

10 Costs and rewards of motor commands

If we were to summarize the concept of internal models thus far, it would go something like this: 1) you begin with a prior belief regarding the state of your body and the environment; 2) you generate some motor command; 3) you predict the resulting change in the state of your body and the environment; 4) you predict the information that should be conveyed by your sensory organs; 5) you combine the sensory feedback with your predictions; and 5) you update your estimate of the state. Internal models are the computational maps that allow you to make the predictions in steps 3 and 4 outlined above: the prediction about the change caused by your motor commands in the state of your body and the environment, and the prediction about the sensory information that should be provided by your sensory afferents. These predictions provide a source of information that is independent of the actual sensory feedback. If your internal models are unbiased, by combining their predictions with the actual sensory feedback you can produce an estimate of state that is better than either source of information alone. Therefore, a fundamental advantage of internal models is that it can improve your ability to sense the world around you.

However, our brain would be nothing more than a fancy surveillance system if all it could do was accurately estimate the state of the environment. We have a central nervous system in order to *do things*, i.e., generate motor commands and perform actions. An often cited example is the sea squirt, a cucumber shaped organism that in its larval form actively swims and is equipped with a brain-like structure of 300 cells (Llinas, 2001). It uses these cells to search for a suitable surface. Once it finds a surface, it buries its head, absorbs most of its brain, and never swims again. So if the purpose of having a nervous system is to move, how does one generate useful motor commands? It turns out that if we have an internal model of our body and the environment, then we can generate motor commands that change the state of our body and the environment to one that is more advantageous for us. That is, the ability to predict the consequences of our motor commands is crucial for generating movements so that we can get *reward*. Without accurate internal models, we would have little ability to perform movements that provide us with rewarding outcomes.

In this chapter and next, we will begin considering the problem of how to produce useful motor commands, i.e., motor commands that can change the state of our body to one that is more rewarding or valuable. We will imagine that there is something of value that we would like to have, and this object is at some position. It is better to receive reward sooner rather than later, so

there will be a cost associated with the time it takes to arrive at this valuable state. We will have to exert some effort in terms of motor commands to get to this rewarding state, and therefore we will have to balance the need for acquiring reward with costs that we are willing to pay in terms of motor commands (or effort). Together, the reward that we expect to get when we arrive at the goal state, and the effort that we expect to expend to arrive at that state, form a cost function. Our objective will be to find a *control policy* that given any current state, produces the ‘best’ motor commands in the sense that the motor commands will bring us to the goal state while minimizing a cost. Internal models allow us to predict how motor commands will change our state, and therefore are critical for forming control policies. Our main theoretical tool will be optimal control theory, a framework in which we will describe the cost of movements. Our ultimate objective is to produce a theory that can not only account for the regularity that exists in movements of people, but also explain why movement patterns change with aging, development, and disease.

10.1 Voluntary eye movements

As you are reading this book, your brain shifts your gaze from one point to another, rapidly moving your eyes. Each movement is a saccade that positions your eyes so that the fovea can sample the visual space. The fovea has densely packed neurons that provide an exquisitely fine resolution of the visual scene. It is like a very high resolution camera. However, the rest of our retina does not have nearly the resolution as the fovea. As we move away from the fovea on the retina, the density of the neurons drops exponentially, and as a result, the visual acuity drops exponentially (Fig. 10.1). It is clear that we make saccades in order to place the image of interest on our fovea. But why don’t we have an eye that has uniformly high resolution capabilities? That way, we wouldn’t have to make saccades. Why is it that our retina is endowed with only one region of high resolution?

In his book *Movements of the Eyes*, Carpenter (1988) wrote:

The reason why the whole visual field is not provided with the luxury grade of vision is presumably [because] the dense packing of the receptors and the one-to-one connections to the optic nerve would lead to the latter becoming unmanageably large. For example, if our entire retina were of the high-quality type, the cross-sectional area of our optic nerve would have to increase by a

factor of over two hundred, and no doubt the size of the blind spot would also have to increase proportionately. (p. 7)

So it would seem that having the very high resolution area limited to a small part of our retina allows us to have a relatively small blind spot, as well as a manageable optic nerve, but at the cost of having to constantly move the eyes so to place the image of interest on our fovea. In fact, voluntary eye movements only occur in vertebrates which have a fovea (Walls, 1962).

In making these saccades, our brain solves two kinds of problems: first, it selects where to look, and next, it programs the motor commands that move the eyes to that location. Let us consider the problem of ‘where to look’. It has long been recognized that the scan sequence, i.e., the locations that people fixate when they are given an image, is not random. The pioneering work on this was performed by Alfred Yarbus (1967), a Russian physicist turned psychologist who invented new methods to record eye movements and wrote the results of his dissertation (which took about 20 years to complete) in a book that was published in 1965 (and translated to English two years later). In thinking about how we direct movements of our eyes to examine a picture, he wrote: “It may be seen to some people that when we examine an object we must trace its outlines with our eye and, by analogy with tactile sensation, ‘palpate’ the object. Others may consider that, when looking at a picture, we scan the whole of its surface more or less uniformly with our eyes.” However, Yarbus performed experiments that showed that motion of the eyes was neither like the motion of the hand in examining a surface, nor uniform like a scanning beam that you find on a copy machine. For example, he presented to his subjects a painting by Shishkin, *Morning in the Pine Forest*, in which four black bears are playing on a fallen tree (Fig. 10.2). He imagined that people might look more at parts of the image that had a lot of detail (like the intricate branches on the trees). He thought that perhaps people might look more at colors that they had previously identified to be ‘their favorite’, but people did not seem to do that. Instead, he found that people tended to move their eyes so that the image of one of the bears fell on the fovea (i.e., fixate one of the bears). When he presented another painting by Shishkin, *Countess Mordvinova’s Forest*, people tended to fixate the hunter, which is only a small figure in a rather large and highly detailed image. For the Petrushavich sculpture *My Child*, the fixations were centered on the faces of the mother and child (especially the eyes), and the hands. From these experiments Yarbus concluded that:

All records of the eye movements show that neither bright nor dark elements of a picture attract the observer’s attention unless they give essential and useful information... When examining complex objects, the human eye fixates mainly

on certain elements of these objects. Analysis of the eye-movement records show that the elements attracting attention contain ... information useful and essential for perception.

