

6 Learning to make accurate predictions

In the previous couple of chapters we suggested that as the brain sends a motor command to the muscles, it also predicts the sensory consequences. When the sensory system reports on the consequences, the brain combines the observations with its own predictions to form an estimate of the state of the body and the environment. However, combining predictions with observations makes sense only if the predictions are unbiased estimates of the observations. If trial after trial there are persistent differences between predictions and observations, that is, the brain's predictions are consistently biased, then there is something wrong in these predictions.

Occasionally, our predictions are quite wrong. For example, one of us (RS) has a couple of teenagers at home who occasionally raid the refrigerator, drinking the milk carton almost dry. On more than one occasion picking up the carton has resulted in his hand jerking upward. The brain predicted the carton to be heavier than reality, and so the motor commands produced an unexpected motion, resulting in a difference between the expected sensory feedback and the observed quantities. Combining a biased prediction with an unbiased sensory feedback results in a biased estimate of state, something that we would generally want to avoid. Therefore, it seems rational that the brain should somehow change its predictions so that eventually the predictions agree with what the sensory system reports. How should it do this? That is, how should the brain go about improving its predictions?

The problem of forming unbiased predictions of sensory observations is one of the fundamental problems of learning (the other is learning how to produce useful motor commands, something that we will tackle in a few chapters). For example, when you have formed an accurate representation, i.e., an *internal model*, of your arm, you can apply motor commands to it and (on average) accurately predict the sensory consequences of those commands in terms of what you will see and what you will feel. An internal model is simply a map that mimics – in its relation between input and output signals - certain aspects of reality in which the brain is immersed. For example, if you hold a golf club in your hand, learning an accurate internal model of your arm and the club allows you to apply motor commands and produce a motion of the club, and the ball, that agrees with your predictions. Professional athletes are particularly good at making predictions about the sensory consequences of their motor commands. In tennis, a good player will know whether the ball will land in or out at the moment he or she hits the ball. However, note that having an accurate internal model does not mean that you can hit the ball well; it only

means that given some motor commands that you produced, you can predict the sensory consequences.

Let us consider the problem of picking up a carton of milk. Our internal model is a *generative model*, i.e., a model that is capable of generating data (things that we can observe, or measure) according to a set of rules, as shown in Fig. 6.1. According to the model in Fig. 6.1, when we send motor command u to our arm, it produces a change in the state of our hand h . This state depends on our motor command, as well as the mass that the hand may be holding, m . The consequence of our motor command is observed via our visual and proprioceptive sensors, y_v and y_p . What we have expressed in Fig. 6.1 is an internal model that describes how motor commands produce sensory consequences. In this internal model there are two hidden states: position of our hand, and mass of the object. Using our observations we can estimate these hidden states. We begin with some prior belief about the position of our hand and mass of the object. We have an objective, e.g., bring the state of the hand to some position. Based on our prior belief regarding state of our hand and the mass of the object, we generate motor command u . For example, if our objective is to lift the object, a prior belief that m is large would dictate that u should be large as well. Using our internal model, we predict the sensory consequences \hat{y}_v and \hat{y}_p . If the mass is less than what we had expected, there will be a difference between the predicted and actually observed sensory consequences. Using this prediction error, we update our belief about the state of our hand and the mass of the object, forming a posterior belief. This posterior becomes the prior for our next attempt to lift the carton. As a result, the motor commands that we will generate to pick up the carton for a second time will be different than the first time (because our prior associated with m will have changed). The trial to trial change in the motor commands will reflect our learning.

The idea is that the brain is constantly trying to predict things that it can observe. To make accurate predictions, it builds an internal model of the process that generates the data that it is observing. We can represent this generative model as a collection of observable quantities as well as a group of hidden states that cannot be observed. The problem of learning can be viewed as the problem of estimating these hidden states. The driving force for learning is the difference between the predicted sensory measurements and the actually measured quantities, i.e., sensory prediction error.

We saw in the last chapter that there are Bayesian algorithms that can help us solve the estimation problem. Therefore, if we frame the problem of learning in terms of state estimation, then we can apply these same algorithms. Here, we are interested in asking whether people and other animals learn internal models in ways that resemble optimal estimation. To help answer our question, it will be useful to consider other algorithms that can solve our learning problem, in particular those that are not Bayesian, so we will have a point of comparison.

To start our discussion, we will begin with some very simple learning problems. These problems come from experiments in which animals learn to predict the relationship between stimuli and reward or punishment. The experiments are called *classical conditioning*. We will see that even in this very simple paradigm, animal behavior exhibits surprisingly complex characteristics. The estimation framework that we developed in the last two chapters does a fairly good job of explaining the behaviors. We will then apply the same framework to experiments in which people learn to control novel tools or adapt to novel distortions on their sensory feedback.

6.1 Examples from animal learning

Let us start with a simple example from the animal learning literature called the *blocking* experiment, introduced by Leon Kamin (1968). In this experiment, the animal (a rat) was trained to press a lever to receive food. After completion of this baseline training, a light was turned on for 3 minutes and during the last 0.5 second of that period, a mild electric shock was given to the foot. In response to the shock, the animal reduced the lever-press activity. As the light-shock pairing was presented again, the animal learned to associate the light with the shock and in anticipation stopped pressing the lever (freezing behavior). After completion of this training, a second stimulus (a tone) was presented with the light. That is, the light and a tone were presented together, followed by the shock. Pairing of light with shock and pairing of light+tone with shock constituted set 1 and set 2 of the training. The experiment concluded with a test period, during which only the tone was presented. The crucial result was that during this test period, the animal did not reduce its lever pressing activity, despite the fact that in set 2 the tone was present when it was being shocked. Somehow, the animal did not fear the tone.

In control animals that had only experienced a tone paired with shock, in the test period there was reduced lever pressing in response to the tone. Therefore, the animal could easily learn to associate the tone with the shock. However, the fact that in set 1 the light was paired with shock,

this prior pairing somehow blocked the ability of the brain to associate at least some of the shock to the tone in the subsequent light+tone training. The effect is called *Kamin blocking*.

Our objective here is to describe this learning, and other examples like it, in an estimation framework, and then test the idea that the animals are implicitly doing optimal estimation: computing uncertainties, forming a learning rate (i.e., the sensitivity to error or the Kalman gain) that varies according to this uncertainty, and then applying this learning rate to their prediction errors. In comparing the experimental data with the Kalman algorithm, it is useful to have an alternate hypothesis regarding how the learning might take place. Let us sketch one such alternate hypothesis, called the *LMS algorithm* (least mean squared).

6.2 The LMS algorithm

Suppose that we have a very simple internal model that relates inputs (stimuli: light and tone) on trial n , written as $\mathbf{x}^{(n)}$ with an output (e.g., shock) $y^{(n)}$. That is, the two input quantities are x_1 and x_2 (representing light and tone). When light is on, $x_1 = 1$ and when it is off $x_1 = 0$.

Similarly, when tone is on $x_2 = 1$. There is one output quantity y^* , representing shock. When we make an estimate $\hat{y}^{(n)} = \mathbf{x}^{(n)T} \hat{\mathbf{w}}^{(n)}$ and observe a prediction error $y^{(n)} - \hat{y}^{(n)}$, we face a credit assignment problem: how should we assign responsibility for the error on our estimate $\hat{\mathbf{w}}^{(n)}$? That is, which element of $\hat{\mathbf{w}}^{(n)}$ should get blamed most for the error, and which should get blamed the least? To answer this question, let us consider a simple learning algorithm called LMS. In our problem, we have \mathbf{x} and $\hat{\mathbf{w}}$ as 2×1 vectors, as shown in Fig. 6.2. When we project vector $\hat{\mathbf{w}}$ onto \mathbf{x} , i.e., the dot product of the two vectors ($\mathbf{x}^T \cdot \hat{\mathbf{w}}$), we get a scalar quantity p such that $p = \|\hat{\mathbf{w}}^{(n)}\| \cos \alpha$, where α is the angle between the vectors $\hat{\mathbf{w}}^{(n)}$ and $\mathbf{x}^{(n)}$, and the double vertical lines refer to the magnitude of the vector. The quantity $\hat{y}^{(n)}$ is related to p :

$$\begin{aligned} \cos \alpha &= \frac{\hat{\mathbf{w}}^{(n)T} \mathbf{x}^{(n)}}{\|\hat{\mathbf{w}}^{(n)}\| \|\mathbf{x}^{(n)}\|} \\ p &= \frac{\hat{y}^{(n)}}{\|\mathbf{x}^{(n)}\|} \end{aligned} \tag{6.1}$$

Eq. (6.1) implies that we want to change $\hat{\mathbf{w}}^{(n)}$ by amount $\Delta^{(n)}$ so that when we project

$\hat{\mathbf{w}}^{(n)} + \Delta^{(n)}$ onto $\mathbf{x}^{(n)}$, we get a length specified by the dashed line in Fig. 6.2, i.e., $y^{(n)} \|\mathbf{x}^{(n)}\|^{-1}$.

