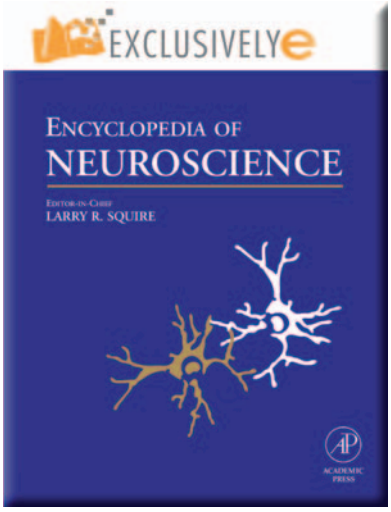

Shadmehr R (2009) Computational Approaches to Motor Control. In: Squire LR (ed.) Encyclopedia of Neuroscience, volume 3, pp. 9-17.
Oxford: Academic Press.

Computational Approaches to Motor Control

R Shadmehr, Johns Hopkins School of Medicine, Baltimore, MD, USA

Introduction

There are an infinite number of movements that we can do at any given moment. How do we select what kind of movement to perform? We perform actions to acquire rewarding states. For example, we may reach in order to pick up a cup of tea and bring it to our mouth. The drinking of the tea is the rewarding state that we have chosen as our goal, and the reach is the action that allows us to experience that rewarding state. The problem of what to do begins with a catalog of all potential rewarding states and proceeds with an evaluation of these possibilities to decide on one that is most valuable.

Our ability to choose from among the infinite number of potentially rewarding states depends on our brain's ability to assign value to those states. This value includes a measure of cost and a measure of reward. The basal ganglia play a fundamental role in this ability to assign value to potentially rewarding states and help in choosing the goal. However, once the goal has been chosen, the problem of how to perform it remains. The cerebellum plays a fundamental role in this ability to control the process of performing actions.

The problem of motor control is summarized in Figure 1. Through our belief about costs and rewards of various states, we choose a goal state. We use the cost and rewards associated with our goal to generate motor commands that we believe will bring us closer to the goal state. Our motor commands cause states of our body to change. We sense these state changes through our sensory system. However, we also predict what the sensory changes should be through internal models (called forward models). We combine what we predicted with what we actually sensed to form beliefs about the state of our body. We compare our belief about our current state with our goal state and continue to produce further motor commands until either the goal state is achieved or we abandon that goal and consider another one.

The Problem of What to Do

We select actions based on cues that are available to us. Some of these cues tell us about the position of our body and objects in the environment with respect to a spatial map. 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 waterline 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. For example, the rat might learn that the platform is near the wall with a particular feature. 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 they learn the map, rats can find the platform regardless of where they are released into the water because the map is related to the cues on the walls. If the platform is removed, the normal animal will spend most of its time searching in the region where the platform used to be.

Learning 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. 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 rat 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, and the flag attached to the small platform is green. The platforms may be positioned in any part of the pool, and their locations 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 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. The two versions of the experiment are depicted in Figures 2(a) and 2(b). There is a natural competition between the learning systems that might be involved in these two conditions: Is the platform 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 in which the flags

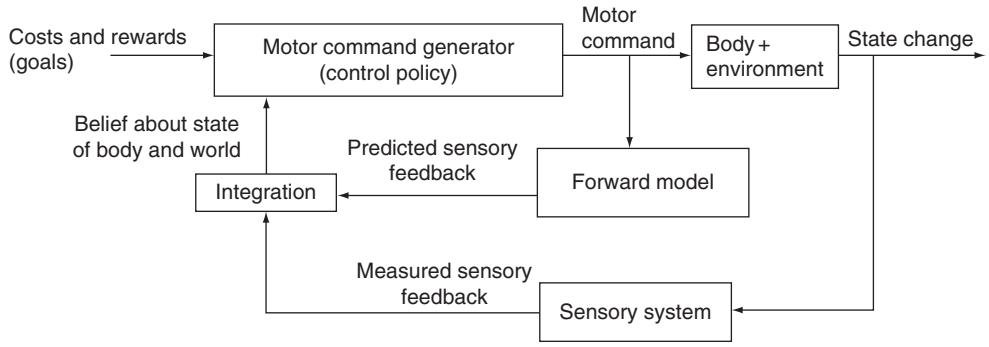

Figure 1 Schematic of motor control for voluntary movements.

move around, in the first experiment the animal should learn that the spatial map is not a good indicator of the platform and that therefore the values associated with places in the spatial map should be near zero. However, 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 that the flag colors are irrelevant.

