

A computational neuroanatomy for motor control

Reza Shadmehr · John W. Krakauer

Received: 11 August 2007 / Accepted: 10 January 2008 / Published online: 5 February 2008
© Springer-Verlag 2008

Abstract The study of patients to infer normal brain function has a long tradition in neurology and psychology. More recently, the motor system has been subject to quantitative and computational characterization. The purpose of this review is to argue that the lesion approach and theoretical motor control can mutually inform each other. Specifically, one may identify distinct motor control processes from computational models and map them onto specific deficits in patients. Here we review some of the impairments in motor control, motor learning and higher-order motor control in patients with lesions of the corticospinal tract, the cerebellum, parietal cortex, the basal ganglia, and the medial temporal lobe. We attempt to explain some of these impairments in terms of computational ideas such as state estimation, optimization, prediction, cost, and reward. We suggest that a function of the cerebellum is system identification: to build internal models that predict sensory outcome of motor commands and correct motor commands through internal feedback. A function of the parietal cortex is state estimation: to integrate the predicted proprioceptive and visual outcomes with sensory feedback to form a belief about how the commands affected the states of the body and the environment. A function of basal ganglia is related to optimal

control: learning costs and rewards associated with sensory states and estimating the “cost-to-go” during execution of a motor task. Finally, functions of the primary and the premotor cortices are related to implementing the optimal control policy by transforming beliefs about proprioceptive and visual states, respectively, into motor commands.

Keywords Optimal control · Computational models · Reaching · Cerebellum · Basal ganglia · Motor cortex · Parietal cortex

Introduction

Over the last 25 years, a large body of experimental and theoretical work has been directed towards understanding the computational basis of motor control, particularly visually guided reaching. Roboticians and engineers largely initiated this work, galvanized by the obvious discrepancy between the clumsy movements of robots and the ease with which healthy people move their limbs. Their aim was to derive from first principles some of the strikingly stereotypical features of movements observed in people and other primates. That is, they aimed to understand why we move the way that we do. For example, the theories began to explain why in reaching to pick up a cup, or in moving the eyes to look at an object, there was such consistency in the detailed trajectory of the hand and the eyes. In many ways, the approach was reminiscent of physics and its earliest attempts to explain regularity in motion of celestial objects. Except here, the regularity was in our movements and the search was for normative laws that governed our behavior (for review see Kording 2007).

The computational models relied on empirical data generated by observation and experiment. Experiments

R. Shadmehr (✉)
Laboratory for Computational Motor Control,
Department of Biomedical Engineering,
Johns Hopkins School of Medicine, 410 Traylor Building,
720 Rutland Ave., Baltimore, MD 21205, USA
e-mail: shadmehr@jhu.edu

J. W. Krakauer
The Motor Performance Laboratory,
Department of Neurology, Columbia University College
of Physicians and Surgeons, New York, NY 10003, USA

differed from observation of natural behavior in that they perturbed the behavior under study. Perturbations led to errors. The nature of these errors and the way that they were corrected (adaptation) provided insight into the computational organization of the motor system. However, the language that was chosen to represent these results consisted of equations that were independent of the properties of the hardware needed to implement them (i.e., neurons). More often than not there was not even a mention of the brain region where these computations might be occurring. In David Marr's (1982) terminology, the scientists were working on a *computational problem*: inquiring about the driving purpose and logic of actions.

This approach stood in contrast to models that one occasionally encounters in neuroscience, where the experiments and mathematics concern specific neuronal machinery. An example of this approach would be the extensive work on the cerebellar circuitry underlying eye-blink conditioning (Ohyama et al. 2003), the vestibulo-ocular reflex (Raymond et al. 1996), or reaching (Kawato and Gomi 1992). Marr called these levels of analysis the *algorithmic problem*, inquiring about how a computation is represented in an input–output form and implemented via an algorithm, and the *implementation problem*, the physical realization of the algorithm.

There is a perspective in neurology that shares characteristics of the computational approach in motor control. In neurology, one begins with the assumption that complete behaviors have a functional architecture, and its organization is revealed when the architecture is damaged by disease or accident. That is to say, the brain's more basic modules or subsystems can be inferred through the effects of focal lesions. In essence, lesions can make humans as clumsy as robots. However, careful observation reveals that the clumsiness has structure and can come in different forms, each carrying a signature that one can use to infer the locus of damage. Unfortunately, inferring normal function from the effects of lesions is difficult for many reasons but one is that neurologists tend to equate lesions of the motor system with deficits in execution rather than seeing the behavior as a possible form of compensation for the change in the state of the nervous system.

One way to evaluate the success of the computational models is to ask whether they help infer normal function of a damaged brain region from the behavioral effects of lesions. We will attempt to link computational theories of the motor system with findings in patients with neurological disease. This review will be organized into the following sections: (1) a computational framework for motor control, (2) the problem of inferring function from lesions, (3) computational steps involved in making a movement, and (4) effects of focal and distributed lesions on these specific computational steps.

The current state of computational motor control (and how we got here)

In 1954, Fitts published a short paper in which he reported that there were regularities in how people planned their movements (Fitts 1954). He asked volunteers to move a pen from one "goal region" to another as fast and accurately as they could. He found that the movement durations grew logarithmically as a function of the distance between the goals (Fig. 1a). This relationship was modulated by two factors. One factor was the size of the goal region. As the goal region became smaller, movements slowed down. Therefore, the speed at which people planned their movements depended on the accuracy requirements of the task. A second factor was the mass of the pen. People slowed their movements when they had to move a heavier pen. This suggested that it was harder to maintain accuracy when one moved a heavier object (it is remarkable that the heaviest object in that paper was less than a pound, yet produced robust changes in behavior). Together, the data suggested that movement planning was affected by accuracy constraints of the task as well as the forces that were required to perform the task.

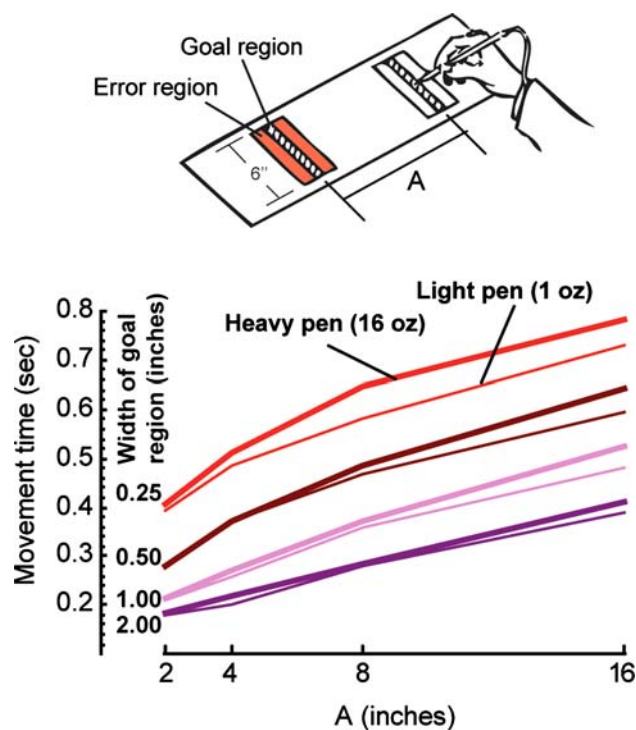


Fig. 1 Accuracy constraints affect control of reaching. Volunteers were instructed to tap the two goal regions with a pen as many times as possible during a 15 s period. Movement time increased as the accuracy requirements increased (width of target region decreased), and as weight of the hand-held pen increased (figure constructed from data in Fitts 1954)

In 1981, Morasso closely examined the trajectory of reaching movements and noted further regularities: the hand moved in roughly a straight line with a bell-shaped speed profile, while the various joints moved in more complicated patterns (Morasso 1981). This was surprising because muscles controlled the joints, and it was unclear why control should be complex at the joint level in order to produce regularity in the hand path. Particularly strong regularities were observed in saccadic eye movements. Bahill et al. (1975) and Collewijn et al. (1988) reported that people moved their eyes to targets of various eccentricities with trajectories that had highly reproducible timing and speed profiles. Was there a unified explanation for these regularities?

To answer this question, let us consider the simple reaching task that Fitts performed in some detail. Suppose you are a subject and are handed a pen. Because the target box is surrounded by two penalty regions, it seems rational to aim for the center of the target box. What if the penalty region was only on one side? Now you should aim for a point farther away from the penalty region and not at the center of the target box. This is because movements have variability, and you will maximize your performance (in terms of sum of hits and misses) if you take into account this variability. Trommershauser et al. (2005) demonstrated this in an experiment illustrated in Fig. 2a. Costs for the penalty region, as well as the variance of the tool affected where subjects aimed their movements. Importantly, if the tool increased its variability (through artificial means), people aimed farther away from the penalty region. Perhaps this variability also explains the speed of movements in Fitt's experiment and its sensitivity to pen weight: rapid movements are more variable than slow movements, so one should slow down if there is a need to be accurate. Moving heavier objects tends to increase movement variability, again requiring a reduced speed to maintain accuracy. Therefore, movement planning takes into account movement variability because variability affects accuracy, which in turn affects our ability to acquire reward.

Harris and Wolpert (1998) began formalizing these ideas by linking variability and movement planning. They noted that larger motor commands required larger neural activity, which in turn produced larger variability due to a noise process that grew with the mean of the signal. Therefore, motor commands carried an accuracy cost because the larger the command, the larger the standard deviation of the noise that rides on top of the force produced by the muscles (Jones et al. 2002). Noise makes movements inaccurate. The link between size of motor commands and accuracy of movements was crucial because in principle it could account for the data that Fitts and Morasso had observed in reaching and some of the stereotypical behaviors that one observes in saccades (Harris and Wolpert 2006).

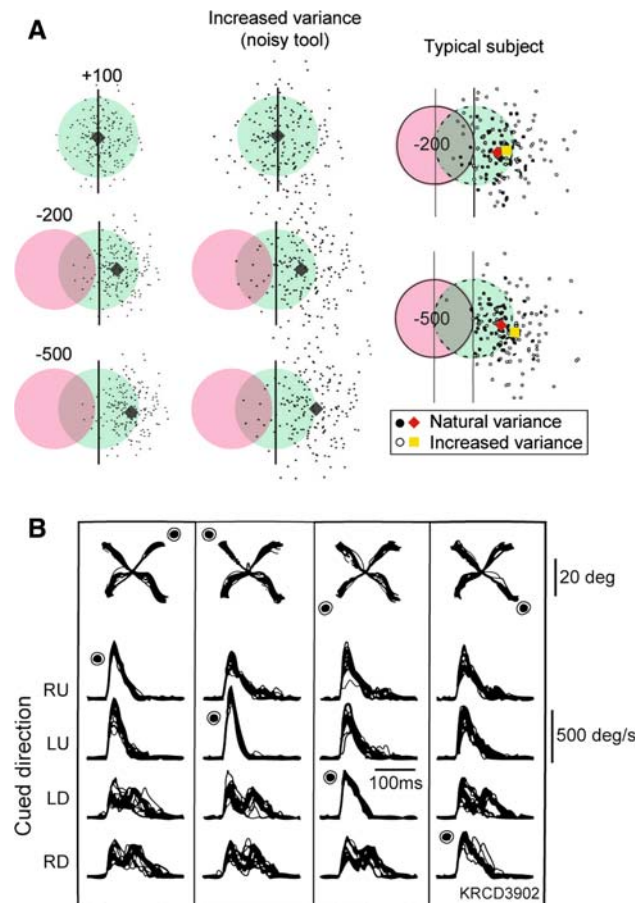


Fig. 2 Reward affects control of movements. **a** The task is to make a rapid pointing movement so to maximize reward. Endpoints in the blue region are rewarded while those in the red region are penalized. Because movements are variable, subjects should plan their movements so that with increased error costs mean of the endpoint distributions shifts away from the penalty region (pink region). If the tool that they are holding increases their endpoint noise, the shift should increase. The left two columns are predicted performances with small and large noise. The right column is data from a typical subject (from Trommershauser et al. 2005). **b** The task for the monkey is to saccade to a visual target that can appear in one of four locations (LU left upper, RD right down, etc.). However, only one target location in a given set is rewarded. Each set is identified by a column. The top row shows indicates the rewarded target location in each set (filled circle). The bottom four rows show saccade speed to each target location under each reward condition. When the target is rewarded, saccades to that location have a higher speed, smaller duration, and less variability (from Takikawa et al. 2002)

In a sense, the theory re-stated the purpose of these tasks using language of mathematics: be as fast as possible, while trying to be as accurate as the requirements imposed by the task. However, by doing so, it forced the theorists to think how one would actually achieve this optimality. Certainly, the solution to the problem could not be “hard wired”: First, costs and rewards of tasks are not constant. Take the simple saccade task where an animal is given more reward for certain visual targets and less for others. Hikosaka and colleagues (Takikawa et al. 2002) examined

eye trajectories when a monkey was asked to make saccades to various target locations (Fig. 2b). They noted that peak speeds tended to be higher and less variable when saccades were to rewarded target locations. Therefore, as expected rewards of the task change, movement planning responds to these changes. Second, the brain alters movement planning as the dynamics of the body or a tool change (Figs. 1, 2a). Even if one considers only the motor plant itself, the dynamics of the plant change over multiple timescales: in the long timescales of development and aging, limbs grow and muscles strengthen, and then muscles age and some motor neurons die. In the short timescale, muscles and neurons fatigue with repeated use and recover with the passage of time. Therefore, the nervous system cannot rely on a motor plant that is time-invariant. Rather, it seems more reasonable that the nervous system should monitor these changes and form an *internal model* of the plant and/or the tool (Shadmehr and Mussa-Ivaldi 1994). Indeed, maintaining performance in something as simple as a saccade or a reach probably requires constant adjustment of this internal model (Smith et al. 2006; Kording et al. 2007).

