The acquistion of motor behavior
Emilio Bizzi; Ferdinando A Mussa-Ivaldi
Daedalus; Spring 1998; 127, 2; Research Library Core
pg. 217

Emilio Bizzi and Ferdinando A. Mussa-Ivaldi

The Acquisition of Motor Behavior

INTRODUCTION

RECENTLY, GREAT STRIDES HAVE BEEN MADE in understanding the neural foundations of motor behavior. Through the combined efforts of biologists, computer scientists, physicists, and engineers, a picture has begun to emerge of the way in which the nervous system regulates movement.

The human body is capable of an extraordinary range of movements. Years of practice shape the complex skills of professional dancers, pianists, and tennis players. But to neuroscientists, even the simplest everyday movements—reaching for a cup, buttoning a jacket, descending a flight of stairs—present a challenge to scientific explanation. We still do not fully understand how the brain controls these actions, nor can the most sophisticated robotics expert create a machine capable of matching the everyday competence of the central nervous system of the bird, the frog, or the cat, much less that of the human being.

The goal of this essay is to explain what neuroscience has established so far about how the central nervous system (CNS) deals with the complex dynamics of our limbs as they interact with a variable and often unpredictable environment. We will review how scientists have approached the study of movement, the problems they have encountered, and the solutions they have proposed.

Emilio Bizzi is E. McDermott Professor in the Brain Sciences and Human Behavior at the Massachusetts Institute of Technology.

Ferdinando A. Mussa-Ivaldi is associate professor of physiology at Northwestern University.

One issue of particular interest to researchers has been the question of how the brain handles the staggering number of mechanical variables involved in even the simplest movement. To illustrate the complexity of this basic problem, consider the analogy of a marionette—a rough imitation of the human body with a head, a trunk, two arms, two hands, two legs, and two feet. Rather than pulling on wires, a modern-day puppeteer uses a computerized control board with a switch connected to each of the marionette's thirteen joints. Each switch can take one of five positions: two for the extreme angles and three for the intermediate values.

To bring the marionette to life, the puppeteer faces the daunting task of mastering and controlling over 5 to the 13th different positions, or approximately one billion. If we now make this simple marionette more like the infinitely more complicated human body—say, by adding ball joints with two angles at the hip, shoulders, hands, and feet—the number of possible positions rises to 5 to the 19th power, or more than ten thousand billion. This analogy gives some sense of the monumental problem handled routinely by the CNS in the ongoing course of motor control.

To further complicate things, there are countless different ways for the CNS to achieve any given goal involving movement. When a reader turns a page, for example, there are a variety of different trajectories the hand could follow, with many combinations of motions at the shoulder, elbow, and wrist. In addition, the single motion of a joint can be "scripted" with numerous patterns of muscle activations. This characteristic of the biological system is called "kinematic redundancy." It means that there is no single solution to a given problem of motor control. How the CNS decides which plan of action to pursue is a difficult and fascinating question for researchers.

Finally, there is the issue of motor learning. In the course of a lifetime, a human being masters a huge repertoire of movements, the memory of which must somehow be stored in the CNS, despite the very real constraints presented by brain anatomy. Even if one were to assume that each of the billions of neurons in the human brain were to represent a posture in the body's repertoire, storage capacity would fall far short of what is needed. How, then, do our brains meet this challenge?

This essay will address the above questions by reviewing some of the experimental findings made over the last few years. First, we will focus on possible ways in which the CNS may produce the forces necessary to generate movements. In this context, we will also consider the problems presented by kinematic redundancy. Second, we will show how motor memories may be represented, stored, and retrieved through the formation of internal models of limb dynamics. Finally, we will review some of the neurophysiological evidence that suggests that motor learning consists of tuning the activity of a relatively small group of neurons. Each of these groups constitutes a "module," which combines with others to produce a vast repertoire of motor behaviors.

