

1

THEORETICAL MODELS OF MOTOR CONTROL AND MOTOR LEARNING

Adrian M. Haith¹ and John W. Krakauer²

¹NEUROLOGY / BIOMEDICAL ENGINEERING, JOHNS HOPKINS UNIVERSITY, BALTIMORE, USA

²DEPARTMENTS OF NEUROLOGY AND NEUROSCIENCE, JOHNS HOPKINS UNIVERSITY, BALTIMORE, USA

Introduction

The ease with which one is able to walk, talk, manipulate objects and play sports belies the fact that generating coordinated movement is a tremendously complex task. We have to control our bodies through a muscular system that is highly redundant, nonlinear and unreliable. Furthermore, we are reliant on sensory feedback that is also unreliable and substantially delayed. Yet many tasks that robotic systems achieve either clumsily or not at all are routine to us. Expert performers push the limits of performance even further. Our advantage over synthetic manipulators and – arguably – a professional sportsperson's advantage over a rookie, lies not so much in the hardware performing the task, but in the way it is controlled.

A theoretical approach to motor control and motor learning seeks to explain regularities in behavior in terms of underlying principles. This typically entails formulating mathematical models that describe the mechanics of the body or task, the way in which appropriate motor commands are selected, or the way in which prior experience influences future behavior. Many theories are mechanistic in nature – appealing to computations or plasticity occurring at the level of individual neurons or synapses in order to explain observations at the behavioral level. More abstract theories may not necessarily refer to any specific neural substrate but instead seek to explain behavior in terms of the way in which information can be represented and transformed. In both approaches, predictions about behavior stem largely from constraints imposed by the assumed circuitry or algorithm. Note that these two modelling approaches are similar to the implementational and algorithmic levels of analysis discussed by Marr (1982).

An alternative approach is to set aside questions about mechanism or algorithms and attempt to characterize and understand motor system function purely at the behavioral level. The sheer flexibility of the motor system makes it seem unlikely that underlying mechanisms place a significant constraint on the kinds of movement that can be generated. Instead, it seems that regularities in behavior are mostly dictated by ... mostly dictated by features of

the task at hand rather than by features of the underlying implementational mechanism. A *normative* modeling approach seeks to explain behavior by first understanding the precise computational problem that the brain faces, and then asking what, theoretically, is the best possible way to solve it (akin to Marr's computational level of analysis). Finding solutions to such problems typically leverages ideas from control theory or machine learning. Mechanistic and normative approaches are far from mutually exclusive endeavors – breakthroughs in normative models of behavior often inspire and help guide mechanistic models. A deeper mechanistic understanding can help to constrain normative models. The normative point of view effectively assumes that the underlying neural mechanisms have omnipotent capacity. Consequently, aspects of the task itself, rather than the underlying mechanisms responsible for implementing the solution, are what primarily dictate our patterns of behavior.

In this chapter, we provide an introduction to the core concepts that underlie most recent theoretical models of motor control, state estimation and motor learning. We examine the assumptions – many of which often go unchallenged – underlying these models and discuss common pitfalls in their application. Finally, we discuss important unanswered questions and consider possible future directions for research.

Theoretical models of control

The fundamental problem the motor system faces is to decide upon appropriate motor commands to bring about a desired outcome in the environment. For example, suppose you want to move your hand to push a button to call an elevator. What makes this problem difficult is that it is not enough to simply know the location of the elevator button in space. Changing the position of the arm can only be done very indirectly by using the muscles to generate forces that cause acceleration about the joints of the arm. Thus the dynamics of our bodies place a fundamental constraint on how we are able to move. Furthermore, these dynamics are highly nonlinear – the exact same motor commands may lead to a very different acceleration depending on the state of the arm and muscles. As a result, even an apparently simple task like a point-to-point movement actually requires a complex sequence of motor commands to achieve success.

Compounding the fact that task goals are distally and nonlinearly related to motor commands, motor execution itself is highly unreliable – forces generated by a muscle are inherently variable. Though there are many potential mechanisms that contribute to this variability (Faisal et al., 2008), the net effect appears to be that force variability grows linearly with force magnitude (Jones et al., 2002) – a phenomenon known as *signal-dependent* or multiplicative noise. Managing this noise to minimize its negative impact is a major theme in models of the motor system. Since execution noise acts at every instant, its effects will, if unchecked, accumulate over the course of a movement so that even moderate variability can end up significantly interfering with task goals.

In addition to the problems of acting through potentially complex dynamics and noise, a further factor that complicates the control problem is *redundancy*. Although a movement goal might be specified by a unique point in space, there is no such unique set of controls to get there. The movement may take many potential paths through space, may take any amount of time and vary in speed during movement in infinitely many ways, all of which must be achieved using very different muscle activations. Even for a given trajectory, many different combinations of muscle activations can lead to exactly the same kinematic outcome, just with varying degrees of co-contraction. The importance of redundancy in the motor

system has long been recognized. Since the time of Bernstein (1967), who dubbed the need to resolve redundancy as ‘the degrees-of-freedom problem’, it has been common to regard redundancy as a nuisance that the motor system has to deal with on top of all the other factors that complicate control. However, far from viewing redundancy as a problem, redundancy should actually be regarded as a positive thing. It makes it easier to find solutions to a given task and allows goals to be achieved more flexibly and robustly. Redundancy, therefore, makes life easier for the motor system to develop adequate means of control and in general enables superior control strategies. However, redundancy complicates life for the motor system in the sense that it leads to a more complex and challenging control problem if one wants to exploit it intelligently.

The considerable redundancy in a task such as a point-to-point reaching movement means that, in principle, the same task could be successfully performed by two different people in totally different ways. Yet experimental data show that point-to-point reaching movements tend to have highly consistent characteristics across individuals. For example, Morasso (1981) found that kinematics of point-to-point movements were similar across individuals as well as across different directions, different amplitudes and in different parts of one’s workspace. All movements followed a more or less Cartesian straight path and showed a characteristic bell-shaped velocity profile. One possible explanation for this regularity is that it is a consequence of a particular, idiosyncratic control mechanism that is common across individuals. In other words, regularities across the population may be arbitrary and purely a consequence of a shared motor heritage. For example, some models have attempted to explain regularities in the way we move as emerging from a simplistic controller coupled to the intrinsic dynamical properties of the musculoskeletal system (Gribble et al., 1998). An alternative view, and one that is adopted by the majority of recent theories of motor control, is that we possess sophisticated controllers that select the particular movements we make because they optimize some aspect of our behavior. Regularity across individuals emerges due to common properties of the underlying task. Explaining features of behavior as being the result of an optimization process has the advantage that, in principle, a range of behaviors can be explained through a single set of principles. Exactly what aspect of behavior should be optimized in such models is difficult to say, since it is something that the motor system – or, at least, evolutionary pressures – dictates. In most cases, one assumes that the motor system aims to minimize a cost function that reflects some combination of effort, variability, or the satisfaction of task goals. As we will see, in many cases it can be shown that a rational underlying principle can offer a parsimonious explanation for observed features of movement and generate novel predictions about features that our movements should possess.

A fundamental concern with normative models of control is that it might be possible to frame any regular behavior as optimizing something. Adhering to cost functions that make some ecological sense provides some protection against this concern. However, it is impossible to say in any principled way what kinds of cost functions should and should not be allowable. And, in any case, this cannot completely avoid the possibility of inferring spurious cost functions from behavior that may not truly be optimizing anything. It is therefore important to – where possible – specify cost functions *a priori* rather than reverse-engineering a cost function based on observed behavior. That said, there is reason to believe that behavior truly does reflect a process of optimization. When subjects are asked to control an unfamiliar object with complex dynamics, each individual initially adopts an idiosyncratic way of solving the task. With extended practice, however, all subjects gradually converged on almost identical patterns of behavior (Nagengast et al., 2009). This convergence is naturally explained by the idea of an optimization process.

