a forward motion towards the target. This finding indicates that the muscular activation does not specify a force or a torque, as suggested by the inverse dynamic models, nor a final target position. Instead, the response to the initial displacement suggests that the activation of the muscles produces a gradual shift of the limb's equilibrium from the start to end location. Accordingly, at all times the limb is attracted by an elastic force towards the instantaneous equilibrium point. If during a goal-directed movement, the limb is forcefully moved ahead towards the target, the elastic force will drive it towards the lagging equilibrium point, as observed in the experiment.

The sequence of equilibrium positions produced during movement by all the muscular activations has been called by Hogan (1985*b*) a 'virtual trajectory'. The virtual trajectory is a sequence of points where the elastic forces generated by all the muscles cancel each other. By contrast, the actual trajectory is the result of the interaction of these elastic forces with other dynamic components such as limb inertia, muscle velocity–tension properties and joint viscosity. To intuitively illustrate this distinction, consider a ball attached to a rubber band. When the band is displaced from its equilibrium position, a restoring force is generated with amplitude proportional to the displacement. If we move the free end of the rubber band, we control the equilibrium position. As we move the rubber band along a trajectory, the ball will follow a trajectory that results from the interaction of the elastic force with the mass of the ball.

The idea of a virtual trajectory provides a new unified perspective for dealing with (i) the mechanics of muscles, (ii) the stability of movement, and (iii) the solution of the inverse dynamic problem. In fact, a strictly necessary and sufficient condition for a virtual trajectory to exist is that the motor commands directed to the muscle define a sequence of stable equilibrium positions. If this requirement is met, then there exists a single well-defined transformation from the high-dimensional representation of the control signal as a collection of muscle activations, to a low-dimensional sequence of equilibrium points. An advantage of this low-dimensional representation is that, unlike muscle activations, the virtual trajectory may be directly compared with the actual movement of the limb.

The relationship between actual and virtual trajectory is determined by the dynamics of the system and by the stiffness, which transforms a displacement from the equilibrium into a restoring force. In the limit of infinite stiffness, the actual trajectory would match exactly the virtual trajectory. On the other end, with low stiffness values, the difference between virtual and actual trajectory may become quite large. In a work that combined observations of hand movements and computer simulations, Flash (1987) tested the hypothesis that multijoint arm movements are obtained by the CNS shifting the

equilibrium position of the hand along a straight and rectilinear motion from the start to end position. As shown by Morasso (1981), approximately straight hand paths characterize planar hand movements between pairs of targets. If the same movements are analysed at a finer level of detail, however, the paths present certain degrees of inflection and curvature, depending on the direction of movement and on the work-space location. In the simulations Flash made the assumption that the hand equilibrium trajectories (but not necessarily the actual trajectories) are invariantly straight. In addition, she assumed that the equilibrium trajectory had a unimodal velocity profile. The results obtained from the simulation captured the subtle inflections and the curvatures of the actual trajectories. Moreover, the direction of curvature in different work-space locations and with different movement directions matched quite closely the observed movements.

It must be stressed that the stiffness values used in this simulation were taken from measurements that had been performed not during movements but while subjects were maintaining their arm at rest in different locations (Mussa-Ivaldi *et al.* 1985). Katayama & Kawato (1993) and then Gomi & Kawato (1997) repeated Flash's simulation using lower values of stiffness and found, not surprisingly, that in order to reproduce the actual trajectory of the hand, the virtual trajectory had to follow a much more complicated pathway. The results obtained by Gomi & Kawato are at variance with those of Won & Hogan (1995), who were able to show that for relatively slow and low-amplitude arm trajectories the virtual equilibrium point was close to the actual trajectory. Clearly, the complexity of the virtual trajectory depends critically upon the elastic field surrounding the equilibrium point. Experimental estimates of the elastic field under static conditions have shown that the local stiffness, i.e. the ratio of force and displacement, changes at different distances from the equilibrium point (Shadmehr *et al.* 1993). Specifically, it was found that the stiffness decreased with this distance. This is a nonlinear feature of the elastic field. Accordingly if, as in Gomi & Kawato (1997), one attempted to derive the equilibrium point using a linear estimate based on the stiffness at the current position, one would overestimate the distance between current and equilibrium position. At present, however, there is not yet an acceptable technique for measuring the elastic force field generated by the muscles during movement. But, if the shape of the virtual trajectory is a complex path, as in Gomi & Kawato's simulations, then the apparent computational simplicity of the earlier formulation of the equilibrium-point hypothesis is lost.

Another challenge to the equilibrium-point hypothesis comes from the work of Lackner & Dizio (1994) who asked subjects to execute reaching hand movements while sitting at the centre of a slowly rotating room. Because of this rotation, a Coriolis force proportional to the speed of the hand perturbs the subject's arm. The Coriolis force acts perpendicularly to the direction of motion. Lackner & Dizio found that, under this condition, there is a systematic residual error at the final position in the direction of the Coriolis force. This finding seems incompatible with the equilibrium-point hypothesis because the

Coriolis force depends upon hand velocity but not upon hand position. Therefore, it should not alter the location of the final equilibrium point. However, the experimental results of Lackner & Dizio are in apparent contrast with other experimental findings obtained with similar force fields. In particular, Shadmehr & Mussa-Ivaldi (1994) used an instrumented manipulandum for applying a velocity-dependent field to the hand of the subjects. In this paradigm the perturbation acted specifically on the arm dynamics and did not affect in any way other systems, such as the vestibular apparatus. Shadmehr & Mussa-Ivaldi, as well as Gandolfo *et al.* (1996) found that the final position of the movement was not substantially affected by the presence of velocity-dependent fields, in full agreement with the equilibrium-point hypothesis. The cause of the discrepancy between these results and those of Lackner & Dizio (1994) has yet to be determined.

7. BUILDING BLOCKS FOR COMPUTATION OF DYNAMICS: SPINAL FORCE FIELDS

Recent electrophysiological studies of the spinal cord of frogs and rats by Bizzi and co-workers (Bizzi *et al.* 1991; Giszter *et al.* 1993; Mussa-Ivaldi *et al.* 1990; Tresch & Bizzi 1999) suggest a new theoretical framework that combines some features of inverse dynamic computations with the equilibrium-point hypothesis. In these studies, electrical stimulation of the interneuronal circuitry in the lumbar spinal cord of frogs and rats has been shown to impose a specific balance of muscle activation. The evoked synergistic contractions generate forces that direct the hindlimb towards an equilibrium point in space (figure 1). To measure the mechanical responses of the activated muscles, Bizzi *et al.* (1991), Giszter *et al.* (1993) and Mussa-Ivaldi *et al.* (1990) attached the right ankle of the frog to a force transducer. To record the spatial variations of the force vectors generated by the leg muscles, they placed the frog's leg at a location within the leg's work-space. Then, they stimulated a site in the spinal cord and recorded the direction and amplitude of the elicited isometric force at the ankle. This stimulation procedure was repeated with the ankle placed at each of nine to 16 locations spanning a large portion of the leg's work-space. The collection of the measured forces corresponded to a well-structured spatial pattern, called a vector field. In most instances, the spatial variation of the measured force vectors resulted in a field that was at all times both convergent and characterized by a single equilibrium point.

