8 The multiple timescales of memory

Despite the modest success of our simple framework for learning, it turns out that there is something seriously wrong with the models that we have considered so far in our discussion of learning and memory. They fail to account for two fundamental properties of memory: *savings*, and *spontaneous recovery*. Let us explain these two features of memory with some examples.

8.1 Savings and spontaneous recovery of memory

A few years ago, Yoshiko Kojima, Yoshiki Iwamoto, and Kaoru Yoshida (2004) took their monkey to a darkened room and measured its eye movements (saccades) in response to display of visual targets. Usually, when a target is shown at 10^o with respect to current fixation, the animal makes a 10^o saccade to foveate the target. If we define *saccadic gain* as the ratio between the displacement of the eyes (about 10^o) and the displacement of the target (also 10^o), then in the healthy monkey (and healthy person) the saccadic gain is approximately one. However, Kojima et al. were interested in having the monkey learn to change this gain. To do this, they tested the monkey in a gain adaptation paradigm (McLaughlin, 1967). In the experiment, a target is shown at 10^o but as soon as the eyes begin moving toward it, the 10^o target is extinguished and a new target is shown at 13.5^o. Saccades are so fast (peak speed of greater than 400 deg/s) and so brief (movements complete within 60ms) that the brain cannot use any sensory feedback during the saccade to help control it. In fact, one is effectively blind during a saccade. So in response to a 10^o target, the monkey makes a 10^o saccade, observes the endpoint error, and follows this with a second saccade. As the trials continue, the brain learns to increase the saccadic gain so that in response to the 10^o target it makes a larger than 10^o saccade. This is called gain-up adaptation.

Kojima and her mentors performed a variant of a classic experiment called *extinction*. They initially trained the monkey on the gain-up task. They then followed this training with a gain-down task in which the 10^o target was extinguished and a new target was shown at 6.5^o. This is called extinction training because the objective is to counter the initial adaptation with a reverse-adaptation protocol until the gain is returned back to baseline (i.e., gain of one) They observed that in the gain-up trials, the gain of saccades increased, and in the following gain-down trials, the gain returned back to near baseline (Fig. 8.1A). The rate of adaptation in the gain-up trials was assayed by fitting a line to the initial 150 trials (this is noted by the slope of the line in Fig. 8.1A, which is about 4.0). The training in the gain-up session was followed by training in a gain-down

session until the gain had returned to one. That is, in the de-adaptation session the performance of the animal returned to baseline. Was the motor memory also at baseline, that is, had the de-adaptation session wiped out the memory that was acquired due to the prior adaptation session? To answer this question, Kojima et al. once again presented the animal with the gain-up paradigm (labeled re-adapt in Fig. 8.1A). They observed that the animal re-learned faster (the slope of the line is now 6.9). This faster re-learning is evidence of *savings*. Savings refers to the fact that despite return of behavior to a baseline condition, some component of the animal's memory still remembered the earlier gain-up training, producing a faster re-learning. Therefore, the gain-down training did not completely wash out the previous gain-up training.

In the next set of experiments (Fig. 8.1B), Kojima and her mentors once again trained the monkey in the gain-up paradigm and followed this with a gain-down paradigm until behavior had returned to baseline. At this point they simply turned the lights off in the room and left the monkey alone for 30 minutes. When they returned and re-tested the monkey in the gain-up paradigm, they noticed that the gain had suddenly increased. This is an example of *spontaneous recovery* of motor memory. With passage of time after extinction, the initial memory of adaptation seemed to spontaneously recover.

Another example of spontaneous recovery is in a form of classical conditioning in honey bees. Nicola Stollhoff, Randolf Menzel, and Dorothea Eisenhardt (Stollhoff et al., 2005) placed a bee in a test chamber and presented it with an odor (oil from a carnation flower) for 5s. Two seconds after the odor onset, they touched the antenna of the bee with a toothpick that was moistened with sucrose. Training consisted of three such odor (conditioned stimulus, CS) and sucrose (unconditioned stimulus, US) pairings. They measured whether the bee extended its proboscis between the onset of the CS and the presentation of the US (a positive score indicates that the bee performed this act). In Fig. 8.2 we see that on the first trial, none of the bees anticipated the sugar, but on the second trial about 60% of the bees put out their proboscis after they sensed the fragrance of the carnation flower. This ratio reached to 85% on the third trial. The next day they brought the bee back to the test chamber and presented it with 5 trials in which the odor was present, but the antenna was not touched (CS-only trials). During this period of extinction training, the percent of bees that responded to the odor gradually dropped so that by the 5^{th} trial, only about 15% of the bees extended their proboscis in anticipation of the sugar. Effectively, performance had returned to near baseline. Finally, they brought back some of the bees at 1hr, 2hr, etc. after completion of the CS-only trials for a test of retention. They observed that with

each passing hour, the percentage of bees that extended their proboscis increased, reaching about 50% by 24 hours. Therefore, with passage of time after extinction, the initial memory that associated the odor with the sucrose seemed to spontaneously recover.

Savings and spontaneous recovery are two fundamental properties of biological learning. Savings suggests that errors that reverse the direction of learning in a behavioral sense may not reverse the state of memory in a neural sense, as evidenced by the observation that re-learning after extinction may be faster than original learning. This implies that de-adaptation may not destroy the memory that was acquired during adaptation. Spontaneous recovery suggests that errors that reverse the direction of learning produce a new memory (call it extinction memory, e.g., CS associated with no US), that competes with the original adaptation memory (CS associated with US). With passage of time after the training episode that produced extinction, the memory acquired during the adaptation training appears to be re-expressed.

8.2 Two state model of learning

Maurice Smith, Ali Ghazizadeh, and Shadmehr (Smith et al., 2006) proposed a simple mathematical model of learning to account for spontaneous recovery and savings. Imagine that when our brain encounters a prediction error, the error affects multiple learning systems, each one with its own internal state: some that are highly sensitive to error but have poor retention, and others that are poorly responsive to error but have strong retention. The key new ideas were: 1) represent learning as a problem in state estimation in which there were multiple hidden states, and 2) associate a different timescale of forgetting to each state.