You will note that there is no unique vector $\Delta^{(n)}$ that will do this for us. As long as we choose $\Delta^{(n)}$ so that $\hat{\mathbf{w}}^{(n)} + \Delta^{(n)}$ touches any point on the dashed line, we have one solution. A reasonable approach would be to add to $\hat{\mathbf{w}}^{(n)}$ a vector parallel to $\mathbf{x}^{(n)}$, with a magnitude that is proportional to $y^{(n)} - \hat{y}^{(n)}$. By making $\Delta^{(n)}$ parallel to $\mathbf{x}^{(n)}$, we are adding the smallest possible vector length to $\hat{\mathbf{w}}^{(n)}$:

$$\begin{aligned}\Delta^{(n)} &= \left(\frac{y^{(n)}}{\|\mathbf{x}^{(n)}\|} - \frac{\hat{y}^{(n)}}{\|\mathbf{x}^{(n)}\|} \right) \frac{\mathbf{x}^{(n)}}{\|\mathbf{x}^{(n)}\|} \\ &= \frac{1}{\|\mathbf{x}^{(n)}\|^2} (y^{(n)} - \hat{y}^{(n)}) \mathbf{x}^{(n)}\end{aligned}\quad (6.2)$$

The quantity $\mathbf{x}^{(n)} \|\mathbf{x}^{(n)}\|^{-1}$ is a unit vector along the direction of $\mathbf{x}^{(n)}$. If we change $\hat{\mathbf{w}}^{(n)}$ by the vector specified in Eq. (6.2), we will completely compensate for the error. In practice, the step size is usually a fraction of this quantity, where $0 < \eta < 1$:

$$\hat{\mathbf{w}}^{(n+1)} = \hat{\mathbf{w}}^{(n)} + \frac{\eta}{\|\mathbf{x}^{(n)}\|^2} (y^{(n)} - \hat{y}^{(n)}) \mathbf{x}^{(n)} \quad (6.3)$$

The algorithm in Eq. (6.3) is the estimation method presented by Widrow and Hoff (Widrow and Hoff, 1960). It is widely known as the Delta rule or LMS (Least Mean Squared) rule in the engineering literature, and Rescorla-Wagner rule in the psychological literature (Rescorla and Wagner, 1972). A less graphical but more immediate way to derive the LMS rule is by taking the gradient of the squared error with respect to the parameter vector and then “descending” by a small increment in that direction. The gradient of the squared error, $\mathcal{E}^{(n)2} = (y^{(n)} - \hat{y}^{(n)}(\mathbf{w}))^2$ is derived by the chain rule:

$$-\nabla_{\mathbf{w}} \mathcal{E}^{(n)2} = -\frac{\partial \mathcal{E}^{(n)2}}{\partial \mathbf{w}} = 2\mathcal{E}^{(n)} \frac{\partial \hat{y}^{(n)}}{\partial \mathbf{w}} = 2(y^{(n)} - \hat{y}^{(n)}) \mathbf{x}^{(n)} \quad (6.4)$$

The quantity of the left-hand side – modulo the step size – is the step of the LMS rule in Eq. (6.3). How does LMS compare to the Kalman approach?

6.3 Learning as state estimation

To solve the same problem using state estimation, the basic idea is to describe our observations y as a function of some hidden states \mathbf{w} , that is, we start with a generative model. The generative model that describes our observations can be as follows:

$$\begin{aligned}\mathbf{w}^{(n+1)} &= \mathbf{w}^{(n)} + \boldsymbol{\varepsilon}_w & \boldsymbol{\varepsilon}_w &\square N(0, Q) \\ y^{(n)} &= \mathbf{x}^{(n)T} \mathbf{w}^{(n)} + \varepsilon_y & \varepsilon_y &\square N(0, \sigma^2)\end{aligned}\quad (6.5)$$

The model is shown graphically in Fig. 6.3, in which the shaded circles are the observed variables. Our objective is to find some estimate of the hidden states $\hat{\mathbf{w}}$ that can predict our observation y . On trial n we start with a prior estimate $\hat{\mathbf{w}}^{(n|n-1)}$ and uncertainty $P^{(n|n-1)}$. We make a prediction $\hat{y}^{(n)} = \mathbf{x}^{(n)T} \hat{\mathbf{w}}^{(n|n-1)}$, and then learn from the sensory prediction error $y^{(n)} - \hat{y}^{(n)}$:

$$\hat{\mathbf{w}}^{(n|n)} = \hat{\mathbf{w}}^{(n|n-1)} + \mathbf{k}^{(n)} (y^{(n)} - \hat{y}^{(n)}). \quad (6.6)$$

The key comparison is between the terms that multiply the prediction error $y^{(n)} - \hat{y}^{(n)}$ in Eq. (6.3) and the Kalman gain in Eq. (6.6). In LMS, the terms that multiply the prediction error depend only on the current input $\mathbf{x}^{(n)}$. In the Kalman gain, we have:

$$\mathbf{k}^{(n)} = \frac{P^{(n|n-1)} \mathbf{x}^{(n)}}{\left(\mathbf{x}^{(n)T} P^{(n|n-1)} \mathbf{x}^{(n)} + \sigma^2 \right)}. \quad (6.7)$$

In effect, in the Kalman gain we weigh the prediction error based on the ratio of the uncertainty in the state that we wish to predict $P^{(n|n-1)}$, and our uncertainty about our measurements (the denominator in Eq. 6.7). An important fact is that the state uncertainty matrix P keeps a history of the variance and covariance in the prior inputs. This uncertainty affects how we will respond to the prediction errors. In contrast, in LMS the history of the prior inputs does not play a role in the problem of learning from error. Furthermore, while LMS leaves the step size somewhat unspecified – by allowing for an arbitrary step parameter η – Kalman offers a rigorous way to determine the step size – as well as its direction – based on the input's statistics.

Now let us simulate the Kamin blocking experiment and compare performance of the two learning algorithms. In the first set of training, suppose that we have 20 trials in which $x_1 = 1$ (light is on), $x_2 = 0$ (tone is off), and $y^* = 1$ (shock is present). In the second set of training

(another 20 trials), x_1 and y^* remain unchanged but $x_2 = 1$. Our objective is to estimate w_1 and w_2 , the weights associated with x_1 and x_2 , and then predict behavior in the test period during which $x_1 = 0$ and $x_2 = 1$ (light is off, tone is on, not shown in Fig. 6.2). Fig. 6.3 shows the simulation results for the Kalman algorithm and LMS. We begin with \hat{w}_1 and \hat{w}_2 set to zero, i.e., no association between inputs and shock. In the first set when $x_1 = 1$, both algorithms learn to set $\hat{w}_1 = 1$. That is, the light fully predicts the shock. In set 2, when light and tone are both on, there is no prediction error to learn from as the light fully predicts the shock, and therefore the weights remain unchanged. So by the end of set 2, the weight associated with tone \hat{w}_2 is unchanged from before the experiment began, i.e., it is still around zero. In the subsequent test period when tone appears alone, its weight is near zero, and the animal does not ‘freeze’. In this case, the two algorithms behave similarly.

A number of observations in this simulation are worth highlighting because, as we will see shortly, they play a crucial role in the success of the Kalman algorithm and the failure of LMS in the next experiment. In the Kalman portion of Fig. 6.3, the terms k_1 and k_2 are the elements of the Kalman gain \mathbf{k} . They specify how much of the prediction error is assigned to \hat{w}_1 and \hat{w}_2 . The terms $P_{1,1}$, $P_{1,2}$, and $P_{2,2}$ are elements of matrix P , describing the variance and covariance of \hat{w}_1 and \hat{w}_2 . In set 1 (the first 20 trials) when $x_1 = 1$ but $x_2 = 0$, uncertainty about w_2 is large because its value does not affect the output, as reflected in the large value of $P_{2,2}$. In set 2, when $x_1 = 1$ and $x_2 = 1$, the value of $P_{1,2}$ becomes negative because of the covariance of \hat{w}_1 and \hat{w}_2 (if \hat{w}_1 increases, then \hat{w}_2 must decrease). That is, when the inputs are ‘on’ together, they produce a negative covariance on the weights associated with them: to maintain a correct output, if a weight increases, the other must decrease. We saw this phenomenon earlier in Fig. 4.13B. The elements of matrix P keep a record of the history of the prior inputs, whereas no such record is kept in the LMS algorithm. Let us now consider an experiment in which this record plays an important role.

The next experiment is called *backward blocking*. In this experiment, during set 1 the animal is trained with both stimuli, i.e., $x_1 = 1$ and $x_2 = 1$. During set 2, only one of the stimuli remains

present, i.e., $x_1 = 1$ but $x_2 = 0$. The Kalman and LMS algorithms now make very different predictions (Fig. 6.4). The Kalman algorithm, like LMS, predicts that during set 1, the system will learn to assign some weight to both stimuli, i.e., $\hat{w}_1 = 0.5$ and $\hat{w}_2 = 0.5$. In set 2 when only x_1 is present, there is a prediction error. Both LMS and Kalman increase the weight of \hat{w}_1 to one. However, because x_2 is not present, LMS does not alter \hat{w}_2 . The Kalman algorithm, however, reduces \hat{w}_2 to zero. Therefore, LMS predicts that at the end of set 2, the animal will show as much freezing in response to x_2 (tone) as it did at the end of set 1. In contrast, Kalman predicts that at the end of set 2, in response to x_2 the animal will show significantly less freezing. Kalman makes the unusual prediction that if light predicts the shock, then tone must be ineffective, despite the fact that earlier the animal was shocked in the presence of both light and tone.

The reason for this is that in set 1, the Kalman algorithm sets the covariance of \hat{w}_1 and \hat{w}_2 to be negative (the term P_{12}). This means that in set 2, when \hat{w}_1 is increased (the weight of light), absent of other evidence \hat{w}_2 (weight of tone) must decrease. This is what covariance implies. In effect, the prior history causes the Kalman algorithm to alter the weight of \hat{w}_2 in set 2, even though in that set x_2 is not present to directly cause learning from the prediction error.