Written more formally, the idea is that our brain continuously assigns a value to every part of the visible space, forming a priority or salience map (Gottlieb et al., 1998; Fecteau and Munoz, 2006). Each saccade is a movement with which the brain directs the fovea to a region where currently, the value is highest. Therefore, the reason why we make a saccade is to place the currently most valuable part of the image on our fovea. The landscape of this value function, i.e., the position of the currently most valuable part of the image, would of course depend on the task: if the *Morning in the Pine Forest* painting is shown and the subject is asked to look for a shoe, they are likely to direct their eyes to the underbrush. Generalizing from this, we can speculate that the reason why we move is because our motor commands are expected to change the state of our body (or state of the environment) to something that is more advantageous or rewarding. This thought was elegantly expressed some 700 years ago by the Persian poet Hafez, who wrote:

Because of our wisdom, we will travel far for love,
As all movement is really a sign of thirst,
And speaking really says "I am hungry to know you."

10.2 Expected reward discounts the cost of the motor commands

Perhaps when we are looking at a visual scene, the decision of where to move our eyes is specified by the peaks and valleys of a landscape of values or expected rewards. For example, in viewing a scene consisting of a face and non-face objects, we are naturally drawn to the face region first and spend a longer time looking at the face compared to the rest of the scene (Cerf et al., 2008). Benjamin Hayden, Purak Parikh, Robert Deaner, and Michael Platt (Hayden et al., 2007) found that the opportunity to look at another person of opposite sex is a valued commodity. They found that this was especially true for men, for whom physical attractiveness of faces of women was one dimension along which value increased rapidly. Indeed, face images activate the reward circuitry of the brain (Bray and O'Doherty, 2007; Kampe et al., 2001). The implication is that faces have an intrinsically higher value than other images. This accounts for the fact that images of attractive faces play a prominent role in advertising.

Until recently, it was thought that the motor commands that the brain generates to move the eyes to a given location are generally independent of the value that the brain assigns to that location. Saccades are so short in duration (50-70ms), and the eyes move with such a high velocity (400-500°/s), that the motor commands that move the eyes during a saccade were thought to be invariant programs: given a desired displacement of the eyes, the program played out the motor commands. Indeed, during a saccade the motion of the eyes tends to be stereotypical, exhibiting a fairly consistent relationship between amplitude, duration, and velocity. Yarbus (1967) noted that as saccade amplitudes increased, durations also increased (Fig. 10.3A). As better recording techniques developed, his observations regarding the approximately linear relationship between saccade duration and amplitude were confirmed (Fig. 10.3B). In addition, it also became apparent that during a saccade, the eye velocity had a symmetric profile for small amplitudes, but an elongated profile for larger amplitudes (Fig. 10.3C and Fig. 10.3D) (Bahill et al., 1975; Collewijn et al., 1988).

This regularity is important because in principle, each person could have chosen a different trajectory of moving their eyes: for a given displacement, one subject might have moved their eyes more slowly than the other, with a peak velocity that arrived earlier or later. The fact that we see regularity (in healthy people) is exciting because it encourages us to ask why the brain should control the movements in this way. In this chapter and next, we build a theory that can account for this regularity. Along the way, we will need to establish two ideas: that movements are directed toward rewarding states, and this reward seems to discount the effort it takes to produce that movement. We will attempt to account for the regularity in movements by claiming that the specific way that we move is the best that we could move, given that our goal is to maximize the rate at which we receive reward.

Let us establish the idea that saccade kinematics are affected by the reward that is expected at the end of the movement. We will do this by reviewing data from a few experiments. Yoriko Takikawa, Reiko Kawagoe, Hideaki Itoh, Hiroyuki Nakahara, and Okihide Hikosaka (Takikawa et al., 2002) trained monkeys to make saccades to a remembered target location (as shown in Fig. 10.4A). The targets would appear in one of four positions. In each block of trials only one of the four targets was associated with a juice reward. That is, the monkey would get juice after making a saccade to this baited target. The monkey would nevertheless have to make saccades to the other targets as well (so that the trials would proceed to the baited target), but would not get a reward for the non-baited targets. Examples of the resulting eye positions and velocities are

shown in Fig. 10.4B. The eye velocity is higher, and latency (reaction time) is shorter, when the target is associated with juice as compared to when it is not (Fig. 10.4C). In fact, when the monkey has to make a saccade to a non-rewarding target, the movement has a long duration and the velocity profile occasionally has multiple peaks. Therefore, the relative reward that was associated with the stimulus affected the velocity and latency of the movement. In effect, in programming the motor commands that moved the eyes, the brain was more willing to ‘spend’ the motor commands if the result of the movement was some juice.

This reward dependent variability in saccade velocities was not limited to targets that were baited with juice. For example, consider a scenario in which there are some objects on a table: a cup of tea, a newspaper, a sugar bowl, and some utensils. Suppose you fixate these objects one at a time, making saccades from one to another. Now you decide to reach for the sugar spoon. As you saccade to the spoon, your eye velocity will be higher and the duration of the saccade will be shorter than if you were to make the same amplitude saccade to the spoon but without the intention of reaching for it. That is, saccades that fixate the goal of a reaching movement are faster than saccades that fixate that same image outside the context of reaching (Epelboim et al., 1997; van Donkelaar et al., 2004; Snyder et al., 2002). Perhaps the value of the stimulus that affords the saccade (e.g., the spoon) is higher when that object is also the stimulus that simultaneously affords a reaching movement.

In the laboratory, we like to present the same visual stimulus (often a simple LED) repeatedly, asking the volunteers to fixate it over and over. Just as you might be less interested in looking at a (particularly boring) picture for the 50th time, making a repeating sequence of saccades to a small point of light might also reduce your interest regarding the visual target. Indeed, repeatedly making saccades to the same visual stimulus produces eye movements with smaller velocities and longer durations (Montagnini and Chelazzi, 2005; Chen-Harris et al., 2008; Golla et al., 2008). An example of this is shown in Fig. 10.5A, in which subjects made 60 saccades from one target to another (target LEDs were positioned 15deg apart and displayed at a rate of about once per 1.5sec). The targets appeared in a short, fixed sequence that repeated many times in each set. Haiyin Chen-Harris, Wilsaan Joiner, Vincent Ethier, David Zee, and Shadmehr noted that saccade velocities dropped as the set proceeded, and then recovered after a short break in which the subjects simply closed their eyes. In almost lock-step fashion, durations increased during each set and then dropped after the short break (Fig. 10.5B). Saccade amplitudes, however, remained fairly stable and showed little correlation to changes in duration and velocity (Fig. 10.5C). In

other words, the eye arrived at the target, yet with a high velocity and short duration at the start of each set, and with a lower velocity and longer duration at the end of each set. Were these changes due to some form of fatigue in the muscles of the eyes?