To illustrate the idea, suppose that we make a prediction $\hat{y}^{(n)}$. Say that this prediction is the sum of two internal states, one a 'fast' state $x_f^{(n)}$, and the other a 'slow' state $x_s^{(n)}$:

$$\hat{y}^{(n)} = x_f^{(n)} + x_s^{(n)} \tag{8.1}$$

Both the fast and the slow states learn from the prediction error $\tilde{y}^{(n)} = y^{(n)} - \hat{y}^{(n)}$, but from one trial to the next partially forget what they have learned. The forgetting in each state is specified by parameter a, and sensitivity to error is specified by parameter b. The important assumptions are that the fast system's learning rate b_f is larger than the slow system's learning rate b_s, while the fast system's retention rate a_f is smaller than the slow state's retention rate a_s:

$$
\begin{aligned}
x_f^{(n+1)} &= a_f x_f^{(n)} + b_f \left(y^{(n)} - \hat{y}^{(n)} \right) \\
x_s^{(n+1)} &= a_s x_s^{(n)} + b_s \left(y^{(n)} - \hat{y}^{(n)} \right) \\
0 &< a_f < a_s < 1 \\
0 &< b_s < b_f < 1
\end{aligned}
\tag{8.2}
$$

Let us use this model to consider the experiment by Kojima and colleagues. In that experiment (Fig. 8.1A), the animal adapted to a perturbation, and this period of adaptation was followed by a period of de-adaptation until performance returned to baseline. The de-adaptation was followed by re-adaptation, during which the animal exhibited savings, i.e., a faster re-learning. Fig. 8.3A shows a simulation of this scheme. A perturbation of magnitude one is imposed for a few hundred trials (adaptation period) and then the perturbation reverses sign (de-adaptation period) until performance returns to baseline. At the onset of the adaptation period the prediction errors $\tilde{y} = y - \hat{y}$ are large, causing rapid changes in the fast state x_f. As the training continues, $\tilde{y}$ becomes smaller and the forgetting term a_f in the fast state becomes relatively more important. As a result, near the end of the adaptation period most of the prediction $\hat{y}$ is due to the slow state x_s. Now in the de-adaptation period, the prediction errors are once again very large, causing a rapid change in the fast state. By the end of the de-adaptation period, the prediction $\hat{y}$ is at zero, i.e., the system appears to be at baseline. Yet, this is only because the fast and slow states are in competition, effectively canceling each other. Now the de-adaptation period is followed by another period of re-adaptation. In this re-learning period, the system adapts faster than before, exhibiting savings (compare dashed line with the solid line in the lower sub-plot of Fig. 8.3A). The reason for the faster re-learning is that the prediction errors during the de-adaptation period hardly affect the slow state. Savings is due to resistance of the slow state to change, enhancing the ability of the system to return to its previously adapted state in the re-adaptation period.

The basic idea of this two-state model is that during the brief de-adaptation period, the system's predictions can change rapidly and return to baseline, but not through complete unlearning of what it has learned in the past. Rather, a fast state of memory is put in competition with a slow state. The system's output $\hat{y}$ exhibits washout only because the two states happen to balance each other.

While there are some similarities between the simulations of extinction (Fig. 8.3A) and the behavioral data (Fig. 8.1A), there are also some critical differences. In both the simulations and

the behavioral data, the de-adaptation period continues until behavior has reached baseline. In the behavioral data, however, the length of de-adaptation training is fairly similar to the length of adaptation training (about 700 trials of adaptation followed by about 500 trials of de-adaptation), whereas in the simulations the adaptation period is much longer than the de-adaptation period (about 400 trials of adaptation followed by 20 trials of de-adaptation). As Eric Zarahn and colleagues pointed out (Zarahn et al., 2008), the two-state model would show complete unlearning if the de-adaptation period was as long as the adaptation period. This is inconsistent with the behavioral data, and something that we will return to later in this chapter.

Now let us consider the spontaneous recovery phenomenon. In Kojima's experiment (Fig. 8.1B), the monkey was left alone in the dark for 30 minutes. When he was once again asked to make a saccade, he made hyper-metric saccades, as he had initially learned to make in the gain-up paradigm. What happened during this period of darkness? The monkey surely moved his eyes, but had little or no sensory feedback to tell him if the saccades that he made were accurate. In effect, darkness reduced his ability to measure sensory consequences of his motor commands. Let us imagine that this is equivalent to making movements, but having no errors to learn from, that is, the 'darkness' period is equivalent to setting the error term to zero:

$$y^{(n)} - \hat{y}^{(n)} = 0 \tag{8.3}$$

(Having zero error is not the same as having no sensory feedback, but this is a useful initial step. We will return to the question of no sensory feedback once we re-cast our problem using a generative model.) So in effect, the 'dark' period is like an error-clamp period in which movements are made, but there is zero error on each trial.

In the re-adaptation period that comes after the error-clamp period, the model exhibits a jump in performance with respect to end of the de-adaptation period (Fig. 8.3B). Therefore, the model accounts for the Kojima data in Fig. 8.1B by suggesting that during the 'dark' period, the animal made saccades but had no visual errors to learn from. During this 'dark' period, the fast state rapidly declined, the slow state slowly declined, and the sum of the two produce the spontaneous recovery. The states decline toward zero because there were no errors to learn from. In this way, changes in the memory states were only due to the forgetting terms a_f and a_s. The rapid decay of the fast state but the gradual decay of the slow state made $\hat{y}$ rise up from zero, resulting in spontaneous recovery of the previously learned behavior.

The strongest prediction of the model is regarding how the behavior should evolve during the error-clamp period (Fig. 8.3B): there should be a very rapid rise of output followed by a gradual decline. A number of recent experiments have directly tested this prediction by measuring motor output during error-clamp trials. In a saccade task, Vince Ethier, David Zee, and Shadmehr (Ethier et al., 2008) examined saccadic gain changes in people (Fig. 8.4A). As before, adaptation was followed by de-adaptation until gain had returned to near baseline. The training was then followed by a block of error-clamp trials. Each error-clamp trial began with a visual target. As soon as the saccade was initiated, the target was turned off. After completion of the saccade the target re-appeared at the current location of fixation, effectively minimizing endpoint errors. As the model had predicted, saccade gain after the de-adaptation period exhibited spontaneous recovery: a rapid change from baseline followed by a gradual change (Fig. 8.4A, the final period of error-clamp trials).