What is the nature of this adjustment? Todorov and Jordan (2002) recognized that a key component of the optimization problem was presence of feedback. One type of feedback is from sensory receptors that monitor the state of the body and the world. Another type of feedback is from internal models that monitor the motor output and predict their sensory consequences, effectively providing a form of internal feedback. Internal predictions can be made long before sensory feedback, making some very rapid movements like saccades depend entirely on internal feedback. However, for longer movements the two kinds of information would need to be combined to form a belief about the state of the body. Todorov and Jordan (2002) suggested that a more appropriate mathematical approach was to first describe the constraints of the task in terms of a function that included explicit terms for gains and losses, and then maximize that function in the framework of *feedback control*. This new formulation was a breakthrough because it formally linked motor costs, expected rewards, noise, sensory feedback, and internal models into a single, coherent mathematical framework.

We have summarized this framework in Fig. 3a. The theory explains that in order to make a movement, our brain needs to solve three kinds of problems: we need to be able to accurately predict the sensory consequences of our motor commands (this is called *system identification*), we need to combine these predictions with actual sensory feedback to form a belief about the state of our body and the world (called *state estimation*), and then given this belief about the state of our body and the world, we have to adjust the gains of the sensorimotor feedback loops so that

our movements maximize some measure of performance (called *optimal control*).

At the heart of the approach is the idea that we make movements to achieve a rewarding state. This crucial description of why we are making a movement, i.e., the rewards we expect to get and the costs we expect to pay, determines how quickly we move, what trajectory we choose to execute, and how we will respond to sensory feedback.

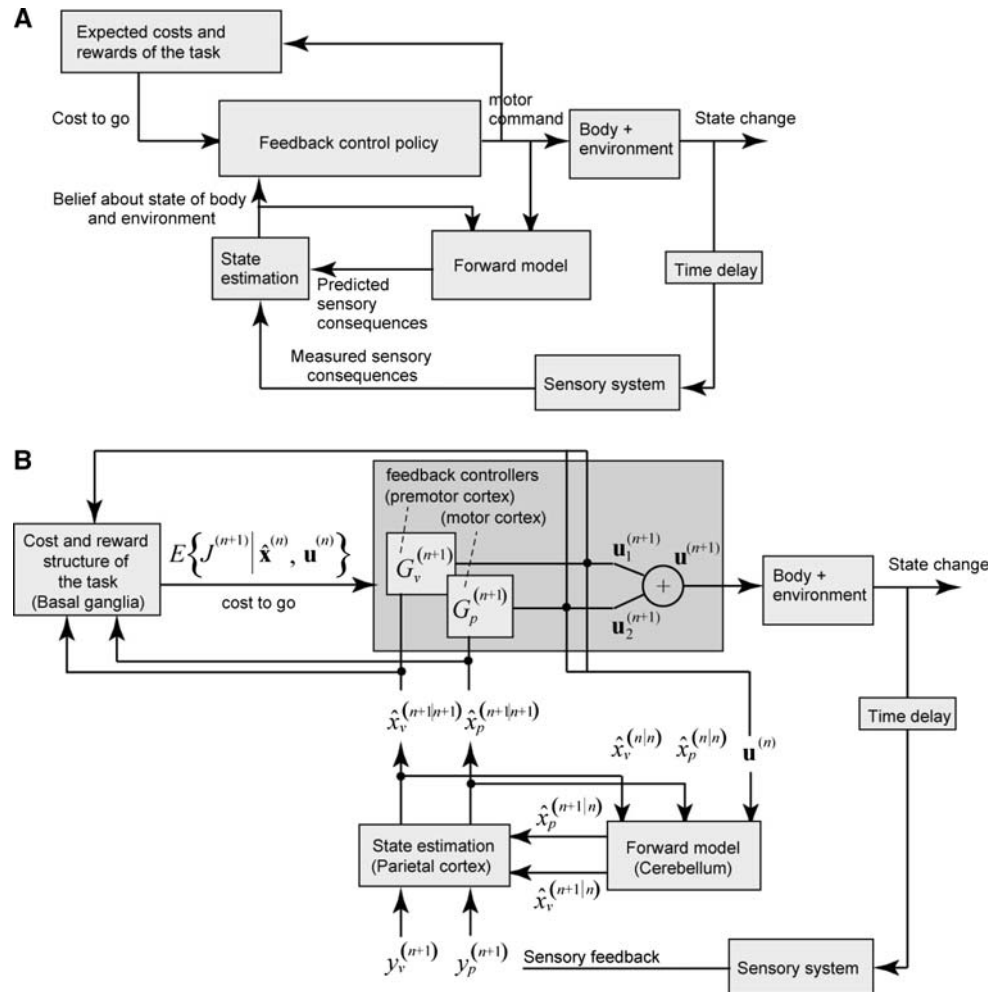
Here, we will use this framework to examine a number of disorders in movement control. We will suggest a specific computational neuroanatomy of the motor system (Fig. 3b). In this framework, the basal ganglia help form the expected costs of the motor commands and the expected rewards of the predicted sensory states. The cerebellum plays the role of predicting the sensory consequences of motor commands, i.e., the expected changes in proprioceptive and visual feedback. The parietal cortex combines the expected sensory feedback with the actual sensory feedback, computing a belief about the current proprioceptive and visual states. Given the motor costs and expected rewards of the sensory states, the premotor and the primary motor cortex assign “feedback gains” to the visual and proprioceptive states, respectively, resulting in sensorimotor maps that transform the internal belief about states into motor commands.

The computational problem in reaching

Let us use the well-studied reach adaptation paradigm to formulate the motor control problem in the framework outlined in Fig. 3. In this framework, our problem can be posed as a series of questions: what are the costs and rewards of the task, i.e., why should a person volunteer to do this task? What is the “best” way, i.e., the best set of motor commands, to maximize the rewards and minimize the costs? Finally, what kinds of computations are required to achieve this goal?

Let us begin with the costs and rewards of a reaching task. For simplicity, allow us to focus on the version that we typically perform in a laboratory. Suppose we are instructed to hold the handle of a tool and move the tool so that a cursor displayed on a monitor arrives at a target. If we accomplish this in a specific time period, we are provided a monetary reward, juice, an “explosion” on the video monitor, or simply an encouraging comment from the experimenter. We can sense the position of the cursor y_v and the target r via vision, and position of our arm y_p via proprioception. Through experience in the task we learn that the objective is to minimize the quantity

Fig. 3 A schematic model for generating goal directed movements. Please see the text for explanation of variables and box labels



$$(\mathbf{y}_v^{(t)} - \mathbf{r})^T (\mathbf{y}_v^{(t)} - \mathbf{r})$$

at time $t = N$ after the reach starts (e.g., this is the time that the movement is rewarded if the cursor is in the target). Superscript T is the transpose operator. To denote the fact that this cost is zero except for time N , we write it as

$$\sum_{t=1}^N (\mathbf{y}_v^{(t)} - \mathbf{r})^T Q^{(t)} (\mathbf{y}_v^{(t)} - \mathbf{r})$$

where the matrix Q is a measure of our cost at each time step (which may be zero except at time N). That is, matrix Q specifies how important it may be for us to put the cursor in the target. If we value the reward, then we set this variable to be large. There is also a cost associated with motor commands \mathbf{u} , which here we assume to grow as a quadratic function. This cost may reflect a desire to be as frugal as possible with our energy expenditure, or it may reflect the fact that the larger the motor commands, the larger the noise in the forces that are produced by the muscles, resulting in variability. This variability increases the difficulty in

controlling the movement. As a result, we want to produce the smallest amount of motor commands possible. Now the total cost becomes:

$$J = \sum_{t=1}^N (\mathbf{y}_v^{(t)} - \mathbf{r})^T Q^{(t)} (\mathbf{y}_v^{(t)} - \mathbf{r}) + \mathbf{u}^{(t)T} L \mathbf{u}^{(t)} \tag{1}$$

where matrix L is a measure of the costs associated with the motor commands. The relative weight of Q and L is an internal measure of expected value of achieving the goal vs. expected motor costs.

To be successful in this task (consistently arrive at the target in time), we need to find the motor commands that on the one hand, are as small as possible, and on the other hand, are large enough to get the cursor to the target in time. To do so, we need some way to relate motor commands to their outcomes. This is called an internal model. For example, through observation, we learn that moving the tool moves the cursor on the screen. In particular, motor commands $\mathbf{u}^{(t)}$ are expected to produce proprioceptive and visual feedback $\hat{\mathbf{y}}^{(t)} = [\hat{\mathbf{y}}_v^{(t)}, \hat{\mathbf{y}}_p^{(t)}]$. These are the expected sensory consequences of our action. Simplifying the

problem, here we write this “internal model” of the dynamics as a linear function of motor commands:

$$\begin{aligned} \hat{\mathbf{x}}^{(t+1|t)} &= \hat{A}\hat{\mathbf{x}}^{(t)} + \hat{B}\mathbf{u}^{(t)} \\ \hat{\mathbf{y}}^{(t)} &= \hat{H}\hat{\mathbf{x}}^{(t)} \end{aligned} \tag{2}$$

where $\hat{\mathbf{x}}^{(t)}$ represents the predicted state (of our body and the world) at time t given the sensory feedback up until that time, H is a transformation of those states to expected sensory feedback $\hat{\mathbf{y}}^{(t)}$ (i.e., proprioception and vision), and $\hat{\mathbf{x}}^{(t+1|t)}$ is predicted state at time $t + 1$ given the state and motor command at time t .

Equation (2) describes an internal model of the dynamical system that we are trying to control. The actual dynamics of that system may be more complicated. For example, the motor commands may carry signal-dependent noise. To solve the optimization problem, it is convenient to represent the stochastic variable representing signal dependent noise in the motor commands $\epsilon_u^{(t)}$ as a linear combination of Gaussian noise stochastic variables $\phi_i \sim N(0, 1)$ with zero mean and variance one (Todorov 2005):

$$\begin{aligned} \epsilon_u^{(t)} &\equiv \begin{bmatrix} c_1 u_1^{(t)} \phi_1^{(t)} & 0 & 0 \\ 0 & c_2 u_2^{(t)} \phi_2^{(t)} & 0 \\ 0 & 0 & \ddots \end{bmatrix} \\ C_1 &\equiv \begin{bmatrix} c_1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \ddots \end{bmatrix} \quad C_2 \equiv \begin{bmatrix} 0 & 0 & 0 \\ 0 & c_2 & 0 \\ 0 & 0 & \ddots \end{bmatrix} \\ \epsilon_u^{(k)} &= \sum_i C_i \mathbf{u} \phi_i^{(k)} \end{aligned}$$

And so we see that the standard deviation of the noise in the motor commands grows linearly with the motor commands with a slope of c , and therefore the variance grows with the “square” of the motor commands:

$$\text{var}[\epsilon_u^{(t)}] = \sum_i C_i \mathbf{u}^{(t)} \text{var}[\phi_i^{(t)}] \mathbf{u}^{(t)T} C_i^T = \sum_i C_i \mathbf{u}^{(t)} \mathbf{u}^{(t)T} C_i^T$$

In general, there may be similar signal-dependent noises on our sensory system, as well as Gaussian type noises on the motor and sensory systems. In sum, a reasonable representation of the stochastic system that we are trying to control might be written as

$$\begin{aligned} \mathbf{x}^{(t+1)} &= A\mathbf{x}^{(t)} + B\left(\mathbf{u}^{(t)} + \sum_i C_i \mathbf{u}^{(t)} \phi_i^{(t)}\right) + \epsilon_x^{(t)} \\ \mathbf{y}^{(t)} &= H\left(\mathbf{x}^{(t)} + \sum_i D_i \mathbf{x}^{(t)} \mu_i^{(t)}\right) + \epsilon_y^{(t)} \end{aligned} \tag{3}$$

$\mu_i, \phi_i \sim N(0, 1) \quad \epsilon_x \sim N(\mathbf{0}, Q_x) \quad \epsilon_y \sim N(\mathbf{0}, Q_y)$

As motor commands are generated, we receive a continuous stream of sensory feedback \mathbf{y} . We combine

the predicted sensory feedback with the observed quantities to form a *belief* about states:

$$\hat{\mathbf{x}}^{(t+1|t+1)} = \hat{\mathbf{x}}^{(t+1|t)} + K^{(t+1)}\left(\mathbf{y}^{(t+1)} - \hat{\mathbf{y}}^{(t+1)}\right) \tag{4}$$

In this equation, the term $\hat{\mathbf{x}}^{(t+1|t+1)}$ is the belief state at time $t + 1$, given that one has acquired sensory information at that time. K is a mixing gain (or a Kalman gain) that determines how much we should change our belief based on the difference between what we predicted and what we observed. Therefore, Eq. (2) describes how we make predictions about sensory feedback, and Eq. (4) describes how we combine the actual sensory observations with predictions to update beliefs about states.

Our task is to perform the movement in a way that maximizes our chances for reward. If Eq. (2) is an accurate model of how motor commands produce changes in the states, then we can use it as a set of constraints with which to minimize Eq. (1). (Equation 2 is called a “constraint” because our solution must lie on the manifold specified by this equation, i.e., our belief about the relationship between the two variables in our cost, \mathbf{y} and \mathbf{u} .) Because there is noise in our system, the cost J in Eq. 1 is a stochastic variable. At each time point during a movement, the best that we can do is minimize the expected value of this cost, given the state that we believe to be in, and the motor commands that we have produced: $E\{J^{(t)} | \hat{\mathbf{x}}^{(t-1)}, \mathbf{u}^{(t-1)}\}$. The term $E\{J^{(t)}\}$ reflects the expected value of the “cost to go”, i.e., the total cost remaining in the current trial. Technically, the cost is minimized by starting at the time when the movement is supposed to end N and then walking back in time toward the current time t . The result is a feedback control law that specifies the motor commands that we should produce at this moment in time as a function of our belief about proprioceptive and visual states:

$$\begin{aligned} \mathbf{u}^{(t)} &= -G^{(t)}\hat{\mathbf{x}}^{(t|t-1)} \\ &= -G_p^{(t)}\hat{\mathbf{x}}_p^{(t|t-1)} - G_v^{(t)}\hat{\mathbf{x}}_v^{(t|t-1)} \end{aligned} \tag{5}$$

The new variable G is a matrix that changes with time during a movement. It tells us how at time t we can transform beliefs in sensory states (in terms of proprioception and vision) into motor commands so that we maximize performance in the remaining task time. The G matrices are basically feedback gains that we compute from our knowledge about the costs/rewards of the task (cost to go) and our knowledge about the dynamics of the task (the internal model). The time-sequence of the feedback gains G is called a control policy. For completeness, we provide the recursive equations that compute these gains:

$$\begin{aligned}
 W_x^{(N)} &= Q^{(N)} \\
 W_e^{(N)} &= 0 \\
 W_x^{(t)} &= Q^{(t)} + A^T W_x^{(t+1)} A - A^T W_x^{(t+1)} B G^{(t)} \\
 &\quad + \sum_i D_i^T H^T K^{(t)T} A^T W_e^{(t+1)} A K^{(t)} H D_i \\
 W_e^{(t)} &= A^T W_x^{(t+1)} B G^{(t)} + (A - A K^{(t)} H)^T W_e^{(t+1)} \\
 &\quad \times (A - A K^{(t)} H) \\
 G^{(t)} &= (L^{(t)} + \sum_i C_i^T B^T W_x^{(t+1)} B C_i + \sum_i C_i^T B^T W_e^{(t+1)} B C_i \\
 &\quad + B^T W_x^{(t+1)} B)^{-1} B^T W_x^{(t+1)} A
 \end{aligned}$$

The W matrices are intermediate variables that are used for defining the feedback gain G . The proof that this algorithm minimizes the cost is provided by Todorov (2005). A tutorial and step-by-step derivation is available as lecture notes on one of the author’s web page (RS). The point to note is that all of the noises, as well as the parameters of the forward model and the Kalman gain, affect the feedback gain. Therefore, the solution is strongly dependent on the specific plant that one is trying to control.