To test for this, Minnan Xu-Wilson, Haiyin Chen-Harris, David Zee, and Shadmehr (Xu-Wilson et al., 2009a) repeated this experiment but now suddenly changed the sequence of stimuli in the middle of a set. The idea was that if the effect on velocities was due to a muscular or neural fatigue, then not much should happen if one makes a small change in the sequence of targets. However, if the effect was due to devaluation of the stimuli because of their predictability, then a sudden but small change in the sequence should produce recovery. That is, novelty might produce an increase in value, and therefore an increase in velocities. They found that on average, during a sequence of 60 saccades the velocities dropped by around 15% (Fig. 10.6B, 'set break'). This decrease was not due to fatigue of the oculomotor muscles or circuitry, because when the sequence of targets suddenly changed in the middle of a set, the velocities once again recovered in the very next saccade (Fig. 10.6B, 'sequence change'). These results suggest, but certainly do not prove, that when a visual stimulus repeats, the internal value that the brain associates with it declines, resulting in a relative increase in the cost of the motor commands, and therefore reduced saccade velocities.

Together, these results suggest that certain manipulations (food, repetition, etc.) alter the implicit value that the brain assigns to the target of the saccade, and that in turn affect the motor commands that move the eyes. To test this idea more directly, Minnan Xu-Wilson, David Zee, and Shadmehr (Xu-Wilson et al., 2009b) asked people to make saccades to foveate a point of light in a dimly lit room. Instead of giving people reward in terms of juice or money, the experiment provided them with an image of a face, an inverted face, an object, or simply noise. The experiment began with fixation of a point of light, and then flash of an image to one side of fixation (Fig. 10.7A). The eyes had to maintain fixation and could not move to look at the image (they of course could identify the image as being face or not, but could not saccade to it). After the image disappeared, the light moved, causing the subject to make a saccade. After completion of the saccade, the subjects were provided with the image that had been previously flashed (Fig. 10.7B). In this way, subjects made a saccade in anticipation of being 'rewarded' by an image. The results showed that saccades that were made in anticipation of seeing the image of a face were faster than inverted faces, objects, and noise (Fig. 10.7D). The implication is that because

faces have an intrinsically higher value than other images, this increased value discounts the motor commands, resulting in faster saccades.

In summary, we perform saccades to change the state of our eyes. Usually, this makes it so that a highly valued part of the visual scene falls on our fovea. Therefore, the motor commands that move the eyes during a saccade change the state of the eyes to one that is more valuable. If we consider saccades as an example of the simplest and most stereotypical movements that our brain controls, we see that even here there is no ‘hard wired’ sequence of motor commands that alters the state of our body. Rather, it appears that the relative gain that we expect from the change in state (the expected value of the final state vs. the value of the current state) affects the strength of motor commands that our brain generates to produce this change in state.

10.3 Movement vigor and encoding of reward

Writing instruments are one of the most common tools that we use in our daily lives. A striking feature 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). However, it can also occur with focal lesions. Fig. 10.8 provides an example of micrographia in a patient who suffered an ischemic stroke in the right basal ganglia, in the head of the caudate nucleus (Barbarulo et al., 2007). When he was asked to copy a four- or eight-letter string of characters with his left hand, he wrote much smaller than with his right hand.

Micrographia reflects an abnormal choice of speed and amplitude and is one manifestation of generalized slowing of movement, called bradykinesia. Bradykinesia is most prevalent in Parkinson’s disease (PD), a disease that arises from loss of neurons in the substantia nigra (a part of the basal ganglia). These neurons release the neurotransmitter dopamine in various parts of the brain, and their activity is strongly tied to expectation of reward (Fiorillo et al., 2003). For example, when animals are presented with a visual stimulus that predicts the amount and probability of reward (juice), after a period of training the dopaminergic cells in the substantia nigra fire in proportion to the value of the stimulus (amount of reward times its probability). That is, the greater expected value of a stimulus, the greater the release of dopamine.

There is evidence that if an animal has to perform a movement to get reward, in conditions for which the movement is expected to be rewarded a larger dopamine response occurs, and this response coincides with a generally faster movement. Tomas Ljungberg, Paul Apicella, and Wolfram Schultz (Ljungberg et al., 1992) trained a monkey to put its hand on a table and wait for opening of a small door that hid a food box (Fig. 9). In some block of trials the box occasionally contained a food reward (called the intermittent task). In the other block of trials, the box always contained a food reward (called the full task). In all trials after the door opened the monkey made a saccade to the box, and when the box was full, the monkey also reached and picked up the food. Interestingly, in the intermittent task in trials in which the box held a piece of food the saccades had a shorter latency (107ms, on average, vs. 140ms in the trials in which the box was empty). After the door opened, dopamine neurons fired, and the firing rate was higher in the trials in which the box contained food vs. for when it was empty. In these trials in which the box contained food, the reach was a bit slower and had a longer reaction time as compared to the reaches that were performed in the full task. The discharge of dopamine neurons was also smaller in the intermittent task with the full box vs. the full task with the full box. Therefore, dopamine response seems to indicate some aspect of the value associated with a stimulus, and on the occasion in which a movement is performed, greater dopamine release often coincides with a more brisk movement (faster, with a shorter reaction time).

In PD, the dopaminergic cells in the substantia nigra gradually die, which presumably results in a reduced valuation of stimuli. Recently, Pietro Mazzoni, Anna Hristova, and John Krakauer (2007) put forth the idea that in PD, bradykinesia is not due to some inherent inability of these patients to make accurate and fast movements, but it is due to a general devaluation of the state changes caused by the movements, which in turn increases an internal measure of motor costs. They wrote: “The motor system has its own motivation circuit, which operates analogously to but separately from explicit motivation... We propose that striatal dopamine energizes action in a more literal sense, namely by assigning a value to the energetic cost of moving.” In effect, they suggested that in PD, there is a high cost associated with motor commands because the stimuli that afford movements have an unusually low value assigned to them by the brain.