The same pattern appeared in adaptation of reaching movements (Smith et al., 2006;Criscimagna-Hemminger and Shadmehr, 2008). For example, Sarah Criscimagna-Hemminger and Shadmehr (2008) asked volunteers to hold the handle of a robotic arm and reach to a target. Forces that were perpendicular to the direction of motion perturbed the reach. With training, the brain learned to predict and compensate for these forces. To measure the motor output that represented this prediction, in some trials the robot produced a virtual channel consisting of two stiff walls that surrounded a straight line between the start and the target positions (Scheidt et al., 2000). In these error-clamp trials, it was possible to measure both the motor output (i.e., what the brain predicted about the perturbation, reflected in the forces that the subject produced against channel walls), and minimize any errors in the movement. In the baseline trials before imposition of the perturbations, people simply moved their hand straight to the target, producing little or no forces against the channel walls (Fig. 8.4B, 'null' condition). Two groups of subjects were then considered. One group experienced only a field that perturbed the hand to one direction (Field *A*). Another group was first exposed to *A*, and then to Field *B*. *B* had a perturbation that was equal but opposite in direction to *A*. Effectively, in the A+B group the experiment replicated the adapt/de-adapt paradigm that we saw earlier in the saccades. In Group A, a block of error-clamp trials immediately followed training in *A*. In Group A+B, a block of error-clamp trials followed training in *B*. While Group A showed a motor output that gradually declined from the adapted state, Group A+B demonstrated spontaneous recovery, i.e., motor output rose from baseline toward the previously adapted levels. However, something interesting happened: spontaneous recovery of motor output was so strong in the A+B group that motor output converged to the A

group. This could only happen if de-adaptation training produced *little or no unlearning* of previous adaptation. That is, it appeared that de-adaptation training instantiated an independent fast memory, and this fast memory rapidly faded in the subsequent error-clamp trials.

The spontaneous recovery patterns in Fig. 8.4B suggest that a long period of training is sufficient to produce a slow motor memory that is resistant to subsequent performance errors (e.g., in Field B). When the field changes and induces large performance errors, the errors do not erase the memory of *A*, but rather install a competing memory of *B*. The general implication of this work is that once a memory has been acquired, all further learning may be instantiation of competing memories.

8.3 Timescales of memory as a consequence of adapting to a changing body

While the two-state model of Eq. (8.2) is certainly more capable of accounting for the experimental data than a single-state model, we need to ask why the nervous system should learn in this way. That is, what might have given rise to this kind of learning? What is this kind of learning good for?

The main new idea in this section is that animals exhibit patterns of learning and forgetting, i.e., the multiple timescales that we alluded to earlier, because these timescales describe perturbations that naturally affect their own body. There are many natural perturbations that can affect the body, some with rapid timescales and others with slower timescales. Perhaps when we perform a task and observe an error, our brain tries to estimate the source of the error with the same generative model that represents natural perturbation to the body.

To illustrate the point, let us suppose that you are designing the control mechanism of an autonomous robot. You recognize that the input-output properties of the motors in various limbs will change with use and with passage of time. For example, with repeated use over a short period, a motor may experience heating and change its response transiently until it cools. On the other hand, with repeated use over a long period of time the motors may decline in function due to wear and tear. Both of these conditions will produce movement errors, i.e., the inputs to the motor will not produce the expected sensory feedback. The errors will require your controller to adapt and send updated commands to the motors to produce the desired actions. However, your controller should interpret these errors differently: errors that have a fast time scale should result

in rapid adaptive changes, but should be quickly forgotten (because a heating motor will cool simply by passage of time). Errors that persist for extended periods of time should result in slow adaptive changes that should be remembered (because a worn out motor will not fix itself no matter how long we wait). Now if a mischievous child were to attach a small weight to one of the arms, the sudden drop in performance of the arm might be interpreted as a motor heating up, and therefore produce a rapid adaptation followed by rapid forgotten if the arm is held still for a while (during which time we would expect the motor to cool, removing the potential source of performance error). In this way, we might view the problem of learning as one of credit assignment to various timescales of memory. To solve the credit assignment problem, we need a generative model that describes the likely sources of perturbations and their timescales.

Konrad Kording, Joshua Tenenbaum, and Shadmehr (Kording et al., 2007) proposed that our nervous system faces similar problems in controlling our body. Properties of our muscles change due to a variety of disturbances, such as fatigue, disease, exercise, and development. The states of these disturbances affect the motor gain, i.e., the ratio of movement magnitude relative to the input signal. States of disturbances unfold over a wide range of timescales. Therefore, when the nervous system observes an error in performance, it faces a credit assignment problem: given that there are many possible perturbation timescales that could have caused the error, which is the most likely? It seems reasonable that the solution to the credit assignment problem should dictate the temporal properties of the resulting memory. That is, adaptation in response to things that are likely to be permanent (slow states) should be remembered, while adaptation in response to things that appear transient (fast states) should be forgotten.

Now this kind of thinking does not explain all forms of perturbations. For example, a sudden injury may last a long time to heal. In this case, a perturbation that comes on rapidly does not go away rapidly. Similarly, certain low frequency events like a car accident or child birth leave memories that are never forgotten. In these scenarios, there are neuro-modulators that are released due to the heightened emotional state of the learner, making memories 'stick'. But for the motor system, the most likely source of a natural perturbation is fatigue. For example, you decide to walk up a few flights of stairs. Step after step, the leg muscles fatigue, altering their force production capabilities. To maintain your stride, the brain must keep an estimate of the state of the muscles so that it can deliver the right amount of activation. This state estimate should change rapidly as the leg muscles fatigue. However, when you arrive at your floor and get to the lab and sit down, the fatigue state of the muscles naturally fades away. Your estimate of

the state of the muscles should also rapidly fade away, without requiring errors to drive changes in your estimate. You do not want a brain that assumes that the muscles are still fatigued after sitting for 10 minutes just because during that period you have not taken a step and therefore had not had the opportunity to observe a prediction error. Perhaps a generative model that describes dynamics of these natural, internally driven sources of perturbation in our motor system acts as a generative model for external perturbations as well. That is, the way animals adapt to artificially induced perturbations on their movements may be based on a generative model of perturbations that affect their own body.