Some examples

Let us consider some examples to illustrate this framework. A first example is a simple task described by Uno et al. (1989) and shown in Fig. 4a. The objective is to reach from point T1 to T2. In one condition, the subject is holding a lightweight tool that moves freely in air. In a second condition, the tool is attached to a spring that pulls the hand to the right. Without the spring, people reach in a straight line. This is the path that minimizes the cost. However, once the spring is attached, the straight path incurs substantially more motor costs than a curved path. The curved path is the one that subjects choose (Uno et al. 1989).

In our second example, the task is to move one’s hand from one point to another in a given amount of time (450 ms), but now instead of a spring, there is a velocity dependent field that pushes the hand perpendicular to its direction of motion (Shadmehr and Mussa-Ivaldi 1994). Before the field is imposed, the motion that minimizes the cost (and maximizes probability of reward) is simply a straight line with a bell-shaped velocity profile. However, when the field is imposed, the solution is no longer a straight line (Izawa et al. 2008). For example, if the field pushes the hand to the left, the policy that produces the least cost in terms of Eq. (1) is one that moves the hand slightly to the right of a straight line, resulting in a curved movement that appears to over-compensate for the forces (Fig. 4b). As subjects train, their hand paths converge to

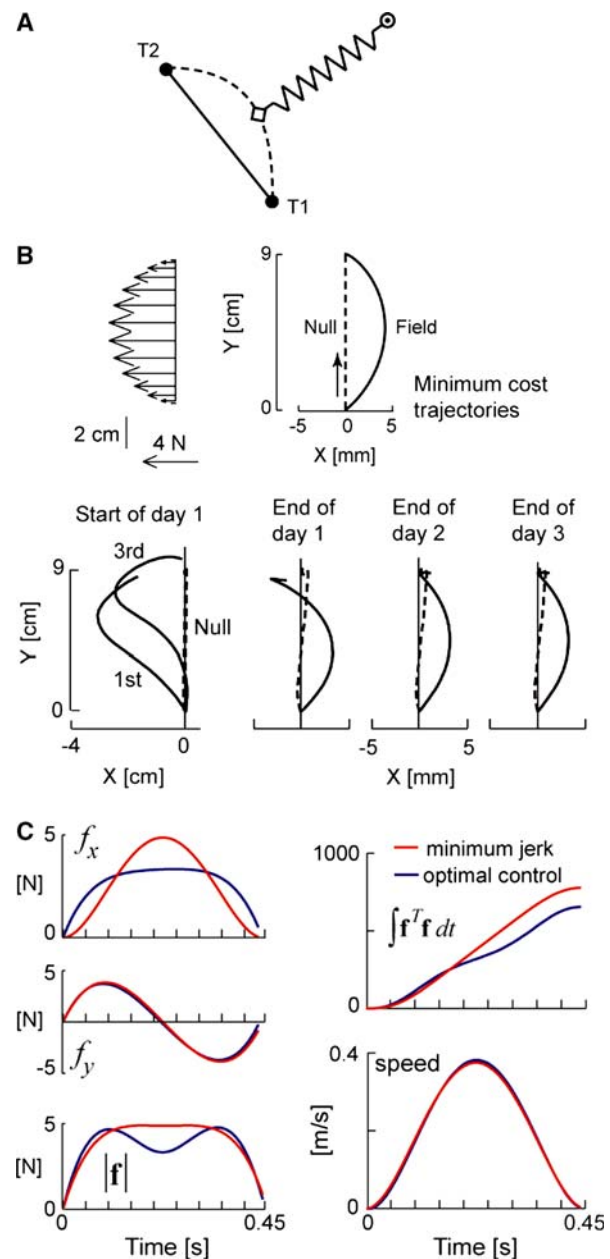


Fig. 4 Task dynamics affect reach trajectories. **a** The task is to reach from point T1 to T2. In one condition, the reach takes place in free space (*straight line*). In another condition, a spring is attached to the hand. In this case, the subject chooses to move the hand along an arc (redrawn from Uno et al. 1989). **b** A velocity dependent force field pushes the hand perpendicular to its direction of motion. For example, for an upward movement the forces push the hand to the left. The motion that minimizes cost of Eq. (1) is not a straight line, but one that has a curvature to the right. The data shows hand paths for a typical subject at start of training on day 1, and then end of training each day. Except for the first and third trials, all other trajectories are average of 50 trials. **c** A rationale for why a curved movement is of lower cost. The *curves* show simulation results on forces that the controller produces, and speed of movement, in the optimal control scenario of Eq. (1) and in a scenario where the objective is to minimize jerk (data in **b** and **c** from Izawa et al. 2008)

this curved trajectory. To see the rationale for this behavior, Fig. 4c plots the forces produced by the optimal controller and compares it to forces that must be produced if a mass is moving along a “minimum-jerk” trajectory. By moving the hand along a curved path, the optimal controller produces less total force: It over-compensates early into the movement when the field is weak, but under-compensates at peak speed when the field is strongest. Therefore, the curved path actually produces less total force than a straight trajectory. People produce similarly curved trajectories when they move in such fields (Thoroughman and Shadmehr 2000).

These examples demonstrate that kinematic costs like maximizing smoothness (Flash and Hogan 1985) are limited in their ability to explain the diversity of motor control because these costs are invariant to the task’s dynamics: in the above examples the prediction would be to keep moving along a straight line. However, we do not need optimal feedback control (OFC) to explain the above two examples. Any cost structure that penalizes motor commands would probably be sufficient and can be used in an open-loop scheme in which a sequence of motor commands are pre-programmed and produce a curved “desired trajectory” (Thoroughman et al. 2007), as in minimum torque-change (Uno et al. 1989) and minimum end-point variance schemes (Harris and Wolpert 1998). So what do we gain with the OFC framework?

All movements that we make include some sensory states that are crucial to the success of the task, some sensory states that are less important, and some sensory states that are practically irrelevant. OFC suggests that movements are planned as a sequence of feedback gains, and these gains take into account how the states or combination of states contribute to the cost. In effect, the gains define how at any point in time, each state or combination of states matters. Consider a condition in which both arms are involved in performing a task. In this experiment (Diedrichsen 2007), the objective is to move a cursor to a target position (Fig. 5, right column). However, the position of the cursor is a weighted sum of the left and the right hand: $1/2(\mathbf{x}_L + \mathbf{x}_R)$. If target position is denoted by \mathbf{r} , then the cost to go can be written as

$$J^{(t)} = \sum_{i=t}^N \left(\mathbf{r} - \frac{\mathbf{x}_L^{(i)} + \mathbf{x}_R^{(i)}}{2} \right)^T Q^{(t)} \left(\mathbf{r} - \frac{\mathbf{x}_L^{(i)} + \mathbf{x}_R^{(i)}}{2} \right) + \mathbf{u}^{(i)T} L \mathbf{u}^{(i)} \tag{6}$$

where the vector \mathbf{u} now includes the motor commands to the left and the right arms. The idea is that motion of the cursor can be influenced by motor commands of either arm, and so the arms should cooperate to control the cursor in order to minimize the cost. Intuitively, the cooperation means that the motor commands that move the right arm

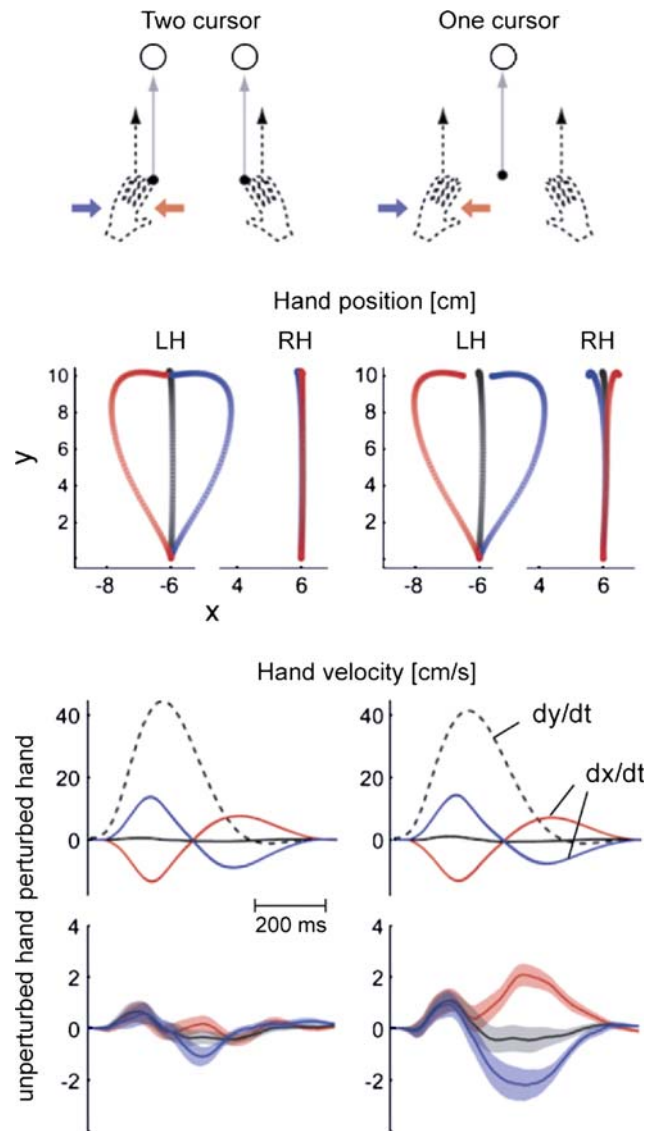


Fig. 5 Task constraints affect feedback response to perturbations. In this bi-manual task, there are either two cursors visible on a screen or a single cursor. In the two-cursor condition, each hand controls one cursor. In the one-cursor condition, the average position of each hand is reflected in cursor position. In the *top row* of this figure, the *blue* and *red arrows* show that a perturbation is applied to the left hand. The *middle row* shows the hand paths in each condition (*black trace* is the condition without a perturbation). The *bottom row* shows the hand velocities. In the two-cursor condition, perturbation to the left hand is compensated by the left hand only. In the one-cursor condition, the same perturbation is compensated by motion of both hands (from Diedrichsen 2007)

should depend not only on the state of the right arm, but also on the state of the left arm. Mathematically, this means that when we minimize Eq. (6) with respect to \mathbf{u} , we find that the resulting G_v in Eq. (5) has off-diagonal terms that produce motor commands in the left arm that depend on the visual state of both the left and the right arms. As a result, when the left arm is perturbed during the reach, the

feedback gain will produce a response both in the left and the right arms.

The experimental results in Fig. 5 (right column) show that in the one-cursor condition, perturbation of the left hand produces a compensatory response in both the left and the right arms. Indeed, simple variability due to noise in one arm affects the motor commands in the other arm. In contrast, when the task involves moving two cursors to two distinct targets, the states of the cursors are now independently controlled by the motor commands that move each arm. The motor command to the right arm can no longer influence the left cursor, and therefore the feedback gains become decoupled. The data (Fig. 5, left column) shows that a perturbation to the left hand now produces a much smaller response from the contralateral arm. Therefore, OFC predicts that in the one cursor condition, the state of one arm is relevant to the motor commands that move the other arm. In the two-cursor conditions, that same state is irrelevant.

This idea that the nervous system controls states that are relevant to the task is called an “uncontrolled manifold hypothesis” (Scholz and Shoner 1999). The idea has been used to account for complex movements that offer clear redundancies, like bimanual pointing (Domkin et al. 2005) and Frisbee throwing (Yang and Scholz 2005). OFC basically states in mathematics what some of these ideas have been implying.

To summarize, the theory states that:

1. At any given time there are many actions possible. For any possible action, we need to know the costs that are associated with it as well as the rewarding nature of the sensory states that it might achieve (Eq. 1).
2. Once we choose to acquire a rewarding state, we need to know how our motor commands produce changes in things that we can observe (Eq. 2). The problem of predicting the consequences of our actions is called system identification.
3. In order to be optimal, we need to produce those motor commands that minimize the costs and maximize the reward (Todorov and Jordan 2002; Todorov 2005). This is the constrained minimization problem in optimal control (minimize Eq. 1 under the constraints of Eq. 2). The result of the minimization is a feedback control policy that specifies the feedback gain that we should apply to sensory states (Eq. 5).
4. As we generate motor commands, we make predictions about the sensory consequences (Eq. 2). When actual sensory feedback arrives, we integrate our observations with our predictions and form a belief about how our motor commands have affected the state of our body and the world around us (Eq. 4) (Wolpert et al. 1995; Vaziri et al. 2006). This is called state estimation.