THE FORCES THAT DRIVE OUR LIMBS

In the last eighty years, biologists, engineers, and computer scientists have proposed theories to explain how the CNS may produce the forces necessary to generate movements. Generating movements in biological or robotic systems is computationally complex because of the large number of mechanical degrees of freedom of the body. In this section, we will review two sets of ideas: those derived from the field of biology, and those derived from the field of robotics.

A simple yet very common task for our brain is to generate a trajectory of the forearm, involving a temporal sequence of elbow angles from an initial value to a final one. This movement is produced by muscles that together must apply a net force on the elbow joint. From Newton's equation (force = mass × acceleration), we know that the acceleration of an object is proportional to the applied force. Thus in order to move the forearm, the brain must solve a specific problem in physics—that of determining which force must be applied by the muscles in order to produce movement through the desired sequence of angles. Roboticists have called this an "inverse dynamics problem" to distinguish it from the direct dynamics problem, whose goal is to find the trajectory that would result from the application of a known force.

In solving this inverse dynamics problem for a simple movement of the forearm, the brain faces a complex computational challenge. The net force at the elbow is the sum of the forces exerted by all the muscles around a joint; there is thus a degree of arbitrariness in the choice of each muscle's contribution. This situation reflects the ubiquitous "redundancy" that characterizes the motor system and makes inverse dynamics an "ill-posed" problem.[1]

Of course, in everyday life we deal with more challenging tasks than moving the elbow between two angles. Accordingly, the computational problems that the brain must face are more complex than solving Newton's equation for a single joint. For instance, the inertia of the arm, something that our puppeteer must know in order to program the marionette's motion, depends in complex ways upon the angles of the joints. In addition, for purely physical reasons it so happens that the motion of one joint causes a force to be exerted on the neighboring joints. These are factors that the puppeteer must consider. Furthermore, there is the additional issue of redundancy, which while providing flexibility poses a difficult problem for motor control. If we were to ask you to touch a word on this page, you may do so in an infinite variety of ways, each of which involves a different posture for your shoulder, elbow, and wrist. While you make your choice effortlessly, the presence of multiple solutions and the necessity of selecting one among them poses a significant computational challenge for the brain. How does the central nervous system solve all these problems?

Several possible explanations have emerged from studies in robotics and computational neuroscience. At the beginning of this century, Sir Charles Sherrington proposed feedback as a way for the CNS to control a limb's motion.[2] In a feedback system, sensory signals would provide information to the CNS about the position and velocity of the controlled limb at each point in time. If a subject's goal was to reach a desired position with the arm, a feedback-control system would compare the arm's current position with the one desired. The difference between the two positions would serve as a measure of the error at any given time. Once the error was computed, all the brain would need to do is to produce a force directed toward the desired position with an amplitude proportional to the amplitude of the error. This theory of control had the appeal of simplicity.

Sherrington observed that when a muscle is extended, 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 fibers. Muscle spindles are well suited for feedback control because they provide direct information on a muscle's length to the CNS. Sherrington hypothesized that voluntary movements were accomplished by combining stretch reflexes with other reflexes in a continuous chain. The theory proposed that movement patterns as complex as walking could be generated by local reflexes without central supervision.

The idea that all movements can be set up by the brain as a chain of reflexes was later found to be simplistic and incompatible with experimental results. If movements were pure reflexes, then we would be paralyzed in the absence of feedback information. In fact, we now know that monkeys and humans can execute various limb movements even after the complete surgical interruption of the pathways that convey sensory information from the limb to the nervous system.

Once it became clear that experimental facts did not support the idea that reflexes alone generated movements, investigators began to search for more effective explanatory alternatives. In recent years, an important contribution to research in biological motor control originated in the field of robotics.