To make our generative model, we will make two assumptions: first, we will assume that there are many sources of perturbations that can affect our performance, with some perturbation states changing rapidly, while others changing slowly. Second, we will assume that states that change rapidly are affected by greater noise and variability than those that change slowly.

Such a generative model might look something like Fig. 8.5. We represent the state of each potential source of perturbation by x_i. Some of these perturbation states decay quickly with time, while others decay slowly. This is reflected in the elements of the diagonal state transition matrix A. When we make an observation y, we see the combined effects of all the states:

$$\begin{aligned} \mathbf{x}^{(n+1)} &= A\mathbf{x}^{(n)} + \boldsymbol{\varepsilon} \\ y^{(n)} &= \mathbf{c}^T\mathbf{x}^{(n)} + \eta \end{aligned} \tag{8.4}$$

To note the different timescales of change for each state x_i, we can manipulate the elements of the matrix A to reflect our assumption that perturbation state x_1 decays more rapidly than perturbation state x_2, etc.

$$0 < A_{1,1} < A_{2,2} < \cdots < A_{m,m} < 1$$

(The subscripts in the above expression refer to the elements of the matrix.) For simplicity, we will assume that A is diagonal. Furthermore, it seems rational that we should be most uncertain about states that change most quickly, an idea that we can incorporate by making the fast states in Eq. (8.4) have the most amount of noise.

$$\begin{aligned} &\boldsymbol{\varepsilon} \sim N(\mathbf{0}, Q) \\ &Q = \begin{bmatrix} \sigma_1^2 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \sigma_m^2 \end{bmatrix} \qquad \sigma_1^2 > \cdots > \sigma_m^2 \end{aligned} \tag{8.5}$$

Finally, while we cannot observe these states individually, what we can observe is their combined effects; so we set vector $\mathbf{c}$ in Eq. (8.4) to have elements that are all one:

$$\mathbf{c} = \begin{bmatrix} 1 & 1 & \cdots & 1 \end{bmatrix}^T$$

In summary, when we observe a disturbance in our performance, our job is to estimate which state is the most likely candidate. Is the disturbance due to a fast changing state? In that case, we should learn a lot from our performance error (because we are least certain about the fast states), but also quickly forget it (because we assume that these states quickly decay, e.g., $A_{1,1} \square 1$). Is the error due to a slow changing state? In that case, we should change these states by a relatively small amount (because we are most certain about the slow states), but remember what we learned (because we assume that these states decay slowly).

With this generative model, Kording et al. used the Kalman algorithm to make predictions about learning and memory. Their specific model had 30 hidden states, with timescales that were distributed exponentially between 2 and 10^5 trials. The simulation result for the spontaneous recovery experiment is shown in Fig. 8.6A. To simulate the darkness period, rather than setting errors to zero, the sensory noise η in Eq. (8.4) was set to near infinite variance. As a consequence, during the darkness period the Kalman gain became zero, and the performance was dictated by the terms in matrix A.

Because the timescales ranged from very short to very long, the model was also able to replicate results from a long-term adaptation experiment. An example of a long-term adaptation experiment is shown in Fig. 8.6B. In this experiment, Rick Robinson, Robijanto Soetedjo, and Christopher Noto (2006) trained a monkey on a gain-down saccade adaptation protocol for 22 days. During each training session the gain decreased, but some of this learning was forgotten during the time between sessions (Fig. 8.6B). On the 23rd day and beyond, the targets were no longer moved during the saccades, encouraging the animal to return the saccadic gain back to one (called the washout period). Interestingly, between the sessions in the washout period the forgetting now reversed directions. That is, whereas during adaptation forgetting was toward a gain of one, now during washout forgetting was toward a gain of 0.5.

The simulations with the generative model of Eq. (8.4) produced a similar result (Fig. 8.6C). The simulations explained that during initial days of gain-down adaptation, most of the changes were in the fast states (they became negative, where zero indicates a baseline condition). These states

exhibited forgetting between the sessions, and so behavior reverted back toward gain of one between training days. By the end of training on day 22, only the slowest of the slow states had changed (became negative) and there was little or no contribution to behavior from the faster states (which were all near zero). When the washout trials begin on day 23, all the changes were once again in the fast states (which now became positive). Between the training days, the fast states once again forgot back to zero, but because the slow states had not returned to baseline, between day forgetting in the washout session appeared to be in the opposite direction of adaptation.

The reason why the simulations needed 30 states rather than two is because of the very long length of the training protocol. Whereas from day 3 to day 22 there was little change in performance (as shown by the generally flat lines in Fig. 6B), in the model there were significant changes as the performance shifted from a reliance on some of the faster states in the early days of training to the slowest states as the weeks wore on. The effect of this very long-term training could only be recorded by states that changed on a similar long-term timescale. As a result, very long-term training produced memories that decayed extremely slowly.

In summary, when adaptation is followed by de-adaptation, during the following period of error-clamp trials there is a spontaneous recovery of the adapted state. This pattern is consistent with a state estimation process in which our brain assumes that the observations that it makes are driven by a generative model in which multiple hidden states contribute; some states change slowly with time, while others change rapidly.

8.4 Passive and active meta-states of memory

When you learn to use a tool (for example, the robot in Fig. 6.10A), the motor memory can be re-activated when you return a day later and are asked to use it again. For example, if during the experience with the robot you felt a force that pushed your hand to the right on day 1, upon your return on day 2 on your very first movement you will produce a motor output that pushes the robot handle to the left, in apparent anticipation of the force field you experienced yesterday. However, your experience with the robot on day 1 was rather brief, perhaps an hour long. For the rest of the day you continued to make reaching movements without the tool. Presumably, the motor memory that you acquired regarding the novel tool during those brief training trials was ‘de-activated’ for the rest of the day, and then ‘re-activated’ when on day 2 you needed to move

the handle of the robot again. Therefore, it seems rational that motor memory should exist in two meta-states: an active meta-state in which the motor memory is being used to perform an action, and a passive meta-state in which the memory is not engaged.