In general, the activation of a region within the spinal cord does not produce a stationary force field. Instead, following the onset of stimulation, the force vector measured at each limb location changes continuously with time (figure 2). As the force vectors elicited by a stimulus change, so does the equilibrium position: the sites occupied by the equilibrium position at subsequent instants of times define a spatial trajectory. The time-varying field is the expression of a mechanical wave that summarizes the combined action of the muscles that are affected by the stimulation. Mechanical waves of the same kind can be used to describe the operation of central pattern generators and of other natural structures




Figure 1. Force fields induced by microstimulation of the spinal cord in spinalized frogs. (From Bizzi *et al.* 1991.) (a) The hindlimb was placed at a number of locations on the horizontal plane (indicated by the dots). At each location a stimulus was derived at a fixed site in the lumbar spinal cord. The ensuing force was measured by a six-axes force transducer. (b) Peak force vectors recorded at the nine locations shown in (a). (c) The work-space of the hindlimb was partitioned into a set of non-overlapping triangles. Each vertex is a tested point. The force vectors recorded on the three vertices are used to estimate, by linear interpolation, the forces in the interior of the triangle. (d) Interpolated force field.




Figure 2. Temporal evolution of a spinal force field. Following the stimulation of a site in the spinal cord, the force vectors change in a continuous fashion. The result is a mechanical wave, described here by a sequence of frames ordered by increasing latency from the onset of the stimulus. The frames are separated by intervals of 86 ms. The dot indicates the location of the static equilibrium point (where the estimated force vector vanishes) in each frame. (From Mussa-Ivaldi *et al.* 1990.)

involved in the control of motor behaviour. At all latencies after the onset of stimulation, the force field converges towards an equilibrium position. The temporal sequence of these equilibrium positions provides us with an image of a virtual trajectory, as in the sequence of frames of figure 2. Sometimes we found that the virtual trajectories observed after electrical stimulation followed circular pathways starting and ending at the same point (Mussa-Ivaldi *et al.* 1990). In contrast, the virtual trajectories inferred by Flash (1987) and Won & Hogan (1995) from reaching arm movements followed rectilinear and smooth pathways, from start to final position of the hand. This is not a surprising discrepancy given the great difference in experimental conditions, limb mechanics and neural structures involved in these studies. Despite these differences, however, it is remarkable that the essential biomechanics of the moving limb is the same for the hindlimb of the spinalized frog and for the arm of the human subject. In both cases, movement is described as a smooth temporal evolution of a convergent force field produced by the spring-like properties of the neuromuscular apparatus.

In the spinal frog, different groups of leg muscles were activated as the stimulating electrodes were moved to different loci of the lumbar spinal cord in the rostro-caudal and mediolateral direction. After mapping most of the premotor regions in the lumbar cord with the technique of electrical microstimulation, Bizzi *et al.* (1991) reached the conclusion that there were at least four areas from which distinct types of convergent force fields were elicited. These results were confirmed by Saltiel *et al.* (1998) with the more selective method of chemical microstimulation. N-methyl-D-aspartate iontophoresis applied to a large number of sites of the lumbar spinal cord revealed a map comparable with that obtained with electrical microstimulation.

Perhaps the most interesting aspect of the investigation of the spinal cord in frogs and rats was the discovery that the fields induced by the focal activation of the cord follow a principle of vectorial summation (figure 3). Specifically, Mussa-Ivaldi *et al.* (1994) developed an experimental paradigm involving the simultaneous stimulation of two distinct sites in the frog's spinal cord. They found that the simultaneous stimulation of two sites led to vector summation at the ankle of the forces generated by each site separately. When the pattern of forces recorded at the ankle following co-stimulation were compared with those computed by summation of the two individual fields, Mussa-Ivaldi *et al.* (1994) found that 'co-stimulation fields' and 'summation fields' were equivalent in more than 87% of cases. Similar results have been obtained by Tresch & Bizzi (1999) by stimulating the spinal cord of the rat. Recently, Kargo & Giszter (2000) showed that force field summation underlies the control of limb trajectories in the frog.

Vector summation of force fields implies that the complex nonlinearity that characterizes the interactions both among neurons and between neurons and muscles is in some way eliminated. More importantly, this result has led to a novel hypothesis for explaining movement and posture based on combinations of a few basic elements. The few active force fields stored in the spinal cord may be viewed as representing motor primitives from which,

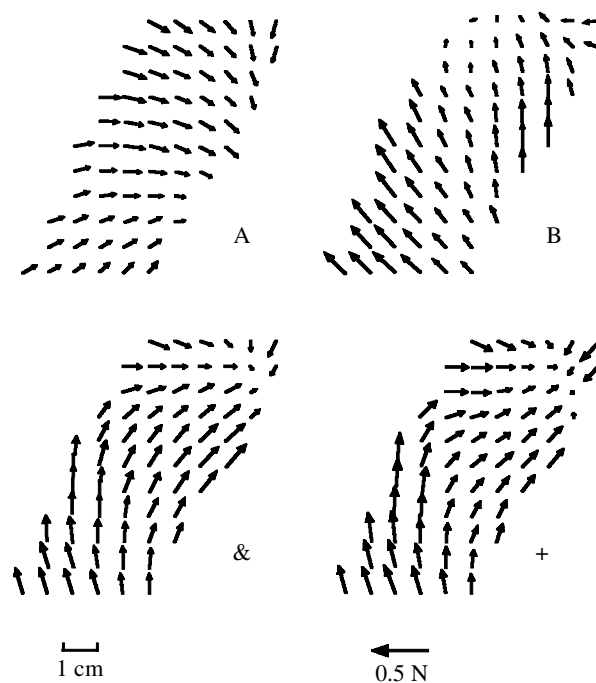


Figure 3. Spinal force fields add vectorially. Fields A and B were obtained in response to stimulations delivered to two different spinal sites. The & field was obtained by stimulating simultaneously the same two sites. It matches closely (correlation coefficient larger than 0.9) the force field in +, which was derived by adding pairwise the vectors in A and in B. This highly linear behaviour was found to apply to more than 87% of dual stimulation experiments. (From Mussa-Ivaldi *et al.* 1994.)

through superposition, a vast number of movements can be fashioned by impulses conveyed by supraspinal pathways. Through computational analysis, Mussa-Ivaldi & Giszter (1992) and Mussa-Ivaldi (1997) verified that this view of the generation of movement and posture has the competence required for controlling a wide repertoire of motor behaviours.