Optimal control

The motor system can be modeled mathematically at many different levels of detail. Most simply, one can model the end-effector as a point mass subject to accelerations in one-, two- or three-dimensional Cartesian space. A more realistic model replaces the point mass dynamics with a multi-link rigid body subject to torques around each joint. More detailed models still replace torques with the combined action of individual muscles generating forces across joints and may encompass intrinsic properties of the muscles themselves, such as the nonlinear relationship between muscle length, muscle velocity and muscle force. There is no single ‘correct’ level of detail to adopt. A point mass serves as an excellent model of the oculomotor system (Robinson et al., 1986), but may be overly simplistic in other settings. A more detailed musculoskeletal model may be unnecessarily cumbersome for modeling some behaviors, but can in certain cases prove to enlighten our understanding of how a task is performed (Todorov, 2000). The appropriate level of modeling detail is largely a matter of judgment. However, for the sake of both parsimony and transparency, it is generally best to work with the simplest model possible that is able to explain a particular phenomenon of interest.

Mathematically, regardless of the level of detail employed, we can represent the state of the body at time t by the vector \mathbf{x}_t . This will typically contain the position and velocity of an end-effector or set of joints (e.g. the shoulder and elbow angles of the arm), but may also include things such as intrinsic states of each muscle. The motor commands themselves, which we denote by the time-varying vector \mathbf{u}_t , may correspond to joint torques, muscle forces or motor neuronal activity that only indirectly leads to changes in muscle force. The dynamics of the system – the way in which motor commands change the state – can be expressed in terms of a forward dynamics equation:

$$\dot{\mathbf{x}}_t = f(\mathbf{x}_t, \mathbf{u}_t). \quad (1.1)$$

This describes how changes in the state, represented as a derivative, depend in a particular way on the current state and on the current outgoing motor commands. The change in state is represented as a derivative – though note that this is the change in state of *all* components of the state. If the vector \mathbf{x}_t contains position and velocity, then $\dot{\mathbf{x}}_t$ contains velocity and acceleration. This equation describes mathematically the properties of the apparatus under control.

The role of the controller is to specify the motor commands \mathbf{u}_t . Typically, the process of motor command selection is described mathematically in terms of a *control policy* – a mapping π from some relevant variable, such as time or state of the body, to controls \mathbf{u}_t . A control policy can be either purely feedforward (open-loop), in which case the control policy is a mapping from time to motor commands

$$\mathbf{u}_t = \pi(t), \quad (1.2)$$

or feedback (closed-loop), in which case the motor commands may depend also on the current state of the plant

$$\mathbf{u}_t = \pi(\mathbf{x}_t, t). \quad (1.3)$$

More generally, feedback control commands may depend directly on sensory feedback rather than on the state of the plant *per se*. The control policy, coupled with the dynamics of the motor apparatus, fully determines the course of behavior (barring the effects of execution

noise). The job of the motor system, therefore, is to employ a control policy that will ensure completion of the task while minimizing some cost J that will in general depend on both the states and the outgoing motor commands.

An early theory, put forward to explain the kind of kinematic regularities observed by Morasso (1981), is that movements are selected to be as smooth as possible. Flash and Hogan (1985) suggested squared jerk – the derivative of acceleration – as an appropriate measure. If the scalar position of the effector at time t is x_t , then the minimum jerk hypothesis states that among all possible ways of moving from $x_0 = 0$ to a goal located at g in time T (i.e. $x_T = g$), the best way to move is the one that minimizes a cost J that is given by the summed squared jerk (the derivative of acceleration):

$$J = \int_0^T |\ddot{x}_t|^2 dt. \quad (1.4)$$

In this case, it is possible to analytically find the unique best sequence of positions (and therefore unique motor commands) that minimize this cost – see Shadmehr and Wise (2005) for details. The result is a smooth trajectory that is best characterized by examining the velocity, which follows a symmetric, bell-shaped profile, matching human movement data quite well.

A more dynamical version of this idea is to replace jerk with the rate of change of torque about a joint (Uno et al., 1989):

$$J = \int_0^T |\dot{\tau}_t|^2 dt. \quad (1.5)$$

Note, since forces and torques are linearly related to accelerations (according to Newton’s second law of motion), rate of change of torque is qualitatively similar to jerk. However, nonlinearities in the mapping from joint angles to end-effector location, and the fact that limb dynamics are different in different directions, mean that these two costs do not generate exactly the same predictions.

This pair of models (minimum jerk and minimum torque change) represents the first time that movement regularity was formally described in terms of an optimization of some feature of the movement. These two theories have been superseded by more recent frameworks and we therefore won’t dwell on the pros and cons of these models. However, they do largely capture many of the important features of more contemporary models. Both minimum jerk and minimum torque change essentially penalize some kind of third derivative of position. More recent formulations mostly penalize motor commands quadratically:

$$J = \int_0^T |u_t|^2 dt. \quad (1.6)$$

In these models motor commands are typically assumed to influence the rate of change of muscle force (Todorov & Jordan, 2002; Diedrichsen, 2007; Izawa et al., 2008). Therefore, these models also effectively penalize the squared third derivative of position.

Why should the squared motor command matter? The quadratic form of this cost is partly chosen for mathematical convenience. Many problems turn out to be straightforward to solve, provided one assumes such a quadratic penalty on motor commands. But the main reason the use of such a cost function persists is that it has proven successful in leading to models that offer a faithful description of behavior. It is important to note that it is not actually chosen based on any principled theoretical rationale. Energy consumption, which can be

quantified in terms of ATP utilization, appears to scale linearly with muscle force (Szentesi et al., 2001), not quadratically. So even if the motor commands themselves represent or are at least well-correlated with muscle force, this cost does not appear to reflect expended energy. Instead, this quadratic motor command penalty should be viewed as a more abstract notion of ‘effort’, which appears to be successful in describing certain aspects of behavior, despite having no clear theoretical foundation.

There is some independent empirical support for the notion that the motor system attempts to minimize some quadratic cost function. In static force generation tasks, subjects must generate a sustained force of a specific amplitude in a specific direction. These tasks are particularly useful for studying how the motor system resolves redundancy. They eliminate the dynamical complexities and kinematic redundancies associated with point-to-point movement, thereby isolating the redundancy associated with coordinating multiple muscles to generate a prescribed force. Because there are multiple muscles spanning each joint, generating a force in a particular direction – e.g. at the wrist – can be achieved through an infinite number of combinations of individual muscle activity (Hoffman & Strick, 1999). It is possible, for instance, to generate the same net force with varying degrees of co-contraction of an agonist/antagonist pair. In reality, muscles at joints with multiple degrees of freedom are not organized as simple agonist/antagonist pairs, but simply as a collection of muscles that can each generate force in different directions. At the wrist, for instance, there are five muscles that can each generate force along a single line of action. Generating a force in a direction that is not aligned with the line of action of any one muscle requires multiple muscles to be recruited. Hoffman and Strick (1999) measured the contributions of individual muscles during such a static force generation task at the wrist and found that individual muscles possessed a tuning curve that was in fact not centered on their anatomical line of action. The preferred pulling directions – i.e. the force direction for which each muscle was most active – were found to be more uniformly distributed than the anatomical pulling directions. Fagg and colleagues (2002) showed that these patterns of muscle activations can be explained through a model in which muscle forces are chosen to minimize the squared force in each muscle. This finding is not limited to the wrist; a similar result was described by Kurtzer et al. (2006) for force generation using the elbow and shoulder. In this case, the force is shared across mono-articular (spanning one joint) and bi-articular (spanning two joints) muscles of the arm. Alternative versions of this cost with exponents other than two can be considered (e.g. penalizing $|u|^3$). Powers below two predict broader tuning with greater sharing of load across muscles, while higher powers predict recruitment of fewer muscles where possible. A quadratic cost appears to describe data about as well as any alternative. In practice, however, dissociating these variants on the notion of an effort cost relies on quantitatively precise predictions that may be more sensitive to other parameters of the model than those in the cost function (Kurtzer et al., 2006).