Until now, we have not considered the dissociation between these two hypothetical meta-states of memory. When we talked about forgetting, e.g., the period termed 'retention' in Fig. 7.5A, we referred to an experiment in which motor output declined during error-clamp trials. This forgetting referred to decay in an activated memory, as the subjects in that experiment were using the tool as we measured their motor output. However, forgetting in the usual sense refers to changes in memory as a function of time, i.e., changes to a passive memory. For example, if you learn to hit a tennis serve as a result of an hour of instruction, you may not be able to reproduce the same performance the next day when you pick up a racquet. This reduction in performance is presumably a forgetting that affects the memory when it is in a passive meta-state.

We can represent the active and passive properties of memory as meta-states: when the memory is activated, the rules that describe the state transitions may be different than when it is not activated. It is easy to see that forgetting in the active memory, i.e., state changes as defined in error-clamp trials, cannot be the same as time-dependent forgetting when the memory is passive. For example, consider the decline in motor output for Group 1 in Fig. 7.5B. These subjects learned to move the robot to a single direction for 60 trials, and were then placed in a block of 60 error-clamp trials. Their motor output declined to near baseline during these 60 post-adaptation trials. The amount of time it takes to perform these trials is about 2-3 minutes. Experiments show that if we waited 3 months after a comparable period of initial training, subjects would show savings (Shadmehr and Brashers-Krug, 1997). That is, the decline in the active meta-state of the memory is not the same as the decline that takes place during the same period of time in a passive meta-state of memory.

Returning to the state estimation framework, this implies that the state transition matrix for the active meta-state A_a must differ from the state transition matrix for the passive meta-state A_p.

$$\begin{aligned} &\text{active meta-state} \begin{cases} \mathbf{x}^{(t+\Delta)} = A_a \mathbf{x}^{(t)} + \boldsymbol{\varepsilon}_a \\ y^{(t)} = \mathbf{c}^T \mathbf{x}^{(t)} + \eta_a \end{cases} \\ &\text{passive meta-state} \begin{cases} \mathbf{x}^{(t+\Delta)} = A_p \mathbf{x}^{(t)} + \boldsymbol{\varepsilon}_p \end{cases} \end{aligned} \tag{8.6}$$

In the active meta-state, we can make observations and use prediction errors (differences between $y^{(t)}$ and $\hat{y}^{(t)}$) to update our estimate of the perturbation state. In the passive meta-state, the states continue to change but we cannot make observations. (We made the superscripts in Eq. 8.6 in terms of time, rather than trial, so that the matrices A_a and A_p could be compared. In effect, we will assume that a 'trial' takes a constant amount of time, and use this calibration to represent changes in the states of the passive meta-state.)

Sarah Hemminger and Shadmehr (Criscimagna-Hemminger and Shadmehr, 2008) performed an experiment that estimated A_a and A_p. They assumed a two-state model of motor memory (fast and slow states) and trained people in a reaching task with a robot that perturbed the limb with force fields. They had two groups of subjects: Group A that trained for hundreds of trials on Field *A*, and Group A+B that had the same training in *A* but then briefly was trained in the opposite field *B* until motor output returned to baseline. We can assume that for Group A, most of the memory by the end of the training is of the slow variety, whereas for Group A+B, there is an equal amount of slow and fast memories (canceling each other, reflected in the fact that by end of training in *B*, motor output was at baseline). They then waited a variable period of time and measured performance in error-clamp trials. They observed that at 0 and 2 minutes after training, motor output of the A+B group rose from zero and converged onto the motor output of the A group (Fig. 8.7A). That is, there was a spontaneous recovery, indicating resistance of the slow state to change. This implies that at 0 and 2min post acquisition of *B*, the 'fast' memory (of *B*) rapidly decayed within 30 trials. However, at 10 minutes and beyond, motor output of the A+B group no longer started at zero, but at a significantly higher value. Importantly, the performance of A and A+B groups no longer converged within 30 trials. That is, within a few minutes after acquisition, the fast memory of *B* no longer decayed away within 30 trials. It had somehow gained resistance to the error-clamp trials.

The resulting generative model that fitted the data well is shown in Fig. 8.8. The state transition matrix in the active meta-state had values of $A_a = [0.984, 0; 0, 0.855]$. The state transition matrix in the passive meta-state had values of $A_p = [0.9999, 0.0043; 0, 0.983]$, where Δ=6sec for a typical inter-trial interval. [The matrix elements may look similar, but keep in mind that 6 hours has 3600 'trials', so these numbers are raised to very large powers and a small difference matters.] The passage of time in the active state was far more destructive to the state of memory

than the same passage of time in the passive meta-state. That is, there was more forgetting in the active meta-state. Furthermore, note that A_p had a non-zero off-diagonal term. This indicates that in the passive meta-state, some of the fast state became a slow state with passage of time. This transition of the fast into a slow state during the time when the memory was not activated is an example of a transformation that makes it resistant to change, a phenomenon that is called consolidation.

In summary, these results suggest that error-driven learning not only engaged processes that adapt with multiple timescales, but that once practice ends and the memory is placed in a passive state, passage of time transforms some of the fast states into slower states that resist change when the task again resumes and the memory is re-activated.

8.5 Protection of motor memories

When a period of adaptation training is followed by de-adaptation, does the brain protect the memory that was acquired during adaptation, or do the prediction errors during de-adaptation destroy the previously acquired memory, producing effective unlearning? This simple question has been surprisingly difficult to answer. While saccade experiments in monkeys such as those shown in Fig. 8.1, and classical conditioning experiments such as those shown in Fig. 8.2, would suggest that adaptation memories are protected during de-adaptation, there is also significant data that does not agree with this view.

For example, Graham Caithness, Rieko Osu, Paul Bays, Henry Chase, Jessica Klassen, Mitsuo Kawato, Daniel Wolpert, and Randy Flanagan (2004) trained volunteers in a reaching task that for some groups involved visuomotor rotation and for other groups involved force fields that perturbed motion of the arm. They adapted the volunteers to a perturbation for *n* trials, and then reversed the perturbation and continued the de-adaptation training for an equal number of *n* trials. When they retested the volunteers in the original perturbation, they found that performance was no different than naïve, i.e., there was no evidence of savings. They wrote: "when people successively encounter opposing transformations (A then B) of the same type (e.g., visuomotor rotations or force fields), memories related to A are reactivated and then modified while adapting to B." They concluded that memory of adaptation was unprotected during de-adaptation, resulting in catastrophic destruction.