Packard and McGaugh 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 these as errors. In the first experiment, where reward was associated with the red flag, healthy animals gradually learned to swim to the red flag (**Figure 2(a)**). It is interesting that animals with damage to the medial temporal lobe (which includes the hippocampus and the fornix, a major afferent–efferent pathway of the hippocampus) learned the task just as well as the healthy controls. However, animals with damage to the caudate nucleus (a part of the basal ganglia) 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 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 (**Figure 2(b)**). 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.

The Major Reward System in the Brain Is the Neurotransmitter Dopamine

How does reward get associated with a stimulus? That is, how do we learn the value of a location on a spatial map or the color of a flag? This issue is still poorly understood, but a major piece of the puzzle appears to be the neurotransmitter dopamine. Neurons that release dopamine have their cell bodies in the midbrain area (an area called the substantia nigra) of the brain stem. They project to three main areas: the striatum (encompassing the caudate and putamen nuclei of the basal ganglia), the hippocampus, and the prefrontal cortex. About 80% of the dopamine in the brain is in the basal ganglia.

The earliest studies of neuronal activity in the dopaminergic cells of the midbrain led to few published reports. The investigators conducting those experiments often reported frustration because the neurons never seemed to 'do anything.' For example, the dopaminergic cells discharged at about 2–4 Hz, but this rate never seemed to change. Although they did not recognize it at the time, the people who made these observations had actually stumbled on a crucial and important property of dopaminergic cells. The earliest recordings were made in monkeys that performed highly overlearned movements – a simple flexion and extension of the arm – in a highly stereotyped and repetitive manner. The monkeys performed the task perfectly and learned nothing during the recording sessions.

When Wolfram Schultz and his colleagues reexamined the activity of midbrain dopaminergic cells years later, they placed their monkeys in a situation involving learning, and the cells showed significant modulation in their activity level. They found that these cells carried a signal that conveyed a measure of the predicted reward: when an unexpected reward occurred, the dopaminergic cells increased their discharge rates. After repeated experience, the cells became unresponsive to rewards *per se* but discharged after signals that predicted reward. The omission of expected reward led to a decrease in the firing rate of dopaminergic neurons. Thus, dopaminergic cells seemed to encode both errors in predicted reward and stimuli that predicted reward.