Lesions and separable brain processes

There is an extensive literature on the complexities of inferring normal function from patient impairments (see Shallice 1988, for review). This topic can only be touched on briefly here, but it is important to note the principal problems in order to avoid the most blatant kinds of error when applying results from patient studies to test theories of normal motor control.

One pitfall is localization: inferring the normal function of a region from the effect that damage has had on behavior. Such inference is problematic because abnormal behavior after a focal lesion, especially when studied at a delay after the insult, is the sum of the direct effect of the lesion plus its effect on regions connected to the damaged region. In addition, the implicit assumption that the patient is using the same brain subsystem as before, just less efficiently or effectively, may be wrong. From an optimization standpoint it could be that after injury a qualitatively distinct subsystem is used to achieve the task goal (compensation). If one part of the brain performs system identification and another performs feedback optimization, then damage to the part that performed system identification might still allow the optimization process to proceed, resulting in *optimization in response to the damage*. For example, in cerebellar patients, reaching slowly may be an optimized response for an inability to predict sensory consequences of movements, which is critical for control during fast movements.

A second concern is disagreement over whether single case reports or averaged group effects yield the most informative results in patient studies. Some psychologists have argued that the only valid form of neuropsychological data comes from single-case studies (Caramazza 1986). The argument is that averaging artifacts may lead to a group result not shown by any individual subject. Although this is a potential concern, certain kinds of group results may be more resistant to this problem and, more importantly, group data are indispensable when parametric relationships need to be explored and learning effects preclude testing more than one value of a variable of interest in any given individual.

A third issue is the question of what form of evidence is the best proof of separate brain subsystems. The simplest kind of evidence for damage to a particular subsystem is to show impairment in a specific task, for example a reaching task. However, as alluded to earlier, the abnormal behavior may not reflect suboptimal functioning of the system used by healthy subjects, but might indicate instead re-optimization in response to a functional loss. A second kind of evidence is showing dissociation—the patient performs normally or close to normal on Task A but below normal on Task B. The problem here is that the same brain

subsystem might be involved in both tasks, but task B is more difficult and so damage makes patients more likely to perform poorly on task B than task A. Conclusions are on firmest ground when a strong double dissociation can be demonstrated. In a double dissociation, one patient performs normally on task A and abnormally on task B, whereas for the second patient the opposite is true.

To better understand the idea of a double dissociation we can consider an example from outside of motor control; the words-and-rules theory from linguistics (Pinker and Ullman 2002). This theory posits that in English, irregular past-tense forms are stored in the lexicon, a division of declarative memory system, whereas the regular past-tense forms are computed with a grammatical processing rule (addition of “-ed”). Finding a double dissociation in patients would lend support to this theory and indeed that is what has been found. A patient with anomia fluent aphasia was able to form the regular past tense but not the irregular, whereas a patient with a grammatical non-fluent aphasia showed the reverse pattern, although the formation of the regular past tense was not quite at control levels. This example serves two purposes. First, it shows the value of single case studies for validating a theory. Second, the demonstration of double dissociation in-and-of-itself supports the theory; the exact location of the lesions in the two patients is not required for this purpose. It should be added that double dissociations are also subject to certain caveats but for the purposes of this review they are considered the strongest evidence for isolable motor subsystems.

Thus, it is a conceptual mistake to consider lesion studies synonymous with studies of localization. Patient studies do not have to be about localization; instead they can serve to identify separable qualitative processes through the demonstration of associations and dissociations. Demonstration of double dissociations in patients, and to a lesser extent associations and single dissociations, will not only allow one to infer the existence of functional specialization, but also determine its architecture, i.e., the anatomical basis for the specialization. This provides a springboard for study of the anatomical subsystems at the algorithmic and biological substrate levels.

The cerebellum: predicting sensory consequences of motor commands

Control policies generate motor commands based on beliefs about the state of the body and the environment (Eq. 5). This state estimate depends on two quantities: a prediction, and an observation. The prediction comes from an internal model that uses a copy of the motor commands to estimate the state change that is expected to occur. The observation comes from the sensory system that provides a

measure of those state changes (Fig. 3). That is, our beliefs are not based on our observations alone. Rather, our beliefs are a combination of what we predicted and what we observed (Kording and Wolpert 2004; Vaziri et al. 2006).

Some movements are so fast that there is no time for the sensory system to play a role. A prominent example is control of saccades (rapid eye movements that move the eyes to a new location typically with in 50–80 ms). Such movements are too brief for visual feedback to influence saccade trajectory. In fact, the brain actively suppresses visual processing during saccades to reduce the perception of motion (Thiele et al. 2002). Furthermore, proprioceptive signals from the eyes do not play any significant role in controlling saccade trajectories (Keller and Robinson 1971; Guthrie et al. 1983). Thus, the brain must guide saccade trajectories in the absence of sensory feedback. How is this accomplished? A plausible solution is for the brain to use an internal estimate of the state of the eye, derived from a copy of ongoing motor commands (Robinson 1975). This internal feedback probably accounts for the fact that variability at saccade initiation is partially corrected as the saccade progresses (Quaia et al. 2000).

What are the neural substrates of this internal feedback? The cerebellum is known to be critical for many aspects of saccade control and adaptation (Hopp and Fuchs 2004; Girard and Berthoz 2005). The projections from the superior colliculus to the cerebellum may provide the efference copy. Indeed, Takeichi et al. (2005) showed that adaptive changes in saccade amplitude are reflected in the nucleus reticularis tegmenti pontis, a major source of input to the cerebellum. Together, the superior colliculus-cerebellar-brainstem side loop seems important for steering saccade trajectories mid-flight (Optican and Quaia 2002; Optican 2005). This side loop is a likely candidate for acting as a forward model of the plant to produce mid-flight corrective feedback.

A simple experiment can test whether the cerebellum plays a role in predicting consequences of self-generated motor commands. Nowak et al. (2007) asked subjects to hold a force transducer that measures grip force (subject's right hand in Fig. 6). Attached to the transducer is a basket. The experimenter drops a ball into the basket. When the ball drops, it exerts a downward force on the hand. The subject responds by squeezing the transducer so that it will not slip out of his/her hand. Because there are delays in sensing the impact of the ball, the grip response comes about 100 ms after the ball's impact. Nowak et al. (2007) described patient HK who did not have a cerebellum due to a very rare developmental condition. When the experimenter dropped the ball in the basket, both the healthy individuals and HK showed the delayed response (data not shown). Therefore, the sensory feedback pathways appeared intact. In a subsequent trial, the subject (rather

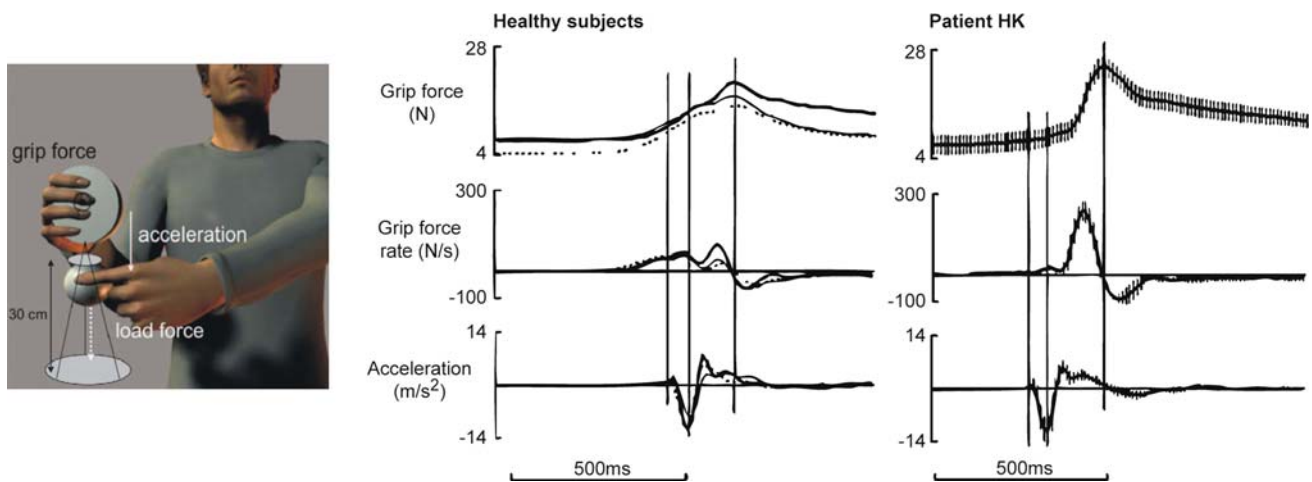


Fig. 6 Predicting and compensating for consequences of motor commands depends on the cerebellum. Subject holds a ball with their *left hand* and releases it into a basket held by the *right hand*. Healthy individuals increase their grip force in anticipation of the ball's impact. Patient HK, who suffered from cerebellar agenesis, has a grip

force that rises in response to the impact but not before it. The *bottom trace* refers to the acceleration of the right hand (holding the basket). The impact of the ball is marked by the first vertical line (from Nowak et al. 2007)

than the experimenter) dropped the ball. In a healthy individual, the brain can predict that the release of the ball will soon result in an impact that will increase the downward load. In anticipation of this event, the healthy individual squeezes the basket's handle harder around the time that the ball is released (grip force and grip force rates, Fig. 6). HK, however, could not make this anticipatory adjustment. Rather, she responded to the perturbation the same way that she responded when the ball was dropped by the experimenter. Therefore, the cerebellum appears to be required for the ability to predict the sensory consequences of motor commands (Wolpert et al. 1998b).

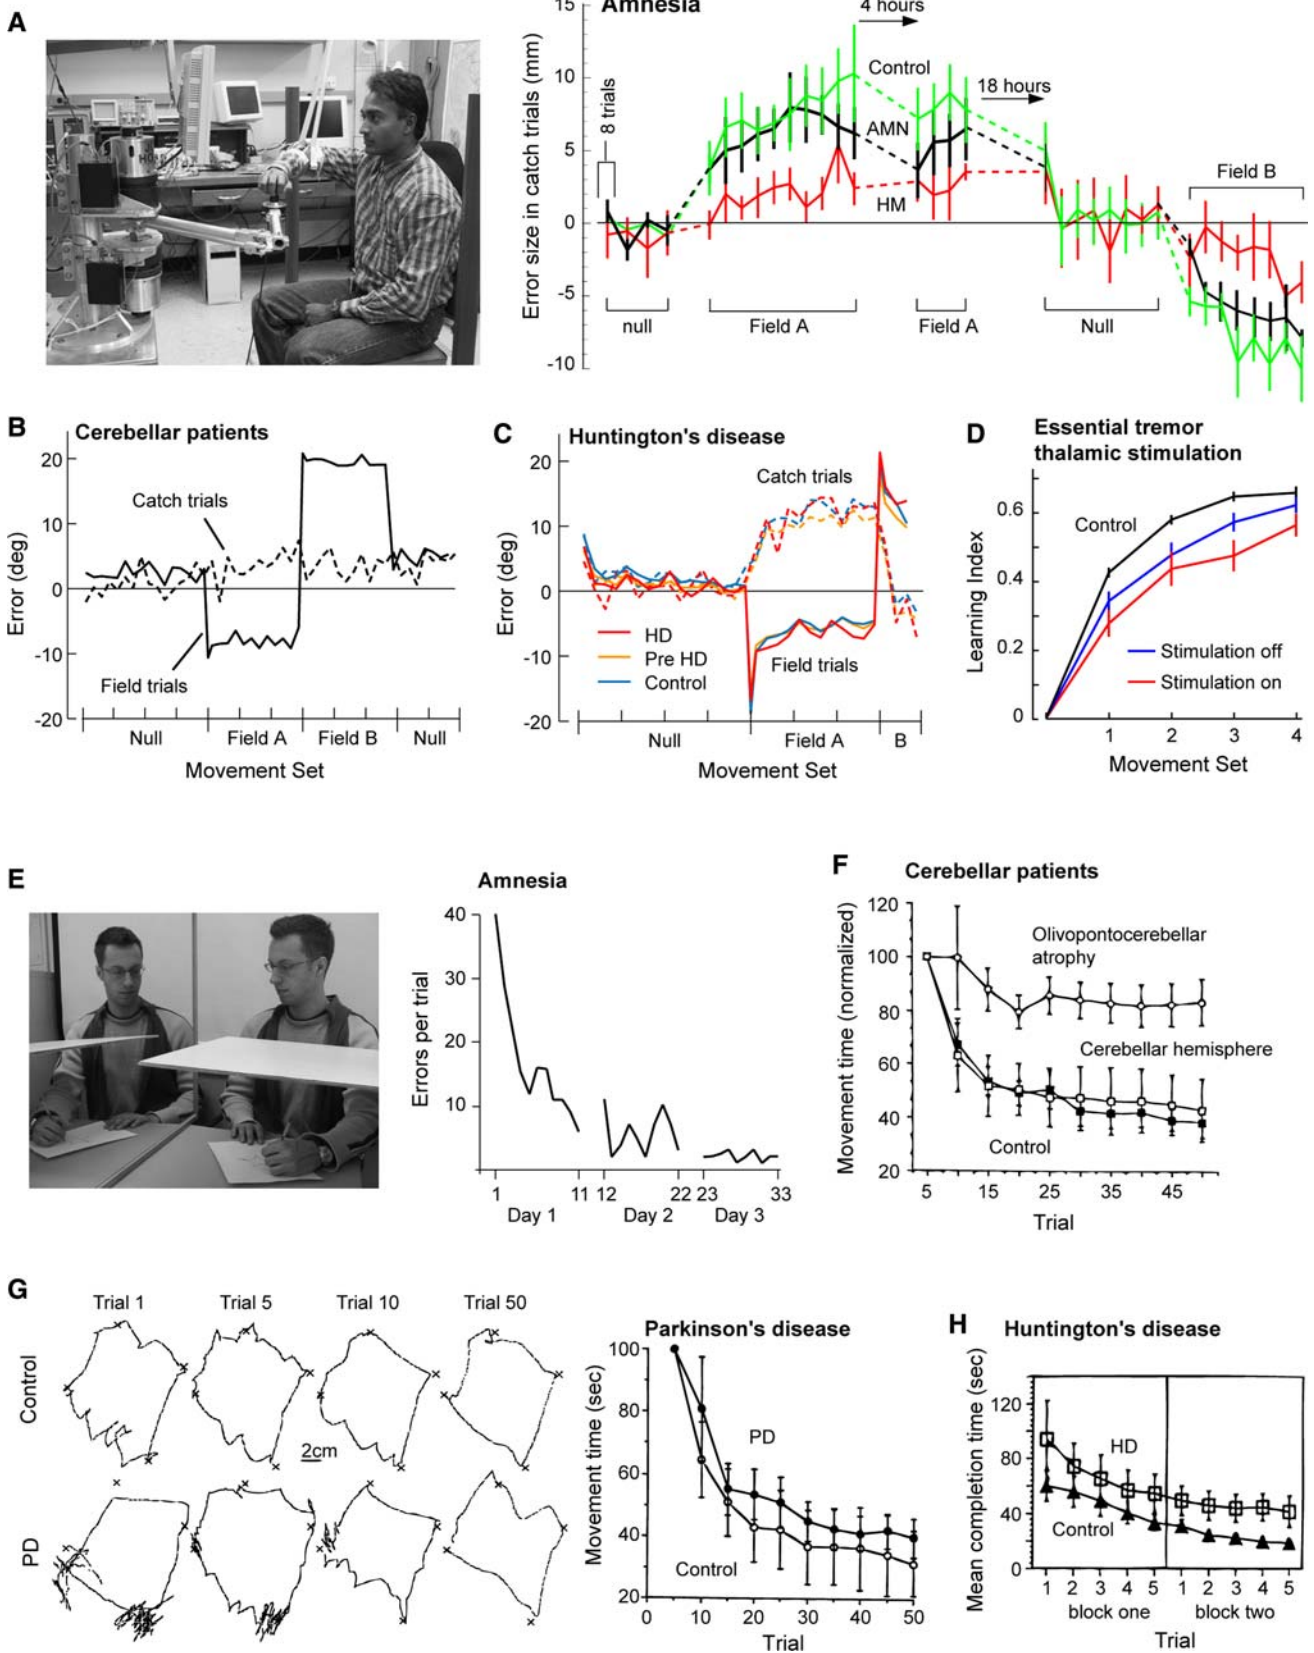
The cerebellum and construction of internal models