An alternative to the notion of feedback control would be to assume that the CNS explicitly solves the inverse dynamics problem. In other words, the brain computes the forces that the muscles must generate in order to move a limb along the desired trajectory. In theory, this dynamic problem can be addressed only after the trajectory of the joint angles has been derived from the trajectory of the endpoint—that is, after an inverse kinematics problem has been solved. Investigations into robot-control done in the late 1970s and early 1980s have shown that both the inverse kinematics and inverse dynamics problem may be efficiently solved on a digital computer for many robot geometries. On the basis of these studies, John Hollerbach and Tamar Flash put forward the hypothesis that the brain may also be carrying out inverse kinematic and dynamic computations when the arm moves in a purposeful way.[3]

Hollerbach's work in robotics was aimed at finding efficient algorithms for calculating the inverse dynamics of artificial arms.[4] His algorithms are well-organized sequences of elementary operations—additions and multiplications—that lead from the desired trajectories of the limb to the needed forces.

A simpler way to compute inverse dynamics was proposed by Marc Raibert in 1977.[5] Raibert started from the observation that the inverse dynamics problem can be represented as the operation of a memory that associates a set of forces with each specific state of motion of the arm. In his approach, the values of the various torques for each possible value of position, velocity, and acceleration of the limb are stored in a computational device that computer scientists call a "look-up table." Unfortunately, the huge demand for memory size makes the look-up table an impractical solution in the biological context.

The work of Raibert and Hollerbach had the merit of showing that the inverse dynamics of limbs may be computed for the robot with a reasonable number of operations and with reasonable memory requirements. However, this work provided no direct evidence that the brain engages in such computation. Furthermore, on a purely theoretical level, explanations based on computing inverse dynamics are unsatisfactory because there is no allowance for the inevitable mechanical uncertainty associated with a limb's interaction with the environment. Living organisms, unlike conventional robots and computers, generally do not operate on the basis of some predefined program. Instead, they learn from experience. As a result, the theories from early robotics, which focused on how a system can be programmed to compute dynamics, did not shed much light on how the brain could learn to deal with the dynamics of limbs operating in the context of a dynamically changing environment.

MOTOR LEARNING: THE ROLE OF INTERNAL MODELS

The focus on learning from experience as a means of acquiring motor skills has gained great strength in recent years. This new approach derives in large part from theoretical and experimental studies on networks of idealized neurons. A number of theoretical studies have shown that when networks of artificial neurons are exposed to repeated motor commands paired with their sensory consequences, learning of fairly complex motor tasks may take place without the need for explicit programming. The learning results from a change in the internal structure of the artificial network, specifically a change in the connectivity among its elements.

On the basis of these results, scientists have proposed that similar processes might be present in the central nervous system. The hypothesis is that learning is the result of repeated exposures to sensory signals coming from the moving limbs as they interact with the environment. The repeated sensory signals are funneled to the motor areas of the central nervous system, where signals that activate the muscles are produced. The actions produced by the activity of the motor areas are initially imprecise, but a feedback mechanism produces a gradual convergence on the correct solution. Ultimately, this iterative process would lead to the establishment of an internal representation of the task through the gradual change in the synaptic strength of the neurons of the motor areas. If the task is that of moving a limb, for example, the outcome of learning would be the formation of an internal model of the limb's dynamics. The internal model, according to this view, is embedded in the newly formed connectivity of a group of neurons. The activity of this group of neurons generates the neural impulses necessary for the execution of the learned motor task.

The experimental results obtained by Reza Shadmehr and Ferdinando Mussa-Ivaldi support the notion of internal models.[6] Their experimental setup was simple: human subjects were asked to make reaching movements in the presence of externally imposed forces. These forces were produced by a robot whose free endpoint was held as a pointer by the subjects. The subjects were asked to move the pointer toward a number of visual targets. Since the forces produced by the robot significantly changed the dynamics of the reaching movements, the subjects' movements were at first grossly distorted when compared to the undisturbed movements. However, with practice, the subjects' hand trajectories in the force field converged to a path similar to that produced in the absence of any force field.

In other words, the subjects learned to compensate for the applied forces. In order to investigate the neural changes underlying this type of motor learning, Shadmehr and Mussa-Ivaldi devised a simple but revealing experimental manipulation. After the subjects had learned to compensate, the researchers removed the perturbing force for the duration of a single movement. The resulting trajectories, named "aftereffects," were approximate mirror images of the distorted movements that were observed when the subjects were initially exposed to the forces.