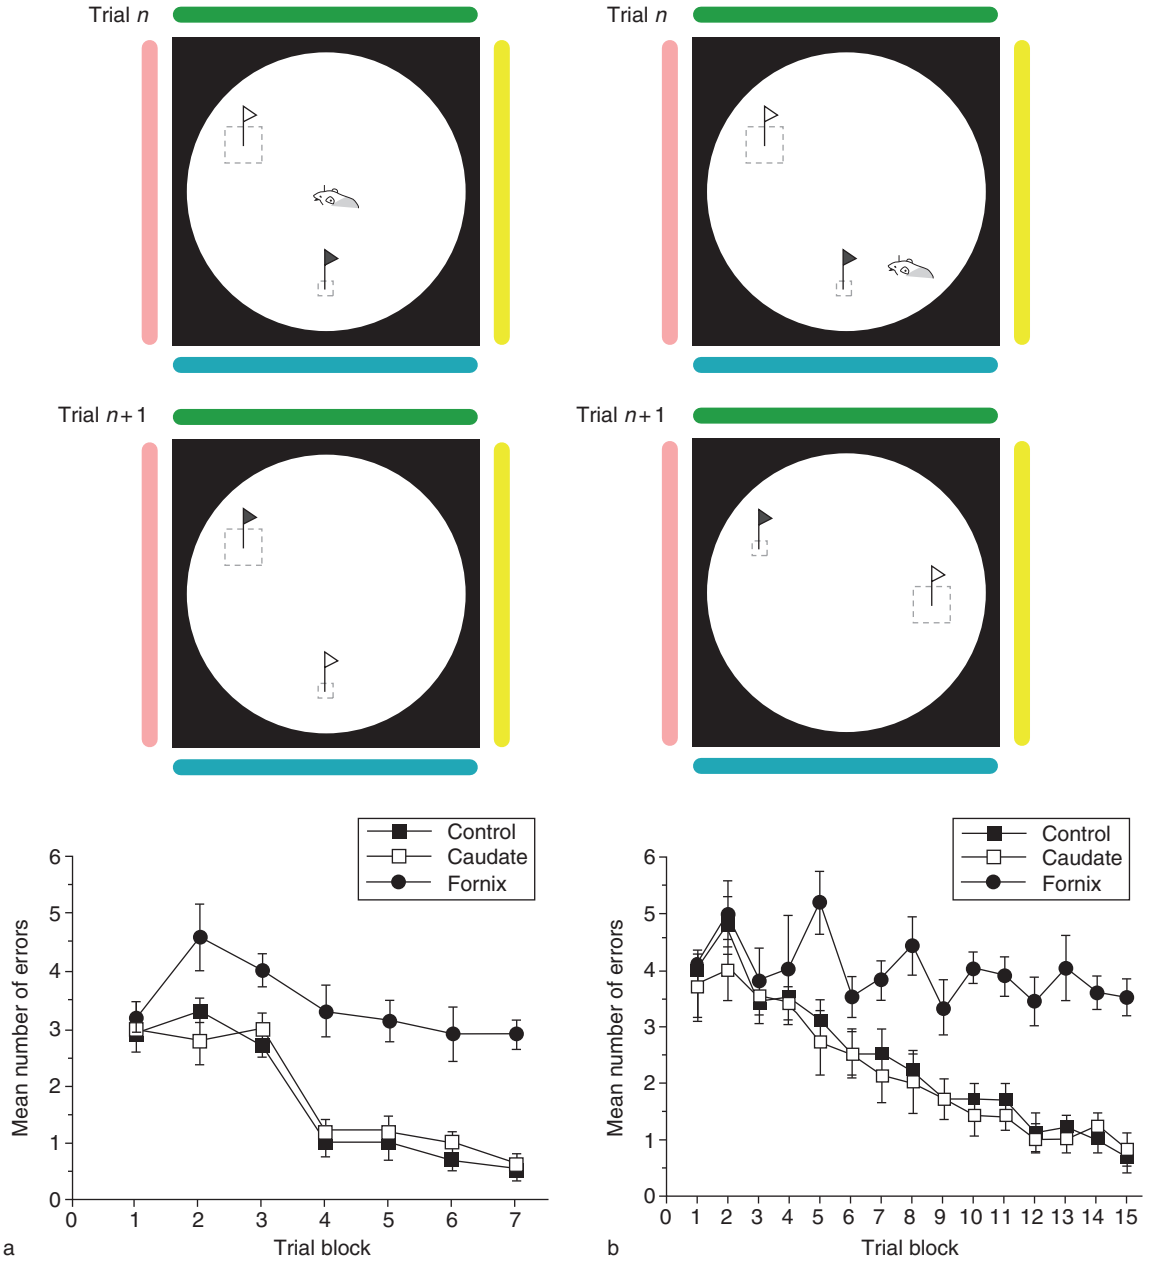

Figure 2 Learning to associate reward to locations on a spatial map depends on the hippocampus, while learning to associate reward to arbitrary visual cues depends on the striatum. (a) Two hidden platforms, one small and one large, are positioned in a constant position with respect to the colored walls. A flag is attached to each platform. On each trial, the flag is randomly changed, but the position of the platforms remains constant with respect to the walls. Each trial block is eight trials. Errors refer to the number of times the animal attempted to mount the small platform. (b) The position of the platforms randomly changes on each trial, but the colors of the flags atop each platform remain constant. Error bars are standard errors of the mean (SEM). Data from Packard MG and McGaugh JL (1992) Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: Further evidence for multiple memory systems. *Behavioral Neuroscience* 106: 439–446.

Okihide Hikosaka and his colleagues have studied the prediction of reward by neurons in the caudate nucleus. In their experiment, the monkey had to make either a left or a right saccadic eye movement, cued by a light spot, in order to keep the trials coming. However, for a block of consecutive trials, only movements to the left target would produce a reward. Then the experimenters reversed that contingency, and after the reversal, only movements to the right target produced reward. The activity of striatal neurons related to leftward saccades increased after just one trial when a leftward target became associated

with reward. In conjunction with that increase in activity, the monkeys responded faster to that cue. Then, after the experimenters switched the reward contingency so that rightward, not leftward, movements produced rewards, the cell slowly returned to its original, lower level of activity as the monkey responded more slowly to the leftward target. Therefore, rewarding behaviors are performed faster, which is certainly quite familiar to parents who request that their child do something. It is important to note that this reduced reaction time is associated with increased activity in the caudate nucleus of the basal ganglia. The basal ganglia seem to be crucial for associating reward to stimuli.

Balancing Expected Rewards with Motor Costs

Associating reward to sensory states is important, but there is another crucial piece of the puzzle: costs of motor commands. Although drinking a cup of tea may provide us with a pleasant sensory state, if that were the only constraint, then we should reach as fast as possible. But people do not reach as fast as possible. Why do we reach with a specific trajectory and speed? What is the cost that prevents us from moving at maximum speed?