To test their idea, they asked subjects to make accurate reaching movements of specified speeds without visual feedback of their hand. Just like healthy controls, PD patients produced movements that had greater endpoint variance when they moved faster. However, these faster movements were simply less probable than in healthy controls. When comparing the

performance of PD to that of control subjects, the 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. Therefore, the PD patients had an increased reluctance to execute movements requiring greater effort, in spite of comparable spatial accuracy with healthy controls. Perhaps the devaluation of the visual stimulus that was produced by loss of dopamine resulted in an unusually large motor costs.

10.4 Motor costs

Let us suppose that the purpose of generating a motor command \mathbf{u} is to change our state \mathbf{x} to one that is more rewarding or valuable. Perhaps the specific motor commands that our brain programs depend on a balance between the expected change in value of state, and the expected effort that it takes to produce the motor commands that cause that change in state. That is, there may be a cost to the motor commands, preventing us from taking actions, unless those actions are discounted by the expected increase in the value of our states. What might be the nature of this motor cost?

Let us imagine a task in which we provide reward for performing an action, and then see how the action that is performed changes as the cost of the motor commands changes. Suppose that you were to have a subject hold onto a rigid bar as shown in Fig. 10.10 (in this case, the subject is a monkey) and provide a task in which the objective is to pull up or down, push left or right, so that a force vector is produced. The handle is attached to a force transducer, and you will reward the subject based on the difference between the goal force \mathbf{f}_g and the force produced \mathbf{f} . Now there are lots of muscles that act on the wrist joint, and each has a pulling direction specified by a unit length vector \mathbf{p}_i . Let us assume that when muscle i is activated by amount $u_i \geq 0$, it will produce a force \mathbf{f}_i such that:

$$\mathbf{f}_i = u_i \mathbf{p}_i \quad (10.1)$$

For the sake of argument, assume that there are 8 muscles that act on the wrist, with pulling directions shown in Fig. 10.10A. The resulting force produced by activation vector

$\mathbf{u} = [u_1, \dots, u_8]^T$ is:

$$\mathbf{f} = P\mathbf{u} \quad (10.2)$$

where $P = [\mathbf{p}_1, \dots, \mathbf{p}_8]$. Given a goal force \mathbf{f}_g , how should these muscles be activated? Let us describe a cost function that penalizes the difference between \mathbf{f}_g and \mathbf{f} , as well as activations \mathbf{u} :

$$J = (\mathbf{f}_g - P\mathbf{u})^T (\mathbf{f}_g - P\mathbf{u}) + \lambda \sum_i u_i^m \quad (10.3)$$

We want to see how the patterns of muscle activations change when we change the motor costs (i.e., the parameters λ and m). We set $|\mathbf{f}_g| = 1$ and vary its direction along a circle and minimize the above cost for various values of λ and m , arriving at $\mathbf{u}^* = [u_1^*, \dots, u_8^*]^T$. For example, for a goal force along direction θ , the optimum activation u_1^* is plotted in Fig. 10.10A as a vector with magnitude u_1^* along direction θ . In Fig. 10.10B, each muscle's activation $u_i^*(\theta)$ is normalized with respect to the maximum activation for that muscle, resulting in a 'tuning function'. The residual force magnitude $|\mathbf{f}_g - \mathbf{f}|$ is plotted in Fig. 10.10C. We note that the muscles are not necessarily most active in their pulling directions. For example, muscles with pulling directions along 270 and 340 degrees are maximally active to force directions other than their pulling directions. We also note that as λ changes, there is little change in the shape of the tuning functions (the effect is mainly on the residuals). Most importantly, as m increases, the tuning functions become broad. That is, when m is small, the muscles do not 'share' the burden of producing force, but rather specialize in their own pulling direction and do little in other directions. In contrast, as m increases, the muscles share the burden, becoming active in many directions. The interesting idea is that the shape of the tuning functions of muscles may be a useful way to infer the nature of the motor costs used by the brain.

The idea that the shape of the muscle tuning functions may be a reflection of the motor cost was put forth by Andrew Fagg, Ashwin Shah, and Andrew Barto (Fagg et al., 2002). They considered the EMG data that was recorded by Donna Hoffman and Peter Strick (Hoffman and Strick, 1999) from the wrist muscles of a monkey who was instructed to produce force in various directions (Fig. 10.11A). Hoffman and Strick had noted that muscle tuning functions were broad, cosine-like, and that several muscles were maximally activated in directions that differed significantly from their pulling directions. [They had established the pulling direction of each muscle by electrically stimulating that muscle and noting the force produced at the handle.] For example, the tuning functions for muscle ECRB (extensor carpi radialis brevis), ECRL (extensor carpi radialis longus), ECU (extensor carpi ulnaris), and FCR (flexor carpi radialis) are shown in Fig.

10.11A. In this figure, the pulling direction of the muscle is noted by the solid line, and the arrow indicates the direction of maximum activation for that muscle during the voluntary force production task. For example, note that for ECRB the pulling direction (as defined by artificial stimulation of the muscle) is about 45 degrees apart from the peak of the tuning function (as defined by voluntary activation of that muscle). The simulation results for increasing values of m are shown in Fig. 10.11B, and for increasing values of various λ are shown in Fig. 10.11C. A reasonable fit to the data is at $m = 2$ and $\lambda = 0.5$, in which the simulations reproduce both the broad tuning and the mismatch between the pulling direction and peak of the tuning function.

Broad tuning functions are a fundamental characteristic of muscles and neurons in the primate motor system. For example, in the primary motor cortex, neurons exhibit broad, cosine-like tuning functions similar to those found in muscles (Georgopoulos et al., 1986). To see why increasing m in Eq. (10.3) broadens the shape of the tuning functions, consider a simplified scenario in which we have two muscles that pull nearly in the same direction. When $m \approx 1$, producing a total force of 1N can be done optimally when one of the muscles produces 1N and the other is inactive. However, when $m = 2$, the smallest motor cost is attained when each muscle produces a force of 0.5N. As Fagg et al. noted: “Under the total squared muscle activation criterion, both muscles would be activated to equal levels, rather than activating one and not the other.”

Quadratic motor costs (i.e., $m = 2$) discourage specialization, and encourage cooperation. The broad tuning of muscles and neurons may arise in part as a reflection of a cost structure in which the motor commands are penalized in a quadratic fashion.

10.5 Motor noise and variability in performance

A significant problem with our approach is that we have not considered the fact that when muscles are activated, they produce forces that have stochastic properties. This stochastic behavior will produce variability in the force, potentially affecting the ability to acquire reward (the first term in Eq. 10.3). Let us show that when we consider the noise properties of muscles, we can no longer infer that their broad tuning is a reflection of motor costs (the second term in Eq. 10.3). Rather, as we will see, the broad tuning can also be associated with maximizing reward.