The fields generated by focal activation of the spinal cord are nonlinear functions of limb position, velocity and time: $\phi_i(q, \dot{q}, t)$ (figure 2). Consistent with the observation that these fields add vectorially, one may modify the formulation of the inverse dynamic problem by replacing the generic torque function, $\tau(t)$, with a superposition of spinal fields:

$$D(q, \dot{q}, \ddot{q}) = \sum_{i=1}^K c_i \phi_i(q, \dot{q}, t). \tag{2}$$

Here, each spinal field is tuned by a (non-negative) scalar coefficient, c_i , that represents a descending supraspinal command. We should stress that in this model, the descending commands do not alter the shape of the fields—that is their dependence upon state and time. This is consistent with the empirical observation that the pattern of force orientation of spinal fields remained invariant in time and with different intensities of stimulation (Giszter *et al.* 1993). Thus, it is plausible to assume that the supraspinal signals select the spinal fields by determining how much each one contributes to the total




Figure 4. A simplified model of limb dynamics. The mechanics of the arm in the horizontal plane are approximated by a two-joint mechanism (a). Shoulder and elbow are modelled as two revolute joints with angles q_1 (with respect to the torso) and q_2 (with respect to the forearm), respectively. (b) The dynamics are described by two nonlinear equations that relate the joint torques at the shoulder (D_1) and at the elbow (D_2) to the angular position velocity and acceleration of both joints. The parameters that appear in these expressions are the lengths of the two segments (l_1 and l_2); their masses (m_1 and m_2); and their moments of inertia (I_1 and I_2). The numerical values used in the simulations are the same as those listed in Shadmehr & Mussa-Ivaldi (1994, table 1) and correspond to values estimated from an experimental subject. The terms β_1 and β_2 describe the viscoelastic behaviour of the resting arm. They are simulated here by linear stiffness and viscosity matrices.

field. The computational model of equation (2) is simply a reformulation of inverse dynamics, with the additional constraint that joint torque is produced by the modulation of a set of pre-defined primitives, the fields $\phi_i(q, \dot{q}, t)$. How does the nervous system derive the tuning coefficients, c_i , from the specification of a desired movement? We do not yet have an answer to this question. However, a simple mathematical analysis demonstrates that the model is competent to generate movements similar to those observed in experimental studies. In particular, the superposition of few stereotyped fields is sufficient to control the movements of the two-joint arm shown in figure 4. To demonstrate this, we begin by defining a set of force fields that capture the main qualitative features of the spinal force fields. Here we focus on two specific features: (i) the convergence of the field towards a single

equilibrium point, and (ii) the tendency of muscle forces to grow, reach a peak and then smoothly decrease when a muscle is stretched. A simple way to capture both features is to represent the force fields as gradients of Gaussian potential functions. Each field in this model (figure 5a) is centred at an arm configuration, q_0 and generates a joint torque that depends upon the distance of the limb from this configuration:

$$\chi(q, \dot{q}) = K(q - q_0)e^{-(q - q_0)^T K (q - q_0)} + B\dot{q}. \quad (3)$$

The exponential term ensures that the joint torques do not keep growing as the limb moves away from the equilibrium point. The last term, $B\dot{q}$, represents a viscous dissipative component in its simplest form.

The field $\chi(q, \dot{q})$ depends upon the state of motion of the limb but not upon time. In contrast, it is reasonable to assume that the modules implemented by the neural circuits in the spinal cord have well-defined timing properties, established for example by recurrent patterns of interconnections. A simple way to introduce stereotyped temporal features in our model is to express each force field as a product of the constant viscoelastic term, χ , and a time function $f(t)$:

$$\phi(q, \dot{q}, t) = f(t) \times \chi(q, \dot{q}). \quad (4)$$

The separation of time and state dependence is also consistent with the observation that the forces generated by electrical stimulation of the spinal cord maintain a relatively constant orientation while the overall field amplitude changes in time following each stimulus (Giszter *et al.* 1993). Always for sake of simplicity, here we consider only time-functions that have the form of a smooth step (figure 5b,c) and its first derivative (figure 5d,e) (a smooth 'pulse').

This model provides us with a way to design a family of stereotyped force fields with features that are qualitatively consistent with empirical observations. Here we have derived a small family of eight fields by combining the four fields of figure 6a with each of the time-functions of figure 5. In the end, we have a model of an arm that may only be operated by specifying eight positive numbers, the coefficients c_i of equation (2). The simulation results in figure 6c show that by modulating these eight numbers it is possible to approximate the minimum-jerk movements of figure 6b. The procedure for determining the coefficients is described in Mussa-Ivaldi (1997). Briefly, for each desired movement in figure 6b, one derives the corresponding joint angle trajectory, $q_D(t)$. Then, the dynamics equation (2) is projected on each field, evaluated along the desired trajectory. The result of this operation is a system of eight algebraic equations in the eight unknowns c_i :

$$\sum_{i=1}^8 \Phi_{j,i} c_i = \Lambda_j \quad (j = 1, \dots, 8), \quad (5)$$

with

$$\begin{cases} \Phi_{l,m} = \int \phi_l(q_D(t), \dot{q}_D(t), t) \bullet \phi_m(q_D(t), \dot{q}_D(t), t) dt \\ \Lambda_j = \int \phi_j(q_D(t), \dot{q}_D(t), t) \bullet D(q_D(t), \dot{q}_D(t), \ddot{q}(t)) dt \end{cases} \quad (6)$$




Figure 5. A simplified model of spinal force fields. The force field in (a) is the gradient of a Gaussian potential function defined over the angular coordinates of the mechanism in figure 4. The force vectors converge towards a stable equilibrium point indicated by the small cross. Gaussian potentials are smooth functions defined over the entire limb work-space. The gradient defines a stable equilibrium and the forces grow in amplitude within a region defined by the variance of the Gaussian potential. This behaviour simulates the tendency of muscle-generated forces to grow until a critical amount of stretch is reached. At that point the forces yield and then begin to decline. It is worth observing that in this mechanical context, the variance of the Gaussian potential has the dimension of compliance (the inverse of stiffness). The functions of time in (b) and (d) are a smooth step and a smooth pulse, respectively. When they multiply the field in (a) they generate the wave functions depicted in (c) and (e). The time corresponding to each frame is indicated by the shaded areas in (b) and (d). The step field enforces a persistent equilibrium position. The pulse field is a transient response that emulates the response to spinal stimulation shown in figure 2.

A standard non-negative least-squares method is used to derive the coefficients with the additional requirement that these are greater than or equal to zero. This is an important condition reflecting the fact that muscles cannot push.