The obvious cost is a metabolic cost of the motor commands. Therefore, in choosing how we should move, we need to consider how much reward we might get when we achieve our goal state and the costs associated with the motor commands that we would produce to get to that state. However, for something as simple as an eye movement, it is hard to imagine that the metabolic costs are very high. Is there some cost other than metabolic cost that constrains our movements?

It turns out that a more compelling motor cost is noise. Large motor commands produce noisy, variable movements. That is, if our objective is to reliably achieve a rewarding state, then large motor commands are a bad idea because they will increase the variance of our movements. These ideas were first described by Harris and Wolpert in a theoretical paper that showed that the cost of minimizing variance of our movements nicely explains the trajectories that animals choose for their eye and arm movements. At the heart of the theory is the assumption that large motor commands make our movements more variable and therefore difficult to control. An experiment by Jones and colleagues has demonstrated this idea (**Figure 3**). In this experiment, participants produced a given force with their thumb flexors (**Figure 3(a)**). To guide force production, the participants viewed a cursor that displayed thumb force, but the experimenters analyzed the data during a 4-s period in which this feedback had disappeared. They observed that variance of force was larger when the force magnitude was larger: standard deviation of the force grew linearly with the size of the force

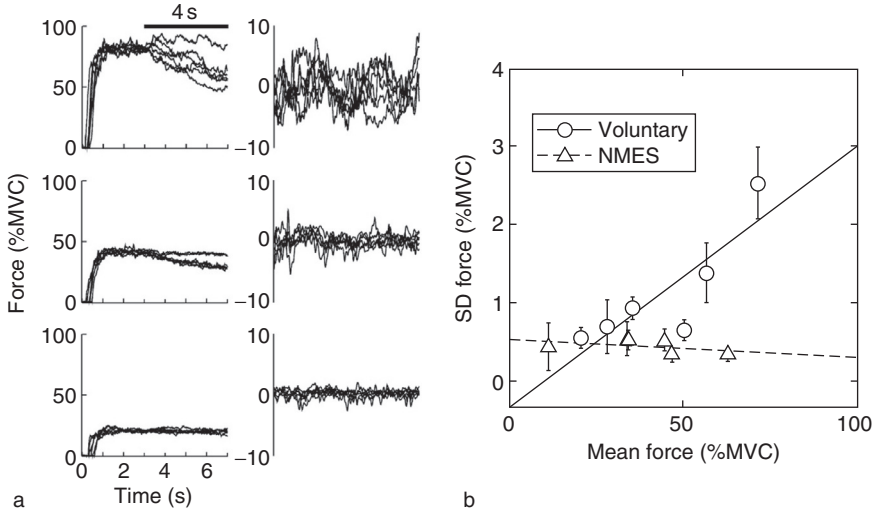

Figure 3 The standard deviation of noise grows with mean force in an isometric task. Participants produced a given force with their thumb flexors. In one condition (labeled 'voluntary'), the participants generated the force, whereas in another condition (labeled neuromuscular electrical stimulation (NMES)), the experimenters stimulated their muscles artificially to produce force. To guide force production, the participants viewed a cursor that displayed thumb force, but the experimenters analyzed the data during a 4 s period in which this feedback had disappeared. (a) Force produced by a typical subject. The period without visual feedback is marked by the horizontal bar on top (4 s), and is expanded in the second column. (b) When participants generated force, noise (measured as the standard deviation (SD)) increased linearly with force magnitude. MVC, maximum voluntary contraction. From Jones KE, Hamilton AF, and Wolpert DM (2002) Sources of signal-dependent noise during isometric force production. *Journal of Neurophysiology* 88: 1533–1544.

(Figure 3(b)). Therefore, large motor commands produce inaccurate movements because the standard deviation of the noise in the signal grows with the size of the signal.

The Problem of How to Do It: Minimizing Costs to Find a Control Policy

Todorov and Jordan have suggested that the movements we make are a reflection of a cost that includes both a measure of expected rewards and a measure of motor costs. They suggested that our movements are optimal with respect to a cost function that includes both motor costs and expected rewards. For example, if we use the vector $\mathbf{y}$ to represent our belief about our current state, vector $\mathbf{r}$ to represent our goal state, and vector $\mathbf{u}$ to represent the motor commands, the objective is to find the motor commands $\mathbf{u}$ that minimize the cost:

$$\sum_{t=1}^{N}(\mathrm{y}^{(t)}-\mathrm{r})^{T}T^{(t)}(\mathrm{y}^{(t)}-\mathrm{r})+\mathrm{u}^{(t)T}L^{(t)}\mathrm{u}^{(t)} \qquad [1]$$

where matrices T and L are measures of the weighted costs associated with achieving the rewarding state and producing the motor commands.