Recall that the noise properties of muscles are not Gaussian, but signal dependent. In particular, the force generated by motor commands has signal dependent noise properties (Fig. 4.19). We can incorporate this fact into Eq. (10.1):

$$\mathbf{f}_i = u_i \mathbf{p}_i (1 + k_i \phi_i) \quad \phi_i \sim N(0,1) \quad (10.4)$$

We note that the variance of force increases as a function of the signal u_i :

$$\text{var}[\mathbf{f}_i] = u_i^2 k_i^2 \mathbf{p}_i \mathbf{p}_i^T \quad (10.5)$$

And so the standard deviation of force increases linearly with u_i , with slope k_i . The cost J in Eq. (10.3) is of course a scalar, but it now becomes a stochastic variable due to the properties of the first term in Eq. (10.3). The best that we can do is to find the motor commands that minimize the expected value of this cost. A useful identity that will expedite our calculations is the expected value of a ‘squared’ random variable \mathbf{x} :

$$E[\mathbf{x}^T \mathbf{A} \mathbf{x}] = E[\mathbf{x}]^T \mathbf{A} E[\mathbf{x}] + \text{tr}[\mathbf{A} \text{var}[\mathbf{x}]] \quad (10.6)$$

The second term in Eq. (10.6) is the trace operator. Noting that $\mathbf{f} = \sum_i \mathbf{f}_i$, the expected value of our cost is:

$$E[J] = E[\mathbf{f}_g - \sum_i \mathbf{f}_i]^T E[\mathbf{f}_g - \sum_i \mathbf{f}_i] + \text{tr}[\text{var}[\mathbf{f}_g - \sum_i \mathbf{f}_i]] + \lambda \sum_i u_i^m \quad (10.7)$$

The second term in Eq. (10.7) is:

$$\begin{aligned} \text{tr}[\text{var}[\mathbf{f}_g - \sum_i \mathbf{f}_i]] &= \text{tr}\left[\sum_i u_i^2 k_i^2 \mathbf{p}_i \mathbf{p}_i^T\right] \\ &= \sum_i u_i^2 k_i^2 \mathbf{p}_i^T \mathbf{p}_i \end{aligned} \quad (10.8)$$

The expected value of our cost becomes:

$$E[J] = (\mathbf{f}_g - P\mathbf{u})^T (\mathbf{f}_g - P\mathbf{u}) + \sum_i u_i^2 k_i^2 \mathbf{p}_i^T \mathbf{p}_i + \lambda \sum_i u_i^m \quad (10.9)$$

When we compare Eq. (10.9) with Eq. (10.3), we see that when the motor commands have noise properties that are signal dependent, in effect we are adding a quadratic motor cost to our cost function (the second term in Eq. 10.9). As we increase the motor commands, we also increase the variance of these commands, affecting the variance of the output force, the task relevant variable. Because we are rewarded based on a measure of accuracy in our force, increasing the motor commands will cost us in terms of accuracy, and this cost appears as a squared term on the motor commands. [This would not be the case if the noise process was Gaussian, i.e., a process in which noise variance was independent of the mean of the signal.]

For example, when the system is noiseless ($k_i = 0$ in Eq. 10.4), setting $m \approx 1$ produces narrowly tuned muscle activation patterns (Fig. 10.12A). However, when the same system has signal dependent noise ($k_i = 1$ in Eq. 10.4), setting $m \approx 1$ now produces broadly tuned muscle activation patterns (Fig. 10.12B). Therefore, the broadly tuned activation patterns (for example, the patterns observed in the wrist muscle in Fig. 10.10) are not necessarily a result of a quadratic motor cost. Rather, the same broadly tuning can result when the system suffers from signal dependent noise. This is because the motor noise affects accuracy, and reward is often a function of accuracy. The problem that we face is that noise in the motor commands affects our accuracy, which in turn affects our ability to get reward. We need a way to dissociate motor costs from the effects that the motor commands have on accuracy.

10.6 Maximizing performance while minimizing effort

How can we dissociate between costs that truly penalize effort, and costs that penalize the inaccuracy that results from their noise? We need to know the noise properties of the system that generates force. With this information, we can estimate how much this noise affects the cost function, and then see if the system still penalizes the motor costs over and above those that arise from signal dependent noise properties of the system.

The first experiment that dissociated performance variability vs. effort costs was performed by Ian O'Sullivan, Etienne Burdet, and Jörn Diedrichsen (2009). In the task, volunteers were asked to put one finger of the right hand on a force transducer, and one finger of the left hand on another force transducer (Fig. 10.12). They pushed with the two fingers (either the index or the little finger on each hand) so that a cursor whose position represented the sum of the two forces was placed in a target area during a 7 second period. The subjects were rewarded inversely proportional to the mean squared error of force during that period. Let us assume that the displacement of the cursor due to the force produced by a finger on the right hand is:

$$x_i \propto N(u_i, k_i^2 u_i^2) \quad (10.10)$$

Further suppose that there are no motor costs, and the only objective is to minimize the squared error between the sum of the forces produced by the two fingers and the goal force:

$$\begin{aligned}
J &= E \left[(x_i + x_j - g)^2 \right] \\
&= E \left[(x_i + x_j - g) \right]^2 + \text{var} \left[(x_i + x_j - g) \right] \\
&= (u_i + u_j - g)^2 + k_i^2 u_i^2 + k_j^2 u_j^2
\end{aligned} \tag{10.11}$$

Minimizing this cost results in the following optimal motor commands:

$$\begin{aligned}
\frac{dJ}{du_i} &= 2(u_i + u_j - g) + 2k_i^2 u_i = 0 & u_i^* &= \frac{g - u_j}{1 + k_i^2} \\
\frac{dJ}{du_j} &= 2(u_i + u_j - g) + 2k_j^2 u_j = 0 & u_j^* &= \frac{g - u_i}{1 + k_j^2}
\end{aligned} \tag{10.12}$$

Simplifying the first equation in Eq. (10.12) we have:

$$u_i^* = \frac{g - \frac{g - u_i}{1 + k_j^2}}{1 + k_i^2} = \frac{gk_j^2 + u_i}{(1 + k_i^2)(1 + k_j^2)} \tag{10.12}$$

We can now view the optimal motor commands by describing the percentage of the force produced by one hand as a function of the total force by the two hands:

$$\frac{u_i^*}{u_i^* + u_j^*} = \frac{k_j^2}{k_i^2 + k_j^2} \tag{10.13}$$

The result in Eq. (10.13) implies that if the noise on the left finger is more than on the right, then the right hand should contribute more to the total force. The implication is that if the cost function depends only on acquisition of reward (Eq. 10.11), then the subject should divide the total force based on the ratio of the noise on each finger. O'Sullivan et al. measured the noise properties of each finger on each hand using the same procedure as in the main task, arriving at an estimate of the rate at which standard deviation of the force for each finger increased as a function of mean force (data for a typical subject is shown in Fig. 10.13A). Using this estimate of k for each of the two fingers of each hand, they then computed the optimal ratio of forces and compared it to the actually measured ratio. If the selection of forces on the right and left hand depended solely on maximizing reward, then the data should fall on the diagonal line of Fig. 10.13B. The actual data did not agree well with this prediction.