Perhaps we should note that the solution given by the Kalman algorithm is not better than the solution of the LMS algorithm in an absolute sense. It is plausible to encounter situations in which either one is more effective than the other. However, the Kalman algorithm is consistent with Bayesian inference of the posterior probability to observe a shock given the history of experienced stimuli, in the context of our generative model. More interestingly, it nicely agrees with behavior of rats. Ralph Miller and Helena Matute (Miller and Matute, 1996) performed the backward blocking experiment in a group of rats (Fig. 6.5). They had three stimuli: A, B, and X (three different kinds of auditory stimuli). In phase 1, they paired A and X with B, i.e., $AX \rightarrow B$. In phase 2, they paired A with B, i.e., $A \rightarrow B$. In phase 3, they paired B with shock. In the test period, they presented X or A alone and tested the time period in which the animal stopped moving (freezing period). They found that the animals froze much longer in response to stimulus

A than stimulus X in the test period. If we label stimulus A as x_1 and stimulus X as x_2 , then their results are consistent with the Kalman algorithm and inconsistent with LMS.

Therefore, in our example the sensitivity to prediction error, i.e., the learning rate, depended on the history of the past observations. In the Kalman framework, this history defined an uncertainty that dictated the sensitivity to prediction error.

6.4 Prediction errors drive adaptation of internal models

In examples above the tasks were forms of classical conditioning. The animals did not have to do anything to get reward or avoid punishment; the best they could do was predict the consequences of an externally controlled set of stimuli. A more interesting scenario is one in which the brain needs to figure out the consequences of a self-generated action and produce actions that maximize some measure of performance. An example of this is tasks in which people learn to control movements of their body.

During development, bones grow and muscle mass increases, changing the relationship between motor commands (e.g., torques) and motion of the limb (position and velocity). In addition to such gradual variations, the arm's dynamics change over a shorter timescale when we grasp objects and perform manipulation. It follows that in order to maintain a desired level of performance, our brain needs to be 'robust' to these changes. This robustness may be achieved through an updating, or adaptation, of an internal model that predicts the sensory consequences of motor commands. Indeed, people appear to excel in the ability to adapt rapidly to the variable dynamics of their arm as their hand interacts with the environment: we easily switch from moving a small pen, to a long stick, to a baseball bat. A task in which people use their hand to interact with a novel environment might be a good candidate to study how the brain learns internal models.

There are two well studied versions of the adaptation paradigm. In one version, called *visuomotor adaptation*, the investigator introduces a perturbation that distorts the visual consequences of the motor commands but leaves the proprioceptive consequences unchanged. This is typically done by wearing prism goggles, or having people move a cursor on the screen in which the relationship between cursor position and hand position is manipulated (Fig. 6.6A). In another version of the adaptation paradigm, called *force field adaptation*, the investigator

introduces a perturbation that alters both the visual and proprioceptive consequences of motor commands. This is typically done by having the volunteer hold the handle of an actuated manipulandum (a robotic arm) that produces forces on the hand (Fig. 6.6B). This type of adaptation can also be done by having people reach in a rotating room (the rotation imposes novel forces on the hand), or even in micro-gravity in which the usual forces are removed. In both the visuomotor and force field experiments, learning depends on sensory prediction errors.

The oldest record of visuomotor adaptation experiment that we know of is an 1867 report by Hermann von Helmholtz. In that work, he asked subjects to point with their finger at targets while wearing prism lenses that displaced the visual field laterally. When the displacement was to the left, subjects initially had errors (an overshoot) in that direction and after some practice, they learned to compensate for the visual displacement. Helmholtz observed that as soon as the prisms were removed, subjects made erroneous movements to the right of the target. This is known as an *after-effect* of adaptation.

Nearly a century later, in the early 1960s there was renewed interest in motor adaptation because of the space program. Astronauts were preparing to leave earth, and there was concern as to whether their motor system could function in zero gravity. Could the brain adapt their motor commands so that they could function in this radically different environment? Were there methods to help speed the process of adaptation? To quantify the ability of the brain to adapt, Richard Held and Sanford Freedman (Held and Freedman, 1963) repeated Helmholtz's experiment with a new twist. They compared the performance of subjects when they actively moved their arm while viewing their finger through prism glasses, versus when they viewed their finger but their arm was passively moved for them (Fig. 6.7). In the experiment, the subject's arm was tied to a board that could move about the elbow. When the subject actively moved her elbow, the result was a circular motion of the finger. When the board was moved by the experimenters, the result was the same circular motion of the subject's finger, but now there were little or no motor commands that were generated by the subject. In both cases, the subject viewed the motion of the finger via a prism that induced a displacement in the visual feedback. After this viewing, the subject was tested in a pointing task (marked Test in Fig. 6.7). Held and Freedman found that in the test session, the subject showed after-effects, but only if during the earlier session she had viewed her hand while actively moving it. She did not have after-effects if she had viewed her hand while it was moved passively. In their words: "Although the passive-movement condition provided the eye with the same optical information that the active-

movement condition did, the crucial connection between motor output and visual feedback was lacking”. In our terminology, sensory prediction error was missing in the passive condition, as the subjects did not actively generate a movement, and therefore could not predict the sensory consequences.

There is a more recent example of visuomotor adaptation that provides striking evidence for the crucial role of sensory prediction errors. Pietro Mazzoni and John Krakauer (Mazzoni and Krakauer, 2006) had people move their wrist so that the position of the index finger was coupled with the position of a cursor on a screen. There were always 8 targets on display, spanning 360° . On a given trial, one of the targets would light up and the subject would move the cursor in an out-and-back trajectory, hitting the target and then returning to the center. After a baseline familiarization period (40 trials), the experimenters imposed a 45° counter-clockwise rotation on the relationship between the cursor and finger position (early adaptation, Fig. 6.8A). Let us represent this perturbation as a hidden state and label it with r . Now, a motor command u that moved the hand in direction θ did not produce a cursor motion in the same direction, but in direction $\theta + r$. If we label the predicted sensory consequences $\hat{y} = \theta$ and the observed consequences $y = \theta + r + \varepsilon_y$, then there is a sensory prediction error $y - \hat{y}$. The objective is to use this prediction error to update an estimate for \hat{r} . With that estimate, for a target at direction θ^* , we can generate a motor command $u = \theta^* - \hat{r}$ to bring the cursor to the target. Indeed, after about 80 trials, in response to target at θ^* people would move their hands to $\theta^* - 40$ so the cursor would land within 5 degrees of the target (as shown in ‘adaptation’ subplot of Fig. 6.8B).

Now, Mazzoni and Krakauer tried something quite clever: they took another group of naïve subjects and after they had experienced a couple of rotation trials, they simply told them: “Look, you made two movements that had large errors because we imposed a rotation that pushed you 45° counter clockwise. You can counter the error by aiming for the neighboring clockwise target.” That is, simply issue the motor command $u = \theta^* - 45$ and as a consequence, the cursor will move at direction θ and land at the target. Indeed, the subjects followed this strategy: on the very next trial, all the error dropped to zero (strategy group, Fig. 6.8B). However, now something very interesting happened: as the trials continued, the errors gradually grew! What’s more, the rate of change in the errors in this ‘strategy’ group was exactly the same as the rate of change in the regular adaptation paradigm.

To explain this, Mazzoni and Krakauer hypothesized that on trial 43, when the subjects in the strategy group were producing the motor commands that brought the cursor to the target, there was still a discrepancy between the predicted and observed sensory consequences of motor commands $y - \hat{y}$. This is because whereas *explicitly* they had been told of the perturbation, *implicitly* their estimate was still around zero, $\hat{r} \approx 0$. The implicit estimate (the motor system's estimate) learned from prediction error, and that learning took many trials (as shown in the adaptation group, Fig. 6.8B).

This experiment hints at the complexity of the problem that we are considering: there are multiple learning systems in the brain. The explicit system is typically associated with our conscious awareness, whereas the implicit system is typically associated with an unconscious process. Both can formulate internal models, learn from prediction errors, and contribute to our motor commands. The data in Fig. 6.8B hints that the explicit system learns extremely quickly, whereas the implicit system learns gradually. In the tasks that we will be considering in this chapter, our assumption will be that we are primarily gauging the influence of the implicit system. The main reason for this assumption is that, as we will see shortly, people with amnesia who have severe deficit in remembering explicit information nevertheless learn to alter their motor commands in the adaptation experiments. This does not mean that in healthy people, internal models learned by the explicit system play no role in adaptation. Indeed, we will return to this topic when we consider the multiple timescales of memory in the next chapter. However, for now, let us imagine that in a typical adaptation experiment involving tens to hundreds of trials, the main contributor is the implicit system.

In summary, an internal model is simply an association between motor commands and their sensory consequences. The driving force in learning an internal model is the sensory prediction error.

6.5 A generative model of sensorimotor adaptation experiments

Whereas in the visuomotor paradigms the visual consequence of motor commands is altered but not the proprioceptive consequences, in a force field paradigm both sensory consequences are altered, resulting in multiple prediction errors. An example of the force field adaptation paradigm is shown in Fig. 6.9. In this experiment, Shadmehr and Mussa-Ivaldi (Shadmehr and Mussa-Ivaldi, 1994) had a volunteer reach to a target while holding a handle attached to a robotic arm

(Fig. 6.10A). When the robot motors were turned off, the hand motion was generally straight (Fig. 6.9A). To induce adaptation, the robot produced a force field that pushed on the hand as a function of hand velocity, as shown in Fig. 6.9B. These forces perturbed the reach trajectories (Fig. 6.9C). With practice, the subject adapted his motor commands and produced forces on his hand that generally compensated for the robot forces, resulting in a restoration of the straight trajectories (Fig. 6.9D). When the forces were unexpectedly removed (in catch trials), the hand trajectory exhibited an after effect.