A year later, John Krakauer and colleagues (2005) repeated some of the experiment performed by Caithness and colleagues and found different results. They trained volunteers in a reaching task that perturbed the relationship between hand position and cursor position via a visuomotor rotation. They observed that when adaptation to a perturbation was followed by adaptation to the reverse-perturbation, in some conditions the re-learning of the initial perturbation was faster (these conditions involved either a large number of trials in which subjects adapted to the perturbation, or passage of time between initial adaptation and the de-adaptation episode). Because the experiment uncovered evidence for savings, they wrote: "the persistence of interference across long intervals … is not definitive proof of erasure of initial learning." Based on their results the authors concluded that the inability to show savings in previous experiments was not due to destruction of the memory, but an inability of the nervous system to express that memory.

The published data demonstrates that in experiments in which adaptation is followed by de-adaptation, some researchers have found evidence of savings while others have not. Indeed, there is currently no consensus on the question of whether memories that are acquired during adaptation are protected during de-adaptation. One possibility is that savings, as defined by a faster re-learning, is a weak assay of memory. Is there a different way to assay motor memory?

A recent experiment approached the issue from a different perspective. The idea was to use the error-clamp trials as a way to probe the contents of motor memory. Recall that in the error-clamp trials, movements are constrained so that no matter the motor output, there is no performance error and the subject is rewarded for her behavior. We saw earlier that when a long period of adaptation is followed by a brief period of de-adaptation, motor output in the following error-clamp trials is a mixture of both the memory acquired during adaptation and the memory acquired during de-adaptation (Fig. 8.4). For example, when people train for ~400 trials to reach in force field *A*, and then are trained for 20 trials in the opposite force field *B*, the motor output in the following error-clamp trials starts at zero, but then rapidly rises toward *A* (Fig. 8.4B). This suggests that during the error-clamp trials the brain expresses both the memory of *A* and the memory of *B* (*B* rapidly decays, *A* gradually decays, resulting in the phenomena of spontaneous recovery). Sarah Pekny, Sarah Criscimagna-Hemminger, and Shadmehr trained people so that a long period of exposure to *B* preceded an equally long period of exposure to *A*. They then briefly re-exposed the subjects to *B*. This paradigm was termed *BAb*. They argued that if *A* destroys *B*, then the brief exposure to *B* after *A* should not produce savings, and more importantly, the pattern

of spontaneous recovery that is expressed in the error-clamp trials show be the same in *BAb* and *Ab*. If, on the other hand, the memory of *B* is protected during adaptation to *A*, then the motor output in the error-clamp trials should show evidence for *B* in the *BAb* group as compared to the *Ab* group.

The results of this experiment are shown in Fig. 8.9. When re-exposed to *B*, the *BAb* group re-learned slightly faster than the *Ab* group (middle column, Fig. 8.9B). More importantly, the motor output of the *BAb* group in the error-clamp trials is clearly biased toward field *B* as compared to the *Ab* group. Because of the large difference between *Ab* and *BAb* groups in the error-clamp trials, it seems likely that in the *BAb* group the memory of *B* was not destroyed by the subsequent training in *A*. Similar results were obtained when training in *B* was followed by training in a null field (*Bnb* group vs. *nb* group, left column of Fig. 8.9). The *Bnb* group displayed a motor output in the error-clamp trials that was biased toward field *B*. Finally, when training in *B* was done gradually (right column, Fig. 8.9A), the resulting memory for *B* was even stronger, producing a large bias in the error-clamp trials toward *B* (right column, Fig. 8.9C). Taken together, the results of this experiment suggest that a long period of adaptation produces a memory that is at least partially protected from the large errors that could produce unlearning during de-adaptation. That is, the performance changes that take place during de-adaptation appear to be due to learning of a new memory.

But if this is the case, how does the brain decide when to spawn a new memory? One possibility is that sudden, large prediction errors are interpreted by the brain as a change in context, resulting in engagement of new resources to learn from prediction errors, rather than un-learn the previously acquired memory. Jeong-Yoon Lee and Nicolas Schweighofer (2009) proposed a model of adaptation in which motor output was supported by a single fast adaptive state (represented by scalar variable x_f), and multiple slow adaptive states (represented by a vector variable $\mathbf{x}_s$). The appropriate slow state was selected based on context $\mathbf{c}$:

$$\begin{aligned} \hat{y}^{(n)} &= x_f^{(n)} + \mathbf{c}^{(n)T}\mathbf{x}_s^{(n)} \\ x_f^{(n+1)} &= a_f x_f^{(n)} + b_f\left(y^{(n)} - \hat{y}^{(n)}\right) \\ \mathbf{x}_s^{(n+1)} &= A_s\mathbf{x}_s^{(n)} + b_s\left(y^{(n)} - \hat{y}^{(n)}\right)\mathbf{c}^{(n)} \end{aligned} \tag{8.7}$$

This model predicts that whereas the fast state is completely erasable (because it is not context selectable), the slow states can be afforded protection if the context changes. The lack of context dependence of fast states accounts for the fact that if training in *A* precedes training in *B*,

performance in *B* is worse than naïve, a phenomenon called *anterograde interference*. Yet, the model can account for savings because context allows protection of the slow states. Lee and Schweighofer (2009) did not describe a method with which one might select the correct context. The mechanisms of spawning new memories and contextual selection of old ones remains poorly understood.

8.6 Multiple timescales of memory in the cerebellum

The adaptation tasks that we have considered here are known to rely on the cerebellum. For example, damage to the cerebellum makes people and other animals generally impaired in their ability to adapt to perturbations in saccadic eye movements (Straube et al., 2001;Golla et al., 2008) or reaching movements (Smith and Shadmehr, 2004;Rabe et al., 2009;Criscimagna-Hemminger et al., 2010). Are the multiple timescales of memory, savings, and spontaneous recovery reflected in the neural activity of the cerebellum? Unfortunately, relatively little is known regarding the neural changes in the cerebellum during saccade and reach adaptation. However, classical conditioning also exhibits savings and spontaneous recovery, and for one form of classical conditioning there is substantial data on the role that the cerebellum plays in encoding the memory. Here we will briefly review this data, as it sheds light on the neural basis of the fast and slow processes that potentially support learning.