As an alternative, O'Sullivan et al. considered a cost that not only penalized a measure of inaccuracy in the task relevant variable, but also a measure of motor costs:

$$J = vE \left[(x_i + x_j - g)^2 \right] + \lambda (u_i^2 + u_j^2) + \mu (u_i^2 / MVC_i^2 + u_j^2 / MVC_j^2) \quad (10.14)$$

The term MVC_i refers to the maximum voluntary contraction of muscle i . In Eq. (10.14), motor commands are penalized in terms of their absolute value, and also in terms of their ratio with respect to MVC. Simplifying the above equation we have:

$$\begin{aligned} J &= v(u_i + u_j - g)^2 + (vk_i^2 + \lambda + \mu / MVC_i^2) u_i^2 + (vk_j^2 + \lambda + \mu / MVC_j^2) u_j^2 \\ &= v(u_i + u_j - g)^2 + a_i u_i^2 + a_j u_j^2 \end{aligned} \quad (10.15)$$

where $a_i = vk_i^2 + \lambda + \mu / MVC_i^2$. Minimizing this cost we have:

$$\begin{aligned} \frac{dJ}{du_i} &= 2v(u_i + u_j - g) + 2a_i u_i \\ u_i^* &= -\frac{v(u_j - g)}{(v - a_i)} = \frac{vga_j}{-va_i - va_j + a_i a_j} \end{aligned} \quad (10.16)$$

Once again we can view the optimal motor commands by describing the percentage of the force produced by one hand as a function of the total force by the two hands:

$$\frac{u_i^*}{u_i^* + u_j^*} = \frac{a_j}{a_j + a_i} \quad (10.17)$$

The authors then fit the model (Eq. 10.17) to the data (actual ratio of the forces) and found the two free parameters. [There were only two free parameters because of the constraint that the sum of the three parameters $v + \lambda + \mu$ should equal one.] Performance of the fitted model and the parameter values are shown in Fig. 10.13C and Fig. 10.13D. The term that penalized the quadratic motor costs had the largest value.

Therefore, in this simple bimanual force production task, the distribution of forces between the two hands was not done to maximize accuracy (despite the fact that maximizing accuracy was the criterion for which the subjects were rewarded). Rather, the cost of maximizing accuracy was discounted by a cost that penalized the motor commands, i.e., a motor cost. The most costs appeared to be approximately associated with the squared force produced by each finger.

Motor costs during a movement

Most of the experiments that we have considered thus far were isometric. Motor costs also appear to influence planning of voluntary movements. Our first example is a simple task described by Y. Uno, Mitsuo Kawato, and R. Suzuki (1989) (shown in Fig. 10.14). The objective for the volunteers was to reach from one point to another. In one condition, the subject holds a light-weight 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. 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.

In our second example, the task is to move one's hand from one point to another in a given amount of time (450ms), but now instead of a spring, there is a velocity dependent force field that pushes the hand perpendicular to its direction of motion (Shadmehr and Mussa-Ivaldi, 1994). The field is shown as a sequence of arrows in Fig. 10.15A. We might consider two kinds of cost functions here: one that minimizes a measure of kinematics (e.g., a measure of smoothness like squared jerk of the hand), and one that attempts to bring the hand to the target while minimizing a measure of effort (e.g., squared force). Before the field is imposed, the motion that minimizes both kinds of costs is simply a straight line with a bell-shaped velocity profile. However, when the field is imposed, the solution for the kinematic cost remains a straight line, while the solution for the effort cost is no longer a straight line. Jun Izawa, Tushar Rane, Opher Donchin, and Shadmehr (Izawa et al., 2008) showed that if the field pushes the hand to the left, the policy that produces the least effort 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. 10.15A). To test this idea, they had subjects train in the field for a long period (3 sessions over 3 days). They found that with training, subjects did not return their hand trajectories to a straight line. Rather, they produced a curved trajectory (Fig. 10.15B). To see the rationale for this behavior, Fig. 10.15C plots the forces produced by a policy that minimizes effort, 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 controller produces less total squared 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 that appears to over-compensate for the forces actually requires less total force than a straight trajectory. If we define effort as the sum of squared force, then this path is the one that minimizes an effort cost.

These examples suggest that movement trajectories are a result of motor commands that produce changes in state that attempt to meet task goals (e.g., bring the hand to the target), while minimizing some measure of effort (e.g., cumulative squared force). The task goals describe how we will be rewarded for our efforts, while task costs describe how much we will work to acquire this reward.

10.7 Summary

The reason why we move may be because our motor commands are expected to change the state of our body (or state of the environment) to something that is more advantageous or rewarding. For example, each saccade is a movement with which the brain directs the fovea to a region where currently, the value is highest. Images of faces tend to have higher value for us humans than most other images, which is why we tend to first look at the face in a portrait.

The relative gain that we expect from the change in state (the expected value of the final state vs. the value of the current state) affects the strength of motor commands that our brain generates to produce this change in state. For example, saccades that are made to stimuli that are expected to be more rewarding have higher speed. That is, the expected reward discounts the motor commands. Diseases that affect the reward prediction mechanism of the brain produce motor disorders that maybe understood in terms of unusually small or large discounting of motor commands.

Discounting of motor commands implies that the brain maintains an implicit cost associated with generating a movement. To explore the nature of these motor costs, a number of studies have considered how the brain solves redundancies in which many muscles contribute to generating a force. Under an assumption of noiseless or Gaussian noise process, penalizing motor commands tends to discourage specialization of muscles and encourage cooperation. The implication is that the broad tuning of muscles and neurons are a reflection of a cost structure in which the motor commands are penalized.