When the subjects are caught by surprise by the field, the hand trajectory is perturbed, but it eventually reaches the target, as shown in Figure 6.8C. This is because the field depends only on the velocity of the hand: at the end of movement the hand is at rest, and there is no perturbing force pushing it away from the target. Reaching the target is all that subjects were asked to do. The experiment shows that subjects are compelled to adjust their motor command so as not only to reach the target, but to do so along a path that is similar to the initial unperturbed path. If asked to describe the forces they experience and what they did to compensate, they would not have a clear answer. Nevertheless, their brain alters the motor commands and attempts to bring the cursor to the target in more or less a straight line.

An important observation is that when the severely amnesic subject H.M. and other amnesic subjects were tested on this task (Fig. 6.10B), they learned and showed after effects during the training session, a few hours after this initial training session, and even the next day. This adaptation took place despite the fact that they could not consciously remember having seen the robot or having done the task before (Shadmehr et al., 1998). Clearly, a crucial aspect of adaptation in the force field paradigm is via an implicit, non-conscious memory system. [It is of course possible that some explicit system was involved in the motor learning but that this explicit component itself was not remembered the next day. What is clear is that it is possible to recall motor memory without a declarative awareness of that memory.]

How do subjects adapt to the visuomotor or force field perturbations? This is a very broad question that might be answered at many levels. For example, our answer might invoke something about the molecular basis of memory in the brain regions involved in motor control. However, let us focus here on the basic computational problem: in principle, what is the problem that is being solved? As we will see in this and the following chapters, the perturbation is producing two kinds of errors: a sensory prediction error (the motor commands did not produce

the expected hand trajectory in visual and proprioceptive coordinates), and a reward prediction error (the motor commands did not get us to the target in time and so we did not get rewarded for our efforts). In this chapter, we will focus on the problem of accurately predicting the sensory consequences of the motor command.

Perhaps the simplest approach is to imagine that our hand did not move as we had expected because there was a force (a hidden state) that we did not account for. Therefore, if we could estimate this force, we could accurately predict the motion that should result. To get the hand to the target, we would produce motor commands that compensate for this perturbing force. In the visuomotor paradigm, the perturbation is a bias that distorts cursor position with respect to hand position (another hidden state). If we had a correct estimate of this perturbation, we could produce motor commands that accounted for it, producing accurate movements of the cursor to the target.

To approach the problem of learning internal models, we need to specify a generative model that relates motor commands with sensory consequences. To motivate our approach, we need to consider a curious fact: people who participate in these experiments do not simply learn to produce motor commands that compensate for a perturbation. Rather, practice produces a sensory illusion.

For example, consider a paradigm in which people use their right arm to adapt to a visuomotor rotation, i.e., they reach to a target and eventually learn to cancel the imposed bias on the visual feedback. It turns out that the training also alters where they think their right arm is located in space (Fig. 6.11). Let us explain this result in some detail. In Figure (6.5), the two setups not only allow the subject to reach to a target, but the setups also allow the subject to point to the position of one hand with the other hand. Now suppose that before the reach task starts, we remove vision of both hands, have the subject move one hand to some location, and then have him point with the other hand to the location of the moved hand. People are usually pretty accurate at this kind of alignment. Now as the visuomotor adaptation experiment begins, the subject reaches to the visual target and she would see her hand (or cursor representing her hand) to the left of the target because of the bias that has been imposed by the visual display (Fig. 6.11, before adaptation). Trial after trial, the seen position of the hand gets closer to the target, and the actual hand position gets farther from the target (Fig. 6.11, after adaptation). However, if we now ask the subject to estimate her hand position, she points to somewhere a bit to the left of where it

actually is (van Beers et al., 2002). Indeed, a similar thing happens when people adapt their reaching movements to force fields (Haith et al., 2008): if a force pushes the hand to the right, people learn to compensate for these forces but afterwards believe that their hand is a little to the right of its actual position. That is, it is not just the motor commands that have adapted. Something has also changed the perception of where the hand is located in space. Adaptation seems to produce a sensory illusion.

We can account for these results by describing the process of adaptation in the framework of state estimation. A reasonable approach was suggested by Adrian Haith, Carl Jackson, Chris Miall, and Sethu Vijayakumar (2008). Their generative model is graphically described in Fig. 6.12. To make things easy, let us represent our problem in a one dimensional world (so our variables are scalars). The position of the hand (or finger) on trial n , represented by variable h , depends on the motor commands u , and is observed by the brain through two independent sensors, vision and proprioception, represented by y_v and y_p . The sensory measurements are affected by noise $\varepsilon_v \sim N(0, s_v)$, and $\varepsilon_p \sim N(0, s_p)$, as well as by two potential sources of perturbation r_v and r_p :

$$\begin{aligned} y_v^{(n)} &= r_v^{(n)} + h^{(n)} + \varepsilon_v \\ y_p^{(n)} &= r_p^{(n)} + h^{(n)} + \varepsilon_p \end{aligned} \quad (6.8)$$

Motor commands produce a change in the state of the finger, and they too are subject to a potential source of perturbation r_u , and noise $\varepsilon_u \sim N(0, q_u)$. Grossly simplifying the dynamics of movements, let us represent this relationship as:

$$h^{(n+1)} = u^{(n)} + r_u^{(n)} + \varepsilon_u. \quad (6.9)$$

Eq. (6.8) is a model of the sensory system: we are assuming that our two sensors, vision and proprioception, can tell us something about the position of our finger. We are also assuming that our model of how the sensors are related to the actual position of the hand may be biased and that bias is a form of perturbation. Similarly, Eq. (6.9) is a model of how the motor commands produce a hand position. This model can also be biased, i.e., our model of the relationship between motor commands and the position of the hand may be wrong. This bias is represented as a perturbation (e.g., a force field is a perturbation that ‘adds’ a bias to our motor commands, moving our hand in an unexpected way).

Now when we see a visual target at y_v^* , our objective is to place the visual representation of our hand y_v at the target. To do that, we need to produce the following motor command:

$$u^{(n)} = y_v^* - \hat{r}_v - \hat{r}_u. \quad (6.10)$$

That is, we produce a motor command that cancels the estimated perturbations \hat{r}_v and \hat{r}_u . When we are asked to indicate the position of our hand by pointing to it, we are being asked about our estimate \hat{h} . Summarizing our model, the estimation problem is as follows: given motor command u and the sensory feedback y_v and y_p , estimate the state of the perturbations r_v , r_p , and r_u , and also estimate the position of the hand h .

Before we solve the problem mathematically, let us get an intuitive feel for what the solution should look like. The purpose of training is to help us form a better estimate of the various perturbations so that by the end of training, given a motor command u , our predictions about sensory feedback (\hat{y}_v and \hat{y}_p) match the actual observations (y_v and y_p). Right after the visual perturbation is introduced (say $r_v = -2$, i.e., we see our hand or cursor 2 cm to the left of where it actually is), \hat{y}_v will differ substantially from y_v , i.e., there will be a prediction error regarding the visual sensory feedback. However, at the same time we will have $\hat{y}_p \approx y_p$, i.e., no prediction error regarding proprioceptive feedback. What can account for this pattern of prediction errors?

One possibility is that all the prediction error is due to perturbation r_v . For example, if $r_v = -2$, then by setting $\hat{r}_v = -2$ and $\hat{r}_p = \hat{r}_u = 0$, Eq. (6.10) will bring the visual representation of the hand to the target as well as maintain $\hat{y}_p \approx y_p$. This is one way with which we can assign credit to the various potential perturbations. However, this is not the only way. If $r_v = -2$, then by setting $\hat{r}_v = -1.5$, $\hat{r}_u = -0.5$, and $\hat{r}_p = +0.5$, once again Eq. (6.10) will be sufficient to make $\hat{y}_v = y_v$ and $\hat{y}_p = y_p$, i.e., eliminate the prediction error. If we assign the credit this way, we interpret part of the visual perturbation in terms of a motor perturbation r_u , and part in terms of sensory perturbations r_v and r_p . Indeed, as long as $\hat{r}_v + \hat{r}_u = -2$ and $\hat{r}_p + \hat{r}_u = 0$, we have a solution to our problem. So we see that in principle, there are multiple solutions to our problem. Which kind of credit assignment is best? The credit assignment will depend on the uncertainty

that we have regarding how the various hidden states (the perturbations) affected our observations.

The key idea is that despite the fact that the subject has experienced a prediction error because the experimenter introduced a bias on the visual feedback, from the point of view of the subject the optimal thing to do may not be an assignment of the error solely to a visual perturbation. The error may also be due to an incorrect internal model that associates motor commands with position of the hand, as well as sensory feedback that measures that position. To make a sensible decision, the subject must take into account the reliability of each source of information. And this can only be done based on past experience. As we will see, it is this ambiguity that can account for the fact that during adaptation to a visual perturbation, subjects acquire a sensory illusion regarding estimate of their hand position.