The same condition is also sufficient to ensure the stability of posture and movement by imposing that the forces generated by each field converge towards the equilibrium point. Another significant issue, from a computational

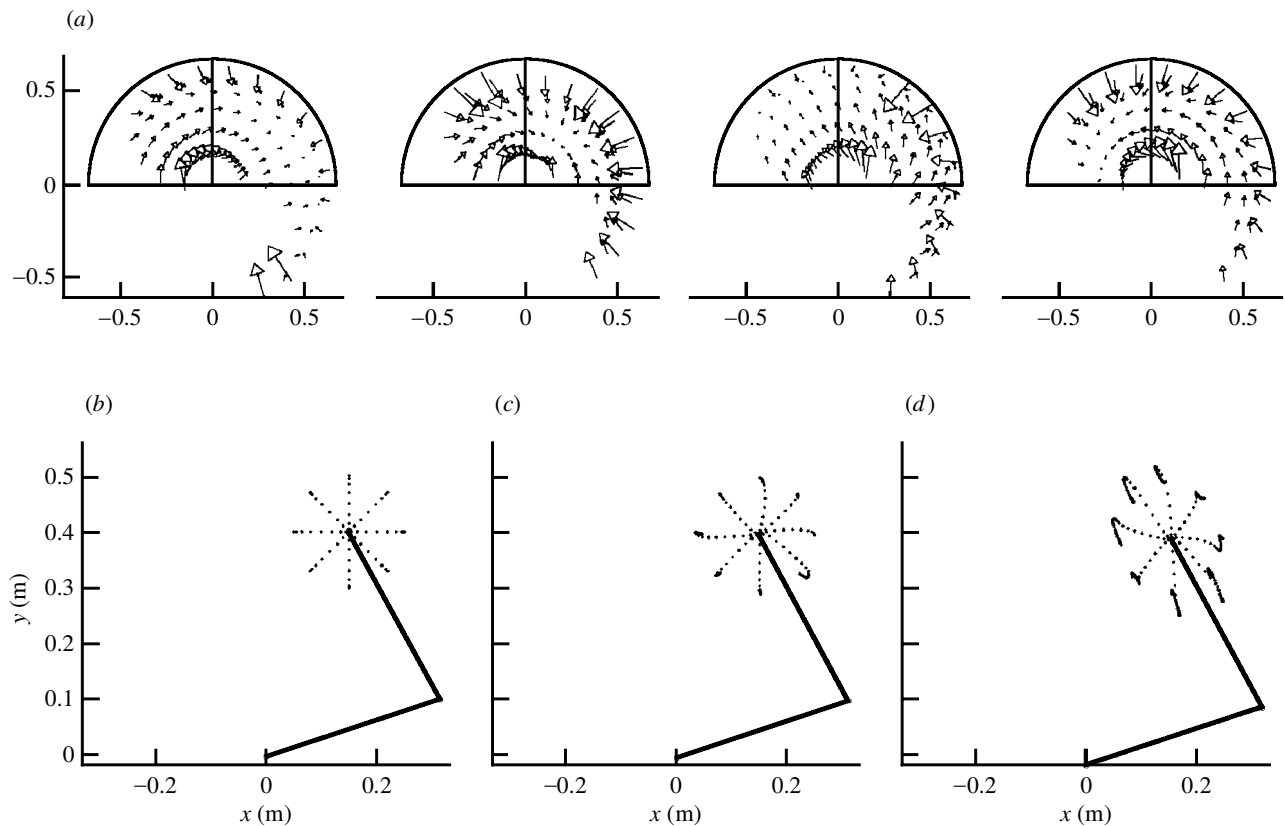


Figure 6. The vectorial superposition of few force fields is competent to reproduce the kinematics and dynamics of arm movements. These movement simulations have been obtained by combining step and pulse fields generated by four Gaussian potentials. The gradients of these potentials are shown in (a). The least-squares procedure defined by equations (3) and (4)—described in more detail in Mussa-Ivaldi (1997)—was used to approximate the desired trajectories in (b). The outcome of the procedure is a set of constant coefficients that modulate a linear combination of the step and pulse fields (equation (2)). The trajectories generated by these linear combinations are shown in (c). When the arm dynamics are perturbed by the application of the force field shown in figure 7*b*, the resulting hand movements are distorted as shown in (d). These trajectories have been obtained by applying the same coefficients as in (c). There is a striking similarity between the simulated perturbations and the experimentally observed responses shown in figure 7*d*.

standpoint, is to ensure that equation (5) may be inverted. We know from elementary algebra that this is contingent upon the matrix Φ being full rank, a condition that is met by the class of nonlinear force fields used here (Poggio & Girosi 1990; Mussa-Ivaldi & Giszter 1992).

Remarkably, the simulation results of this extremely simplified example are not only consistent with the kinematics of reaching, but also with the responses observed (figure 7*d*) when unexpected mechanical perturbations (figure 7*b*) are imposed upon the moving hand. In this case, the trajectories executed by experimental subjects display a distinctive pattern of deflections. The same pattern was produced by the simulation (figure 6*d*) when the same perturbing field was added to the dynamics of the model arm with the same coefficients used to generate the reaching movements of figure 6*c*.

Obviously, the repertoire of behaviours generated by equation (2) depends on the functional form of the fields that, at present, still needs to be accurately determined. In the current model we have strongly simplified the velocity-dependent forces by neglecting the known nonlinear features of muscle force versus velocity dependence. Instead, here we are focusing on the convergent features of the static fields generated by the spinal cord. A particularly significant feature of this field is that

they have a broad but limited region where they exert an influence. This feature is captured by the variance of the Gaussian potentials and may be characterized as the motor counterpart of a receptive field. A computational analysis by Schaal & Atkeson (1998) indicated that on-line learning of complex behaviours is successful only when the receptive fields of the motor primitives are sufficiently small. If each primitive had a large region of influence, the tuning of its parameters might interfere disruptively with neighbouring regions. Remarkably, the force fields elicited by stimulation of muscles and spinal cord have consistently large domains of action.