It is not easy to make accurate predictions about the sensory consequences of motor commands: our muscles respond differently depending on their fatigue state, and our limbs move differently depending on whether we are holding a light or heavy object. To maintain accuracy of the predictions, our brain needs to learn from the sensory feedback and adapt its internal model. This adaptation can be simple like changing parameter values of a known structure (changing A , B , or H in Eq. 2), or complex like identifying the structure de novo (replacing the linear form of Eq. 2 with some nonlinear function). The cerebellum appears to be one of the crucial sites of this process.

Cerebellar damage often prevents individuals from learning how to use novel tools. For example, when subjects are asked to move the handle of a robotic tool to manipulate cursor positions, they may not be able to learn to compensate

for forces generated by the robot (Fig. 7b) (Maschke et al. 2004; Smith and Shadmehr 2005), or to compensate for the novel visual feedback through a mirror (Fig. 7f) (Sanes et al. 1990). If the cerebellum is the crucial site for learning internal models, then it probably makes its contribution to control of reaching via its outputs to the thalamus, which in turn projects to the cerebral cortex. In humans, it is possible to reversibly disrupt this pathway. Essential tremor patients are occasionally treated with deep brain stimulators that artificially disrupt the ventrolateral thalamus, improving their tremor. However, these patients learn the reach task better when the stimulator is turned off (Chen et al. 2006). Therefore, it seems quite likely that the cerebellum is a key structure that allows us to learn tool use.

Another well-studied paradigm is adaptation to laterally displacing prisms that distort the path of light to the eyes. When cerebellar patients don prism glasses, they are severely impaired in adapting their reaching and throwing movements (Martin et al. 1996). Specifically, patients with infarcts in the distribution of the posterior inferior cerebellar artery, in the ipsilateral inferior peduncle, in the contralateral basal pons or in the ipsilateral middle cerebellar peduncle had impaired or absent prism adaptation. Interestingly, patients with infarcts in the distribution of the posterior inferior cerebellar artery usually had impaired or absent adaptation but little or no ataxia. In contrast, patients with damage in the distribution of the superior cerebellar artery or in cerebellar thalamus usually had ataxia but preserved adaptation. Thus, these studies show a dissociation within the cerebellum for anatomical loci that cause abnormal prism adaptation and a double dissociation between prism adaptation and ataxia for one anatomical



◀ **Fig. 7** Examples of motor skill learning in health and disease. **a** Learning to control a tool that has novel force characteristics. Subjects reached to visual targets while the robot perturbed the hand with either no forces (null field), or fields *A* or *B*. In 1/6th of the trials, the field was removed, resulting in catch trials. Learning was measured by the size of errors in catch trials. Data shown is for amnesic subject HM, other amnesic subjects (*AMN*), and normal control subjects (*NCS*). HM learned the task more slowly, but had excellent retention. Data from Shadmehr et al. (1998). **b, c** Performance of cerebellar and Huntington's disease patients on the same task. While cerebellar patients are profoundly impaired in learning, HD patients are normal. Data from Smith and Shadmehr (2005). **d** Effect of stimulation of the cerebellar thalamus of patients with essential tremor on the same task. Stimulation of the cerebellar thalamus reduces tremor (data not shown), but impairs learning. Data from Chen et al. (2006). **e** Learning to control a tool with novel kinematic characteristics: the mirror drawing paradigm. Data is from amnesic subject HM, redrawn from Milner et al. (1998). Each trial requires the subject to trace the star while keeping within the two lines. An error occurs when the pencil goes outside the two lines. **f** Performance of two groups of cerebellar patients on the same task. Data from Sanes et al. (1990). **g** Performance of patients with Parkinson's disease on the same task. *Filled symbols* represent the patient group. Data from Agostino et al. (1996). **h** Performance of patients with Huntington's disease on the same task. *NC* normal controls. Data from Gabrieli et al. (1997)

location, which suggests that these may be distinct abnormalities—one relates to trajectory errors in the extrinsic space whereas the other is caused by errors in anticipation of proprioceptive events (Vilis and Hore 1980). From the perspective of Marr, the identification of the cerebellum as a locus for prism adaptation is already a good demonstration of coarse-grained modularity. However, in the case of prism adaptation it seems that more can be said about the “what, where and how” *within* the cerebellum.

In contrast, patients with damage to the basal ganglia showed little or no deficits in adaptation. Figure 7c shows performance of Huntington's disease patients on the robot task (Smith and Shadmehr 2005), and Fig. 7g, h show performances of Parkinson's and Huntington's disease on the mirror task (Agostino et al. 1996; Gabrieli et al. 1997).

The question that remains is the precise role of the cerebellum in adaptation. Experiments show that the cerebellar damage causes abnormalities in adaptation to both kinematic (Tseng et al. 2007) and force (Smith and Shadmehr 2005) perturbations. One unifying concept is that the cerebellum may be the site of forward models, which predict the sensory consequences of motor commands. The output of the forward model could be used to generate a prediction error that drives adaptation and also be used to update a previous estimate of limb state. Support for this idea comes from a recent experiment in which TMS was used to disrupt the lateral cerebellum in human subjects while they slowly moved their arm in preparation for a making a rapid reaching movements (Miall et al. 2007). Reaching errors in initial direction and final finger position

suggested that the reaching movements had been made from an estimated hand position approximately 140 ms out of date, consistent with a role for the cerebellum in iteratively updating limb state.

The relationship between learning better sensory predictions and learning better motor commands

It is important to note that optimal control can only be “optimal” if the internal models are accurate. That is, we cannot hope to maximize reward unless our predictions about the behavior of our body generally agree with the sensory feedback (otherwise, the predictions would bias observations). Therefore, a strong implication of the OFC theory is that building internal models (a process called system identification) must go hand-in-hand with the process of optimization. This is an interesting prediction that to our knowledge has not been tested. However, it already points out a major problem with the framework presented here: because our body is constantly changing, and internal models are constantly adapting, do the timescales of adaptation match the timescales of optimization? Or is the timescales of optimization much slower? We do not know the answers to these questions. We will return to this question when we examine more thoroughly the limitations of OFC near the end of this review.

Amnesia and learning tool use

In 1997, one of us was involved in a two-day experiment on the severely amnesic patient HM to see how well he could learn a new motor skill and retain the new memory (Shadmehr et al. 1998). The task was the standard reach adaptation task (Fig. 7a) where subjects hold the handle of a robotic tool and learn to use it to guide a cursor to a sequence of targets (Shadmehr and Mussa-Ivaldi 1994). When we seated HM in front of the robot, he, like all naïve volunteers, sat quietly and avoided touching the machine. We asked him to put his hand on the robot's handle and move it around a bit. Naturally, he kept his gaze on his hand as he moved the robot's handle. He was instructed to not look at his hand, but rather at the video monitor, where a cursor was present. After a minute or so of moving the cursor around, a center target was presented and he was asked to move the cursor to that location. Subsequently, another target was shown and he was encouraged to move the cursor there. He was instructed to get the cursor to the target in a given amount of time. If he did so, the target “exploded”. For HM, the target explosion triggered a childhood memory of going bird hunting. As he was performing the task and was able to get a target explosion, he would spend the next few minutes describing the memory in

detail: the type of gun that he used, the porch in the rear of his childhood home, the terrain of the woods in his backyard, and the kinds of birds that he hunted. (He repeated these details many times during the two-day experiment.)

After a few minutes of reaching to targets, the robot began to impose forces on HM's hand, perturbing the path of the cursor. With more practice, he altered his motor commands to predicatively compensate for the forces: we inferred this from the fact that when the forces were suddenly removed (in catch trials), his movements had large errors, in a pattern that was mirror symmetric to errors that he had early in the training trials. We then thanked him for his time and he left to have lunch.

When he came back to the experiment room four hours later, he claimed that he had never seen the robotic device or knew what it was for. We pushed the robotic arm aside and asked him to sit down. He sat down, but then something interesting happened: rather than avoiding touching the machine, he voluntarily reached and grabbed the robot's handle, brought it toward him, and looked at the video monitor, apparently in anticipation of a target. It was clear that despite having no conscious recollection of having done the task before, some part of HM's brain recognized that the contraption was a tool that had a particular purpose: to manipulate cursors on a screen. When a target was presented, he showed strong after-effects of the previous training (Fig. 7b). That is, his brain expected the robot to perturb his movements, and so he generated motor commands in an attempt to compensate for these forces. Therefore, the motor memory was much more than just how to manipulate a tool. Rather, the memory also included information about the rewarding nature of the tool's purpose: the sight and touch of the robot was sufficient to encourage a motor act that was expected to be rewarding. If use of the robot in the first session had been paired with a shock or another noxious stimulus, it seems likely that he would have been reluctant to use the device again.

Brenda Milner had of course made a similar observation in HM some 30 years earlier in a task (Fig. 7e) where he drew on a piece of paper while looking in a mirror (Milner 1962). In the novel visual feedback setting, HM adapted his motor output and learned to draw accurately. When he returned the next day, the visual and/or tactile cues associated with the experimental setup were sufficient to allow him to recall the motor skill that he had learned before. Sue Corkin (1968) noted that after HM was trained on a pursuit rotor task (a task where one is asked to move a joystick so its position matches a moving target), upon return on the next day he not only retained the motor memory, but he had also retained certain "habits" regarding the purpose of the task: how to acquire reward, and what kinds of behaviors to avoid. Over the years, a number of other investigators made similar observations in other amnesic patients (Gabrieli et al. 1993;

Tranel et al. 1994; Yamashita 1993; Cavaco et al. 2004), culminating in the theory that our ability to learn motor skills is independent of the medial temporal lobe (Mishkin et al. 1984).

HM's performance in these tasks demonstrated that the brain could solve three general problems without conscious awareness and without the damaged medial temporal lobes. First, during the initial training session, the brain could learn to use a novel tool in order to achieve an instructed goal. Second, when tested hours later, the sight of the tool was sufficient to produce voluntary use, suggesting that potential rewards associated with the use of the tool had been learned and stored. Third, the sight of the tool and holding it was sufficient to allow recall of both the purpose of the tool (move a cursor) and the motor commands needed to achieve that purpose (the control policy). This is despite the fact that the same visual information was not sufficient to recall conscious memory of having trained on the task.

Learning the rewarding nature of sensory states

We expected that a severely amnesic individual who was performing a novel task would have to be regularly reminded of the task's instructions: "try to move the cursor to the target fast enough so it explodes." However, after HM had exploded a few targets, he no longer needed verbal reminders. The visual appearance of the target was enough for him to initiate a reaching movement. Strikingly, when he returned a few hours later and the next day he voluntarily reached for the robot handle and began preparing for onset of targets by moving the cursor to the center location. This behavior suggests that during the first session, he implicitly learned the reward basis of the task (Eq. 1). What brain regions were involved in learning the rewarding nature of bringing the cursor to the target?

Experiments on action selection in rodents provide important insights into this question. For example, suppose that a rat is released into a pool of water from some random starting point. A platform is positioned in a specific location just below the water line and cannot be seen. The platform is always at the same location in the pool. Rats dislike being wet, and will try to find a way to elevate themselves. The normal rat can learn to locate the platform position by paying attention to the visual cues that surround the pool. This requires learning a spatial map of where the platform is located with respect to the surrounding visual cues. With repeated swims, the animal learns a spatial map. This spatial map is analogous to a reward function that associates places in the pool with the likelihood of the platform (and therefore, the likelihood of not having to be wet).

Once the map is learned, the animal can find the platform regardless of where it is released into the water because the

map is with respect to the cues on the walls. If the platform is removed, the normal animal will spend most of his time searching in the region where the platform should be.

Learning of this sort of a spatial map depends on the hippocampus. If a genetically altered rat with a malfunctioning hippocampus is given the same training, he will not learn the spatial map and will spend equal time in each quadrant (Tsien et al. 1996). Therefore, selecting an action based on a spatial map likely relies on the hippocampus.