6.6 Accounting for sensory illusions during adaptation

Let us represent the learning problem in a visuomotor adaptation experiment via the generative model shown in Fig. 6.12. Set vector \mathbf{x} to be the states that we need to estimate

$\mathbf{x} = [r_v \quad r_p \quad r_u \quad h]^T$ and vector \mathbf{y} to be the sensory observations $\mathbf{y} = [y_v \quad y_p]^T$. The generative model can be written as:

$$\begin{aligned} \mathbf{x}^{(n)} &= A\mathbf{x}^{(n-1)} + \mathbf{b}u^{(n-1)} + \boldsymbol{\varepsilon}_x & \boldsymbol{\varepsilon}_x &\square N(0, Q) \\ \mathbf{y}^{(n)} &= C\mathbf{x}^{(n)} + \boldsymbol{\varepsilon}_y & \boldsymbol{\varepsilon}_y &\square N(0, R) \end{aligned} \quad (6.11)$$

$$A = \begin{bmatrix} a_v & 0 & 0 & 0 \\ 0 & a_p & 0 & 0 \\ 0 & 0 & a_u & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad \mathbf{b} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} \quad C = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 \end{bmatrix}$$

The matrix A has diagonal terms that are all less than one, implying that we believe that all perturbations are transitory and will eventually go away. To estimate the position of the hand and the various perturbations, we apply the Kalman algorithm. We start with a belief $\hat{\mathbf{x}}^{(n|n-1)}$ and uncertainty $P^{(n|n-1)}$ and produce a motor command $u^{(n-1)}$ (using Eq. 6.10) and predict its consequences $\hat{\mathbf{y}}^{(n)}$ (using Eq. 6.11). Our motor command produces a hand position (Eq. 6.9) and its visual and proprioceptive sensory consequences $\mathbf{y}^{(n)}$ (Eq. 6.8, where we apply the experimenter's visual perturbations). Using our prior uncertainty, we form the Kalman gain:

$$K^{(n)} = P^{(n|n-1)} C^T \left(C P^{(n|n-1)} C^T + R \right)^{-1} \quad (6.12)$$

We then update our estimate of the perturbations and hand position by combining our predictions with the observations:

$$\begin{aligned} \hat{\mathbf{x}}^{(n|n)} &= \hat{\mathbf{x}}^{(n|n-1)} + K^{(n)} \left(\mathbf{y}^{(n)} - \hat{\mathbf{y}}^{(n)} \right) \\ P^{(n|n)} &= \left(I - K^{(n)} C \right) P^{(n|n-1)} \end{aligned} \quad (6.13)$$

The final step is to form the prior uncertainty for the next movement, which is computed as:

$$P^{(n+1|n)} = A P^{(n|n)} A^T + Q \quad (6.14)$$

Fig. 6.13A shows a typical simulation result. In this simulation, all noise variances were independent and set to 1 (matrices R and Q), and a_v , a_p and a_u were set at 0.99. The visual target was set at $y_v^* = 1$. For the first 10 trials $r_v = 0$ but for the remaining trials $r_v = -2$, i.e., a visual perturbation was imposed at trial 10 and beyond. All other perturbations were maintained at zero. By end of the adaptation period, the algorithm has assigned a non-zero value to \hat{r}_p and \hat{r}_u . As a result, the estimate of current hand position \hat{h} is closer to the visual target than actual hand position h (Fig. 6.13C). That is, there is a perceptual illusion of where the hand is located.

The reason for this illusion is that the prediction error $y_v - \hat{y}_v$ can be due to a combination of estimation errors in \hat{r}_u and \hat{r}_v . It is possible that some of the prediction error is due to a faulty internal model that associates motor commands with hand position (mis-estimation of \hat{r}_u), while some of it is due to the model of the sensory system that measures hand position via visual feedback (mis-estimation of \hat{r}_v). However, because there is never a bias in the proprioceptive feedback ($y_p - \hat{y}_p$ is zero on average), any change in \hat{r}_u must be balanced with an opposite change in \hat{r}_p . It is the fact that \hat{r}_p is not zero that produces the perceptual illusion of feeling one's hand at a location other than it actually is.

The model explains that the perceptual illusion regarding where the hand is located arises because of the relative uncertainties in the sensory and motor factors that affect the estimate of hand position. To explore this idea further, in Fig. 6.13B we increased the state noise associated with r_v (the noise term for this variable in matrix Q in Eq. 6.11). Now there is little or no perceptual

illusion as the algorithm assigns nearly all of the error to \hat{r}_v . That is, if one is very uncertain about the relationship between hand position and the visual sensory feedback, then most of the prediction error is assigned to \hat{r}_v . As a result, \hat{r}_p remains near zero and there is no perceptual illusion regarding hand position.

In summary, when we generate motor commands, perturbations like force fields or visuomotor rotations produce discrepancies between the predicted and observed sensory consequences. The process of adaptation involves learning an internal model that accurately predicts these sensory consequences. An important consequence of this learning is often a sensory illusion. The generative model in Fig. 6.12 explains the adaptation process as one of estimating the potential contributions of the various factors that could have caused the prediction error. The ambiguity in these potential causes can account for both the adaptive changes in the motor commands and the sensory illusions.

6.7 History of prior actions affects patterns of learning

We started this chapter with backward blocking, an example of learning from the classical conditioning literature. In that example, it seemed that learning depended not only on the prediction errors, but also the history of the previous observations that the animal had made. The Kalman model accounted for the learning by keeping a history of these observations (i.e., the covariance in the inputs) in the uncertainty matrix. Does something like that happen in human learning as well? Do we respond to prediction errors based on the history of our prior observations?

In the backward blocking experiment, it was the history of the previous inputs (light and tone) that affected the learning in the rat. We can think of these two inputs as contextual cues that the rat was trying to associate with the unpleasant shock. We imagined that the rat's objective was to find an appropriate weighting for the two cues. After a training period in which both cues were present, prediction errors during a training period in which only one cue was present affected the state associated with the absent cue. The Kalman algorithm explained that this was because in the initial training period, the history of the two cues led to a negative covariance associated with the uncertainty of the two states. This covariance affected the credit assignment of the prediction errors.

Now suppose that these two cues are contexts in which we perform an action like moving a cursor to a target. In one context, we move our arm and the cursor is associated with the motion of our hand. In another context, we move only our wrist and the cursor is associated with the motion of our index finger. The cursor that we are moving is perturbed by a visuomotor rotation of 30° and we are going to learn to compensate for this perturbation. As John Krakauer, Pietro Mazzoni, Ali Ghazizadeh, Roshni Ravindran, and Shadmehr (Krakauer et al., 2006) observed, people show rather peculiar behaviors during these context-dependent adaptation experiments. First, for a given magnitude of perturbation, they tend to learn much slower in the arm context than in the wrist context. That is, for some reason the wrist context is easier to learn. Second, the arm context generalizes to the wrist context: after volunteers learn the 30° rotation in the arm context, they are much better than naïve when tested in the same 30° rotation in the wrist context. However, the wrist context does not generalize to the arm context: after training in the wrist task, they are about the same as naïve in the arm task. That is, the pattern of generalization is asymmetric between these two contexts. What can account for these curious facts?

It turns out that Krakauer et al. (Krakauer et al., 2006) could account for these observations when they formulated the problem in an estimation framework. Interestingly, they found out that the problem looked a lot like the backward blocking example that we considered earlier.

To start, consider that subjects were trained in two situations: arm-controlled cursor, and wrist-controlled cursor. In each case, the cursor indicated the position of the end-effector (hand or finger). We represent this position in polar coordinates and focus only on its angular component. That is, if in trial n the end-effector angle is $e^{(n)}$ and the imposed perturbation (rotation) is $r^{(n)}$, then the computer displays the cursor at $y^{(n)}$:

$$y^{(n)} = e^{(n)} + r^{(n)} \quad (6.15)$$

Now there are different contexts in which movements are performed. Let $\mathbf{c}^{(n)}$ be a binary vector that specifies this context and $\mathbf{w}^{(n)}$ be the weight vector that specifies the contribution of the context to the perturbation:

$$\begin{aligned} r^{(n)} &= \mathbf{c}^{(n)T} \mathbf{w}^{(n)} \\ y^{(n)} &= e^{(n)} + \mathbf{c}^{(n)T} \mathbf{w}^{(n)} + \varepsilon_y^{(n)} \quad \varepsilon_y^{(n)} \square N(0, \sigma^2) \end{aligned} \quad (6.16)$$

In Eq. (6.16), we have the ‘measurement’ equation, specifying that what the subject observes (cursor position) is affected by some weighted contribution of the two contexts. The problem of learning is to estimate these contributions: the state that we are trying to estimate is specified by vector $\mathbf{x} = [w_1, w_2, e]^T$. The state update equation takes the form:

$$\mathbf{x}^{(n+1)} = A\mathbf{x}^{(n)} + \mathbf{b}u^{(n)} + \boldsymbol{\varepsilon}_x \quad \boldsymbol{\varepsilon}_x \sim N(0, Q) \quad (6.17)$$

So the generative model takes the form specified by Eqs. (6.16) and (6.17). Given a target at $y_t^{(n)}$, we rely on our prior estimate of the perturbation

$$\hat{r}^{(n)} = \mathbf{c}^{(n)T} \hat{\mathbf{w}}^{(n|n-1)} \quad (6.18)$$

to produce a motor command that compensates for the estimated perturbation:

$$u^{(n)} = y_t^{(n)} - \hat{r}^{(n)} \quad (6.19)$$

To explain the peculiar experimental results, Krakauer et al. (Krakauer et al., 2006) made a critical assumption. They assumed that for their tasks, the context was not defined in terms of which muscles or joints moved, but in terms of motion of body parts in extrinsic space. In the wrist context, only the hand moved and not the upper arm. However, in the arm context, both the hand and the upper arm moved. As a result, the wrist context was defined by the context vector $\mathbf{c}^{(n)} = [0 \ 1]^T$, whereas the arm context was defined by $\mathbf{c}^{(n)} = [1 \ 1]^T$. The basic idea then is that the two contexts are not independent, but that one (the arm context) is a super set of the other (the wrist context).