Minimum endpoint variance

In the above examples, movements are considered to be deterministic – i.e. there is no noise associated with executing a movement. However, movements are inherently variable. Importantly, the variability introduced during a movement appears to grow linearly with the size of the motor commands (Jones et al., 2002; Faisal et al., 2008). This leads to faster movements being more variable than slower movements. Harris and Wolpert (1998) proposed that a natural cost on candidate control policies is the endpoint variability. Endpoint variability reflects an accumulation of instantaneous signal-dependent noise that is introduced throughout the course of the movement. The endpoint variance can be expressed as a weighted sum of

squared motor commands (Shadmehr et al., 2010). This quadratic structure is similar to the effort costs considered above, and solution can be found through similar methods. Applying this principle to model fast point-to-point movements of the eye (saccades) leads to saccade velocity profiles that are remarkably similar to those measured experimentally.

Effort versus variability

The fact that endpoint variance is a quadratic function of motor commands means that minimizing endpoint variance can be considered equivalent to minimizing a time-varying cost on squared motor commands, i.e.

$$J = \int_0^T w(t)u_i^2 dt, \quad (1.7)$$

where $w(t)$, illustrated in Figure 1.1c, is a time-varying cost weight that depends on the dynamics of the eye. Motor commands early in the movement are penalized far less heavily than motor commands late in the movement. The reason for this is that the elasticity of the eye acts to dampen out variability over time, making it advantageous to introduce noise earlier in the movement than later on. The trajectories of saccades that minimize endpoint variability have a velocity profile that is skewed towards the beginning of the movement (Figure 1.1a), when the effective cost of motor commands is relatively cheap. The time-varying cost implicit in the minimum endpoint framework can be contrasted with an effort cost that is uniform in




Figure 1.1 Comparison of effort and endpoint variability costs for open-loop arm and eye movements. (a) velocity profile for a 30 degree saccade lasting 100ms, as predicted by effort (solid line) and endpoint variability (dashed line) costs; (c) relative weight of quadratic motor command cost as a function of time within movement; (b) and (d) as (a) and (c) but for a point-to-point reaching movement of 10cm lasting 300ms.

time. The saccades predicted by a uniform cost have symmetric velocity profiles (Figure 1.1a) and do not closely resemble the kinematics of real saccades. The principle of minimizing endpoint variance can also be applied to reaching movements (Figure 1.1b, 1.1d), in which case it also leads to reasonable predictions about hand trajectories during point-to-point reaching movements. Exactly as in the case of saccades, a minimum effort cost leads to symmetric velocity profiles. For minimizing endpoint, the effective time-varying cost is greater for motor commands later in the movement resulting in a velocity profile that is slightly skewed towards being faster later in the movement – opposite from the result for saccades. The key difference between oculomotor control and control of the arm is that the arm does not have a natural equilibrium point in the same way that the eye does. Noise does not, therefore, dissipate in the same way that it does for the eye. The effective cost of motor commands in fact decreases towards the end of the movement because noise introduced by motor commands late in the movement does not have time to significantly affect final position. It should also be noted that the dynamics of the arm are less well specified than those of the eye. Different simulations use a wide range of values for inertia and viscosity of the arm – compare e.g. Harris and Wolpert (1998), Liu and Todorov (2007) – which lead to a fairly broad range of behaviors. We used an intermediate set of parameters to illustrate some of the general characteristics of the predictions of these models (1kg mass, 2Nm/s viscosity, 5ms muscle time constant).

Based on these results, one might be tempted to conclude that noise is in fact the thing of primary interest to the motor system. The effort penalty in models which assume deterministic dynamics (or at least non-signal-dependent noise) might be deemed simply as providing a good approximation to the true cost relating to variability. The main difficulty in validating this theory more rigorously is that it is difficult to directly measure the noise associated with a particular motor command at a given instant in time. The best one can do is to examine patterns of variability and show that they are consistent with a particular signal-dependent noise model coupled with an associated optimal controller (van Beers et al., 2004; van Beers, 2007). It is feasible, however, to measure variability in a static force generation task where motor commands are constant. O’Sullivan et al. (2009) measured variability of forces of various magnitude generated by different digits of the left and right hands. They fitted a model in which force and variability were linearly related, and used this model to predict behavior in a second task in which two fingers, one from each hand, cooperated to achieve a desired force goal. Since the relationship between force and variability was different for each finger, this led to a rich set of predictions for how the force should be divided among varying combinations of fingers. Minimizing the variability of the net force did not turn out to predict behavior well. Rather, in order to completely describe the data, it was necessary to invoke an effort cost penalizing the squared force generated by each finger.

The similar mathematical forms of effort and variability make it difficult to disambiguate one from the other. However, it seems unlikely that variability is the sole factor determining how we move. In addition to the results of O’Sullivan et al., it has been suggested that effort may play an important role in deciding how fast to move to acquire rewards of different magnitude (Shadmehr et al., 2010). Perceived effort leads to robust decisions even in a task in which variability plays little role (Körding et al., 2004). Perceived effort also appears to be awry in patients with Parkinson’s disease (Mazzoni et al., 2007).

Feedback control

All of the theories presented so far consider only feedforward control policies, in which a pre-determined sequence of commands is rolled out at the time of execution. In reality,

however, we are not restricted to being passive spectators to the effects of execution noise or external perturbations. For all but the fastest of movements (such as saccades) we are able to observe deviations from expected behavior and make online adjustments during a movement. A simple feedback control policy can be formed by augmenting a feedforward control policy with a feedback component that tries to cancel out the effects of noise and external perturbations by keeping the movement of the plant close to some planned trajectory. While this kind of strategy will indeed help to negate the impact of noise, it is highly inflexible, attempting to rapidly correct any deviation from desired behavior. These corrections can be quite costly and in many cases largely unnecessary. Better strategies are possible that do not rigidly adhere to a single trajectory, but allow more flexibility in the way errors are corrected.

The key difference between feedforward and feedback control is that a feedback control policy selects motor commands as a function not just of time, but also of the current state:

$$\mathbf{u}_t = \pi(\mathbf{x}_t, t). \tag{1.8}$$

Feedback control thus relies on knowledge of the instantaneous state of the motor apparatus. While this is somewhat unrealistic, given the noisy and delayed nature of available sensory feedback, it serves as a reasonable simplifying approximation to explore the most salient aspects of feedback control. More detailed models replace the exact state of the system \mathbf{x}_t with an estimate of the state $\hat{\mathbf{x}}_t$ that must be generated based on sensory feedback. For now, however, we make the simplifying assumption that if feedback is available at all, the state of the system is known precisely and instantaneously. We focus instead on the question of how this knowledge about the state of the system should be used.

Just as in the feedforward case, the controller coupled with the dynamics of the motor apparatus completely determines the behavior of the system. Optimal feedback control theory addresses the question of what feedback control policy should be used in the same way feedforward models do, by associating a cost with each potential control policy based on the resulting behavior, and identifying the control policy that leads to minimal cost.

Determining the cost associated with each policy and finding the best possible is much harder for feedback control policies, since each policy does not specify motor commands directly, but instead specifies a rule for determining motor commands based on the current state. Fortunately, the general mathematical problem has been extensively studied in control theory (Bertsekas, 1995) and reinforcement learning theory (Sutton & Barto, 1998), and requisite extensions for dealing with features peculiar to the motor system, such as signal-dependent noise, have been developed (Todorov, 2005). The methods used to find the optimal feedback control policies are beyond the scope of this chapter, but a thorough introduction can be found in (Todorov, 2006).