The vector fields generated by the spinal cord offer a clear example of the impedance control that has been discussed in § 6. The experiments suggest that the circuitry in the spinal cord—and perhaps also in other areas of the nervous system—is organized in independent units, or modules. While each module generates a specific field, more complex behaviours may be easily produced by superposition of the fields generated by concurrently active modules. Thus, we may regard these force fields as independent elements of a representation of dynamics. Recent simulation studies (Mussa-Ivaldi 1997) have demonstrated that by using this modular representation, that is by adding convergent force fields, the CNS may learn to




Figure 7. Adaptation to external force fields. (a) Sketch of the experimental apparatus. Subjects executed planar arm movements while holding the handle of an instrumented manipulandum. A monitor (not shown) placed in front of the subjects and above the manipulandum displayed the location of the handle as well as targets of reaching movements. The manipulandum was equipped with two computer-controlled torque motors, two joint-angle encoders and a six-axes force transducer mounted on the handle. (b) Velocity-dependent force field corresponding to the expression

$$\mathbf{F} = B \cdot \mathbf{v} \text{ with } B = \begin{bmatrix} -10.1 & -11.2 \\ -11.2 & 11.1 \end{bmatrix} \text{ N} \times \text{s m}^{-1}.$$

The manipulandum was programmed to generate a force \mathbf{F} that was linearly related to the velocity of the hand, $\mathbf{v} = [v_x, v_y]$. Note that the matrix B has a negative and a positive eigenvalue. The negative eigenvalue induces a viscous damping at 23° whereas the positive eigenvalue induces an assistive destabilizing force at 113° . (c) Unperturbed reaching trajectories executed by a subject when the manipulandum was not producing disturbing forces (null field). (d) Initial responses observed when the force field shown in (b) was applied unexpectedly. The circles indicate the target locations. (Modified from Shadmehr & Mussa-Ivaldi 1994.)

reproduce and control the dynamics of a multijoint limb coupled with the dynamics of the environment.

8. EVIDENCE FOR INTERNAL MODELS

The findings on the spinal cord suggest that the CNS is capable of representing the dynamic properties of the limbs. This representation is an internal model. The term 'internal model' refers to two distinct mathematical transformations: (i) the transformation from a motor command to the consequent behaviour, and (ii) the transformation from a desired behaviour to the corresponding motor command (Jordan & Rumelhart 1992; Kawato & Wolpert 1998; McIntyre *et al.* 1998). A model of the first kind is called a 'forward model'. Forward models provide the controller with the means not only to predict the expected outcome of a command, but also to estimate the current state in the presence of feedback delays (Miall & Wolpert 1996). A representation of the mapping from

planned actions to motor commands is called an 'inverse model'. Studies by Wolpert *et al.* (1998) proposed that the neural structures within the cerebellum perform sensorimotor operations equivalent to a combination of multiple forward and inverse models. Strong experimental evidence for the biological and behavioural relevance of internal models has been offered by numerous recent experiments (Brashers-Krug *et al.* 1996; Flanagan & Wing 1997; Flash & Gurevich 1992; Gottlieb 1996; Sabes *et al.* 1998; Shadmehr & Mussa-Ivaldi 1994). In particular, the experimental results obtained by Shadmehr & Mussa-Ivaldi (1994) demonstrate clearly the formation of internal models. Their experimental subjects were asked to make reaching movements in the presence of externally imposed forces. These forces were produced by a robot whose free end-point was held as a pointer by the subjects (figure 7). The subjects were asked to execute reaching movements towards a number of visual targets. Since the force field produced by the robot significantly changed

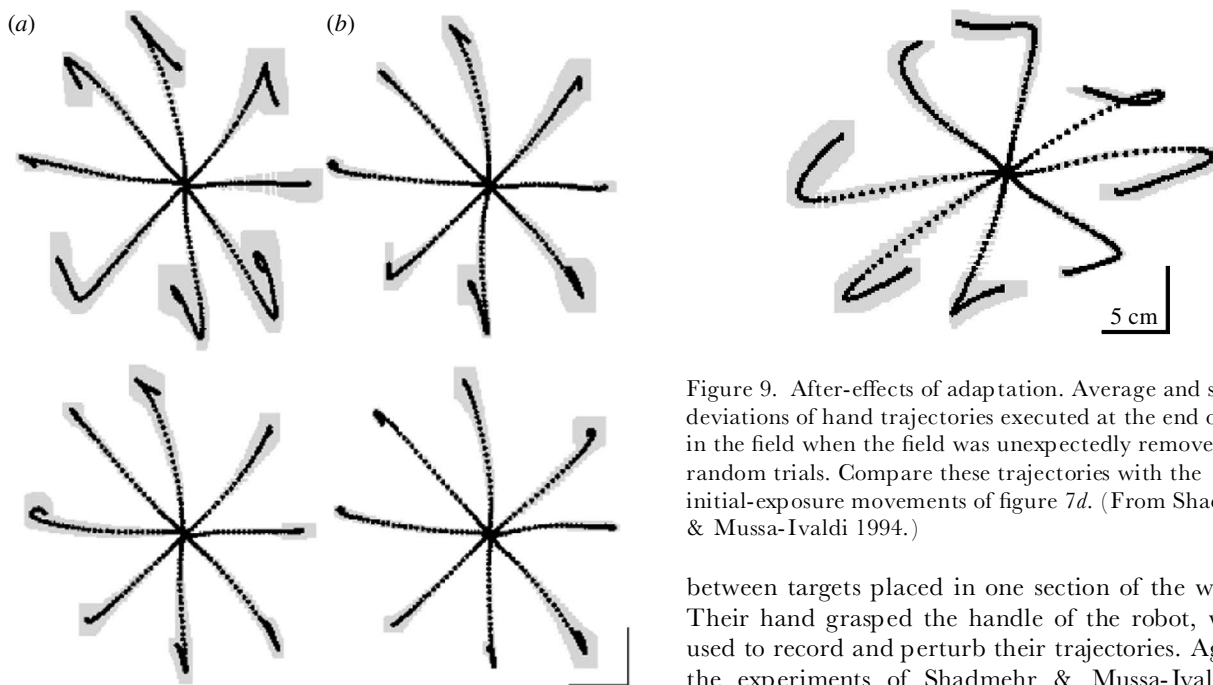


Figure 8. Time-course of adaptation. Average and standard deviation of hand trajectories executed during the training period in the force field of figure 7*b*. Performance is plotted during the (a) first, (b) second, (c) third and (d) final set of 250 movements. All trajectories shown here were under no-visual feedback condition. (From Shadmehr & Mussa-Ivaldi 1994.)

the dynamics of the reaching movements, the subjects' movements, initially, were grossly distorted when compared with the undisturbed movements. However, with practice, the subjects' hand trajectories in the force field converged to a path similar to that produced in absence of any perturbing force (figure 8).

Subjects' recovery of performance is due to learning. After the training had been established, the force field was unexpectedly removed for the duration of a single hand movement. The resulting trajectories (figure 9), named after-effects, were approximately mirror images of those that the same subjects produced when they had initially been exposed to the force field. The emergence of after-effects indicates that the CNS had composed an internal model of the external field. The internal model was generating patterns of force that effectively anticipated the disturbing forces that the moving hand was encountering. The fact that these learned forces compensated for the disturbances applied by the robotic arm during the subjects' reaching movements indicates that the CNS programmes these forces in advance. The after-effects demonstrate that these forces are not the products of some reflex compensation of the disturbing field.