Sometimes, certain cues are rewarding no matter where they are located. Consider a pool where there are two hidden platforms: one that is large enough for the mouse to mount, and one that is too small. Both have a distinct visual cue associated with them: a little flag attached to each platform, each of a different color, sticking out of the water. Suppose that the flag attached to the large platform is red, while the flag attached to the small platform is green. The platforms may be positioned in any part of the pool, and will change from trial to trial. Therefore, in this experiment the animal needs to learn that the red flag indicates the location of the suitable platform and is a rewarding object. In another version of the experiment, the large platform will always be located in a particular spatial location, but the flag atop of it will be a random color. In this version of the experiment, the animal needs to learn that it is not the color of the flag that is important, but the spatial location.

We see that there is a natural competition between the learning systems that might be involved in these two conditions: is the platform in the same “place” as before (where place refers to a location in the spatial map), or is the platform always where the red flag is located? After a few trials where the flags move around, in the first experiment the animal should learn that the spatial map is a not a good indicator of the platform, and therefore the values associated with places in the spatial map should be near zero and the value associated with the red flag should rise. In the second experiment, the animal should learn that it is the spatial location that is of value, and the flag colors are irrelevant.

Packard and McGaugh (1992) performed both experiments by having their animals swim eight times per day for a number of days. They recorded the number of times the animals mounted the small platform and labeled this as errors. In the first experiment, where reward was associated with the red flag, healthy animals gradually learned to swim to the red flag. Interestingly, animals with damage to the medial temporal lobe learned the task just as well as the healthy controls. However, animals with damage to the caudate nucleus were much slower in learning the association. After days of training, they continued to attempt to mount the platform under the green flag. Therefore, it appears that the ability to associate reward to stimuli regardless of its spatial location depends on the basal ganglia.

In the second experiment, where reward was associated with a spatial location, healthy animals gradually learned to swim to that location and ignore the color of the flag. Animals with damage to the caudate nucleus performed similarly to the healthy controls. However, animals with damage to the medial temporal lobe were much slower in learning the association. Therefore, the ability to associate reward to a spatial location depends on the medial temporal lobe.

Returning to our observations in HM, we would speculate that it was his basal ganglia that learned that if he were to place the cursor in the box on the screen, and do so rapidly, a rewarding state would be experienced (explosions, which triggered a pleasant childhood memory). During the later sessions, the visual appearance of the machine, and the act of holding its handle, likely triggered a recall of this reward structure.

Effects of striatal damage on the assessment of movement cost

Writing instruments are one of the most common tools that we use in our daily lives. One of the striking features of damage to the human striatum is micrographia, an impairment of writing where letters become very small and writing speed becomes slow. This condition is most common in degenerative diseases of the basal ganglia like Parkinson’s disease (Van Gemmert et al. 2001), a condition causing progressive dysfunction of the neostriatum. However, it can also occur with focal lesions. Figure 8 provides an example of micrographia in patient FF, an individual who suffered an ischemic stroke in the left basal ganglia, in the head of the caudate nucleus and the anterior part of the putamen (Barbarulo et al. 2007). When asked to copy a



Fig. 8 Writing ability of patient FF, who suffered a lesion in the right caudate nucleus. Four- and eight-letter string copying (models on the upper lines) by the right (middle lines) and the left hand (lower lines). Micrographia was evident only with the right hand (Barbarulo et al. 2007)

four- or eight-letter string of characters, writing with the right hand was much smaller than with the left hand. Micrographia reflects an abnormal choice of speed and amplitude and is one manifestation of generalized slowing of movement (bradykinesia). The basic observation of smallness of writing size and slowness of the writing speed has been puzzling, as is the phenomenon of movement speed selection in general.

In the optimal control framework, there are no desired trajectories for our movements. Rather, the path is a result of a control policy (Eq. 5), which itself is a result of minimization of a cost (Eq. 1). The cost depends on two quantities: spatial accuracy (error cost) and required effort (energy cost). Accuracy requirements influence speed selection, due to the signal-dependent noise property of motor commands. The desired accuracy of a movement sets an upper limit on the maximum speed of a movement. The accuracy term of the cost function offers an explanation for the wealth of experimental data demonstrating speed-accuracy trade-off in reaching movements. Normal movements, however, do not appear to be made at the limits imposed by the speed-accuracy trade-off: we can reach for an object faster than usual without appreciable loss of accuracy. Although very little experimental data exists on spontaneous speed selection, the effort term of the cost function offers a potential explanation for this phenomenon. That is, perhaps micrographia is an indication of an abnormally high motor cost.

One of us recently tested this idea that in Parkinson's disease (PD), there may be an abnormally high cost associated with motor commands (Mazzoni et al. 2007). We required healthy control subjects to make accurate reaching movements of specified speeds. As the required speed increased, subjects took longer (required more trials) to accumulate a set number of movements at the required speed. This reluctance to move faster could be explained by the increase in required energy as well as by the degradation of spatial accuracy, and thus did not disambiguate the contribution of these two costs. We then compared the performance of patients with PD to that of control subjects in this task. PD patients demonstrated normal spatial accuracy in each condition, but required more trials than controls to accumulate the required number of movements in each speed range. The patients' increased reluctance to execute movements requiring greater effort, in spite of preserved spatial accuracy, provided experimental demonstration of the contribution of energy cost to speed selection, independent of spatial accuracy.

In the experiment described, PD patients exhibited increased sensitivity to a movement's effort requirement. Mazzoni et al. (2007) referred to the sensitivity to a movement's energy requirement (the effort term in the optimal control framework's cost function) as a type of

motivation, which they called "motor vigor". They suggested that this interpretation would link the neostriatum's motor control functions to the motivation/reward framework that has implicated the ventral striatum in relationship to behaviors explicitly guided by reward (Schultz 2006). The execution of an accurate movement at a comfortable speed could be the reward that results by selecting a speed that minimizes both error and effort terms in overall cost. This minimization is driven by the speed-accuracy trade-off for the error term, and by motor motivation for the effort term. Bradykinesia results when striatal dysfunction changes the value of effort minimization, (increased sensitivity to effort cost; L in Eq. 1) relative to that of accuracy optimization (error cost; Q in Eq. 1). Thus it appears that the basal ganglia either provides the motor motivation signal, which is then used to compute the "cost-to-go" elsewhere or is where the "cost-to-go" is computed. This cannot be disambiguated at this time.

Parietal cortex damage and state estimation

In a classic series of lesion experiments performed in monkeys, Mathew Rushworth and colleagues demonstrated a double dissociation between areas of posterior parietal cortex required for reaching under visual control and those required for reaching using proprioception (Rushworth et al. 1997). Monkeys were tested on two tasks: reaching in the light to visual targets and reaching in the dark to targets defined by arm position. Lesions to LIP and area 7 in the inferior parietal lobe (IPL) did not affect reaches in which the goal was defined in proprioceptive coordinates (in the dark), but produced mis-reaching in the light. In contrast, monkeys with lesions to area 5 in the superior parietal lobule (SPL) could reach accurately in the light but not in the dark. Therefore, SPL may be involved in proprioceptive estimates of limb position in space. An example of this in humans comes from patient PJ who had an extra-axial cyst encroaching on her left SPL (Wolpert et al. 1998a). Without vision of her right arm, PJ's sense of arm position became increasingly uncertain until she reported that the arm disappeared altogether. Thus, PJ was unable to store a proprioceptively derived estimate of limb state but, presumably because of the intact IPL, could use vision to maintain a sense of the limb's position in space. Similar to Rushworth's monkeys with SPL lesion, PJ made accurate reaches when vision was present.

Sometimes goal states change as the task is being performed. For example, in reaching to pick up a pen, the pen may start rolling away. Healthy individuals have no problems adjusting their movements to compensate for this change. However, parietal patients show particular difficulties with this task. For example, if parietal damage

impairs representation of visual states contralateral to the fixation, then motion of the goal state to this region during a movement impairs the ability to adjust the reach mid-flight. Grea et al. (2002) observed this phenomenon in a patient with bilateral posterior parietal cortex damage. The patient had no problems reaching to targets in central fixation. However, when at reach onset the target shifted to the right, the subject continued to reach to the original location of the target as if the target had not moved.

Disruption of the parietal cortex in healthy individuals can produce a similar phenomenon. Desmurget et al. (1999) provided a single pulse via a transcranial magnetic stimulator as the reach to the target began. On trials in which the target jumped, most of the participants had hand movements that disregarded the shift in the target location.

Let us examine these results in the framework of Fig. 3. The relevant state variables in this task include position of the limb (in proprioceptive and visual coordinates) and the position of the target (in visual coordinates). As motor commands are generated, the forward model should update its predicted state of the limb. Generally, we expect targets to remain stationary, and therefore the output of the forward model should continue to predict the target position. Together, these predictions represent the prior belief about the state of the body and the world. The sensory feedback from proprioception and vision is integrated with this prediction to make a posterior belief. When the target jumps, the novel sensory information needs to be integrated with the output of the forward model. If it is not, the reach will continue to the prior expectation of its location. The results noted above suggest that either this integration step is affected by damage or stimulation of the parietal cortex, or that the sensory information outside the central fixation region cannot reach the integration step.

Sometimes, there is no sensory feedback associated with our intended movements and the only thing that we can rely on is predicted states. An example of this is when we imagine performing a movement. Sirigu et al. (1996) examined the ability of patients with parietal cortex damage to monitor imagined movements of their fingers or arms. In their first task, they compared the speed of imagined movements with the participants' actual movements. The experimenters asked participants to imagine touching the tip of the thumb with the tips of each of the fingers of the same hand, in time to the sound of a metronome. They increased the speed of the metronome slowly until the individual reported that the imagined movement could no longer keep up. The experimenters then measured how fast the participants could actually make the movement. They found that in healthy people, the maximum speed of the imagined movements agreed remarkably well with the maximum speed of the actual ones (within 2%). Patients with unilateral posterior parietal cortex damage

could not accurately estimate their performance with the hand contralateral to the lesioned hemisphere. They could, however, estimate their maximum movement speed for the hand ipsilateral to the lesion. In contrast, a patient with a motor cortex lesion in the right hemisphere could accurately estimate maximum speed with both hands, despite the fact that the left hand moved much more slowly.

Next, Sirigu et al. (1996) had their participants imagine movements that varied from easy to hard. For example, they imagined reaching with a pen in order to place the tip inside a small or a large square. Afterwards, they performed the actual movements. In healthy participants, the time to completion of the imagined and actual movements agreed closely. As the task became more difficult, both the imagined and the actual movements took longer to complete. Patients with motor cortex lesions had trouble making the actual movement with the contralateral arm, but the duration of imagined movements matched that of the actual ones. In contrast, in patients with right PPC lesions, the duration of imagined movements with the left arm was significantly less than the duration of actual movements. Imagined and actual movements of the less affected right arm, however, showed a close correspondence.

To interpret these results, we need to conjecture as to what is the computational equivalent of an imagined movement. Perhaps imagining a movement is equivalent to formulating the cost of the task, minimizing it using a forward model, and arriving at motor commands, but not actually committing the commands to the descending spinal pathways. If producing a movement, imagined or actual, is a feedback control process that depends on the belief about current state, and this belief depends on both prediction and measurement, to imagine a movement we need to change the weight associated with measured sensory feedback to zero, making our control process rely entirely on our predictions. An inability to do so results in mis-estimation of the state of the limb during imagined movements.

Thus, evidence from physiological studies in non-human primates and patient studies suggests that the parietal cortex is involved in state estimation. In the optimal control framework these state estimates are multiplied by a gain to generate motor commands. This gain is in essence a sensory-motor transformation, which is achieved by an interaction between the parietal cortex and the frontal motor areas.

Motor cortex and feedback control

Hemiparesis after a stroke, the result of a lesion of the primary motor cortex and/or the corticospinal tract, is the

most common motor disorder seen by neurologists. In the majority of cases, ischemic lesions not only interrupt the monosynaptic connections to the ventral horn of the spinal cord, but also the projections to spinal cord interneurons, and often, because of their proximity to the internal capsule, projections to brainstem nuclei, which then project down to the spinal cord (Porter and Lemon 1995). The motor deficits most frequently emphasized after stroke include weakness, increased movement variability, the decomposition of smooth multi-joint movements into stereotypical joint combinations (synergies), and spasticity. These abnormalities are thought to result from interruption of monosynaptic cortico-motoneuronal connections to alpha-motor neurons, with a decrease in the amount and rate of their recruitment, reversion to control by spared descending brain stem pathways, and changes in segmental reflex circuits. What do the ideas of OFC and state estimation bring to the understanding of hemiparesis?

It is necessary to concede that the theoretical framework outlined earlier does not speak to deficits in implementation. In other words, the process of ascertaining rewards based on state estimation and the search for optimal control laws all presuppose that the computed commands have access to muscles via the corticospinal tract. Thus, the consequences of interruption of the corticospinal tract can be considered downstream from the computations we have described. However, there are two reasons why the framework has the potential to provide insights into hemiparesis. First, experimental results suggest that interruption of descending pathways leads to impairments in sensorimotor integration that are not reducible to downstream implementation deficits like weakness or impaired inter-joint coordination (see below). Second, it has long been appreciated that patients partially compensate for their deficits, i.e., they find alternative strategies to achieve the task goal. Thus, compensation could be considered a new optimal control law in the setting of an altered system with increased noise or uncertainty.

Is hemiparesis merely a deficit in communicating commands to the contralateral motor apparatus? Recently, one of us was involved in a study of patients with right arm paresis secondary to a subcortical stroke (Raghavan et al. 2006). These patients had impaired anticipatory scaling of grip and load forces during precision lifting of objects. In other words, they were not able to form an internal model of the properties of the object. The critical finding, however, was that if they lifted the object just once with their unaffected left hand they were then able to properly scale the grip and load forces when they lifted with their right hand. This means that their initial failure to scale was not attributable to a downstream implementation deficit, but rather to an inability to map the state of the object into an appropriate control signal. This might occur because after

damage to output from M1, the motor command may emanate from premotor cortex, which, relatively speaking, has less access to proprioceptive state information (Rizzolatti and Luppino 2001). It could be conjectured that the left hemisphere has access to output from the right hemisphere's controller through callosal connections.