To get to the rewarding state, we need to find the motor commands that will minimize this cost (eqn [1]). To find the motor commands, we need to understand how our motor commands produce changes in our states. For example, if we intend to make a reaching movement, we need to know how the motor commands to our arm muscles are expected to produce sensory feedback. The term 'internal model' (or forward model) refers to this ability of the brain to predict the sensory consequences of motor commands. In a gross simplification of the problem, these predicted consequences may be written as a linear function of our motor commands:

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

where $\hat{\mathbf{x}}^{(t)}$ represents the state of the body and the world that we interact with, $\mathbf{u}^{(t)}$ is our motor command, $\hat{\mathbf{y}}^{(t)}$ is our expected sensory feedback, and $\hat{\mathbf{x}}^{(n+1|n)}$ is our predicted state at time $t+1$, given our previous predictions and observations at time t.

Now our body and the world that we interact with have their own dynamics, represented as follows:

$$\begin{aligned}\mathbf{x}^{(t+1)}&=A\mathbf{x}^{(t)}+B\mathbf{u}^{(t)}+\varepsilon_{u}^{(t)}\\ \mathbf{y}^{(t)}&=C\mathbf{x}^{(t)}+\varepsilon_{y}^{(t)}\end{aligned} \qquad [3]$$

These dynamics are affected by noise: $\varepsilon_{u}^{(n)}$ is a stochastic variable representing motor noise, and $\varepsilon_{u}^{(n)}$ is a stochastic variable representing sensory noise. As we send motor commands, we receive a continuous stream of sensory feedback $\mathbf{y}$. We combine what we predicted with what we observed to update our estimate of the state of our body and the world:

$$\hat{\mathbf{x}}^{(t+1|t+1)}=\hat{\mathbf{x}}^{(t+1|t)}+K^{(t)}(\mathbf{y}^{(t)}-\hat{\mathbf{y}}^{(t)}) \qquad [4]$$

In this equation, the term $K^{(t)}$ is a mixing gain (or a Kalman gain) that determines how much we should change our belief based on the errors between what we predicted and what we observed. Therefore, eqn [2] describes how we make predictions about sensory feedback, and eqn [4] describes how we combine the actual sensory observations with our predictions to update our beliefs.

Our problem is to try to get as much reward as possible for as little cost as possible. How do we generate motor commands that produce the greatest amount of reward? If eqn [2] is an accurate model of how motor commands produce changes in the state of the body and the state of the cursor that we are observing on the screen, then we can use it as a set of constraints with which to minimize eqn [1]. This is the classic linear quadratic problem in optimal control. Solving this problem yields a linear feedback control law that specifies the motor commands as a function of our belief about our state with respect to our goal:

$$\mathbf{u}^{(t)}=G^{(t)}\hat{\mathbf{x}}^{(t|t-1)} \qquad [5]$$

This feedback control law is called a control policy. The new variable G in a time-dependent matrix that represents the gain of the sensorimotor pathway. This way of formulating the problem of motor control is extremely powerful because, in principle, it can account for behavior whether we look at a single act, like a reaching movement, or our actions over the course of a day.

The theory states that our aim is specified as a long-term goal, and our actions try to achieve that goal by implementing a control policy that relies on expected costs and rewards and on our belief about where we are at any given time with respect to that goal. Therefore, the computational problem of motor control may be described as having four components:

1. We live in a world where, at any given time, many actions are possible. For any possible action, we need to know the costs associated with it as well as the sensory states that are rewarding (eqn [1]). We choose to perform an action because we believe it will be more rewarding than other potential actions. The way that we perform that action will depend on our internal estimates of costs and rewards.

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 (eqn [2]). This is a system identification problem associated with our body or tool that we may be trying to control. Learning of this map is called forming a forward model.

3. We must learn how to actually produce the motor commands that are needed so that we minimize the costs and maximize the reward. That is, we need to figure out the 'best' motor commands that bring the cursor to the target and get it to explode. This is the constrained minimization problem (minimize eqn [1] under the constraints of eqn [3]). The result of the minimization is a feedback control law, or a control policy, that specifies our motor response to the sensory states that we observe in our body and the environment (eqn [5]). Learning of this feedback control law is sometimes called forming an inverse model.

4. As we generate motor commands, we make predictions about the sensory consequences (eqn [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 (eqn [4]). Combining predictions with observations is sometimes referred to as Bayesian integration.

At the heart of the approach is the idea that we make movements to acquire rewards. The way that we make a movement, that is, how quickly we move, what trajectory we choose to execute, and how we respond to sensory feedback, depends on our control policy, which itself depends on our belief about the costs and rewards associated with that movement.