An intuitive way to understand optimal feedback control without delving into technical details is as a feedforward control policy that is continuously re-planned based on the latest state information. For example, if the hand is perturbed from its initial path, the best thing to do is forget about the original intended movement and plan again from scratch given the new, perturbed state as a starting point. This picture provides a reasonable way to intuit the properties of optimal feedback control policies with a few caveats. Firstly, movements are not actually re-planned. The optimal feedback control policy (Equation 1.8) implicitly encodes the optimal course of action starting from any state at any time. Secondly, there is no need to wait for a large perturbation to prompt an adjustment of one's movement. Even small deviations from expected trajectories should prompt a flexible change in motor commands. Thirdly,

although in simple cases there is a precise equivalence between optimal feedback control and the idea of continually re-planning an open-loop controller, they do not always lead to the same predicted behavior. Knowing that one will be able to make feedback corrections in the future can influence one's control strategy in the present. For instance, acting under signal-dependent noise, it can be beneficial to introduce noise early in a movement in the knowledge that one will have time to correct it later on. Such 'strategies' emerge naturally within the theoretically optimal policy.

Minimum intervention principle

The critical difference at the behavioral level between optimal control policies and other potential feedback control policies is the way in which the optimal feedback control policy flexibly exploits redundancy in order to minimize costs. This is perhaps best understood through the notion of the minimum-intervention principle: an optimal control policy should only correct perturbations that interfere with the achievement of task goals. If a perturbation is irrelevant to a task, for instance, if your elbow is knocked during a reaching movement without affecting your hand position, then there is no need to correct for it – just maintain the new elbow posture during the rest of the movement. The same applies to deviations occurring because of noise. Not making unnecessary corrections allows one to be more sparing with motor commands, which helps reduce both effort costs and the impact of signal-dependent noise. This gives rise to the more general prediction that movement variability will be greatest far away from task-critical periods of a movement. This is true in the case of point-to-point reaching movements, where variability is highest midway through the movement (Liu & Todorov, 2007) and after striking a target object such as when striking a ping-pong ball with a bat (Todorov & Jordan, 2002).

While any motor task naturally contains redundancies, the characteristics of optimal feedback control policies are most striking in tasks where redundancy is exaggerated. Diedrichsen (2007) demonstrated this in a bimanual task in which subjects controlled a cursor that appeared at the average location of their two hands. Subjects had to move the cursor from an initial starting location (with specific initial positions for each hand) to a goal location. On select trials, one hand experienced a force that perturbed it perpendicular to its movement direction, causing a corresponding perturbation to the cursor. According to a desired trajectory hypothesis, subjects should have corrected for this error by returning the right hand back towards its original path. If, however, subjects performed the task according to an optimal control policy, they should exploit the redundancy afforded by the bimanual nature of the task and recruit the other hand to help steer the cursor towards the target. This is precisely what subjects were found to do.

Taking the idea of a redundant task to extreme, Todorov and Jordan (2002) recorded movements of the fingers while subjects scrunched a piece of paper into a ball using one hand. Subjects do not adopt a fixed method, but show huge variability from trial to trial in how they manipulate the object to achieve their goal. Todorov and Jordan interpreted this variability as evidence that subjects were implementing a flexible feedback control policy of the kind predicted by optimal feedback control. In this case, the variability of actions stemmed directly from the variability in the state of the paper in the hand. Although it is difficult to prove that in this case the high variability had anything to do with performing the task well, this experiment serves well as an intuitive example to illustrate the fact that acting flexibly in order to be effective – the key principle underlying optimal feedback control – is something that comes very naturally to us.

Scope and limitations of optimal control models

The basic idea behind optimal control theory is actually a fairly simple one – among the many ways one could accomplish a task, pick the one which is the best according to some cost measure. The power of optimal control theory lies in its ability to explain a wide range of behaviors through a small set of widely applicable cost functions. Whereas early motor control research focused on identifying invariants in kinematics, optimal control focuses on identifying invariants at the level of cost functions. Obtaining the ‘best’ control policy is often the most technically challenging aspect of applying optimal control theory to model motor behavior. In many ways, however, it is the easy part, since the problems one encounters have unambiguous solutions. For a given dynamics model (Equation 1.1) and cost function, various methods exist – albeit in some cases computationally demanding ones – for obtaining the optimal control policy. The real challenges in testing optimal control as a hypothesis about motor control are that it is very difficult to know the dynamics of the motor apparatus precisely and one doesn’t know exactly what kind of costs the motor system uses to evaluate one movement relative to another. Different assumptions about the dynamics of the motor apparatus and different kinds of cost functions can lead to very different control policies. If predictions of a model fail to match those of data, there is no way of knowing if one used a poor dynamics model or the wrong cost function, or whether the motor system simply doesn’t behave in a way consistent with the premise of optimal control. Conversely, and much more problematically, the optimal control framework is sufficiently flexible that it is often possible to find some combination of dynamics and cost function that will give rise to any given behavior. Suppose that subjects are seen to always accomplish a particular task by moving in a straight line. This phenomenon can be easily ‘modeled’ by imposing a cost function that penalizes deviations away from a straight line. This does not, of course, constitute a theory about why subjects move in a straight line, nor does it provide any support for the premise of optimality. In general, optimal control theories of the motor system are most compelling when using cost functions decided *a priori* and generate qualitative predictions about behavior, rather than when they are used to provide quantitatively precise descriptions of data.

Optimal control theory has proven a valuable modeling tool that has both explained previously characterized aspects of motor behavior and generated a host of novel predictions that have been demonstrated experimentally. However, not all behavior is consistent with optimal control predictions. For instance, human subjects seem to have a preference for reaching in straight lines, even in the presence of perturbations that make it sub-optimal to do so (Wolpert et al., 1995). It is difficult to explain this tendency in terms of some underlying normative principle. In reality, it is likely that relies heavily on approximations and heuristics in order to achieve behavior that is near-optimal in a variety of contexts. Outside of these contexts, the particular control strategies that are habitually used may no longer be optimal. This appears to be true of the way in which multiple muscles are coordinated across a joint (de Rugy et al., 2012) and a similar idea may explain the tendency for people to want to reach in straight lines.

State estimation: Making sense of sensory feedback

Any movement we make necessarily depends on incoming sensory information. Knowing exactly what action to take will depend on sensory input to specify the goals of the task and the initial state of the body. Once a movement is underway, feedback control requires ongoing monitoring of the state of the body through vision and proprioception. Individual sensory modalities provide imperfect information. Although we may possess high visual acuity, using visual

information to specify the location of a point in space additionally requires knowledge of the orientation of the eyes within the head and of the head relative to the body, all of which introduce noise and uncertainty. Proprioceptive estimation of hand location is notoriously imprecise – try aligning your two index fingers above and below an opaque tabletop. Given the noisy nature of sensory feedback, two modalities may often report conflicting estimates of the same thing. How should one estimate the position of one’s hand in order to make decisions about movement?

This general problem can be formalized by supposing that our senses yield noisy readouts of the true state of our arm. If the true location of the hand is x , then we can model vision and proprioception as providing independent, unbiased observations of that location, $v \sim N(x, \sigma_v^2)$ and $p \sim N(x, \sigma_p^2)$. A principled way to estimate the hand position x , given these observations, is to ask what hand location is most likely to have led to these observations. This approach is known as Maximum Likelihood Estimation (MLE). The maximum likelihood estimate of hand position can be shown to be equal to a linear combination of the two unimodal observations, with each observation weighted inversely to its variability:

$$\hat{x}_{MLE} = \frac{\sigma_v^{-2}}{\sigma_v^{-2} + \sigma_p^{-2}} v + \frac{\sigma_p^{-2}}{\sigma_v^{-2} + \sigma_p^{-2}} p. \quad (1.9)$$

Thus, the combined estimate should lie somewhere in between the two unimodal estimates. This principle can be applied directly to any pair of modalities, not just vision and proprioception, and can be generalized to incorporate any number of modalities. Maximum Likelihood Estimation provides a normative rationale for how to optimally combine two sensory observations into a single estimate. Individuals’ sensory integration strategies can be tested by providing subjects with conflicting cues about a single object and asking them to judge the location of the common cause, provided the cues are not too discrepant (Körding et al., 2007). In order to test whether people use an optimal sensory integration strategy, it is not enough to simply observe that a subject’s combined estimate lies somewhere in between the two uni-modal estimates. Maximum likelihood integration provides a quantitatively precise prediction about where the integrated observation should lie. In order to properly test this hypothesis, one needs to measure the reliability (i.e. the variance) associated with each individual sensory modality, yielding a quantitative prediction that can be compared with subject behavior. This approach has been successfully applied in a number of contexts, including integrating visual and haptic estimates of the size of objects (Ernst & Banks, 2002).