If the two contexts are viewed in this way, then the arm context is identical to the situation in our earlier classical conditioning experiment in which both light and tone were present (Fig. 6.4). Recall that the presence of both cues made the uncertainty matrix P acquire negative off-diagonal elements, implying that when one weight goes up, the other weight should go down. Krakauer et al. assumed that in the course of daily activity, most of our movements are arm movements, and few are just wrist movements in which the upper arm does not move. In effect, they assumed that the initial uncertainty matrix $P^{(1|0)}$ had negative covariance in its off-diagonal elements. With this prior uncertainty, learning will indeed be slower in the arm context than the wrist context. The reason is the negative covariance.

This idea is flushed out in a simulation. Let us consider wrist training. The experimenter sets $r^{(n)} = 30$ and asks the learner to move the cursor with the wrist. The learner assumes that the

context is $\mathbf{c}^{(n)} = [0 \ 1]^T$. Fig. 6.14A shows the two components of the vector $\hat{\mathbf{w}}^{(n|n-1)}$, i.e., the weight associated with the upper arm \hat{w}_1 and the weight associated with the wrist \hat{w}_2 . With each trial, \hat{w}_2 increases toward 30° . However, despite the fact that the context is wrist only, \hat{w}_1 becomes negative, resulting in an estimate for the whole arm ($\hat{w}_1 + \hat{w}_2$) that is only slightly positive. Therefore, the model reproduces the result that wrist training will not have a significant impact on subsequent training with the arm (i.e., wrist training does not generalize to arm training). The prior uncertainty matrix with negative off-diagonal elements is directly responsible for this generalization pattern.

Next consider arm training. Fig. 6.14B shows the simulation results when we set $r^{(n)} = 30$ and train in the arm context by setting $\mathbf{c}^{(n)} = [1 \ 1]^T$. The observed errors produce changes in \hat{w}_1 and \hat{w}_2 , but because the covariance in the uncertainty matrix is negative, the sensitivity to error (Kalman gains) are much smaller in the arm context than when the task is performed in the wrist context. Consequently, the arm context is learned more slowly than the wrist context. Despite the fact that the uncertainty matrix $P^{(1|0)}$ and the initial estimate $\hat{\mathbf{w}}^{(1|0)}$ are identical in the two simulations of Fig. 6.14A and Fig. 6.14B, the errors decline about twice as slowly in the context of the arm as compared to the wrist. Furthermore, the same uncertainty matrix dictates a generalization from arm to wrist, as the Kalman gain is positive for both the upper arm and wrist. As a consequence, arm training results in the estimate for the wrist \hat{w}_2 to increase to about 10° . If we now test the system in the wrist context, it has already learned much of the perturbation and will show better performance than naïve. Therefore, the simulations also reproduce the asymmetric patterns of generalization (from arm to wrist, but not from wrist to arm).

6.8 Source of the error: my body or my tool?

Suppose that you are a tennis player and after reading the latest tennis magazine you are tempted to change your racquet. You go to the pro-shop, pickup a racquet, step into the court, and hit a few balls. The results look pretty bad. The balls are flying all over the place. What is the problem? Is the problem something to do with the internal model of your own arm, e.g., the muscles have gotten a bit weaker since you last played, or is the problem with the internal model of the racquet, e.g., is the center of mass different than what you are used to? In principle, when

one interacts with a novel tool and experiences an error in performance, the learning problem is one of credit assignment. Should one adjust the parameters of the internal model of their own arm, or the parameters of the novel tool?

The framework that we have built extends naturally to encompass this problem of credit assignment. The generative model of Fig. 6.15 shows that the parameters that we wish to estimate can be classified into parameters associated with our own body (i.e., internal parameters \mathbf{x}_{in}), and parameters associated with the tool (i.e., external parameters \mathbf{x}_{ext}). To learn optimally, the uncertainty that we have regarding each set of parameters should describe how we should divide up the error and update our estimates of the parameters.

Max Berniker and Konrad Kording (2008) used the model in Fig. 6.15 to explain experimental results regarding how people generalize their adaptation. Their idea was that if the brain adjusted the parameters associated with the body, then adaptation should generalize when the body is used in another configuration. For example, if I am using my right arm to hit serves with a new tennis racquet and I use the errors to update parameters associated with the internal model of my arm (\mathbf{x}_{in}), then after adaptation if I use my right arm and the same racquet to hit a forehand, I should be able to generalize. However, if the brain adjusted the parameters associated with the tool, then the generalization patterns should be different: adaptation should generalize to other limbs when those limbs are holding the tool. For example, if during hitting serves with my right arm I update parameters associated with the internal model of the racquet (\mathbf{x}_{ext}), then I should be able to control the racquet even when I am using my left arm.

Consider the task pictured in Fig. 6.16. In this task, a volunteer holds the handle of a robot and reaches to a target position. The robot produces forces on the hand, perturbing its motion. If we interpret these perturbations as having something to do with the object that we are holding in hand, we can write the forces in terms of the velocity of that object:

$$\mathbf{f} = B\dot{\mathbf{x}}_o \quad (6.20)$$

Now if we make a reaching movement and observe a difference between what we observed and what we predicted (in terms of motion of the object), we would use the prediction error to update our estimate of the parameter B . Alternatively, suppose that we interpret the perturbations as having something to do with the motion of our own body, i.e., torques $\boldsymbol{\tau}$ that depend on velocity of our arm:

$$\boldsymbol{\tau} = W\dot{\mathbf{q}} \quad (6.21)$$

In Eq. (6.21), the vector \mathbf{q} refers to angular configuration of the shoulder and elbow joints and $\dot{\mathbf{q}}$ refers to joint angular velocities. If we interpret the perturbations as having something to do with motion of our arm, then when we make a movement and observe a prediction error, we would be updating our estimate of the parameter W . How can we tell whether the brain is interpreting the prediction errors in terms of the external object, or internal body?

Suppose that we learn to move the object with our hand in the right-workspace (as shown in Fig. 6.16A), and then we test for generalization by making movements in the left-workspace (as shown in Fig. 6.16B). In the left-workspace we observe one of two different kinds of force fields. One of the fields is the same as in the right-workspace when we represent it in terms of the velocity of the object. This is called an ‘extrinsic field’. The other is the same as in the right-workspace when we represent it in terms of velocity of our arm. This is called an ‘intrinsic field’. If we learned by updating B in Eq. (6.20), then we should do well when our arm is in the left-workspace and we encounter an extrinsic field, but do poorly if we encounter an intrinsic field. On the other hand, if we learned by updating W in Eq. (6.21), then we should do well in the left-workspace only if we encounter an intrinsic field.

Shadmehr and Mussa-Ivaldi (Shadmehr and Mussa-Ivaldi, 1994) performed this experiment and found that in the left-workspace, performance was much better for the intrinsic field than the extrinsic field. Fig. 6.16E (top) shows reaching movements in the left-workspace in the intrinsic field, and Fig. 6.16E (bottom) shows reaching movements in the left-workspace in the extrinsic field. This suggests that subjects predominantly interpreted the prediction errors in terms of an internal model in Eq. (6.21), i.e., something about their own arm. However, the subjects still learned at least a little about the dynamics of the object (i.e., Eq. 6.20). This conclusion was inferred because of the results of a later experiment by Sarah Criscimagna-Hemminger, Opher Donchin, Michael Gazzaniga, and Shadmehr (2003). In that experiment, volunteers learned to control the robot with their right arm and then were tested with their left arm (Fig. 6.17). Performance on the extrinsic field during test of the left arm was better than naïve (naïve being defined as performance of people who had not previously learned any tasks with the robot), and also better than performance of people who were given the intrinsic field on their left arm. Together, the results of within-arm generalization (Fig. 6.16) and between arms generalization (Fig. 6.17) suggested that during adaptation, the subjects mostly updated a model similar to Eq.

(6.21), but still learned a little about the model in Eq. (6.20). That is, in learning to control the robot, the subjects appeared to be learning mostly in the intrinsic coordinate system of their own body, and a significant but smaller amount in the extrinsic coordinate system of the object.