It is of interest to ask what are the properties of the internal model, and in particular whether the model could generalize to regions of the state space where the disturbing forces were not experienced. Recent experiments by Gandolfo *et al.* (1996) were designed to test the generalization of motor adaptation to regions where training had not occurred. In these experiments, subjects were asked to execute point-to-point planar movements

Figure 9. After-effects of adaptation. Average and standard deviations of hand trajectories executed at the end of training in the field when the field was unexpectedly removed on random trials. Compare these trajectories with the initial-exposure movements of figure 7*d*. (From Shadmehr & Mussa-Ivaldi 1994.)

between targets placed in one section of the work-space. Their hand grasped the handle of the robot, which was used to record and perturb their trajectories. Again, as in the experiments of Shadmehr & Mussa-Ivaldi (1994), adaptation was quantified by the amount of the after-effects observed when the perturbing forces were discontinued.

As a way of establishing the generalization of motor learning, Gandolfo *et al.* (1996) perturbed only the trajectories made to a subset of the targets and searched for after-effects in movements that had not been exposed to perturbations. The amount of the after-effect made it possible to quantify the force field that the subjects expected to encounter during their movements in the trained as well as in the novel directions. The same investigators found that the after-effects were present, as expected, along the trained directions, but the magnitude of the after-effects decayed smoothly with increasing distance from the trained directions. This finding indicates that the ability of the CNS to compensate for external forces is restricted to those regions of state space where the perturbations have been experienced by the moving arm. However, most importantly, subjects were also able to compensate to some extent for forces experienced at neighbouring work-space locations. There is a remarkable degree of consistency between these results on dynamic adaptation and some studies of the responses to perturbations in the perceived kinematics. For example, Ghahramani *et al.* (1996) exposed subjects to a localized shift in the visual presentation of a target and observed the adaptive changes in reaching movements of the hand induced by this shift at a number of surrounding locations. They found that the adaptive responses decayed smoothly with distance from the training location, where the visual information was presented. In a different set of experiments, Martin *et al.* (1996) trained subjects to throw a ball at a visual target, while wearing prism spectacles that displayed the visual field. They found that learning did not generalize between right and left hand. However, they could occasionally, although rarely, observe generalization across different throwing patterns executed with the same hand. A somewhat contrasting result was recently obtained by Vetter *et al.* (1999), who did not observe a decay in generalization after remapping

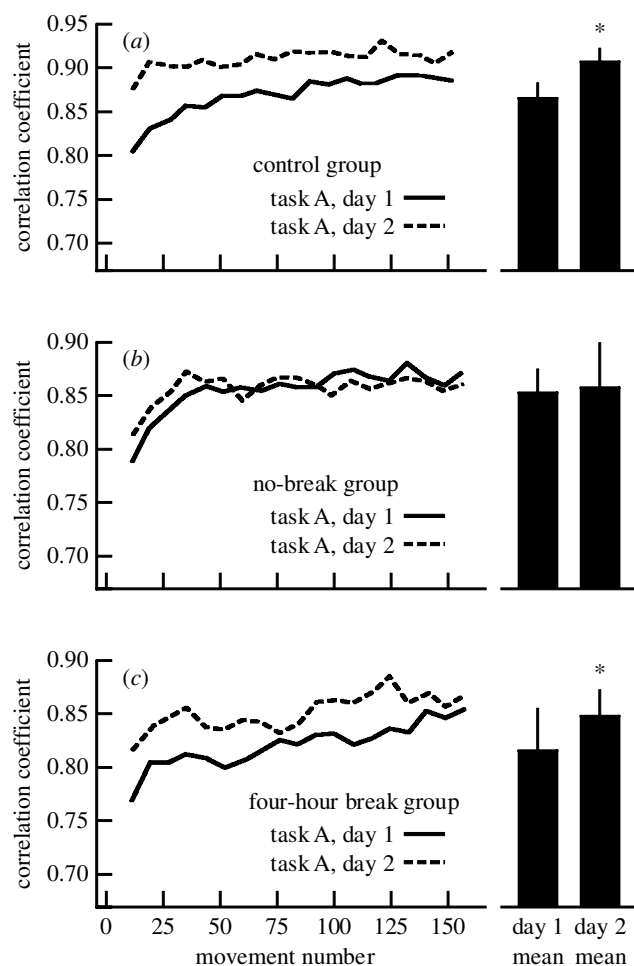


Figure 10. Evidence of motor memory consolidation. The left panels show the learning curves for three groups of subjects. Learning in a perturbing force field was quantified by a correlation coefficient between the trajectories in the field and the average trajectory before any perturbation had been applied. On the right are the mean performances in experiment days 1 and 2. Subjects in the control group (a) practised reaching movements against a force field (task A) in the first day and then were tested again in the same field during the second day. Subjects in the no-break group (b) during the first day practised movements in the field of task A. Then they immediately practised movements in a different field (task B). On the second day they practised again in the field of task A. Finally, subjects of the 4 h break group (c) during the first day were exposed to the fields of tasks A and B but with a breaking interval of 4 h between the two. Their performance was measured on task A in day 2. Learning curves and mean performance were significantly higher in day 2 both for the control group and for the 4 h break group. In contrast, subjects in the no-break group did not display any difference in performance from day 1 to day 2. (From Brashers-Krug *et al.* 1996.)

of the target location in a pointing paradigm similar to that of Ghahramani *et al.* (1996).

The experiments on dynamic adaptation have shown that subjects adapt to a new environment by forming a representation of the external force field that they encounter when making reaching movements. Does this representation form an imprint in long-term memory? Brashers-Krug *et al.* (1996) investigated this question by exposing their subjects to perturbing force fields that

interfered with the execution of reaching movements (figure 10). After practising reaching movements, these subjects were able to compensate for the imposed forces (task A) and were able to guide the cursor accurately to the targets despite the disturbing forces. This group of subjects, which was tested 24 h later with the same disturbing forces, demonstrated not only retention of the acquired motor skill, but also additional learning. Surprisingly, they performed at a significantly higher level on day 2 than they had on day 1. A second group of subjects was trained on day 1 with a different pattern of forces (task B), immediately after performing task A. In task B the manipulandum produced forces opposite in direction to those applied during task A. When this second group of subjects was tested for retention of task A on day 2, the investigators found that the subjects did not retain any of the skills that had been learned earlier. This phenomenon is known as retrograde interference. Next, Brashers-Krug *et al.* (1996) investigated whether the susceptibility to retrograde interference decreased with time. They found that retrograde interference decreased monotonically with time as the interval between task A and B increased (figure 10). When 4 h passed before task B was learned, the skill learned in task A was completely retained—the initial learning had consolidated. What is remarkable in these results is that motor memory is transformed with the passage of time and in absence of further practice, from an initial fragile state to a more solid state.