The emergence of these aftereffects suggests that the central nervous system composes an internal model of the external force field, a model that generates patterns of compensating forces that anticipate the forces that had perturbed the moving hand.

It is of interest to ask what the properties of the internal model might be, and whether the model could generalize to regions of the work space where the perturbing forces had not been experienced. Recent experiments by Francesca Gandolfo and colleagues were designed to test whether motor adaptation generalized to regions of the work space where no training had occurred.[7] In these experiments, subjects were asked to execute point-to-point planar hand movements between targets placed in one section of the work space. The subject's hand grasped the handle of a robot, which was used both to record and disturb their trajectories. Again, as in the experiments of Shadmehr and Mussa-Ivaldi, the adaptation was quantified by the degree of aftereffect observed when the perturbing forces were discontinued.

Gandolfo found that aftereffects were present, as expected, along the directions where subjects had been trained, but the magnitude of the aftereffects diminished smoothly with increasing distance from the trained locations. This finding indicates that the ability of the CNS to compensate for external forces is restricted to those spatial locations where perturbations have been experienced by the moving arm.

In summary, the work of Shadmehr and Mussa-Ivaldi and of Gandolfo and his collaborators has 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? Recently Thomas Brashers-Krug and his coworkers investigated this question by exposing their subjects' movements to forces that interfered with the execution of reaching to a target.[8] After some practice, these subjects were able to guide the cursor accurately to the targets despite the interfering forces.

Twenty-four hours after learning the task, one group of subjects was tested with the same disturbing forces and demonstrated not only retention of the acquired motor skill but also additional learning. Surprisingly, they performed at a significantly higher level the second day than they had the first.

A second group of subjects was trained on day one, like the first group, to execute reaching movements with a perturbing field (task A). Immediately afterwards, on the same day, these subjects were trained to execute the same movements with perturbing forces in the opposite direction (task B). When these subjects were tested on a subsequent day, Brashers-Krug's team found that retention of task A had been significantly impaired by exposure to task B. This phenomenon is known as "retrograde interference." In a later experiment, the same researchers found that retrograde interference decreased monotonically with time as the interval between task A and B increased. When four hours passed before task B was learned, the skill learned in task A was completely retained; apparently, the initial learning had consolidated. What is remarkable in these results is that motor memory was transformed, with the passage of time and in the absence of further practice, from an initial fragile state to a more solid state.

Taken together, the experiments just described indicate several things: 1) There was a certain degree of specificity in the learning of a simple motor task. The internal model that the subjects learned was restricted to that part of the space where interference had been experienced. The same external forces could not be handled in a different part of the work space. 2) There was an enhancement of the learned task that did not depend upon practice, but only on the passage of time. 3) There was a process of consolidation of learning that took four hours at a minimum. The consolidation was not dependent upon practice; it was an internally generated event.

We conceive of the internal model as a newly formed rearrangement of synaptic contacts among a group of neurons. It is theoretically possible that a given neuron may participate in a number of different groups, each supporting different internal models. Given the large number of synapses on the surface of neurons, this sharing could sustain a large number of internal models. As an alternative, the internal models could be conceived not as independent monads but as entities that can be combined into bigger assemblies when more demanding motor tasks are faced by the body. The study of the brain circuitry at the cellular level will undoubtedly provide new evidence on these issues in the near future.

THE NEURAL SUBSTRATE OF INTERNAL MODELS

In the previous section, we outlined the concept of an internal model for the dynamics of a moving limb. Now, we will describe the physiological evidence supporting the theory that the brain areas responsible for generating motor commands also serve as the sites for the storage and retrieval of motor memory. This linkage is consistent with the view that the brain circuit that has learned a task becomes the command center for expressing that task.