Effects of Damage to the Basal Ganglia on Representation of the Cost Function

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 characteristics in which letter sizes become very small and writing speed becomes slow. This condition is most common in degenerative diseases of the basal ganglia such as Parkinson's disease. However, it can also occur in the case of focal lesions. For example, consider 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. When FF was asked to copy a four- or an eight-letter string of letters, the writing with the right hand was much smaller than with the left hand.

Micrographia can be a complex disorder, manifesting itself sometimes in self-generated writing and not in copying of sentences or present in copying letters but not copying geometric shapes. However, the basic observation of smallness of writing size and slowness of writing speed has been puzzling. Current theories suggest that the disorder may be related to a function of the striatum in forming complex repetitive and sequential motor actions. But if we assume that a fundamental role of the striatum is to associate sensory states with their expected reward properties, then the optimal control formulation of motor control provides a different perspective.

In the optimal control framework, there are no desired trajectories for our movements. That is, we do not move our hand along a particular path because that path is somehow hardwired in our brain. Rather, the path is a result of a control policy (eqn [5]), which itself is a result of minimization of a cost function (eqn [1]). The cost function depends on two quantities: an expected reward and a motor cost. What matters is the ratio between these two costs. In theory, movements become small and slow if the cost associated with the motor commands becomes large compared with expected rewards. Therefore, micrographia may be a reflection of a relative loss of expected reward, or increased motor costs, associated with the use of the writing instrument.

If we assume that the control policies that govern our movements (eqn [5]) require an internal measure of costs and rewards (eqn [1]), and if we further assume that the striatum is a crucial node where the brain forms associations between visual cues and expected rewards, then perhaps we can understand bradykinesia (slowness of movements) in Parkinson's disease as a reflection of a disorder in this implicit measure of costs and rewards. Mazzoni and colleagues recently reported that although patients with Parkinson's disease reach more slowly than age-matched controls, they are nevertheless able to make fast movements without a loss in accuracy if forced to move fast by the experimental task. However, the patients take longer (require more trials) to accumulate a set number of movements at the required speed. Thus bradykinesia (slowness of movements), one of the hallmarks of Parkinson's disease, may represent decreased intrinsic value of a given motor task because of an imbalance between an estimate of an implicitly determined cost for making a fast movement and the expected rewards.

Cerebellar Damage and the Ability to Predict Sensory Consequences of Motor Commands

In the optimal control framework, control policies generate motor commands on the basis of beliefs

about the state of the body and the environment (eqn [5]). In eqn [4], we see that the state estimate $\hat{\mathbf{x}}^{(t)}$ depends on two quantities: a prediction $\hat{\mathbf{y}}^{(t)}$ and an observation $\mathbf{y}^{(t)}$. 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 (eqn [2]). The observation comes from the sensory system. The predicted value acts as a prior belief that 'colors' the observed value. 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.

An important part of the optimal control framework is that motor commands continuously depend on predictions about state, which in turn continuously depend on predictions about sensory consequences of motor commands (the internal feedback pathway that relies on the forward model in Figure 2). 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 within 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. Furthermore, proprioceptive signals from the eyes do not play any significant role in controlling saccade trajectories. 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. This idea is supported by the observation that both natural and drug-induced variability of saccade velocity and duration have little influence on saccade amplitude. That is, variability of the saccade's trajectory is partially corrected as the saccade progresses.

What are the neural substrates of this internal feedback? The cerebellum is known to be critical for many aspects of saccade control and adaptation. The projections from the superior colliculus to the cerebellum may provide the efference copy. Indeed, Takeichi and colleagues have shown that adaptive changes in saccade amplitude are already reflected in the nucleus reticularis tegmenti pontis, a major source of input to the cerebellum. Together, the superior colliculus–cerebellar–brain stem side loop seems important for steering saccade trajectories midflight. This side loop is a likely candidate for the role of forward model of the plant to produce midflight corrective feedback.

A simple experiment can test whether the brain has the capacity to predict consequences of self-generated motor commands. In an experiment described by Nowakand colleagues, the participant holds a force transducer that measures grip forces (participant's right hand in Figure 4). Attached to the transducer is a basket. The experimenter drops a ball into the basket. When the ball drops, it exerts a downward