In reality, we are rarely interested in one-dimensional quantities. Position, for instance, is three-dimensional. Extending the MLE framework to higher dimensions actually leads to richer predictions about how information should be combined across modalities than in the one-dimensional case. In higher dimensions, each observation is a vector and the variances associated with each modality become covariance matrices: $\mathbf{y}_i \sim N(\mathbf{x}, \Sigma_i)$. The covariance matrix encodes the relative confidence in a given modality along different dimensions. If two different modalities have different covariance matrices, i.e. if they have different relative reliabilities along different dimensions in space, then combining these observations yields a MLE that, perhaps counter-intuitively, does not lie along the straight line between the individual uni-modal estimates (consider performing independent one-dimensional integrations along each dimension to see how this can be the case – see Figure 1.2b). This qualitative behavior is seen when subjects are asked to integrate visual and proprioceptive estimates of hand location in two dimensions (van Beers et al., 1999).

A similar kind of integration also occurs when estimating the state of the limb for feedback control. In this case, rather than combining two sensory modalities, one combines sensory




Figure 1.2 Maximum Likelihood Estimation of hand position from visual and proprioceptive estimates. (a) One-dimensional integration of discrepant visual and proprioceptive observations of hand position. Dashed lines indicate single-modality likelihoods—i.e. probability of hand location given observation. Solid line indicates joint likelihood integrating both visual and proprioceptive observations. Note that the maximum likelihood estimate is closer to the visual observation due the visual observation having lower variability. (b) Integrating two sensory estimates in two dimensions. Anisotropic uncertainties lead to an integrated estimate that does not lie on a straight line between individual observations.

feedback with an internally-generated prediction about the new state of the limb based on previously issued motor commands. The Kalman filter, a standard approach to state estimation in optimal control models, is essentially an iterative version of this idea.

Box 1.1: Forward and inverse models

In addition to using sensory information to estimate the location of our limbs, we can also use prediction based on known outgoing motor commands. A *forward model* predicts the consequences of outgoing motor commands on the state of our bodies. For example, maintaining an internal representation of the dynamics of the body given by Equation 1.1 would constitute an internal forward model. Forward models allow us to overcome the delays inherent in sensory feedback, enabling early corrections for aberrant outgoing motor commands or for anticipating unwanted consequences of actions in order to prevent them from occurring.

A forward model is truly a model in the sense that it is an internal representation of a process that occurs extrinsically. Historically, motor learning theory has also extensively appealed to the notion of an inverse model that maps desired changes of state to motor commands. Inverse models are not really models in the true sense of the word and, in any case, have been superseded by the more general notion of a control policy.

Motor learning

Optimal control theory, coupled with state estimation, provides a powerful framework for describing coordination in over-learned tasks in which we are already experts. However, most movements require extensive practice before we consider ourselves to have reached any

kind of proficiency. The capacity to learn new patterns of movement and improve and adapt existing ones is arguably the most fundamentally important facet of the human motor system. The vast majority of research in motor learning studies this capacity through adaptation paradigms in which a systematic perturbation is introduced to disrupt a well-practiced behavior, such as point-to-point reaching. The imposed perturbation generally causes significant errors at first, however, subject' performance generally returns to near baseline levels within tens of trials. A classic example is adaptation of reaching movements while wearing prism goggles that shift the visual field to the left or right (von Helmholtz, 1962). Other common examples include rotations of visual feedback on a virtual display (Krakauer et al., 2000) or force fields that perturb the dynamics of the arm (Shadmehr & Mussa-Ivaldi, 1994).

In all cases, the course of learning is qualitatively similar. Error measures, such as directional error, tend to decline approximately exponentially across trials. This exponential decline is consistent with the idea that learning between trials is proportional to error size. Formalizing this insight mathematically leads to a so-called 'state-space model' of motor learning (Thoroughman & Shadmehr, 2000; Donchin et al., 2003; Cheng & Sabes, 2006), which we illustrate through the example of adapting reaching movements under a rotation of visual feedback. Suppose that, on a given trial a subject aims their reach in a direction u_i towards a goal located at 0 degrees. In this case, error can be quantified as the directional error at the start of the movement. The error subjects experience on a given trial will be determined by a combination of the subject's chosen action and the perturbation – in this case, a rotation of visual feedback by r degrees:

$$e = u + r. \tag{1.10}$$

The ideal action u^* is the one that generates zero error:

$$u^* = -r. \tag{1.11}$$

The state-space model assumes that, if an error e_i is experienced on trial i , subjects will adapt their reach angle on the next trial by a proportional amount, i.e.:

$$u_{i+1} = Au_i + \alpha e_i. \tag{1.12}$$

Here α is the learning rate determining the sensitivity to error, while A , which is some number close to but less than 1, captures the natural tendency to return to baseline behavior in the absence of errors (Galea et al., 2011) – often described as 'forgetting'. This summarizes the basic premise of a state-space model. The 'state-space' nomenclature arises from parallels with dynamical systems theory – the trial-to-trial reach errors can be viewed as the output of a simple discrete-time linear dynamical system. This is not, however, a particularly helpful analogy unless one happens to be already familiar with such frameworks. Instead, it is best to think about these models as simply mathematically capturing the idea that changes in behavior are proportional to error size. This example is perhaps the simplest possible such model. Alternative forms of learning, such as adaptation to dynamical perturbations, can be modeled in an identical way with the caveat that one requires an additional parameter to translate between a quantitative measure of the perturbation (typically a force, or viscosity) and a kinematic measure of error (typically a distance or angle) (Donchin et al., 2003; Thoroughman & Shadmehr, 2000).

An important feature of this model is the so-called 'forgetting rate' A in equation 1.12. The existence of this term predicts that, in the absence of observed error, motor commands will gradually return to their original baseline values. For visual perturbations errors can be

removed simply by removing visual feedback, in which case behavior does indeed return to baseline. The fact that there is always some forgetting between trials also means that learning can never completely reduce error to zero. Subjects reach an adaptation asymptote at a particular error where there is an equilibrium between trial-to-trial forgetting and the amount of learning due to error. This incomplete adaptation is precisely what is observed experimentally. Mathematically, the asymptotic behavior u_∞ is given by

$$u_\infty = \frac{-\alpha r}{\alpha - (1 - A)}. \quad (1.13)$$

Note that, as A gets closer to 1, u_∞ approaches the ‘correct’ value of $-r$. The asymptotic error also grows with the size of a learned rotation, suggesting that incomplete adaptation is not simply a reflection of adapting just enough to reach the outside of a target.

Various extensions of this basic model are possible. Most commonly, the basic SSM presented here can be applied in the context of multiple targets. In this case, errors at a given target can also influence learning at neighboring targets (Donchin et al., 2003; Thoroughman & Shadmehr, 2000). Despite the apparently simple assumptions, state-space models describe adaptation behavior remarkably well, even when the imposed perturbation varies randomly from trial to trial (Donchin et al., 2003; Huang et al., 2011).