9. CORTICAL PRIMITIVES

While the internal representation of limb's dynamics based on modules is of central importance for the execution of motor tasks, voluntary movements are often specified and planned in terms of goals. Recordings of cell activity from primates' premotor areas of the frontal lobe have revealed the presence of neurons active during various forms of grasping. Each neuron is selectively active for a specific type of grasping. Rizzolatti *et al.* (1990) interpreted their findings as an indication of a 'vocabulary of actions'. The words of the vocabulary are represented by neuronal populations, each of which specifies a given motor act. It is of interest that these neurons are active not only during the act of grasping, but also when the primate simply looks at the objects that, eventually, will be picked up. Neurons with similar visuomotor properties have been found in the parietal lobe where neurons selectively active during manipulations are present in the anterior intraparietal area (Sakata *et al.* 1995).

Cells active prior and during reaching moments were also found in the parietal lobe by Mountcastle *et al.* (1975) and in the frontal motor area by Georgopoulos *et al.* (1988). However, unlike the cells representing grasping, directionally tuned arm-reaching neurons display continuous parameterization of directional movements.

While the significance and the functional role of distributed and categorical cortical codes remains to be investigated, a question of great importance is to understand how the codes representing reaching and manipulation may be combined with each other by the brain to span a repertoire of purposeful behaviours. At present, we know that spinal force fields implementing

the execution of motor commands are combined by vectorial superposition. However, we do not know the rules that govern the combination of reaching and manipulation goals. If there is a system of high-order primitives that code for goals, then it remains to be established how these goals may be combined and translated into movements so that their concurrent activation leads to meaningful results.

10. CONCLUSION

In this paper we have shown that the problem of planning and execution of a visuomotor task can be divided into a set of subprocesses. Actions are first planned in reference to the objects and the geometry of the surrounding environment. Then, once a movement is specified in the environment, it must be translated into motions of multiple body segments. Finally, the execution phase requires the solution of an inverse dynamic problem. Various schemes have been proposed in order to represent and solve the complex dynamics of the multi-joint apparatus: look-up tables, equilibrium-point trajectory, combination of spinal cord modules and the formation of internal models of dynamics.

Motor patterns come into fragments or modules. These modules find their ultimate expression in the force fields generated by the concurrent activation of multiple muscles. Our current understanding of the spinal cord suggests that this structure provides the brain with a first vocabulary of such synergistic force fields. What we found to be remarkable is that there seems to be only a handful of words in this vocabulary in spite of all the muscle combinations that could be realized. It will certainly be important to understand what are the origin and the rationale for this particular choice of spinal force fields. By focusing on the mechanics of force fields we have not only found a system of modules but also a very simple syntax: fields can be literally added with each other to generate a rich repertoire of behaviours. This additive property is likely to be the basis for our ability to compensate complex patterns of force disturbances, as it has been seen in many of the experiments that we have reviewed. And, ultimately, the internal model of a limb's dynamics is nothing else than another field which relates the forces generated by the muscular apparatus to the state of motion of the limb.

This work was supported by National Institute of Health grants NS3567 to F.A.M.-I., NS 09343 to E.B. and 5 P50 MH48185 to both authors.

REFERENCES

- Albus, J. 1971 The theory of cerebellar function. *Math. Biosci.* **10**, 25–61.
- Bastian, A. J., Martin, T. A., Keating, J. G. & Thach, W. T. 1996 Cerebellar ataxia: abnormal control of interaction torques across multiple joints. *J. Neurophysiol.* **76**, 492–509.
- Bizzi, E., Accornero, N., Chapple, W. & Hogan, N. 1984 Posture control and trajectory formation during arm movement. *J. Neurosci.* **4**, 2738–2744.
- Bizzi, E., Mussa-Ivaldi, F. & Giszter, S. 1991 Computations underlying the execution of movement: a biological perspective. *Science* **253**, 287–291.
- Brady, M., Hollerbach, J., Johnson, T., Lozano-Perez, T. & Mason, M. 1982 *Robot motion: planning and control*. Cambridge, MA: MIT Press.
- Brashers-Krug, T., Shadmehr, R. & Bizzi, E. 1996 Consolidation in human motor memory. *Nature* **382**, 252–255.
- Feldman, A. G. 1966 Functional tuning of the nervous system with control of movement or maintenance of steady posture. II. Controllable parameters of the muscles. *Biophysics* **11**, 565–578.
- Flanagan, J. & Wing, A. 1997 The role of internal models in motion planning and control: evidence from grip force adjustments during movements of hand-held loads. *J. Neurosci.* **17**, 1519–1528.
- Flash, T. 1987 The control of hand equilibrium trajectories in multi-joint arm movements. *Biol. Cybernet.* **57**, 257–274.
- Flash, T. & Gurevich, I. 1992 Arm stiffness and movement adaptation to external loads. *Proc. A. Conf. IEEE Engng Med. Biol. Soc.* **13**, 885–886.
- Gandolfo, F., Mussa-Ivaldi, F. & Bizzi, E. 1996 Motor learning by field approximation. *Proc. Natl Acad. Sci. USA* **93**, 3483–3486.
- Georgopoulos, A., Kettner, R. & Schwartz, A. 1988 Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Coding for the direction of movement by a neuronal population. *J. Neurosci.* **8**, 2913–2927.
- Ghahramani, Z., Wolpert, D. M. & Jordan, M. I. 1996 Generalization to local remappings of the visuomotor coordinate transformation. *J. Neurosci.* **16**, 7085–7096.
- Giszter, S., Mussa-Ivaldi, F. & Bizzi, E. 1993 Convergent force fields organised in the frog's spinal cord. *J. Neurosci.* **13**, 467–491.
- Gomi, H. & Kawato, M. 1997 Human arm stiffness and equilibrium-point trajectory during multi-joint movements. *Biol. Cybernet.* **76**, 163–171.
- Gottlieb, G. 1996 On the voluntary movement of compliant (inertial-viscoelastic) loads by parcellated control mechanisms. *J. Neurophysiol.* **76**, 3207–3229.
- Hogan, N. 1985a Impedance control: an approach to manipulation. *ASME J. Dynamic Syst. Measurement Control* **107**, 1–24.
- Hogan, N. 1985b The mechanics of posture and movement. *Biol. Cybernet.* **52**, 315–331.
- Hogan, N., Bizzi, E., Mussa-Ivaldi, F. & Flash, T. 1987 Controlling multi-joint motor behavior. *Exerc. Sport Sci. Rev.* **15**, 153–190.
- Hollerbach, J. M. 1980 A recursive formulation of Lagrangian manipulator dynamics. *IEEE Trans. Syst. Man, Cybernet.* **SMC-10**, 730–736.
- Hollerbach, J. M. & Flash, T. 1982 Dynamic interactions between limb segments during planar arm movements. *Biol. Cybernet.* **44**, 67–77.
- Jordan, M. & Rumelhart, D. 1992 Forward models: supervised learning with a distal teacher. *Cogn. Sci.* **16**, 307–354.
- Kargo, W. J. & Giszter, S. F. 2000 Rapid correction of aimed movements by summation of force-field primitives. *J. Neurosci.* **20**, 409–426.
- Katayama, M. & Kawato, M. 1993 Virtual trajectory and stiffness ellipse during multijoint arm movement predicted by neural inverse models. *Biol. Cybernet.* **69**, 353–362.
- Kawato, M. & Wolpert, D. 1998 Internal models for motor control. *Novartis Found. Symp.* **218**, 291–304.
- Lackner, J. R. & Dizio, P. 1994 Rapid adaptation to Coriolis force perturbations of arm trajectory. *J. Neurophysiol.* **72**, 299–313.
- McIntyre, J., Berthoz, A. & Lacquaniti, F. 1998 Reference frames and internal models. *Brain Res. Brain Res. Rev.* **28**, 143–154.
- Marr, D. 1969 A theory of cerebellar cortex. *J. Physiol.* **202**, 437–470.