There are several examples of the intermingling between control functions and motor memory in the cortex. Brian Benda and his colleagues have reported some of the most direct evidence for the development of new patterns of activity in the cells of the motor area of the frontal lobe, an area named M1.[9] It should be pointed out that M1 is a key motor area, and damage to it profoundly disturbs the ability to produce voluntary movements. Benda's most striking result showed the gradual appearance of activity in cortical neurons of the M1 area in monkeys practicing arm movements against disturbing forces. These neurons displayed activity related to the production of forces that compensated for externally imposed interference. Remarkably, the same neurons were inactive before the application of the disturbance, but they remained active after the disturbance was removed. This effect is consistent with the hypothesis that neurons in M1 operate as memory elements. Similar results have been reported by Steven Wise, who used the technique of single-cell recordings but with different behavioral paradigms.[10]

John Martin and Claude Ghez reached similar conclusions with the use of a pharmacological ablation of M1.[11] They demonstrated that after the inactivation of this area, their experimental animals could not learn to correct the trajectories of perturbed limbs. In recent years, investigators have demonstrated learning in M1 with imaging techniques, namely positron emission tomography and functional magnetic resonance.[12]

In addition to the primary motor cortex (M1), other cortical areas of the frontal lobe, namely the premotor cortex and the supplementary motor areas, have been found to be involved in motor learning, either in conjunction with M1 or in isolation. In particular, the premotor cortex specializes in the learning and retention of visuo-motor tasks.[13] The supplementary motor areas seem predominantly concerned with sequence learning and conditional learning.[14]

Imaging studies have indicated that the prefrontal areas are also involved in motor learning. In particular, the experiments of Shadmehr and Henry Holcomb used the Brashers-Krug paradigm to gather evidence suggesting that the formation of an internal model of a perturbing force is associated with increased activity in the prefrontal cortex.[15] However, towards the completion of the learning task, recall of the learned internal model became correlated with increased blood flow in other cortical and subcortical areas, such as the premotor cortex and the cerebellum, and with decreased blood flow in the prefrontal cortex. A possible interpretation of this shift in blood flow is that the prefrontal cortex is a temporary storage area for sensory-motor associations.[16]

Taken together, these studies show that the motor cortical areas are linked to processes involved in motor learning—a result that implies that the circuitry of these areas may have the capacity to reorganize its functional properties. This capacity to reorganize depends upon the formation of new synapses. William Greenough reported that the dendritic branches of cortical neurons in M1 increase in number with motor training.[17] Presumably new synapses are formed on these branches. This possibility was confirmed by the recent report from Asanuma's group that electrical stimulation of the thalamus increases the density of synapses in the motor cortex.

MODULAR ORGANIZATION OF THE MOTOR SYSTEM

The evidence discussed in the previous sections suggests that the central nervous system is capable of representing the dynamic properties of limbs as well as the environment with which our limbs interact. Presumably, a representation is built upon some elementary building block, or "module," in the same way that sentences are composed of words. How is the representation accomplished? Recent electrophysiological studies of the spinal cord by Emilio Bizzi and his coworkers suggest how the CNS transforms the internal model into action.

The spinal cord is the final output stage of the motor system. Every muscle is innervated by motoneurons located in the ventral portion of the spinal gray matter. We may regard this system of motoneurons as comparable to the switchboard that drives a marionette. But there is more than this switchboard in the spinal cord. In addition to the motoneurons, the spinal gray matter contains a large population of nerve cells, called interneurons, whose functions are not yet fully understood. We know that these interneurons are capable of forming connections with motoneurons that innervate several different muscles.

In experiments performed by Bizzi, Mussa-Ivaldi, and Simon Giszter, the activity induced by the electrical stimulation of the spinal interneurons of the frog was found to spread to several groups of motoneurons.[18] This distribution of activity was not random but imposed a specific balance of muscle contractions. The mechanical outcome of the evoked synergistic contraction of multiple muscles was captured by a *force field*; the activation of a group of muscles generated a force that was recorded by a sensor at the endpoint of the limb. This force vector changed in amplitude and direction according to the position of the limb. Following stimulation of the spinal cord, the resulting force field converged toward a location in the reachable space of the limb—a stable equilibrium point. At this location, the force vanished and a small displacement of the endpoint in any direction caused a restoring force to appear. Thus this location acted as an attraction point for the limb in the same way as the bottom of a teacup is an attraction point for a marble rolling inside the cup. The analysis of the force field induced by the stimulation of the spinal interneurons revealed that such activation leads to a stable posture of the limb.