The success of state-space models in accounting for trial-to-trial behavior in adaptation paradigms tells us little about the underlying mechanisms mediating that behavior. State-space models are thus best viewed as purely phenomenological descriptions of trial-to-trial learning behavior in adaptation paradigms. Nevertheless, state-space models are variously couched in terms of either learning a forward model of the limb (along with any potential external perturbations) or simply directly learning appropriate motor commands. It is difficult to dissociate these interpretations based on trial-to-trial adaptation data, since they lead to identical sets of equations. One important difference, however, is that errors in learning a forward model are prediction errors, whereas errors in learning actions are task errors. In most circumstances, task errors and prediction errors are identical. In certain circumstances they can be dissociated (Mazzoni & Krakauer, 2006; Wong & Shelhamer, 2011), in which case it becomes clear that the motor system is sensitive to *prediction errors*. These findings support the idea that motor learning proceeds by updating an internal forward model that is then subsequently used to guide planning of future movements.

Regardless of the underlying representation, a state-space model offers a principled way to estimate trial-to-trial learning rates. A model that fits the data well can be thought of as a means to compactly summarize the salient features of the data. Comparing parameter fits across conditions then enables one to make inferences about the effect of manipulations on learning rate (Diedrichsen et al., 2005).

Closer inspection of trial-to-trial learning curves reveals that there are typically two phases of learning. An initial, rapid decrease in error is typically by a more prolonged, gradual reduction in error towards asymptote. Consequently, learning curves are far better fit by a sum of two exponentials – one with a fast time constant and one with a slower time constant. This two-rate behavior can naturally be accommodated within the state-space model framework by allowing motor commands to be comprised of two distinct components that learn from the same error, one quickly and one slowly. In addition, the component that learns quickly from error also forgets quickly (i.e. $A \ll 1$), while the slow learning component forgets slowly. Although this two-rate model is motivated by describing data rather than by theory, this model gives rise to an interesting prediction: if subjects are rapidly adapted and then de-adapted, subjects will ‘spontaneously recover’ previous

learning, as the fast component is forgotten to reveal the slow component that still partially reflects the previous learning. This spontaneous recovery has been demonstrated to occur exactly as predicted for both force field learning (Smith et al., 2006) and saccadic gain adaptation (Ethier et al., 2008). Furthermore, components identified through such a model are predictive of long-term retention (Joiner & Smith, 2008) and anterograde interference (Sing & Smith, 2010).

Beyond state-space models: Learning as estimation

Ultimately, working with purely descriptive models is unsatisfying. We really would like unifying concepts to explain why learning should be proportional to error, not just keep a close track of the behavioral consequences of this kind of assumption. The sensitivity to error is also an arbitrary parameter. What determines how fast one learns? Is learning constrained by the mechanics of synaptic plasticity, or are features of the task more important?

Bayes' Rule: Existing beliefs about some variable x in the environment can be expressed as a probability distribution over possible values $p(x)$, with greater probabilities indicating a stronger belief. This distribution is referred to as the prior distribution. A new observation y may prompt us to revise our beliefs about x . If y depends on the true value of x through a probability distribution $p(y|x)$ (the likelihood of x), then Bayes' rule tells us precisely how to revise our existing beliefs in the light of the new evidence:

$$p(x|y) = \frac{p(y|x)p(x)}{p(y)},$$

where $p(x|y)$ (the posterior distribution) represents our revised belief about x in the light of y . $p(y)$ is the overall probability of our observation for any possible underlying x – otherwise known as the marginal probability of y . $p(y)$ is simply a constant if one views $p(x|y)$ as a function of x and serves to normalize the posterior distribution, ensuring that all beliefs sum to 1. In many cases, one will see Bayes' rule written simply as

$$p(x|y) \propto p(y|x)p(x).$$

A critical idea in models of motor learning is that motor learning in adaptation paradigms amounts to estimating the properties of the perturbation. We have seen how a forward model of the dynamics of the motor apparatus can be useful for estimating the state of the limb. In principle, the same forward model can also be used to plan movement. Motor learning research in recent decades has been dominated by this idea.

Korenberg and Ghahramani (2002) pointed out that changing behavior proportionally to the errors one experiences is consistent with the idea of updating one's beliefs about the perturbation from trial to trial in a statistically optimal manner. A basic version of the idea of iteratively revising one's beliefs in this way leads to the commonly employed Kalman filter, which yields updates of one's mean belief that are proportional to prediction errors – precisely as stated by the state-space model framework. The critical difference between this idea and the state-space model formulation is that, in the Kalman filter equations, the learning rate, i.e. the amount one revises one's beliefs in the light of new observations, is not an arbitrary parameter of the model but emerges as a consequence of the statistics of the task. If observations are typically unreliable and

the perturbation is expected to remain quite stable over trials, one should be cautious about revising beliefs based on a single large prediction error. On the other hand, if observations are typically reliable and the underlying perturbation is liable to rapid fluctuations, one should trust one's immediate observations more than past experience and employ a high learning rate.

In order to frame this problem mathematically, we require a more precise statement of the variability or confidence associated with observations and of the expected variability in the perturbation over time. Formally, such a model is described as a generative model, since in principle it can be used to generate a realistic sequence of observations. Figure 1.3 illustrates this




Figure 1.3 Adaptation as Bayesian estimation. (a) Subject's beliefs about the task represented as a graphical model. Shaded circles represent observed variables. Unshaded circles represent unobserved variables to be inferred. Arrows indicate probabilistic dependencies between variables. A rotation r_i is assumed to vary from trial to trial (arrows indicate) following a Gaussian random walk. Each trial yields an observation γ that derives from the perturbation and the motor command u . The subject's goal is to infer the rotation on each trial and use this to predict rotations on future trials. (b) and (c) Illustration of the time-course of learning in response to a step-perturbation predicted by the Kalman filter model—in this case also equivalent to a state-space model. Note that asymptotic error is not equal to zero.

model graphically. There are two components that need to be specified. Firstly, there must be a model of how a given perturbation gives rise to observations. In specific cases, this can be as simple as saying the observation is a noisy, unbiased version of the perturbation, i.e. $y \sim N(x, \sigma^2)$, though in general there could be a more complex relationship between the two. Secondly, we require a model of how the perturbation is liable to vary over time. Typically, one assumes that the values of a perturbation follow a random walk over trials, i.e.

$$x_{i+1} \sim N(Ax_i, q^2), \quad (1.14)$$

where $A < 1$ is some ‘forgetting factor’, exactly like the forgetting factor in state-space models. In this context, however, A reflects a property of the perturbation itself, rather than the learner’s capacity to retain this information. Here, it additionally has the effect of making sure that the perturbation is assumed to not stray too far from zero – the variance remains bounded.

Qualitatively, this model predicts that one should adjust future movements by a fixed proportion of error size – basically recovering the premise of the state-space model from a theoretical grounding. The Bayesian viewpoint is, however, much deeper in its implications than a state-space model. In particular, the Bayesian point of view implies that the learning rate is dictated by particular aspects of the task – the observation noise σ^2 and the perturbation volatility q^2 . In practice, estimating these variables is difficult since they reflect implicit beliefs of the subject. In particular, q^2 corresponds to the subject’s own estimate of how variable a particular perturbation is from trial to trial. Sensory noise, however, is more amenable to experimental quantification. Burge and colleagues (2008) had subjects adapt to a shift in visual feedback while making point-to-point reaching movements. The cursor representing subjects’ hand position was a small circle that was blurred to varying degrees. Consistent with the qualitative Bayesian prediction, the more blurry cursor, which presumably increased observation noise, led to a significant decrease in the rate of learning. This change in adaptation rate was even direction specific – blurring the cursor only along the x-axis led to slower adaptation to shifts in this direction, but not a simultaneously imposed vertical shift. Similar findings have also been reported elsewhere (Wei & Körding, 2010).