Mike Mauk and his students have been systematically identifying the neural basis of a form of classical conditioning in rabbits. In a typical experiment, a rabbit is presented with a tone for 550ms (the conditioned stimulus, CS). The rabbit does not respond to this tone. Next, they apply a small current to a region around one eye of the rabbit for 50ms at the end of the tone (the unconditioned stimulus, US). In response to the mild shock, the rabbit closes its eyes. With repeated trials, the rabbit learns to associate the CS with the US and closes its eyes at around 500ms, right before the US is given. This change in behavior is mediated via changes in two places in the cerebellum: the cerebellar cortex, and the interpositus nucleus (one of the deep cerebellar nuclei). In the cerebellar cortex, the tone (CS) is conveyed via activity that reaches the Purkinje cells via mossy fibers (Fig. 8.10). The US is conveyed via activity that reaches the Purkinje cells via climbing fibers from the inferior olive. The coincidence of these two signals produces reduction in the strength of the mossy fiber-Purkinje cell synapse, inducing plasticity in the cerebellar cortex. In the cerebellar nuclei, the changes in the Purkinje cell discharge combine with input from the mossy fibers to produce plasticity in the interpositus nucleus.

Tatsuya Ohyama and Mike Mauk (2001) performed an experiment which suggests that the cerebellar cortical plasticity takes place fairly rapidly after start of training, whereas interpositus nucleus plasticity takes place slowly and follows the changes in the cerebellar cortex. In their experiment, in phase 1 the rabbits were trained briefly with a long-duration CS. The training was too brief to produce a learned eye-blink response. However, Ohyama and Mauk posited that synaptic plasticity took place in the cerebellar cortex (reduction in the granule cell-Purkinje cell synaptic strength, Fig. 8.11B), but because there was still no plasticity at the nucleus, this change in the cerebellar cortex could not be expressed in terms of change in behavior. In phase 2, the rabbits were trained for a many trials with a short-duration CS. This training produced robust eye-blink response to the CS, which was posited to include plasticity in both the cerebellar cortex and the nucleus (reduction in the granule cell-Purkinje cells synaptic strength, increase in the mossy fiber-nucleus cell strength, Fig. 8.11C). In the test period, the animals were presented with a very long duration CS. The model predicted that the animals would produce a double blink (Fig. 8.11D), and this was indeed confirmed in the behavioral data. The double blink response suggests that the brief training (to the long duration CS, Fig. 8.11B) was sufficient to produce fast adaptation in the cerebellar cortex, which could subsequently be expressed when adaptation took place in the cerebellar nucleus. Based on this data, it would appear that the neural basis of fast and slow adaptive processes may be distinct, with a fast-like process residing in the cerebellar cortex and a slow-like process residing in the cerebellar nucleus.

To investigate the neural basis of savings after extinction, Javier Medina, Keith Garcia, and Mike Mauk (Medina et al., 2001) trained their rabbits for 5 days, ensuring that they produced a robust eyeblink response to the CS (A5, Fig. 8.12). At this point, they injected a drug (picrotoxin, PTX) into the deep nucleus that was an antagonist to GABA, the neurotransmitter that Purkinje cells release to communicate with the cells in the deep nuclei. This effectively disconnected the cerebellar cortex from the cerebellar nucleus. They noticed that the rabbit still produced a response to the tone. However, this response was at the start of the tone and not near its end (PTX on A5, Fig. 8.12). Therefore, there was a memory in the deep cerebellar nucleus that had formed due to training, but this memory did not have the temporal specificity that was present when the whole network was intact. They now trained the rabbits for 45 days of CS alone. That is, they performed extinction training. After 15 days of extinction training, the rabbits no longer responded to the CS (E15, Fig. 8.12). However, when they injected the GABA antagonist into the cerebellar nucleus, the animals once again started responding to the tone. That is, despite the

extinction of behavior, a component of the original CS-US memory was present in the deep nucleus. This component appeared to be unaffected by the 15 days of extinction training. After 45 days of extinction training, still a small amount of response to CS remained in the cerebellar nucleus (E45, Fig. 8.12). After these 45 days of extinction, Medina and colleagues retrained the animals with the CS-US pair and observed that re-learning was faster than naïve (REACQ, Fig. 8.12). The improvement in re-learning with respect to naïve was correlated with the magnitude of the response that was present in the cerebellar nucleus, i.e., the greater the memory in the nucleus, the greater the behavioral measure of savings.

Together, these results demonstrate that the multiple timescales of memory during classical conditioning are partly due to different rates of learning in the cerebellar cortex and nuclei. The faster rate of learning appears in the cerebellar cortex. When adaptation is followed by extinction, the memory in the cerebellar nucleus shows very strong resistance to change, and appears to be responsible for the savings that is exhibited during relearning.

Summary

When people and other animals adapt their movements to a perturbation, removal or reversal of that perturbation may return behavior to baseline, but this does not imply that the change in the direction of errors erases the acquired memory. The evidence for this comes from savings (faster re-learning), and spontaneous recovery of the acquired memory in the time period after reversal of the perturbation. One way to account for these behaviors is to assume that the learner responds to the prediction errors with adaptive processes that have multiple timescales. Some of the processes have fast timescales, learning rapidly from error but also exhibiting rapid forgetting. Some of the processes have slow timescales, learning gradually from error but exhibiting strong retention and resistance to forgetting. Sudden large errors in performance may act as a cue that context of the task has changed, facilitating spawning of new memories and protection of old ones. Much of the ability to adapt our movements to perturbations is due to the cerebellum. The learning in this structure exhibits a faster rate of plasticity in the cerebellar cortex than the cerebellar nuclei. During extinction, the slowly acquired memories in the nuclei exhibit strong resistance to change so that despite return of behavior to baseline, the memory in the cerebellar nucleus is maintained and appears to form the basis for savings during re-learning.