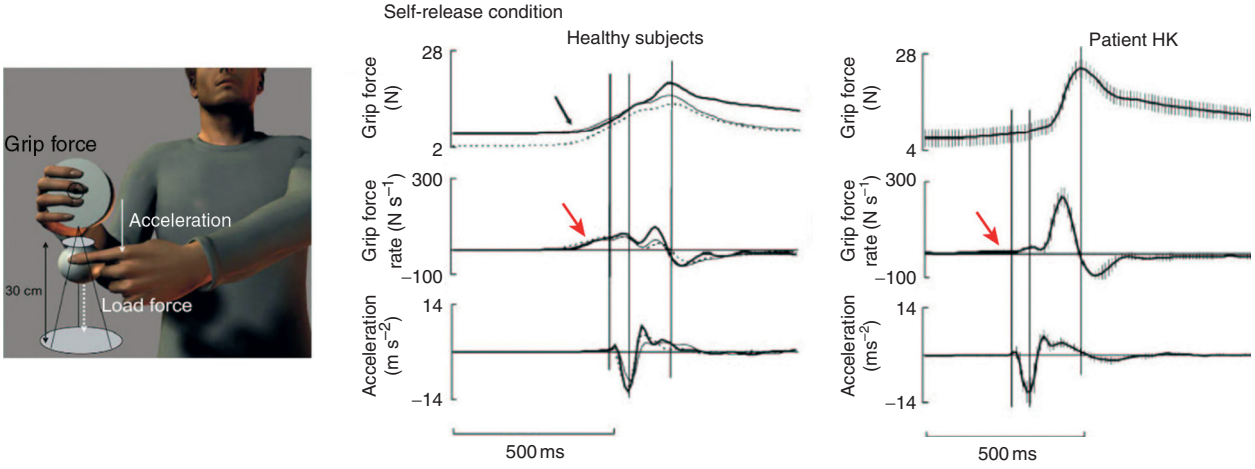

Figure 4 The cerebellum participates in predicting the sensory consequences of motor commands. In an experiment described by Nowak and colleagues, the subject holds a force transducer that measures grip forces. Attached to the transducer is a basket. When a healthy individual drops the ball, the brain predicts the consequences: on impact, the ball will exert a downward force. The traces indicate the grip force of the hand holding the transducer (right hand in the picture), rate of change in the grip force, and acceleration of that hand along the vertical dimension. The first vertical line indicates time of ball impact on the basket. The second vertical line is time of maximum hand acceleration. The third vertical line is time of maximum grip force. To maintain grip, in the healthy subject the right hand exerts a greater force right before the ball hits the basket (blue and red arrows, 50 ms before ball impact). Patient HD, who does not have a cerebellum, does not show this predictive response. Rather, she responds to the sensory feedback from the ball impact by increasing the grip force. From Nowak DA, Timmann D, and Hermsdorfer J (2007) Dexterity in cerebellar agenesis. *Neuropsychologia* 45: 696–703.

load force on the hand. The participant responds by squeezing the transducer so that it will not slip out of the hand. Because there are delays in sensing the drop of the ball, the grip response by the participant is delayed with respect to the drop of the ball (about 100 ms after ball hits the basket, the grip force increases).

Nowak and colleagues have described patient HK, who did not have a cerebellum because of a very rare developmental condition. When the experimenter dropped the ball into the basket, both the healthy individual and HK showed the delayed response. Therefore, the sensory feedback pathways appeared intact in HK. In a subsequent trial, the participant (rather than the experimenter) dropped the ball (**Figure 4**). In a healthy individual, the cerebellum predicts that release of the ball will result in an increased downward load when the ball hits the basket. In anticipation of this event, the healthy individual squeezes harder (with the right hand holding the force transducer) around the time that the ball is released. HK, however, could not make this prediction. 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; that is, it participates in representing a forward model.

Individuals with cerebellar damage are often profoundly impaired in learning to control tools. If the cerebellum is the crucial site for learning forward models (eqn [2]), then it probably makes its contribution to control of reaching via its outputs to the thalamus, which in turn projects to the cerebral cortex. Indeed, artificial stimulation of the ventrolateral thalamus impairs the ability of essential tremor patients to learn the reach adaptation task, even though the stimulation improves the patient's tremor. These results suggest that a crucial function of the cerebellum is to learn forward models that predict sensory consequences of motor commands.

See also: Basal Ganglia and Oculomotor Control; Basal Ganglia: Motor Functions; Bayesian Models of Motor Control; Cerebellar Lesions and Effects on Posture, Locomotion and Limb Movement; Cerebellum and Oculomotor Control; Cerebellum: Clinical Pathology; Cerebellum: Models; Finger Movements: Control.

Further Reading

Barbarulo AM, Grossi D, Merola S, Conson M, and Trojano L (2007) On the genesis of unilateral micrographia of the progressive type. *Neuropsychologia* 45: 1685–1696.

Chen H, Hua SE, Smith MA, Lenz FA, and Shadmehr R (2006) Effects of human cerebellar thalamus disruption on adaptive control of reaching. *Cerebral Cortex* 16: 1462–1473.