However, motor commands tend to be affected by a noise process that is signal dependent, in which the standard deviation of the signal increases with the mean. In this more biologically plausible scenario, the same broad tuning of muscles and neurons arises not from a motor cost, but a cost associated with minimizing variance (i.e., a task cost that attempts to maximize

reward). Recent experiments, however, suggest that motor commands are penalized over and above the cost associated with task performance. That is, our motor commands are due to a cost structure that includes two components: attain the greatest reward possible, while being as lazy as possible. In the next chapter we will consider the third component of this cost, time.

Figure legends:

Figure 10.1. Spatial acuity as a function of distance from the center of fovea. The line represents the function $y = 0.8(1 - \log_{10} x)$, indicating that acuity declines exponentially as a function of distance from the fovea. (Plot constructed from Carpenter (1988), using data from Green (Green, 1970), and Jacobs (Jacobs, 1979)).

Figure 10.2. Saccade scan paths for viewing of images. Top image: Reproduction of Shishkin's painting *Morning in the Pine Forest*. The records of the eye movements were made during free examination of the picture during 2 minutes. Middle image: Shishkin's painting *Countess Mordvinova's Forest*. The records were made during 10 minutes of free examination. Bottom image: Petrushavich's sculpture *My Child*. The records were made during 2 minutes of free examination. (Data reproduced from Yarbus (1967)).

Figure 10.3. Kinematic characteristics of human saccades. **A.** Duration and amplitude relationship for small to medium size saccades of two subjects (Figure from Yarbus (1967)). **B.** Duration and amplitude relationship for a larger range of saccade amplitudes from three subjects. The error bars are 1SD. Note that the variance of the saccade durations increases with saccade amplitude. **C & D.** Amplitude and velocity properties of saccades for one subjects (average of four trials). (Figures B, C, and D are from Collewyn et al. (1988) with permission).

Figure 10.4. Expected reward affects control of movements. **A.** The task for the monkey is to saccade to the remembered location of 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. **B.** The location of the rewarded target is identified by the label of the column. The top row 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. **C.** Comparison of saccade peak velocities and latencies for rewarded and unrewarded movements. Peak velocities are higher and latencies are shorter for movements that are expected to be rewarded. (From Takikawa et al. (2002) with permission).

Figure 10.5. Repeatedly making saccades to the same target reduces saccade velocities. **A.** In a darkened room, volunteers made saccades to targets that appeared in a short, fixed sequence, at a

rate of about one saccade per 1.5 seconds. The figure shows saccade velocities for targets that appeared at 15 degrees displacement. Set breaks were about 30 seconds in duration. **B.** Saccade durations increased as velocities decreased. **C.** Eye displacement showed little or no correlation to changes in duration and velocity. (From (Xu-Wilson et al., 2009a) with permission).

Figure 10.6. Changes in saccade velocities may be related to predictability of the stimulus. **A.** In a darkened room, volunteers made saccades to targets that appeared in a short, fixed sequence. The figure shows saccade velocities for targets that appeared at 15 degrees displacement. In the 10th and 11th set, the target sequence suddenly changed. The dashed lines indicate brief set breaks. **B.** The data was aligned to the saccade following set breaks and to the saccade following sequence change. (From (Xu-Wilson et al., 2009a) with permission).

Figure 10.7. The intrinsic value of visual information affects saccade velocities. **A.** Study protocol. Volunteers fixated an LED in a darkened room while the picture was flashed at 15 degrees. Subjects maintained fixation until the LED moved. After completion of a saccade, the volunteers were shown the same picture. Therefore, saccades were made in anticipation of viewing a face, an object, inverted face, or noise. **B.** Examples of pictures used in the experiment. **C.** Subjects made saccades with higher velocities in anticipation of seeing a face. **D.** Relative decline in peak speed of saccades with respect to face. (From (Xu-Wilson et al., 2009b) with permission).

Figure 10.8. 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).

Figure 10.9. Behavioral and dopamine response to a food reward. In animal faces a wall on which there is a food box. The door of the box opens rapidly upward and the movement is visible and audible to the animals. The animal usually makes a saccade to the box. In the intermittent task, the box occasionally contains food, whereas in the full task, the box always contains food. When the box contains food, the animal reaches into the box and collects the reward. The behavioral performance table shows values in units of ms with ZSE. Reaction and movement times are significantly longer in the intermittent than in the full task. Saccade latencies are significantly longer in trials without food in the box. The neural discharge table shows magnitude

of dopamine firing after door opening in a standard time window. (From (Ljungberg et al., 1992)).

Figure 10.10. Minimization of the cost specified by Eq. (10.3) with the muscle model of Eq. (10.2). Broad tuning curves result from increased motor costs. **A.** Muscle activation patterns. **B.** Normalized activation patterns, i.e., tuning curves. **C.** Force residuals.

Figure 10.11. Wrist activation patterns during an isometric force production task. **A.** Normalized EMG activation patterns for four wrist muscles: ECRB (extensor carpi radialis brevis), ECRL (extensor carpi radialis longus), ECU (extensor carpi ulnaris), and FCR (flexor carpi radialis). The line without an arrow indicates the pulling direction of the muscle when it is electrically stimulated, i.e., the direction of force that the muscle produces when it is stimulated. The line with an arrow indicates preferred direction of the muscle, i.e., direction of force for which the muscle is maximally activated during voluntary recruitment. **B** and **C.** Muscle tuning curves that minimize the cost specified by Eq. (10.3). (The data in part A from (Fagg et al., 2002).)

Figure 10.12. Motor commands that suffer from signal dependent noise result in broad tuning functions of muscles. **A.** The noiseless system. The pattern of motor commands that minimize Eq. (10.9) with the constraint given in Eq. (10.4), with $k_i = 0$, $m = 1.1$, and $\lambda = 0.5$. **B.** A system with signal dependent noise. The pattern of motor commands that minimize Eq. (10.9) with the constraint given in Eq. (10.4), with $k_i = 1$, $m = 1.1$, and $\lambda = 0.5$. The format of this figure is similar to Fig. 10.8.