The key predictions of the Bayesian interpretation of motor learning are that, 1) as uncertainty in sensory feedback increases, learning should become slower. This is directly analogous to the fact that sensory modalities that are uncertain have less influence over the maximum likelihood estimate of state (Fig. 1.3a). 2) As uncertainty in the perturbation increases, learning should become faster. While reducing certainty in sensory feedback appears to reliably influence learning in the way predicted by the Bayesian framework, manipulating certainty in the perturbation is much harder to achieve. One method that has showed partial success is to simply leave subjects sitting idly in the dark. A period of such inactivity leads to a brief increase in learning rates (Wei & Körding, 2010). A similar increase in learning rates is also observed in monkeys during saccade adaptation paradigms, if the monkey is left in the dark for a period of time (Kojima et al., 2004). In both cases, this faster re-learning has been construed as reflecting an increase in uncertainty of an internal model that guides movement planning (Wei & Körding, 2010; Körding et al., 2007). In the case of manipulating uncertainty in visual feedback, it is possible to empirically determine the extent to which a given manipulation of visual feedback affects confidence in visual observations – through the just-noticeable difference, for instance. Unfortunately, there is no analogy of this approach for measuring a subject’s confidence. In their estimate, rather than try to measure

confidence in the state mapping, one can try to manipulate it. However, attempts to decrease certainty in the value of a perturbation – usually by having the perturbation itself follow a random walk with varying degrees of volatility – have yielded little success. It is impossible to know whether this is due to a failure on the subject's part to take known variability in the perturbation into account during learning or due to a failure to learn about the variability of the perturbation. In any case, it appears that the motor system is equipped to respond to some forms of variability better than others.

The basic Kalman filter model of adaptation, that we have outlined here, can be extended to include multiple underlying causes for each potential error, such as world versus body (Berniker & Körding, 2008) or deciding whether a given error has been caused by motor execution or sensory miscalibration (Haith et al., 2008). In general, however, such models display a qualitative consistency with patterns of behavior, but are not typically subject to the level of scrutiny that is found in testing models of statistically optimal multi-sensory integration. The Bayesian framework does, however, provide a rich and rigorous framework within which to frame theories of adaptation and generate new hypotheses.

Issues and outlook

We have introduced the main theoretical concepts underlying many recent models of motor control and motor learning. These models, which for the most part adopt a normative approach, have enjoyed considerable success in recent years. Though we have concentrated on theory and have not discussed the underlying neural mechanisms (either at the systems level or neuronal level), bridging the gap between normative and mechanistic models is an important direction for future research. It appears that the cerebellum is of critical importance in motor learning (Bastian, 2006; Tseng et al., 2007; Taylor et al., 2010), likely subserving adaptation through an internal forward model that is updated by sensory prediction errors. The cerebellum also may contribute a forward model that mediates a state estimate used in control (Xu-Wilson et al., 2009). The control policy itself, however, seems likely to reside in primary motor cortex.

While control and learning have been extensively examined individually, surprisingly little work has addressed the intersection of the two. Multiple studies have suggested that control policies are re-optimized after learning novel dynamics (Izawa et al., 2008; Nagengast et al., 2009). What is unclear at present, however, is how knowledge of dynamics stored in a forward model, presumably in the cerebellum, becomes translated into a control policy in motor cortex. An alternative way in which a control policy could be learned or adjusted is through reinforcement of actions that lead to success. Indeed, it appears that such model-free learning is responsible for savings that occur during adaptation paradigms (Huang et al., 2011) and may be responsible for learning in more complex motor tasks (Hosp et al., 2011). Such forms of learning can also in principle be characterized from a normative point of view.

While normative models offer a potentially powerful lens through which to examine the motor system, it is important to bear in mind that the omnipotent brain is only a simplifying approximation. There is unlikely to be any single principle that will account for all behavior. If anything, such models serve to highlight instances where behavior of the motor system deviates from supposed optimal behavior. Ultimately, the true utility of a normative model is not to provide an overarching theory of everything akin to a fundamental law of physics, but to constrain, inform and inspire new ways of thinking about the motor system across multiple levels of analysis.

References

- Bastian, A.J. (2006) 'Learning to predict the future: the cerebellum adapts feedforward movement control', *Current Opinion in Neurobiology*, 16(6): 645–9.
- Berniker, M. and Kording, K. (2008) 'Estimating the sources of motor errors for adaptation and generalization', *Nature Neuroscience*, 11(12): 1454–61.
- Bernstein, N. (1967) *The Co-ordination and Regulation of Movements*, Pergamon Press, New York.
- Bertsekas, D.P. (1995) *Dynamic Programming and Optimal Control*, Athena Scientific, Belmont, MA.
- Burge, J., Ernst, M.O. and Banks, M.S. (2008) 'The statistical determinants of adaptation rate in human reaching', *Journal of Vision*, 8(4): 20, 1–19.
- Cheng, S. and Sabes, P.N. (2006) 'Modeling sensorimotor learning with linear dynamical systems', *Neural Computation*, 18(4): 760–93.
- Diedrichsen, J. (2007) 'Optimal task-dependent changes of bimanual feedback control and adaptation', *Current Biology* 17(19): 1675–9.
- Diedrichsen, J., Hashambhoy, Y., Rane, T. and Shadmehr, R. (2005) 'Neural correlates of reach errors', *The Journal of Neuroscience*, 25(43): 9919–31.
- Donchin, O., Francis, J.T. and Shadmehr, R. (2003) 'Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control', *The Journal of Neuroscience*, 23(27): 9032–45.
- Ernst, M.O. and Banks, M.S. (2002) 'Humans integrate visual and haptic information in a statistically optimal fashion', *Nature*, 415(6870): 429–33.
- Ethier, V., Zee, D.S. and Shadmehr, R. (2008) 'Spontaneous recovery of motor memory during saccade adaptation', *Journal of Neurophysiology*, 99(5): 2577–83.
- Fagg, A.H., Shah, A. and Barto, A.G. (2002) 'A computational model of muscle recruitment for wrist movements', *Journal of Neurophysiology*, 88(6): 3348–58.
- Faisal, A.A., Selen, L.P.J. and Wolpert, D.M. (2008) 'Noise in the nervous system', *Nature Reviews Neuroscience*, 9(4): 292–303.
- Flash, T. and Hogan, N. (1985) 'The coordination of arm movements: an experimentally confirmed mathematical model', *The Journal of Neuroscience*, 5(7): 1688–703.
- Galea, J.M., Vazquez, A., Pasricha, N., Orban de Xivry, J. and Celnik, P. (2011) 'Dissociating the roles of the cerebellum and motor cortex during adaptive learning: the motor cortex retains what the cerebellum learns', *Cerebral Cortex*, 21(8): 1761–70.
- Gribble, P.L., Ostry, D.J., Sanguineti, V. and Laboisière, R. (1998) 'Are complex control signals required for human arm movement?', *Journal of Neurophysiology*, 79(3): 1409–24.
- Haith, A.M., Jackson, C.P., Miall, R.C. and Vijayakumar, S. (2008) 'Unifying the sensory and motor components of sensorimotor adaptation', In *Proc. Advances in Neural Information Processing Systems*.
- Harris, C.M. and Wolpert, D.M., (1998) 'Signal-dependent noise determines motor planning', *Nature*, 394(6695): 780–4.
- Hoffman, D.S. and Strick, P.L. (1999) 'Step-tracking movements of the wrist. IV. Muscle activity associated with movements in different directions', *Journal of Neurophysiology*, 81(1): 319–33.
- Hosp, J.A., Pektanovic, A., Rioult-Pedotti, M.S. and Luft, A.R. (2011) 'Dopaminergic projections from midbrain to primary motor cortex mediate motor skill learning', *The Journal of Neuroscience*, 31(7): 2481–7.
- Huang, V.S., Haith, A.M., Mazzoni, P. and Krakauer, J.W. (2011) 'Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models', *Neuron*, 70(4): 787–801.
- Izawa, J., Rane, T., Donchin, O. and Shadmehr, R. (2008) 'Motor adaptation as a process of reoptimization', *The Journal of Neuroscience*, 28(11): 2883–91.
- Joiner, W.M. and Smith, M.A. (2008) 'Long-term retention explained by a model of short-term learning in the adaptive control of reaching', *Journal of Neurophysiology*, 100(5): 2948–55.
- Jones, K.E., Hamilton, A.F. and Wolpert, D.M. (2002) 'Sources of signal-dependent noise during isometric force production', *Journal of Neurophysiology*, 88(3): 1533–44.
- Kojima, Y., Iwamoto, Y. and Yoshida, K. (2004) 'Memory of learning facilitates saccadic adaptation in the monkey', *The Journal of Neuroscience*, 24(34): 7531–9.
- Körding, K., Fukunaga, I., Howard, I.S., Ingram, J.N. and Wolpert, D.M. (2004) 'A neuroeconomics approach to inferring utility functions in sensorimotor control', *PLoS Biology*, 2(10): e330.