To explain these results, Berniker and Kording (2008) modeled dynamics of the two-degree of freedom human-like arm using the following:

$$I(\mathbf{q})\ddot{\mathbf{q}} + C(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + W_m\dot{\mathbf{q}} = \boldsymbol{\tau}_m + \boldsymbol{\tau}_p + \boldsymbol{\tau}_n \quad (6.22)$$

Here, the 2x1 vector \mathbf{q} represents joint angles, $I(\mathbf{q})$ is the inertia matrix, $C(\mathbf{q}, \dot{\mathbf{q}})$ is the coriolis/centripetal matrix, W_m is joint viscosity, $\boldsymbol{\tau}_m$ is the commanded torque, $\boldsymbol{\tau}_p$ is a perturbation torque, and $\boldsymbol{\tau}_n$ is noise. Like earlier works (Shadmehr and Mussa-Ivaldi, 1994), Berniker and Kording assumed that the commanded torque was the sum of three components: a feedback torque that stabilized the trajectory, an internal model that compensated for inertial dynamics, and an internal model that estimated the perturbation from external sources:

$$\begin{aligned} \boldsymbol{\tau}_m = & \hat{I}\dot{\mathbf{q}}^*(t) + \hat{C}\dot{\mathbf{q}}^*(t) + \hat{W}_m\dot{\mathbf{q}}^*(t) - \hat{\boldsymbol{\tau}}_p \\ & + K_1(\mathbf{q}^*(t) - \mathbf{q}) + K_2(\dot{\mathbf{q}}^*(t) - \dot{\mathbf{q}}) \end{aligned} \quad (6.23)$$

K_1 and K_2 represent stiffness and viscosities that stabilize the motion around desired trajectory $\mathbf{q}^*(t)$. The term $\hat{\boldsymbol{\tau}}_p$ is an estimate of the external perturbation:

$$\hat{\boldsymbol{\tau}}_p = J^T \hat{B}\dot{\mathbf{x}} \quad (6.24)$$

The matrix J is the Jacobian that relates endpoint displacements and joint displacements. The parameters that the system tried to estimate were W_m (joint viscosity, reflecting a belief that the errors are due to internal perturbations), and B (reflecting a belief that the errors are due to external perturbations). The eight parameters, were collected in vector \mathbf{p} . The output of the system $\mathbf{y}(t)$ was hand position and velocity $[\mathbf{x} \quad \dot{\mathbf{x}}]$. The estimate of this output was $\hat{\mathbf{y}}(\hat{\mathbf{p}}, t)$.

The generative model took the form:

$$\begin{aligned} p_i(t + \Delta) &= ap_i(t) + \varepsilon_i(t) & \varepsilon_i &\square N(0, q_i) \\ \mathbf{y}(t) &= H(t)\mathbf{p} + \mathbf{w}(t) & \mathbf{w} &\square N(0, R) \end{aligned} \quad (6.25)$$

Here, p_i represents the elements of the vector \mathbf{p} (the parameters that described the velocity dependent perturbations), and the matrix $H(t)$ is a linear approximation of the non-linear dynamics in $\mathbf{y}(\mathbf{p}, t)$. $H(t)$ was approximated as:

$$H(t)_{i,j} = \frac{\hat{y}_i(\hat{\mathbf{p}}, t) - \hat{y}_i(\hat{\mathbf{p}} + \delta_j, t)}{\delta} \quad (6.26)$$

In Eq. (6.26), the term $H_{i,j}$ refers to the element (i, j) of matrix H and δ_j is a small perturbation on the j -th element of the \mathbf{p} vector. The credit assignment between parameters representing internal perturbations W_m and external perturbations B depends solely on our assumption regarding the relative variance of the noise term ε_i in Eq. (6.25). To explain the predominately intrinsic patterns of generalization, Berniker and Kording inferred that the noise in the internal parameters W_m must be larger than external parameters B . That is, the model that explained the data was one in which the brain was more uncertain about the parameters that described the viscosity of the arm as compared to the dynamics of the object in hand. In this way, the model explained both the within-arm patterns of generalization (due to changes in the internal parameters W_m), and the across-arm patterns of generalization (due to changes in the external parameters B).

This inference may seem a little odd. How could it be that when the hand is holding a weird, unfamiliar tool (the robot), movement errors are (at least partially) associated with the internal model of the arm, and not the tool? Doesn't the brain 'understand' that it is the tool that is disturbing the hand's motion? Is there another way that we can test the theory? Indeed there is. If the theoretical framework is right and indeed the brain is learning to update an internal model of its own arm, then something strange should happen: when people let go of the robot and make reaches in free air, they should have after-effects of the previous learning. That is, despite the fact that the brain 'knows' that the hand is no longer holding on to the novel tool, because it associated the errors partially to the internal model of its own arm, it should have after-effects of the previous learning. To test for this, JoAnn Kluzik, Jorn Diedrichsen, Shadmehr, and Amy Bastian (2008) trained people to reach with the robot and then tested them in two conditions (Fig. 6.18A). In one condition (called 'free'), the experimenters removed the handle from the robot (the volunteers could see this) and then asked the volunteers to simply reach in free air. In the second condition (called 'robot-null'), the handle remained on the robot but the forces were

turned off. They found that indeed, reaching movements in free-air had significant after-effects (Fig. 6.18B), and the size of these after-effects in free-air were about half the size of the robot-null condition. The results are fairly consistent with the theoretical framework, suggesting that about half of the errors were assigned to the internal model of the arm, and the rest assigned to the internal model of the tool.

Summary

In this chapter, we considered the problem of learning to make accurate predictions. We framed this problem in terms of state estimation. We imagined that in making a prediction about a variable that we can observe, there are some set of hidden states that affect that variable, making it so that sometimes our predictions are wrong. When we observe a prediction error, we try to assign credit to these hidden states, updating our estimate of these states so that we can improve our predictions.

In the Bayesian way of learning, this credit assignment problem is solved using methods that maximize the probability of a posterior distribution about our hidden variables. In the non-Bayesian methods of learning, as in LMS, this credit assignment is solved based on the contribution of each hidden state to the prediction (the greater the contribution, the greater the credit).

In some forms of biological learning, as in backwards blocking, animals appear to learn in a way that resembles the Bayesian method and not LMS. In this example, the learner appears to keep a history of the prior observations, resulting in a credit assignment that resembles the uncertainties accumulated in the Kalman gain. These uncertainties are the covariance between the hidden states.

Sensorimotor adaptation is a widely studied example of biological learning. The key requirement for adaptation is sensory prediction error: adaptation occurs when there is a discrepancy between the predicted and observed sensory consequences of motor commands. There are two widely studied paradigms: visuomotor adaptation, in which the visual consequences of a motor command is manipulated, and force field adaptation, in which both the visual and proprioceptive consequences are manipulated.

An internal model is simply a generative model that relates motor commands with sensory consequences. The hidden states are the sensory and motor perturbations. A Bayesian process of estimation of these hidden states accounts for numerous observations in how people adapt their movements to perturbations, including sensory illusions, and asymmetric patterns of generalization.

Figure Legends

Figure 6.1. A graphical model that associates motor commands, hand position, and mass of an object, with sensory observations via vision and proprioception.

Figure 6.2. A graphical interpretation of the LMS algorithm.

Figure 6.3. A comparison of the Kalman framework and LMS for the Kamin blocking experiment. Training data (top plot): x_1 and x_2 represent light and tone, and y^* represents shock. In the first 20 trials, x_1 is paired with shock. During the last 20 trials, both x_1 and x_2 are paired with shock. Kalman estimator: y and \hat{y} are observed and output quantities, \hat{w}_1 and \hat{w}_2 are weights associated with x_1 and x_2 . The Kalman gain and uncertainty are also plotted. During the last 20 trials, \hat{w}_2 does not change significantly because there are little or no prediction errors. LMS: the LMS algorithm also produces little or no changes in \hat{w}_2 .

Figure 6.4. A comparison of the Kalman framework and LMS for the backward blocking experiment. During the first 20 trials of training, both x_1 and x_2 are paired with shock y^* . During the last 20 trials, only x_1 is paired with shock. In the Kalman framework, the covariance component of the uncertainty P_{12} becomes negative in the first 20 trials. This then affects the Kalman gain k_2 in the subsequent trials, causing ‘unlearning’ in \hat{w}_2 despite the fact that during the last 20 trials x_2 is not present. LMS, on the other hand, does not alter the weight associated with x_2 during the last 20 trials.

Figure 6.5. The ‘backward blocking’ experiment in rats. The stimuli A, B, C, and X represent different kinds of sounds (A=buzzer, B=tone, C=white noise, X=click). In phase 1, A and X are simultaneously presented and paired with B. In phase 2, A is paired with B. In phase 3, B is paired with shock. The prediction of the Kalman framework is that as the animal learns that A fully predicts B, the prior association between X and B should be weakened. As a result, the animal should show reduced freezing in response to X than in control animals who had not learned A-B association (control group). In the experimental results, the bars depict mean freezing time in log seconds. The solid bars are the freezing time in the presence of stimulus X,

and the striped bars are freezing time in presence of stimulus A. Longer times suggest a stronger association. Error bars are SEM. CS = conditioned stimulus. US = foot shock. BB = the backward blocking group. CON = the control group. (From (Miller and Matute, 1996).)

Figure 6.6. Two common paradigms to investigate mechanisms of adaptation in human motor control. **A)** Visuomotor rotation paradigm. People reach with the hand above the tabletop to a target projected by a video projector and viewed through a mirror. In this way, they cannot see their hand, but they see the target in the same plane as their hand. They are provided with feedback regarding their hand position via a cursor (labeled finger on the projection screen). During the adaptation component of the experiment, the location of the projected cursor is shifted with respect to hand position, requiring the subject to alter the motor commands to the hand in order to place the cursor in the target. Occasionally, the subject is also asked to point to their right hand with their left hand. (From (van Beers et al., 2002).) **B)** Force field adaptation paradigm. People reach with their right hand while holding the handle of a robotic manipulandum. The aim is to place a cursor representing handle position inside a target. The robot produces forces that perturb the movement. Occasionally, the subject is also asked to point to their right hand with their left hand. (From (Haith et al., 2008).)

Figure 6.7. Adaptation requires sensory prediction errors. During the exposure part of the experiment, the subject's forearm was tied to a board that could move about the elbow joint. The subject either actively moved her elbow, or her elbow was passively moved by the board. In both cases, she viewed the hand through prism glasses. After a period of self-generated or passive movements, she was tested in a paradigm in which she pointed to targets without the prisms. The targets were displayed through a mirror, preventing her from seeing her hand. Only the active exposure condition resulted in after-effects. (From (Held and Freedman, 1963).)