Similarly, it has been shown that after subcortical stroke patients make trajectory errors consistent with failure to compensate for interaction torques, which suggests an unreliable model of arm dynamics (Beer et al. 2000). Failure to compensate for interaction torques has previously been described in deafferented patients who lack proprioception (Sainburg et al. 1995) and patients with cerebellar lesions (Bastian et al. 2000). All told these results suggest that the abnormalities in precision grip and reaching movements in patients with lesions of the corticospinal tract are partly due to a suboptimal control signal that results from suboptimal proprioceptive state estimation. However, it is also possible that the main problem in hemiparesis may be an inability to implement a correctly determined control law. At the current time determining which particular computational steps are affected in hemiparesis, if any, remains a significant challenge for us.

Distinct feedback gains for visual and proprioceptive states of the task

Selecting a goal state may be the first step in performing a task. However, the goal of a task can be represented at different levels of abstraction: Is the goal of reaching for a glass to sate one's thirst, to get the hand to the glass, to rotate the shoulder and elbow joints in order to get the hand to the glass, or to activate the muscles to rotate the joints? When the question is posed this way it becomes apparent that a task can be described hierarchically, which in turn implies the existence of high and low level feedback controllers. Prior to the OFC framework, the hierarchy was thought of in terms of planned trajectory representations in the brain in either intrinsic or extrinsic space (Cartesian, joint, and torque spaces). In the current framework, the hierarchy may be a set of feedback controllers that act on distinct state estimates provided by the various sensors in our body.

Sensory signals largely project to separate areas of the motor system by distinct routes. Visual signals reach the motor system via posterior parietal cortex to premotor areas. Proprioceptive states reach primary motor cortex from projections from S1 and thalamus. Accordingly, tuning functions of cells in the primary motor cortex are strongly dependent on the proprioceptive state of the arm (Sergio and Kalaska 2003), whereas in the premotor cortex, tuning is more dependent on the visual state of the task

(Takei et al. 2001). Thus, it can be conjectured that there are separate feedback controllers in premotor and motor cortices, each relying on state estimates coming from separate parts of parietal cortex in visual and proprioceptive modalities, respectively.

Healthy people adapt their reaching movements to kinematic perturbations and force perturbations without interference (Krakauer et al. 1999). Notably, subjects learn novel inertial dynamics just as well in the absence of visual feedback. It is possible that adaptation to a kinematic perturbation is driven primarily by visual errors (Krakauer et al. 2000) and adaptation to changes in inertial dynamics is driven primarily by proprioceptive errors (Hwang and Shadmehr 2005). Thus within the framework of OFC, it could be envisaged that because in tasks like reaching there are distinct sensory states (proprioception and vision), there are also distinct feedback control gains associated with those states. Where are these controllers? As outlined above, the control signal almost certainly originates from motor areas of the frontal lobe. Perhaps the feedback gains associated with proprioception dominate the primary motor cortex and feedback gains associated with vision dominate the premotor cortex.

In a recent study, it was shown that a patient with deafferentation due to large fiber sensory neuropathy, which resulted in loss of proprioception, was able to adapt his reaching to kinematic perturbation as well as age-matched controls (Bernier et al. 2006). Conversely, patients with optic ataxia, a disorder that is predominantly associated with lesions of the SPL and intraparietal sulcus and is characterized by directional errors in visually guided reaching, typically show no deficits in movements guided by proprioception, in force production, or visual perception (Perenin and Vighetto 1988). Their deficit is most apparent for non-foveal targets, which suggests that reaching errors relate to incorrect computation of target location in fixation coordinates. This possibility was explored in a recent study of patients with optic ataxia, who were required to make reaches to various targets (Dijkerman et al. 2006). For example, a patient with a right parietal damage fixated a center target and reached a target to either left or right of fixation. The patient had substantially more errors to the left target than right. However, when the subject fixated to the left, targets to the right of fixation were in precisely the same allocentric spatial location as before, yet reaches were markedly improved. One way to interpret this is that damage to the right parietal cortex impaired representation of visual space to the left of fixation. Goal states in this visual space could be as precisely specified. Thus, the results in the deafferented patient suggest that imprecision in the representation of a proprioceptively coded state estimate had no impact on adaptation to visually induced directional errors, but for the patient with optic ataxia,

imprecision in visually coded state estimate caused directional errors in visual space. Although a deafferented patient can adapt to visuomotor rotation, they show marked reach errors, indicative of impaired control of joint interaction torques and inertial anisotropy. Vision can only partially compensate for these errors, but importantly directional biases can be corrected. In summary, the effects of degraded visual and proprioceptive state estimates affect different types of adaptation and different components of reaching trajectories.

However, this simplistic framework is not consistent with what is seen after stroke: patients with a lesion in primary motor cortex or the internal capsule have hemiparesis with or without visual feedback! As an aside, there is a surprising paucity of human studies in the literature on the consequence of lesions isolated to lateral premotor regions (although see Kunesch et al. 1995). One reason is that hemiparesis after a premotor lesion, in the anecdotal experience of one of the authors (JWK), tends to resolve rapidly. A possible explanation for this comes from a recent study in which it was demonstrated that electrical stimulation of ventral premotor cortex did not lead to detectable corticospinal output itself but robustly facilitated corticospinal output from M1 (Shimazu et al. 2004). Thus after a stroke it might be that abrupt loss of this facilitation causes a short-lived hemiparesis. That is, control signals generated from visual and proprioceptive estimates may channel through M1 (Fig. 3a, b).

Limitations in applying the theory to biological motor control

This review of motor control has been written within the framework of OFC. At the heart of the theory is the conjecture that animals make voluntary movements in order to acquire the most reward while expending the least effort. However, the theory cannot make a behavioral prediction unless we can specify three kinds of information: (1) what are the costs and rewards, (2) what are the constraints, i.e., dynamics of the task, and (3) what are the mechanisms of state estimation. In this review we have chosen a specific set of equations to represent each kind of information. However, it is not difficult to find examples of behavior that are inconsistent with our formulation.

The cost that we wrote in Eq. (1) is perhaps the simplest possible cost function for goal directed movements. How seriously can we take this specific representation? As demonstrated by attempts to reverse-engineer the cost (Kording et al. 2004a), the quadratic cost function should not be taken too seriously.

Let us re-consider the task where one uses both arms to control a single cursor. In the right column of Fig. 5b, we

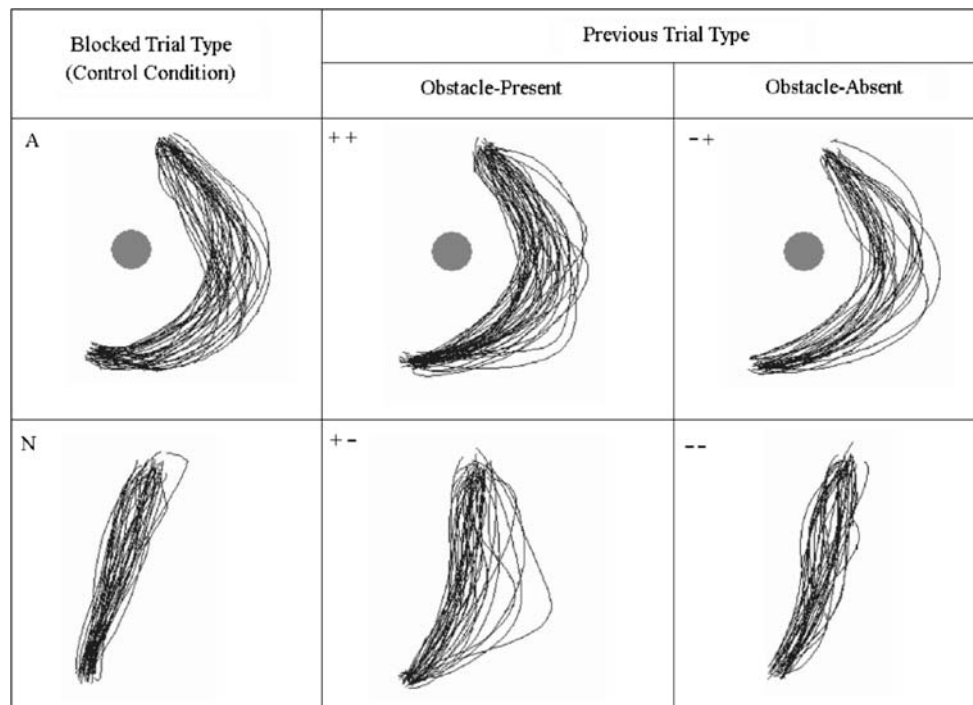


Fig. 9 Reaching around an obstacle affects the subsequent trial when there is no obstacle. On any given trial, the target of the movement and the obstacle, if any, are present before the reach starts. *Left column* shows data from two control groups in which either all (*A*) or none of the movements (*N*) had an obstacle. The data displayed in the *middle and right columns* are for a group where the probability of an obstacle on any given trial was 50%. The marking ++ indicates that

two consecutive movements had an obstacle, and the data shown is for the second of these two trials. The marking +- indicates that the last trial had an obstacle but the current trial does not. No two consecutive trials were to the same direction. The movement directions have been rotated for ease of comparison (from Jax and Rosenbaum 2007)

see that a force field that perturbed the left arm was corrected almost entirely by that arm, with a significant but much smaller contribution from the unperturbed arm. The troubling fact is that if our objective was to minimize a cost like that of Eq. (1), then we should be much more willing to utilize the services of the unperturbed arm, as its actions are not impeded by a resisting force field and therefore carry much smaller motor costs. The theory correctly predicted that there should be a change in perturbation response when the task changed from two-cursor to one-cursor condition, but the amount of change observed appeared to be grossly sub-optimum. Is this a problem with our cost function, a problem with our constraint equations, or a problem with our state estimation process? We do not know the answer.

Consider a set of experiments that highlighted the importance of costs associated with postural stability, a quantity that we did not include in Eq. (1). Scheidt and Ghez (2007) explored a task where continuous random noise perturbed the hand at rest. This constraint encouraged increasing the co-contraction levels of muscles. However, the noise was present only during the postural phase of the task and disappeared when subjects made a reaching movement. They found that if a kinematic perturbation

required adaptation of the movement, the learning did not generalize to the postural phase at the end of the movement. They suggested that the control processes that moved the limb appeared distinct from control processes that set muscle activity levels during posture. If so, do these processes have separate costs? A recent study suggests that the answer is yes, the weighting of postural cost is flexible and can be determined by task context (Liu and Todorov 2007).

Finally, consider an experiment by Jax and Rosenbaum (2007) in which they asked subjects to make arm movements to an array of 12 targets positioned in a 16 cm radius circle on a vertical screen. Targets were presented randomly and in some trials an obstacle was presented halfway between the start and the target (Fig. 9). The same target was never shown twice in a row. Interestingly, whenever a no-obstacle trial followed an obstacle trial, subjects made curved rather than straight trajectories. However, the movements straightened out when a no-obstacle trial followed another no-obstacle trial. Why make a suboptimal curved trajectory when you see that there is no obstacle?

These results highlight a number of important problems with our framework. First, without knowing precisely the costs and rewards of a movement, it will not be possible to

make quantitatively reliable predictions of behavior. Without a priori predictions, how can the theory be falsified? That is to say, if we have experimental results and are allowed to tweak the costs or their weightings until we get a good fit then what have we learned? We would suggest that the best way to proceed is to either specify the cost function before experiments are conducted and make predictions, or fit the costs to data to find the best parameter fit, i.e., derive Q and L from Eq. (1) rather than pre-specify them.

Second, what are the timescales of optimization? Is optimization computed in the reaction time of each trial de novo? The timescale appears to be longer than a single trial because in the experiment shown in Fig. 9, making a curved movement in the previous trial trumped what should be optimal for the subsequent no-obstacle trial. Certainly, new costs can be conjured up—for example, in this case we can assume that finding feedback control gains that minimize a cost requires neural processing that itself has a cost and so it might be more efficient to allow the solution in one trial to linger on to influence the solution in the next trial. Or perhaps there is a cost in switching control policies?

Third, what is the timescale of system identification? Our body changes over multiple timescales. Muscles fatigue and recover quickly, objects are lifted and replaced rapidly, yet aging can produce gradual loss of motor neurons and transformation of muscle fibers. In other words, the parameters of the constraint equation, and perhaps its structure, are changing over multiple timescales. Unfortunately, we cannot make optimized movements unless we have an accurate set of constraint equations, i.e., an accurate internal model. When we see a sub-optimum movement, can we dissociate the effects of an inaccurate internal model from effects of an inaccurate cost function?

Finally, what is the alternate hypothesis to this theory? At this time, the alternative is another cost or constraint, not a fundamentally distinct theory. However, formalization of a theory is the key step that accelerates its evolution toward acceptance or rejection. Experimental data that do not fit with the theory are cause for celebration because only these data encourage progress.

Conclusions

The relationship between theories and the neural machinery that implements them is still in the courtship stage, but despite the separation, it has begun to bear modest fruit: theories have informed the neural basis of motor control in patients while lesion studies have informed the algorithms and representations that implement the computational theories. The result is the functional anatomy of voluntary

movements outlined in Fig. 3b. In this framework, a role for the cerebellum is system identification, i.e., predicting the changes in state that arise as a result of motor commands. A role for the parietal cortex is state estimation, where predictions about sensory feedback are integrated with visual and proprioceptive observations to form beliefs about states of our selves and objects/people around us. The basal ganglia may play a role in computing a “cost-to-go” function, estimating value of states and costs of motor commands. Finally, once a goal state is selected, motor cortical areas minimize this cost function and transform state estimates into motor output by formulating a feedback control policy.

Acknowledgments The work was supported by National Institutes of Health (NIH) grants K02-048099 and R01-052804 to JWK, R01-037422 to RS, and a grant from the Human Frontiers Science Foundation to RS.