- Martin, T. A., Keating, J. G., Goodkin, H. P., Bastian, A. J. & Thach, W. T. 1996 Throwing while looking through prisms. II. Specificity and storage of multiple gaze-throw calibrations. *Brain* **119**, 1199–1211.
- Merton, P. 1972 How we control the contraction of our muscles. *Sci. Am.* **226**, 30–37.
- Miall, R. & Wolpert, D. 1996 Forward models for physiological motor control. *Neural Net.* **9**, 1265–1279.
- Morasso, P. 1981 Spatial control of arm movements. *Exp. Brain Res.* **42**, 223–227.
- Mountcastle, V., Lynch, J., Georgopoulos, A., Sakata, H. & Acuna, C. 1975 Posterior parietal association cortex of the monkey: command functions for the operations within extrapersonal space. *J. Neurophysiol.* **38**, 871–908.
- Mussa-Ivaldi, F. A. 1997 Nonlinear force fields: a distributed system of control primitives for representing and learning movements. In *Proceedings of the 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation*, pp. 84–90. Los Alamitos, CA: IEEE Computer Society Press.
- Mussa-Ivaldi, F. & Giszter, S. 1992 Vector field approximation: a computational paradigm for motor control and learning. *Biol. Cybernet.* **67**, 491–500.
- Mussa-Ivaldi, F., Hogan, N. & Bizzi, E. 1985 Neural, mechanical and geometrical factors subserving arm posture in humans. *J. Neurosci.* **5**, 2732–2743.
- Mussa-Ivaldi, F., Giszter, S. & Bizzi, E. 1990 Motor-space coding in the central nervous system. *Cold Spring Harb. Symp. Quant. Biol.* **55**, 827–835.
- Mussa-Ivaldi, F., Giszter, S. & Bizzi, E. 1994 Linear combinations of primitives in vertebrate motor control. *Proc. Natl Acad. Sci. USA* **91**, 7534–7538.
- Poggio, T. & Girosi, F. 1990 A theory of networks for learning. *Science* **247**, 978–982.
- Polit, A. & Bizzi, E. 1979 Characteristics of motor programs underlying arm movements in monkeys. *J. Neurophysiol.* **42**, 183–194.
- Raibert, M. 1978 A model for sensorimotor control and learning. *Biol. Cybernet.* **29**, 29–36.
- Raibert, M. & Horn, B. 1978 Manipulator control using the configuration space method. *Industr. Robot* **5**, 69–73.
- Rizzolatti, G., Gentilucci, M., Camarada, R., Gallese, V., Luppino, G., Matelli, M. & Fogassi, L. 1990 Neurons related to reaching-grasping arm movements in the rostral part of area 6 (area 6a beta). *Exp. Brain Res.* **82**, 337–350.
- Sabes, P., Jordan, M. & Wolpert, D. 1998 The role of inertial sensitivity in motor planning. *J. Neurosci.* **18**, 5948–5957.
- Sainburg, R. L., Poizner, H. & Ghez, C. 1993 Loss of proprioception produces deficits in interjoint coordination. *J. Neurophysiol.* **70**, 2136–2147.
- Saltiel, P., Tresch, M. & Bizzi, E. 1998 Spinal cord modular organization and rhythm generation: an NMDA iontophoretic study in the frog. *J. Neurophysiol.* **80**, 2323–2339.
- Sakata, H., Taira, M., Murata, A. & Mine, S. 1995 Neural mechanisms of visual guidance of hand action in the parietal cortex of the monkey. *Cerebr. Cortex* **5**, 429–438.
- Schaal, S. & Atkeson, C. 1998 Constructive incremental learning from only local information. *Neural Comput.* **10**, 2047–2084.
- Shadmehr, R. & Holcomb, H. H. 1997 Neural correlates of human memory consolidation. *Science* **277**, 821–825.
- Shadmehr, R. & Mussa-Ivaldi, F. A. 1994 Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* **14**, 3208–3224.
- Shadmehr, R., Mussa-Ivaldi, F. A. & Bizzi, E. 1993 Postural force fields of the human arm and their role in generating multi-joint movements. *J. Neurosci.* **13**, 45–62.
- Sherrington, C. 1910 Flexion-reflex of the limb, crossed extension reflex and reflex stepping and standing. *J. Physiol.* **40**, 28–121.
- Taub, E. & Berman, A. 1968 Movement and learning in the absence of sensory feedback. In *The neuropsychology of spatially oriented behaviour* (ed. S. Freeman), pp. 173–192. Homewood, IL: Dorsey.
- Toni, I., Krams, M., Turner, R. & Passingham, R. E. 1998 The time course of changes during motor sequence learning: a whole brain fMRI study. *NeuroImage* **8**, 50–61.
- Tresch, M. & Bizzi, E. 1999 Responses to spinal microstimulation in the chronically spinalized rat and their relationship to spinal systems activated by low threshold cutaneous stimulation. *Exp. Brain Res.* **129**, 401–416.
- Vallbo, A. 1970 Slowly adapting muscle receptors in man. *Acta Physiol. Scand.* **78**, 315–333.
- Vetter, P., Goodbody, S. J. & Wolpert, D. M. 1999 Evidence for an eye-centered spherical representation of the visuomotor map. *J. Neurophysiol.* **81**, 935–939.
- Wolpert, D., Miall, R. & Kawato, M. 1998 Internal models in the cerebellum. *Trends Cogn. Sci.* **2**, 338–347.
- Won, J. & Hogan, N. 1995 Stability properties of human reaching movements. *Exp. Brain Res.* **107**, 125–136.