Figure 6.8. An example of learning a visuomotor rotation. **A)** Subjects were asked to make an out-and-back motion with their hand so a cursor was moved to one of 8 targets. In the baseline condition, hand motion and cursor motion were congruent. In the adaptation condition, a 45° rotation was imposed on the motion of the cursor and the hand. In the adaptation group (top two plots), the subjects gradually learned to move their hand in a way that compensated for the rotation. In the strategy group (bottom two plots), after two movements subjects were told about the perturbation and asked to simply aim to the neighboring target. **B)** Endpoint errors in the adaptation and strategy groups. The strategy group immediately compensated for the endpoint

errors, but paradoxically, the errors gradually grew. The rate of change of errors in the strategy and adaptation groups was the same. The rapid initial improvement is due to learning in the explicit memory system, whereas the gradual learning that follows is due to an implicit system. (From (Mazzoni and Krakauer, 2006).)

Figure 6.9. Example of adaptation to a force field. **A)** Volunteers made reaching movements while holding the handle of a robotic manipulandum in order to place a cursor in a target. When the robot did not produce a force field (termed null field), hand trajectories were straight. All trajectories shown here were made without visual feedback during the movement (about 1/4th of the trials were without visual feedback). **B)** The force field produced by the robot. **C)** Hand trajectories early in the training protocol (first 250 trials). **D)** Hand trajectories late in the training protocol (trials 750-1000). (From (Shadmehr and Mussa-Ivaldi, 1994).)

Figure 6.10. Example of adaptation to a force field. **A)** Experimental setup. **B)** Size of after-effects in two groups of people: healthy controls, and amnesics (AMN), including subject HM. (From (Shadmehr et al., 1998).)

Figure 6.11. Adapting to a visuomotor perturbation induces an illusion regarding location of the trained hand. **A)** Before adaptation, when the unseen right hand is placed at a target location, the estimated position of that hand corresponds to the actual location. **B)** After adaptation to the visuomotor perturbation, the subject estimates their right hand to be somewhere between the actual position and the seen position.

Figure 6.12. A generative model for motor adaptation (Haith et al., 2008). Motor commands u , affect the position of the hand h , which is observed by the brain through two independent sensors, vision y_v and proprioception y_p . There are three potential sources of perturbation: perturbations r_u that alter the relationship between motor commands and hand position, and perturbations r_v and r_p that alter the relationship between hand position and sensory measurements.

Figure 6.13. Simulations for a reaching experiment with a visuomotor displacement. On trial 10, a perturbation r_v is imposed so that the visually observed hand position y_v is 2 cm to the left of the target, as shown in part C. Trial after trial, the motor commands adapt to place the hand h so

that y_v coincides with the target $y_v^* = 1$. **A)** With equal uncertainties regarding the source of the perturbation, the simulation estimates the perturbation as a combination of \hat{r}_p , \hat{r}_v , and \hat{r}_u . Because \hat{r}_u is not zero, estimate of hand position \hat{h} is different than actual hand position h , as shown in part C. **B)** The effect of increased state noise in the variable r_v (the noise term for this variable in matrix Q in Eq. 6.11). If one is most uncertain about the state of a visual perturbation, then the sensory errors are assigned almost entirely to this modality, resulting in little difference between \hat{h} and h (that is, no sensory illusion). **C)** Estimated hand position and sources of perturbation for part A.

Figure 6.14. Simulation results for the wrist and arm context experiments. All simulations begin with the same initial conditions. **A)** Simulation of a 30 deg rotation in the wrist context. With each trial, the estimate for the wrist increases toward 30 deg. Despite the fact that only the wrist context is present, the estimate for the upper arm becomes negative. This is because the uncertainty matrix has negative off-diagonal elements, which arises from the prior assumption that motion of the upper arm usually results in motion of the wrist. **B)** Simulation of a 30 deg rotation in the arm context. Errors produce changes in the estimates of both the upper arm and the wrist, resulting in transfer to the wrist. Despite identical initial conditions, learning with the arm is slower than learning with the wrist. (From (Krakauer et al., 2006).)

Figure 6.15. A generative model of the source of error in a tool use task. Parameters associated with the tool \mathbf{x}_{ext} and the body \mathbf{x}_m contribute to observations \mathbf{y} .

Figure 6.16. Within-limb generalization patterns after learning to reach with a novel tool. Experimental data from Shadmehr and Mussa-Ivaldi (Shadmehr and Mussa-Ivaldi, 1994). Simulation data from Berniker and Kording (2008). **A)** Volunteers reached while holding the handle of a robotic arm. The training took place in the ‘right-workspace’ (as shown). **B)** After completion of training, they were tested in the ‘left-workspace’ as shown. **C)** Simulation and experimental data are shown for motion in the left-workspace, i.e., the test of generalization. Two conditions were considered: in the intrinsic disturbance condition, the robot produced a force field that was defined in terms of motion of the subject’s joint angles, i.e., $\boldsymbol{\tau} = W\dot{\boldsymbol{\theta}}$. In another condition, the robot produced a field that was defined in terms of the motion of the subject’s hand, i.e., $\mathbf{f} = B\dot{\mathbf{x}}$. The simulated data here shows that if the learning assigned all the errors to

the model of the tool, performance in the intrinsic disturbance would be poor, and performance in the extrinsic disturbance would be excellent. **D)** The simulated data here shows that if the learning assigned all the errors to the model of the subject's own arm, performance in the intrinsic disturbance would be excellent, and performance in the extrinsic disturbance would be poor. **E)** Experimental results. **F)** Simulation results in which most but not all of the error assignment was to the model of the subject's own arm.

Figure 6.17. Between-limb generalization after learning to reach with a novel tool. Experimental data from Criscimagna-Hemminger et al. (2003). Simulation data from Berniker and Kording (2008). **A)** Volunteers reached while holding the handle of a robotic arm. The training took place with their right arm to targets shown. **B)** Testing took place with the left arm in either an intrinsic field in which the forces were defined in terms of motion of the joints, or extrinsic field in which the forces were defined in terms of motion of the hand. **C)** Subjects trained with their right hand in either an extrinsic field or an intrinsic field. The early and late training data shows trajectories of the right hand. After completion of training, they were tested with the left hand on the same field. Performance was much better for the extrinsic field. **D)** Simulation data from a system that assigned most of the error to the body (parameters of the right arm), but some of the error to the tool.

Figure 6.18. Learning to reach while holding a novel tool alters the internal model of both the subject's arm and the tool. **A)** Subjects trained to make reaching movements while holding the handle of a robot. After training, one group of subjects observed as the robot was disconnected from the handle (free), while another group was simply told that the robot motor were turned off (robot null). The reason for holding on to the handle in the free condition was to ensure that the arm/wrist was kept precisely in the same configuration as in training. **B)** Transfer index quantified the size of the errors in reaching in each of the two conditions with respect to the size of the error in catch trials, i.e., trials in which the force field was unexpectedly turned off. (From (Kluzik et al., 2008).)

Reference List

- Berniker M, Kording K (2008) Estimating the sources of motor errors for adaptation and generalization. *Nat Neurosci* 11:1454-1461.
- Criscimagna-Hemminger SE, Donchin O, Gazzaniga MS, Shadmehr R (2003) Learned dynamics of reaching movements generalize from dominant to nondominant arm. *J Neurophysiol* 89:168-176.
- Haith A, Jackson C, Miall C, Vijayakumar S (2008) Interactions between sensory and motor components of adaptation predicted by a Bayesian model. *Advances in Computational Motor Control* 7.
- Held R, Freedman SJ (1963) Plasticity in human sensorimotor control. *Science* 142:455-462.
- Kamin LJ (1968) Attention-like processes in classical conditioning. In: *Miami symposium on the prediction of behavior: aversive stimulation* (Jones MR, ed), pp 9-33. University of Miami Press.
- Kluzik J, Diedrichsen J, Shadmehr R, Bastian AJ (2008) Reach adaptation: what determines whether we learn an internal model of the tool or adapt the model of our arm? *J Neurophysiol* 100:1455-1464.
- Krakauer JW, Mazzoni P, Ghazizadeh A, Ravindran R, Shadmehr R (2006) Generalization of Motor Learning Depends on the History of Prior Action. *PLoS Biol* 4:e316.
- Mazzoni P, Krakauer JW (2006) An implicit plan overrides an explicit strategy during visuomotor adaptation. *J Neurosci* 26:3642-3645.
- Miller RR, Matute H (1996) Biological significance in forward and backward blocking: resolution of a discrepancy between animal conditioning and human causal judgment. *J Exp Psychol Gen* 125:370-386.
- Rescorla RA, Wagner AR (1972) A theory of Pavlovian conditioning: the effectiveness of reinforcement and non-reinforcement. In: *Classical Conditioning II: Current Research and Theory* (Black AH, Prokasy WF, eds), pp 64-69. New York: Appleton Century Crofts.
- Shadmehr R, Brandt J, Corkin S (1998) Time-dependent motor memory processes in amnesic subjects. *J Neurophysiol* 80:1590-1597.
- Shadmehr R, Mussa-Ivaldi FA (1994) Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14:3208-3224.
- van Beers RJ, Wolpert DM, Haggard P (2002) When feeling is more important than seeing in sensorimotor adaptation. *Curr Biol* 12:834-837.
- Widrow B, Hoff ME (1960) Adaptive switching circuits. In: *WESCON Convention Record Part IV* pp 96-104.