After the force field was identified, the stimulating electrodes were placed in different loci of the lumbar spinal cord, which activated a number of groups of leg muscles. After mapping most of the premotor regions in the lumbar cord, Bizzi, Giszter, and Mussa-Ivaldi reached the conclusion that there were at least four areas from which distinct types of convergent force fields were elicited.

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. When two separate sites in the spinal cord were simultaneously active, the resulting force was the sum of the forces induced by the separate activation of each site. This discovery led to a novel hypothesis for explaining movement and posture based on combining a few basic elements. The few force fields stored in the spinal cord may be viewed as representing motor primitives from which, through superposition, a vast number of movements can be formed by impulses conveyed by supraspinal pathways. According to this view, the supraspinal signals would establish the level of activation for each module. By means of mathematical modeling, Mussa-Ivaldi and Giszter along with subsequent work by Alexander Lukashin verified that this view of the generation of movement and posture is capable of accounting for a wide repertoire of motor behaviors.[19]

These 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 behaviors may be produced by superposition of the fields generated by concurrently active modules. We may therefore regard these force fields as independent elements of an internal model of dynamics. Recent simulation studies by Mussa-Ivaldi have demonstrated that by using this modular representation—that is, by adding convergent force fields—the central nervous system may learn to reproduce and control the dynamics of multijoint limbs in the context of a dynamic environment.

CONCLUSION

In this review, we have explained how the brain deals with the complex dynamics of our limbs as they interact with a variable and often unpredictable environment. We have shown that it is computationally difficult to produce the forces that drive our limbs because of the staggering number of mechanical variables involved in even the simplest movement. The kinematic redundancy of our motor system means that there are many different ways for the central nervous system to achieve an intended motor goal. While providing flexibility of motion, the redundancy creates for the CNS the difficult problem of deciding which plan of action to pursue.

We have furthermore proposed a theory based on internal models to explain how the CNS controls limb dynamics. Through repeated exposure to sensory signals coming from the moving limb during the acquisition of a motor task, there is a gradual change in the synaptic strength of the neurons of the motor areas. The outcome of this process is the formation of an internal model of limb dynamics. We have presented experimental evidence demonstrating that the formation of internal models as a means of acquiring motor skills is a more plausible hypothesis than those proposed in the past.

Finally, we have stressed the modular organization of the motor areas of the CNS. At the cortical level, we have demonstrated the tuning of small groups of neurons during motor learning; at the spinal cord level, our work has shown the existence of modules, which can be combined to produce different motor behaviors.

ENDNOTES

[1] Problems that are encountered in physics may be classified into well-posed and ill-posed problems. A well-posed problem is one for which there exists a unique solution that depends continuously upon the data. By contrast, ill-

posed problems may have either no exact solution or a multiplicity of solutions, and these solutions may change abruptly for certain values of the data.

[2] Charles Sherrington, "Flexion-reflex of the Limb, Crossed Extension Reflex and Reflex Stepping and Standing," *Journal of Physiology* 40 (1910): 28–121.

[3] John Hollerbach and Tamar Flash, "Dynamic Interactions between Limb Segments during Planar Arm Movements," *Biological Cybernetics* 44 (1982): 67–77.

[4] John Hollerbach, "A Recursive Formulation of Lagrangian Manipulator Dynamics," *IEEE Transactions on Systems, Man, and Cybernetics* SMC-10 (11) (1980): 730–736.

[5] Marc Raibert, "Analytical Equations versus Table Look-up for Manipulation: A Unifying Concept," in proceedings of the IEEE Conference on Decision and Control, New Orleans, 1977.