- Körding, K.P., Beierholm, U., Ma, W.J., Quartz, S., Tenenbaum, J.B. and Shams, L. (2007) 'Causal inference in multisensory perception', *PLoS One*, 2(9): e943.
- Körding, K.P., Tenenbaum, J.B. and Shadmehr, R. (2007) The dynamics of memory as a consequence of optimal adaptation to a changing body', *Nature Neuroscience*, 10(6): 779–86.
- Korenberg, A.T. and Ghahramani, Z. (2002) 'A Bayesian view of motor adaptation', *Current Psychology*, 21(4–5): 537–64.
- Krakauer, J.W., Pine, Z.M., Ghilardi, M.F. and Ghez, C. (2000) 'Learning of visuomotor transformations for vectorial planning of reaching trajectories', *The Journal of Neuroscience*, 20(23): 8916–24.
- Kurtzer, I., Pruszynski, J.A., Herter, T.M. and Scott, S.H. (2006) 'Primate upper limb muscles exhibit activity patterns that differ from their anatomical action during a postural task', *Journal of Neurophysiology*, 95(1): 493–504.
- Liu, D. and Todorov, E. (2007) 'Evidence for the flexible sensorimotor strategies predicted by optimal feedback control', *The Journal of Neuroscience*, 27(35): 9354–68.
- Marr, D. (1982) *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*, Freeman, New York.
- Mazzoni, P. and Krakauer, J.W. (2006) 'An implicit plan overrides an explicit strategy during visuomotor adaptation', *The Journal of Neuroscience*, 26(14): 3642–5.
- Mazzoni, P., Hristova, A. and Krakauer, J.W. (2007) 'Why don't we move faster? Parkinson's disease, movement vigor, and implicit motivation', *The Journal of Neuroscience*, 27(27): 7105–16.
- Morasso, P. (1981) 'Spatial control of arm movements', *Experimental Brain Research*, 42(2): 223–7.
- Nagengast, A.J., Braun, D.A. and Wolpert, D.M. (2009) 'Optimal control predicts human performance on objects with internal degrees of freedom', *PLoS Computational Biology*, 5(6): e1000419.
- O'Sullivan, I., Burdet, E. and Diedrichsen, J. (2009) 'Dissociating variability and effort as determinants of coordination', *PLoS Computational Biology*, 5(4): e1000345.
- Robinson, D.A., Gordon, J.L. and Gordon, S.E. (1986) 'A model of the smooth pursuit eye movement system', *Biological Cybernetics*, 55(1): 43–57.
- Shadmehr, R. and Mussa-Ivaldi, F.A. (1994) 'Adaptive representation of dynamics during learning of a motor task', *The Journal of Neuroscience*, 14: 3208–24.
- Shadmehr, R. and Wise, S.P. (2005) *The Computational Neurobiology of Reaching and Pointing: A Foundation for Motor Learning*, MIT Press, Cambridge, MA.
- Shadmehr, R., Orban de Xivry, J., Xu-Wilson, M. and Shih, T. (2010) 'Temporal discounting of reward and the cost of time in motor control', *The Journal of Neuroscience*, 30(31): 10507–16.
- Sing, G.C. and Smith, M.A. (2010) 'Reduction in learning rates associated with anterograde interference results from interactions between different timescales in motor adaptation', *PLoS Computational Biology*, 6(8): 1–14.
- Smith, M.A., Ghazizadeh, A. and Shadmehr, R. (2006) 'Interacting adaptive processes with different timescales underlie short-term motor learning', *PLoS Biology*, 4(6): e179.
- Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning: An Introduction*, Cambridge University Press, Cambridge, UK.
- Szentesi, P., Zaremba, R., van Mechelen, W. and Stienen, G.J. (2001) 'ATP utilization for calcium uptake and force production in different types of human skeletal muscle fibres', *The Journal of Physiology*, 531: 393–403.
- Taylor, J. A., Klemfuss, N.M. and Ivry, R.B. (2010) 'An explicit strategy prevails when the cerebellum fails to compute movement errors', *Cerebellum*, 9: 580–6.
- Thoroughman, K.A. and Shadmehr, R. (2000) 'Learning of action through adaptive combination of motor primitives', *Nature*, 407(6805): 742–7.
- Todorov, E. (2000) 'Direct cortical control of muscle activation in voluntary arm movements: a model', *Nature Neuroscience*, 3(4): 391–8.
- Todorov, E. (2005) 'Stochastic optimal control and estimation methods adapted to the noise characteristics of the sensorimotor system', *Neural Computation*, 17(5): 1084–108.
- Todorov, E. (2006) 'Optimal control theory', *Bayesian Brain: Probabilistic Approaches to Neural Coding*, 269–98.
- Todorov, E. and Jordan, M.I. (2002) 'Optimal feedback control as a theory of motor coordination', *Nature Neuroscience*, 5(11): 1226–35.
- Tseng, Y., Diedrichsen, J., Krakauer, J.W., Shadmehr, R. and Bastian A.J. (2007) 'Sensory prediction errors drive cerebellum-dependent adaptation of reaching', *Journal of Neurophysiology*, 98(1): 54–62.

- Uno, Y., Kawato, M. and Suzuki, R. (1989) 'Formation and control of optimal trajectory in human multijoint arm movement. Minimum torque-change model', *Biological Cybernetics*, 61(2): 89–101.
- van Beers, R.J. (2007) 'The sources of variability in saccadic eye movements', *The Journal of Neuroscience*, 27(33): 8757–70.
- van Beers, R.J., Sittig, A.C. and Gon, J.J. (1999) 'Integration of proprioceptive and visual position-information: An experimentally supported model', *Journal of Neurophysiology*, 81(3): 1355–64.
- van Beers, R.J., Haggard, P. and Wolpert, D.M. (2004) 'The role of execution noise in movement variability', *Journal of Neurophysiology*, 91(2): 1050–63.
- von Helmholtz, H. (1962). *Treatise on Physiological Optics*, The Optical Society of America.
- Wei, K. and Körding, K. (2010) 'Uncertainty of feedback and state estimation determines the speed of motor adaptation', *Frontiers in Computational Neuroscience*, 4: 1–9.
- Wolpert, D.M., Ghahramani, Z. and Jordan, M.I. (1995) 'Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study', *Experimental Brain Research*, 103(3): 460–70.
- Wong, A.L. and Shelhamer, M. (2011) 'Sensorimotor adaptation error signals are derived from realistic predictions of movement outcomes', *Journal of Neurophysiology*, 105(3): 1130–40.
- Xu-Wilson, M. Chen-Harris, H., Zee, D.S. and Shadmehr, R. (2009) 'Cerebellar contributions to adaptive control of saccades in humans', *The Journal of Neuroscience*, 29(41): 12930–9.