References

- Agostino R, Sanes JN, Hallett M (1996) Motor skill learning in Parkinson's disease. *J Neurol Sci* 139:218–226
- Bahill AT, Clark MR, Stark L (1975) Dynamic overshoot in saccadic eye movements is caused by neurological control signed reversals. *Exp Neurol* 48:107–122
- Barbarulo AM, Grossi D, Merola S, Conson M, Trojano L (2007) On the genesis of unilateral micrographia of the progressive type. *Neuropsychologia* 45:1685–1696
- Bastian AJ, Zackowski KM, Thach WT (2000) Cerebellar ataxia: torque deficiency or torque mismatch between joints? *J Neurophysiol* 83:3019–3030
- Beer RF, Dewald JP, Rymer WZ (2000) Deficits in the coordination of multijoint arm movements in patients with hemiparesis: evidence for disturbed control of limb dynamics. *Exp Brain Res* 131:305–319
- Bernier PM, Chua R, Bard C, Franks IM (2006) Updating of an internal model without proprioception: a deafferentation study. *Neuroreport* 17:1421–1425
- Caramazza A (1986) On drawing inferences about the structure of normal cognitive systems from the analysis of patterns of impaired performance: the case for single-patient studies. *Brain Cogn* 5:41–66
- Cavaco S, Anderson SW, Allen JS, Castro-Caldas A, Damasio H (2004) The scope of preserved procedural memory in amnesia. *Brain* 127:1853–1867
- Chen H, Hua SE, Smith MA, Lenz FA, Shadmehr R (2006) Effects of human cerebellar thalamus disruption on adaptive control of reaching. *Cereb Cortex* 16:1462–1473
- Collewijn H, Erkelens CJ, Steinman RM (1988) Binocular coordination of human horizontal saccadic eye movements. *J Physiol* 404:157–182
- Corkin S (1968) Acquisition of motor skill after bilateral medial temporal-lobe excision. *Neuropsychologia* 6:255–265
- Desmurget M, Epstein CM, Turner RS, Prablanc C, Alexander GE, Grafton ST (1999) Role of the posterior parietal cortex in updating reaching movements to a visual target. *Nature Neurosci* 2:563–567
- Diedrichsen J (2007) Optimal task-dependent changes of bimanual feedback control and adaptation. *Curr Biol* 17:1675–1679

- Dijkerman HC, McIntosh RD, Anema HA, de Haan EH, Kappelle LJ, Milner AD (2006) Reaching errors in optic ataxia are linked to eye position rather than head or body position. *Neuropsychologia* 44:2766–2773
- Domkin D, Laczko J, Djupsjobacka M, Jaric S, Latash ML (2005) Joint angle variability in 3D bimanual pointing: uncontrolled manifold analysis. *Exp Brain Res* 163:44–57
- Fitts PM (1954) The information capacity of the human motor system in controlling the amplitude of movement. *J Exp Psychol* 47:381–391
- Flash T, Hogan N (1985) The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci* 5:1688–1703
- Gabrieli JDE, Corkin S, Mickel SF, Growdon JH (1993) Intact acquisition and long-term retention of mirror-tracing skill in Alzheimer's disease and in global amnesia. *Behav Neurosci* 107:899–910
- Gabrieli JDE, Stebbins GT, Singh J, Willingham DB, Goetz CG (1997) Intact mirror-tracing and impaired rotary-pursuit skill learning in patients with Huntington's disease: evidence for dissociable memory systems in skill learning. *Neuropsychology* 11:272–281
- Girard B, Berthoz A (2005) From brainstem to cortex: computational models of saccade generation circuitry. *Prog Neurobiol* 77:215–251
- Grea H, Pisella L, Rossetti Y, Desmurget M, Tilikete C, Grafton S, Prablanc C, Vighetto A (2002) A lesion of the posterior parietal cortex disrupts on-line adjustments during aiming movements. *Neuropsychologia* 40:2471–2480
- Guthrie BL, Porter JD, Sparks DL (1983) Corollary discharge provides accurate eye position information to the oculomotor system. *Science* 221:1193–1195
- Harris CM, Wolpert DM (2006) The main sequence of saccades optimizes speed-accuracy trade-off. *Biol Cybern* 95:21–29
- Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394:780–784
- Hopp JJ, Fuchs AF (2004) The characteristics and neuronal substrate of saccadic eye movement plasticity. *Prog Neurobiol* 72:27–53
- Hwang EJ, Shadmehr R (2005) Internal models of limb dynamics and the encoding of limb state. *J Neural Eng* 2:S266–S278
- Izawa J, Rane T, Donchin O, Shadmehr R (2008) Motor adaptation as a process of reoptimization. *J Neurosci* (in press)
- Jax SA, Rosenbaum DA (2007) Hand path priming in manual obstacle avoidance: evidence that the dorsal stream does not only control visually guided actions in real time. *J Exp Psychol Hum Percept Perform* 33:425–441
- Jones KE, Hamilton AF, Wolpert DM (2002) Sources of signal-dependent noise during isometric force production. *J Neurophysiol* 88:1533–1544
- Kakei S, Hoffman DS, Strick PL (2001) Direction of action is represented in the ventral premotor cortex. *Nature Neurosci* 4:1020–1025
- Kawato M, Gomi H (1992) A computational model of four regions of the cerebellum based on feedback-error learning. *Biol Cybern* 68:95–103
- Keller EL, Robinson DA (1971) Absence of a stretch reflex in extraocular muscles of the monkey. *J Neurophysiol* 34:908–919
- Kording K (2007) Decision theory: what “should” the nervous system do? *Science* 318:606–610
- Kording KP, Fukunaga I, Howard IS, Ingram JN, Wolpert DM (2004) A neuroeconomics approach to inferring utility functions in sensorimotor control. *PLoS Biol* 2:e330
- Kording KP, Tenenbaum JB, Shadmehr R (2007) The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat Neurosci* 10:779–786
- Kording KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–247
- Krakauer JW, Ghilardi MF, Ghez C (1999) Independent learning of internal models for kinematic and dynamic control of reaching. *Nature Neurosci* 2:1026–1031
- Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J Neurosci* 20:8916–8924
- Kunesch E, Binkofski F, Steinmetz H, Freund HJ (1995) The pattern of motor deficits in relation to the site of stroke lesions. *Eur Neurol* 35:20–26
- Liu D, Todorov E (2007) Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J Neurosci* 27:9354–9368
- Marr D (1982) Vision: a computational investigation into the human representation and processing of visual information. Henry Holt
- Martin TA, Keating JG, Goodkin HP, Bastian AJ, Thach WT (1996) Throwing while looking through prisms. I. Focal olivocerebellar lesions impair adaptation. *Brain* 119:1183–1198
- Maschke M, Gomez CM, Ebner TJ, Konczak J (2004) Hereditary cerebellar ataxia progressively impairs force adaptation during goal-directed arm movements. *J Neurophysiol* 91:230–238
- Mazzoni P, Hristova A, Krakauer JW (2007) Why don't we move faster? Parkinson's disease, movement vigor, and implicit motivation. *J Neurosci* 27:7105–7116
- Miall RC, Christensen LOD, Owen C, Stanley J (2007) Disruption of state estimation in the human lateral cerebellum. *PLoS Biol* 5:e316
- Milner B (1962) Les troubles de la memoire accompagnant des lesions hippocampiques bilaterales. In: *Physiologie de l'Hippocampe, Colloques Internationaux No. 107*. CNRS, Paris, pp 257–272
- Milner B, Squire LR, Kandel ER (1998) Cognitive neuroscience and the study of memory. *Neuron* 20:445–468
- Mishkin M, Malamut B, Bachevalier J (1984) Memories and habits: two neural systems. In: Lynch G, McGaugh J (eds) *Neurobiology of learning and memory*. Guilford Press, pp 65–77
- Morasso P (1981) Spatial control of arm movements. *Exp Brain Res* 42:223–227
- Nowak DA, Timmann D, Hermsdörfer J (2007) Dexterity in cerebellar agenesis. *Neuropsychologia* 45:696–703
- Ohyama T, Nores WL, Murphy M, Mauk MD (2003) What the cerebellum computes. *Trends Neurosci* 26:222–227
- Optican LM (2005) Sensorimotor transformation for visually guided saccades. *Ann NY Acad Sci* 1039:132–148
- Optican LM, Quaia C (2002) Distributed model of collicular and cerebellar function during saccades. *Ann NY Acad Sci* 956:164–177
- Packard MG, McGaugh JL (1992) Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: further evidence for multiple memory systems. *Behav Neurosci* 106:439–446
- Perenin MT, Vighetto A (1988) Optic ataxia: a specific disruption in visuomotor mechanisms. I. Different aspects of the deficit in reaching for objects. *Brain* 111(Pt 3):643–674
- Pinker S, Ullman MT (2002) The past and future of the past tense. *Trends Cogn Sci* 6:456–463
- Porter R, Lemon R (1995) *Corticospinal function and voluntary movement*. Clarendon Press, Oxford
- Quaia C, Pare M, Wurtz RH, Optican LM (2000) Extent of compensation for variations in monkey saccadic eye movements. *Exp Brain Res* 132:39–51
- Raghavan P, Krakauer JW, Gordon AM (2006) Impaired anticipatory control of fingertip forces in patients with a pure motor or sensorimotor lacunar syndrome. *Brain* 129:1415–1425

- Raymond JL, Lisberger SG, Mauk MD (1996) The cerebellum: a neuronal learning machine? *Science* 272:1126–1131
- Rizzolatti G, Luppino G (2001) The cortical motor system. *Neuron* 31:889–901
- Robinson DA (1975) Oculomotor control signals. In: BachyRita P, Lennerstrand G (eds) *Basic mechanisms of ocular motility and their clinical implications*. Pergamon, Oxford, pp 337–374
- Rushworth MF, Nixon PD, Passingham RE (1997) Parietal cortex and movement. I. Movement selection and reaching. *Exp Brain Res* 117:292–310
- Sainburg RL, Ghilardi MF, Poizner H, Ghez C (1995) Control of limb dynamics in normal subjects and patients without proprioception. *J Neurophysiol* 73:820–835
- Sanes JN, Dimitrov B, Hallett M (1990) Motor learning in patients with cerebellar dysfunction. *Brain* 113:103–120
- Scheidt RA, Ghez C (2007) Separate adaptive mechanisms for controlling trajectory and final position in reaching. *J Neurophysiol* 98:3600–3613
- Scholz JP, Schoner G (1999) The uncontrolled manifold concept: identifying control variables for a functional task. *Exp Brain Res* 126:289–306
- Schultz W (2006) Behavioral theories and the neurophysiology of reward. *Annu Rev Psychol* 57:87–115
- Sergio LE, Kalaska JF (2003) Systematic changes in motor cortex cell activity with arm posture during directional isometric force generation. *J Neurophysiol* 89:212–228
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208–3224
- Shadmehr R, Brandt J, Corkin S (1998) Time dependent motor memory processes in H.M. and other amnesic subjects. *J Neurophysiol* 80:1590–1597
- Shimazu H, Maier MA, Cerri G, Kirkwood PA, Lemon RN (2004) Macaque ventral premotor cortex exerts powerful facilitation of motor cortex outputs to upper limb motoneurons. *J Neurosci* 24:1200–1211
- Sirigu A, Duhamel J-R, Cohen L, Pillon B, Dubois B, Agid Y (1996) The mental representation of hand movements after parietal cortex damage. *Science* 273:1564–1568
- Smith MA, Shadmehr R (2005) Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. *J Neurophysiol* 93:2809–2821
- Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. *PLoS Biol* 4:e179
- Takeichi N, Kaneko CR, Fuchs AF (2005) Discharge of monkey nucleus reticularis tegmenti pontis neurons changes during saccade adaptation. *J Neurophysiol* 94:1938–1951
- Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O (2002) Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res* 142:284–291
- Thiele A, Henning P, Kubischik M, Hoffmann KP (2002) Neural mechanisms of saccadic suppression. *Science* 295:2460–2462
- Thoroughman KA, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. *Nature* 407:742–747
- Thoroughman KA, Wang W, Tomov DN (2007) The influence of viscous loads on motor planning. *J Neurophysiol* 98:870–877
- Todorov E (2005) Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Comput* 17:1084–1108
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235
- Tranel D, Damasio AR, Damasio H, Brandt JP (1994) Sensorimotor skill learning in amnesia: additional evidence for the neural basis of nondeclarative memory. *Learn Mem* 1:165–179
- Trommershauser J, Gepshtein S, Maloney LT, Landy MS, Banks MS (2005) Optimal compensation for changes in task-relevant movement variability. *J Neurosci* 25:7169–7178
- Tseng YW, Diedrichsen J, Krakauer JW, Shadmehr R, Bastian AJ (2007) Sensory prediction errors drive cerebellum-dependent adaptation of reaching. *J Neurophysiol* 98:54–62
- Tsien JZ, Huerta PT, Tonegawa S (1996) The essential role of hippocampal CA1 NMDA receptor-dependent synaptic plasticity in spatial memory. *Cell* 87:1327–1338
- Uno Y, Kawato M, Suzuki R (1989) Formation and control of optimal trajectory in human multijoint arm movement. Minimum torque-change model. *Biol Cybern* 61:89–101
- Van Gemmert AW, Teulings HL, Stelmach GE (2001) Parkinsonian patients reduce their stroke size with increased processing demands. *Brain Cogn* 47:504–512
- Vaziri S, Diedrichsen J, Shadmehr R (2006) Why does the brain predict sensory consequences of oculomotor commands? Optimal integration of the predicted and the actual sensory feedback. *J Neurosci* 26:4188–4197
- Vilis T, Hore J (1980) Central neural mechanisms contributing to cerebellar tremor produced by limb perturbations. *J Neurophysiol* 43:279–291
- Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensorimotor integration. *Science* 269:1880–1882
- Wolpert DM, Goodbody SJ, Husain M (1998a) Maintaining internal representations: the role of the human superior parietal lobe. *Nature Neurosci* 1:529–533
- Wolpert DM, Miall RC, Kawato M (1998b) Internal models in the cerebellum. *Trends Cog Sci* 2:338–347
- Yamashita H (1993) Perceptual-motor learning in amnesic patients with medial temporal lobe lesions. *Percept Mot Skills* 77:1311–1314
- Yang JF, Scholz JP (2005) Learning a throwing task is associated with differential changes in the use of motor abundance. *Exp Brain Res* 163:137–158
- Shallice T (1988) *From neurobiology to mental structure*. Cambridge University Press.