[6] Reza Shadmehr and Ferdinando Mussa-Ivaldi, "Adaptive Representation of Dynamics during Learning of a Motor Task," *Journal of Neuroscience* 14 (1994): 3208–3224.

[7] Francesca Gandolfo, Ferdinando Mussa-Ivaldi, and Emilio Bizzi, "Motor Learning by Field Approximation," *Proceedings of the National Academy of Sciences* 93 (1996): 3843–3846.

[8] Thomas Brashers-Krug, Reza Shadmehr, and Emilio Bizzi, "Consolidation in Human Motor Memory," *Nature* 382 (1996): 252–255.

[9] Brian Benda, Francesca Gandolfo, Chiang-Shan Li, Matthew Tresch, Daniel DiLorenzo, and Emilio Bizzi, "Neuronal Activities in M1 of a Macaque Monkey during Reaching Movements in a Viscous Force Field," *27th Annual Society of Neuroscience Abstracts* 23 (1997): 607.12.

[10] Steven Wise, S. Moody, K. Blomstrom, and A. Mist, "Changes in Motor Cortical Activity during Visuomotor Adaptation," *Experimental Brain Research* (forthcoming).

[11] John Martin and Claude Ghez, "Task-related Coding of Stimulus and Response in Cat Red Muscles," *Experimental Brain Research* 85 (1991): 373–388.

[12] S. Grafton, J. Mazziotta, S. Presty, K. Friston, R. Frackowiak, and M. Phelps, "Functional Anatomy of Human Procedural Learning Determined with Regional Cerebral Blood Flow and PET," *Journal of Neuroscience* 12 (7) (1992): 2542–2548. See also A. Karni, G. Meyer, P. Jezzard, M. Adams, R. Turner, and L. Ungerleider, "Functional MRI Evidence for Adult Motorcortex Plasticity during Motor Skill Learning," *Nature* 377 (1995): 155–158.

[13] Richard Passingham and Ulrich Halsband, "Premotor Cortex and the Conditions for Movement in Monkeys (Macaca fascicularis)," *Behavioral Brain Research* 18 (1985): 269–277. See also Michael Petrides, "Motor Conditional Associative-learning after Selective Prefrontal Lesions in the Monkey," *Behavioral Brain Research* 5 (1985): 407–413.

[14] Hajime Mushiake, Masahiko Inase, and Jun Tanji, "Neuronal Activity in the Primate Premotor, Supplementary, and Precentral Motor Cortex during Visually Guided and Internally Determined Sequential Movements," *Journal of Neurophysiology* 66 (1991): 705–718.

[15] Reza Shadmehr and Henry Holcomb, "Changing Activations in the Prefrontal Cortex during the Time-dependent Phases of Motor Memory Formation," *Journal of Neuroscience* (forthcoming).

[16] Joaquin M. Fuster, *The Prefrontal Cortex: Anatomy, Physiology, and Neuropsychology of the Frontal Lobe* (New York: Raven Press, 1989). See also Patricia S. Goldman-Rakic, "Topography of Cognition: Parallel Distributed Networks in Primate Association Cortex," *Annual Review of Neuroscience* 11 (1988): 137–156.

[17] William Greenough, "Structural Correlates of Information Storage in the Mammalian Brain: A Review and Hypothesis," *Trends in Neuroscience* 7 (1984): 229.

[18] Emilio Bizzi, Ferdinando Mussa-Ivaldi, and Simon Giszter, "Computations Underlying the Execution of Movement: A Biological Perspective," *Science* 253 (1991): 287–291.

[19] Ferdinando Mussa-Ivaldi and Simon Giszter, "Vector Field Approximation: A Computational Paradigm for Motor Control and Learning," *Biological Cybernetics* 67 (1992): 491–500; Alexander V. Lukashin, Bagrat R. Amirikian, and Apostolos Georgopoulos, "Neural Computations Underlying the Exertion of Force: A Model," *Biological Cybernetics* 74 (1996): 469–487.