Derkinderen P, Dupont S, Vidal JS, Chedru F, and Vidailhet M (2002) Micrographia secondary to lenticular lesions. *Movement Disorders* 17: 835–837.

Fiorillo CD, Tobler PN, and Schultz W (2003) Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299: 1898–1902.

Girard B and Berthoz A (2005) From brainstem to cortex: Computational models of saccade generation circuitry. *Progress in Neurobiology* 77: 215–251.

Guthrie BL, Porter JD, and Sparks DL (1983) Corollary discharge provides accurate eye position information to the oculomotor system. *Science* 221: 1193–1195.

Harris CM and Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394: 780–784.

Hopp JJ and Fuchs AF (2004) The characteristics and neuronal substrate of saccadic eye movement plasticity. *Progress in Neurobiology* 72: 27–53.

Jones KE, Hamilton AF, and Wolpert DM (2002) Sources of signal-dependent noise during isometric force production. *Journal of Neurophysiology* 88: 1533–1544.

Jurgens R, Becker W, and Kornhuber HH (1981) Natural and drug-induced variations of velocity and duration of human saccadic eye movements: Evidence for a control of the neural pulse generator by local feedback. *Biological Cybernetics* 39: 87–96.

Keller EL and Robinson DA (1971) Absence of a stretch reflex in extraocular muscles of the monkey. *Journal of Neurophysiology* 34: 908–919.

Kording KP and Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427: 244–247.

Lauwereyns J, Watanabe K, Coe B, and Hikosaka O (2002) A neural correlate of response bias in monkey caudate nucleus. *Nature* 418: 413–417.

Maschke M, Gomez CM, Ebner TJ, and Konczak J (2004) Hereditary cerebellar ataxia progressively impairs force adaptation during goal-directed arm movements. *Journal of Neurophysiology* 91: 230–238.

Nowak DA, Timmann D, and Hermsdorfer J (2007) Dexterity in cerebellar agenesis. *Neuropsychologia* 45: 696–703.

Optican LM (2005) Sensorimotor transformation for visually guided saccades. *Annals of the New York Academy of Sciences* 1039: 132–148.

Optican LM and Quaia C (2002) Distributed model of collicular and cerebellar function during saccades. *Annals of the New York Academy of Sciences* 956: 164–177.

Packard MG and McGaugh JL (1992) Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: Further evidence for multiple memory systems. *Behavioral Neuroscience* 106: 439–446.

Robinson DA (1975) Oculomotor control signals. In: BachyRita P and Lennerstrand G (eds.) *Basic Mechanisms of Ocular Motility and Their Clinical Implications*, pp. 337–374. Oxford, UK: Pergamon.

Smith MA and Shadmehr R (2005) Intact ability to learn internal models of arm dynamics in Huntington's disease but not cerebellar degeneration. *Journal of Neurophysiology* 93: 2809–2821.

Takeichi N, Kaneko CR, and Fuchs AF (2005) Discharge of monkey nucleus reticularis tegmenti pontis neurons changes during saccade adaptation. *Journal of Neurophysiology* 94: 1938–1951.

Thiele A, Henning P, Kubischik M, and Hoffmann KP (2002) Neural mechanisms of saccadic suppression. *Science* 295: 2460–2462.

Todorov E (2005) Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system. *Neural Computation* 17: 1084–1108.

Todorov E and Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nature Neuroscience* 5: 1226–1235.

Troyer AK, Black SE, Armilio ML, and Moscovitch M (2004) Cognitive and motor functioning in a patient with selective infarction of the left basal ganglia: Evidence for decreased non-routine response selection and performance. *Neuropsychologia* 42: 902–911.

Tsien JZ, Huerta PT, and Tonegawa S (1996) The essential role of hippocampal CA1 NMDA receptor-dependent synaptic plasticity in spatial memory. *Cell* 87: 1327–1338.

VanGemmert AW, Teulings HL, and Stelmach GE (2001) Parkinsonian patients reduce their stroke size with increased processing demands. *Brain and Cognition* 47: 504–512.

Vaziri S, Diedrichsen J, and Shadmehr R (2006) Why does the brain predict sensory consequences of oculomotor commands? Optimal integration of the predicted and the actual sensory feedback. *Journal of Neuroscience* 26: 4188–4197.

Wolpert DM, Ghahramani Z, and Jordan MI (1995) An internal model for sensorimotor integration. *Science* 269: 1880–1882.

Wolpert DM, Miall RC, and Kawato M (1998) Internal models in the cerebellum. *Trends in Cognitive Sciences* 2: 338–347.