Figure 10.13. Volunteers were asked to push on two force transducers so that the sum of the two forces, as displayed by a cursor, was equal to a goal. They were rewarded inversely proportional to the mean squared error of the force. **A.** When subjects were asked to perform the task with a single finger, the standard deviation of the force (as shown for one subject) increased as a function of the mean force. This exemplifies the signal dependent noise properties of the motor commands. **B.** If the cost minimized by the subjects depended only on a measure of accuracy (i.e., reward), then the optimal motor commands for each condition and each hand/finger combination should be along the dashed line. **C.** When the costs explicitly penalize motor commands, as well as inaccuracy (Eq. 10.14), there is a better fit between the predicted motor

commands and the actually observed values. **D.** Parameter estimate for the cost function in Eq. (10.14). (From (O'Sullivan et al., 2009) with permission.)

Figure 10.14. Motor costs affect reach trajectories. **A.** The task is to reach from one point to another. In one condition, the reach takes place in free space (straight line). In another condition, a spring is attached to the hand. In this experiment, the maximum force produced by the spring was 10.4N and the minimum was 3.3N. **B.** Data from four reaching movements of a single subject. The subject chooses to move the hand along an arc. (Data redrawn from Uno et al. 1989).

Figure 10.15. During adaptation, the motor commands attempt to minimize motor costs. **A.** 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 a motor costs composed of the sum of squared forces 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 1st and 3rd trials, all other trajectories are average of 50 trials. **B.** A rationale for why a curved movement is of lower cost. The curves show simulation results on forces $\mathbf{f} = \begin{bmatrix} f_x & f_y \end{bmatrix}^T$ that the controller produces, and speed of movement $\sqrt{\dot{x}^2 + \dot{y}^2}$, in an optimal control scenario in which the objective is to minimize total motor costs $J = \int_0^p \mathbf{f}^T \mathbf{f} dt$, and in a scenario in which the objective is to minimize squared jerk $J = \int_0^p (\ddot{x}^2 + \ddot{y}^2) dt$. (Data from Izawa et al. 2006).

Reference List

- Bahill AT, Clark MR, Stark L (1975) The main sequence: a tool for studying human eye movements. pp 191-204.
- 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.
- Bray S, O'Doherty J (2007) Neural Coding of Reward-Prediction Error Signals During Classical Conditioning With Attractive Faces. *J Neurophysiol* 97:3036-3045.
- Carpenter RHS (1988) *Movements of the Eyes*. London: Pion Ltd.
- Cerf M, Harel J, Einhasuer W, Koch C (2008) Predicting human gaze using low-level saliency combined with face detection. In: *Advances in Neural Information Processing Systems* Cambridge, MA: MIT Press.
- Chen-Harris H, Joiner WM, Ethier V, Zee DS, Shadmehr R (2008) Adaptive control of saccades via internal feedback. *J Neurosci* 28:2804-2813.
- Collewijn H, Erkelens CJ, Steinman RM (1988) Binocular co-ordination of human horizontal saccadic eye movements. *J Physiol* 404:157-182.
- Epelboim J, Steinman RM, Kowler E, Pizlo Z, Erkelens CJ, Collewijn H (1997) Gaze-shift dynamics in two kinds of sequential looking tasks. *Vision Res* 37:2597-2607.
- Fagg AH, Shah A, Barto AG (2002) A computational model of muscle recruitment for wrist movements. *J Neurophysiol* 88:3348-3358.
- Fecteau JH, Munoz DP (2006) Saliency, relevance, and firing: a priority map for target selection. *Trends Cogn Sci* 10:382-390.
- Fiorillo CD, Tobler PN, Schultz W (2003) Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299:1898-1902.
- Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neural population coding of movement direction. *Science* 233:1416-1419.
- Golla H, Tziridis K, Haarmeier T, Catz N, Barash S, Thier P (2008) Reduced saccadic resilience and impaired saccadic adaptation due to cerebellar disease. *Eur J Neurosci* 27:132-144.
- Gottlieb JP, Kusunoki M, Goldberg ME (1998) The representation of visual saliency in monkey parietal cortex. *Nature* 391:481-484.
- Green DG (1970) Regional variations in the visual acuity for interference fringes on the retina. *J Physiol* 207:351-356.

- Hayden BY, Parikh PC, Deaner RO, Platt ML (2007) Economic principles motivating social attention in humans. *Proc Biol Sci* 274:1751-1756.
- Hoffman DS, Strick PL (1999) Step-tracking movements of the wrist. IV. Muscle activity associated with movements in different directions. *J Neurophysiol* 81:319-333.
- Izawa J, Rane T, Donchin O, Shadmehr R (2008) Motor adaptation as a process of reoptimization. *J Neurosci* 28:2883-2891.
- Jacobs RJ (1979) Visual resolution and contour interaction in the fovea and periphery. *Vision Res* 19:1187-1195.
- Kampe KKW, Frith CD, Dolan RJ, Frith U (2001) Psychology: Reward value of attractiveness and gaze. *Nature* 413:589.
- Ljungberg T, Apicella P, Schultz W (1992) Responses of monkey dopamine neurons during learning of behavioral reactions. *J Neurophysiol* 67:145-163.
- Llinas RR (2001) *I of the vortex: From neurons to sefl*. Cambridge, MA: MIT Press.
- 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.
- Montagnini A, Chelazzi L (2005) The urgency to look: Prompt saccades to the benefit of perception. *Vision Research* 45:3391-3401.
- O'Sullivan I, Burdet E, Diedrichsen J (2009) Dissociating variability and effort as determinants of coordination. *PLoS Comput Biol* 5:e1000345.
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208-3224.
- Snyder LH, Calton JL, Dickinson AR, Lawrence BM (2002) Eye-hand coordination: saccades are faster when accompanied by a coordinated arm movement. *J Neurophysiol* 87:2279-2286.
- 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.
- 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 Donkelaar P, Siu KC, Walterschied J (2004) Saccadic output is influenced by limb kinetics during eye-hand coordination. *J Mot Behav* 36:245-252.
- Van Gemmert AW, Teulings HL, Stelmach GE (2001) Parkinsonian patients reduce their stroke size with increased processing demands. *Brain Cogn* 47:504-512.
- Walls GL (1962) The evolutionary history of eye movements. *Vision Res* 2:69-80.
- Xu-Wilson M, Chen-Harris H, Zee DS, Shadmehr R (2009a) Cerebellar contributions to adaptive control of saccades in humans. *J Neurosci* 29:12930-12939.

Xu-Wilson M, Zee DS, Shadmehr R (2009b) The intrinsic value of visual information affects saccade velocities. *Exp Brain Res* 196:475-481.

YARBUS AL (1967) *Eye movements and vision*. New York: Plenum Press.