Figure 8.1. Savings and spontaneous recovery of motor memory. **A**) A monkey made saccades to a target that at 10^o. In the adaptation period, the 10^o target was displaced to 13.5^o as the saccade initiated, resulting in an endpoint error. Slowly, the monkey learned to respond to the 10^o by making a larger than 10^o saccade. That is, the gain of the saccade increased. Saccadic gain is defined as the ratio between the displacement of the eyes and the displacement of the target. In the de-adapt period, the 10^o target was displaced to 6.5^o. The de-adapt period continued until the saccade gain returned back to baseline. However, upon re-adaptation, the rate of re-learning was faster than learning during the initial adaptation period. This phenomenon is called savings. The slope of the first 150 trials is noted by the line. **B**) Adapt and de-adapt periods were followed by 30 minutes of darkness in which the animal made saccades but had no stimuli or other visual information. In the re-adapt period, saccade gain appeared to spontaneously change from the end of the de-adapt period. This phenomenon is called spontaneous recovery. (From (Kojima et al., 2004))

Figure 8.2. Acquisition, extinction, and spontaneous recovery of memory in a honey bee. In the three acquisition trials, bees were provided with an odor (CS, conditioned stimulus), and sugar (US, unconditioned stimulus). The y-axis records the percentage of animals that extended their proboscis in response to the odor (percent conditioned response). The next day, some of the bees were provided with five extinction trials (CS-only). At time periods thereafter (retention test), different bees were tested for CS-only. The black bars denote the performance of animals that did not receive extinction training. In the hours that followed extinction training, animals exhibited spontaneous recovery of the initial learning. (From (Stollhoff et al., 2005))

Figure 8.3. Simulation of a two-state system in learning from prediction error (Eqs. 8.1 and 8.2). **A**) Savings experiment. The disturbance $y^{(n)}$ induces adaptation, which initially causes a large change in the fast state. With more trials, the slow state rises while the fast state falls (as the errors decline and forgetting becomes a stronger factor than prediction error). In the de-adaptation period, the training continues until $\hat{y}$ returns to zero. The fast state rapidly changes, setting up a competition between the fast and slow states. In re-adaptation, the slow state is already in an adapted position, which makes re-learning faster than original learning (dashed line vs. solid line, bottom figure). **B**) Spontaneous recovery experiment. Adaptation is followed by de-adaptation until performance returns to baseline. During the error-clamp period, prediction errors are set to zero on every trial. During this period the fast state rapidly decays while the slow

state slowly decays. The sum of the two produces recovery of the output $\hat{y}$ to near a previously adapted state despite the fact that there are no errors to learn from.

Figure 8.4. Spontaneous recovery following adapt/de-adapt training. **A**) Saccade experiment. Top figure shows the state of the perturbation, and the bottom figure shows motor output (eye displacement during each saccade). Target appeared at 15^o. Adaptation began with a gain-down block. Saccade amplitudes are displayed in the lower plot. Note the rapid initial adaptation that was followed by gradual adaptation. Also note the forgetting that took place at set breaks (gray vertical lines). De-adaptation (gain-up block) returned motor output back to near baseline. During the subsequent error-clamp trials, motor output returned toward the previously adapted state. (From (Ethier et al., 2008)) **B**) Reach experiment. People held the handle of a robotic arm and reached to a target. The robot perturbed the arm perpendicular to its direction of motion with a force field. With practice, people learned to expect a force perturbation and compensated. The gray vertical lines indicate error-clamp trials, trials in which the robot produced a stiff channel that guided the hand straight to the target, removing error and allowing the experimenters to measure how hard subject pushed against the wall. This was a proxy for the motor output (plotted in the bottom figure). One group adapted to field *A* and was then presented with a block of error-clamp trials (Group A). Another group adapted to field *A* and then was presented with *B* until performance returned to baseline, at which point a block of error-clamp trials was presented (Group A+B). Motor output of A+B group spontaneous recovered from baseline and converged to output of group A, i.e., the effect of field *B* completely washout within 30 trials after de-adaptation. (From (Criscimagna-Hemminger and Shadmehr, 2008))

Figure 8.5. A graphical representation of the generative model of Eq. (8.4).

Figure 8.6. Simulations of the model in Eq. (8.4). **A**) Simulations results for the spontaneous recovery experiment (Fig. 8.1B). **B**) Experimental data for a multi-day experiment. A monkey trained in a gain-down saccade adaptation paradigm for 22 days (data for 5 days are shown in the adaptation sub-plot). Between days the animal was kept in the dark. Note that the forgetting is toward gain of one. On day 23 and beyond, the learning was washout (target remained stationary, encouraging the gain to return to one). Note that during washout, the forgetting is toward a gain of 0.5. (From (Robinson et al., 2006)) **C**) Simulation results for the model of Eq. (8.4). (From (Kording et al., 2007))

Figure 8.7. Effect of time passage on the spontaneous recovery of motor memory. The task was a reaching movement while holding the handle of a robotic arm, as in Fig. 8.4B. The data describes force produced by subjects during reaching movements as a percentage of the ideal force perturbation. All data are measurements in error-clamp trials. Group A trained on field A for 384 trials, and was then presented with 30 error-clamp trials at a particular time after completion of training. Group A+B trained in field A and then immediately for a brief period in field B (field B had forces that were opposite that of field A). **A**) The 0min group: immediately after completion of training in A or A+B, performance was assayed in error-clamp trials. Group A+B exhibited spontaneous recovery. In the 2min group, two minutes of delay was added to the end of completion of A or A+B training before performance was assayed in error-clamp trials. Similar to the 0min interval, the 2min interval also showed spontaneous recovery that started near zero and rose to meet the A only group. **B**) As the interval between completion of training and error-clamp trials increased, the pattern of spontaneous recovery changed. The changed pattern suggested that whereas at 0 and 2min post acquisition the memory of B was fragile and declined to zero within 30 trials, with further passage of time some memory of B decayed away but the rest gained resistance and no longer decayed to zero within the same number of trials. (From (Criscimagna-Hemminger and Shadmehr, 2008))

Figure 8.8. A model of the effect of motor memory under passive (the task is not being performed) and active (the task is being performed) processes, as in Eq. (8.6). In the passive meta-state, some of the fast state transitions into a slow state.

Figure 8.9. De-adaptation does not erase adaptation memories, as evidence by patterns of spontaneous recovery. The task was reaching in a force field. **A**) Perturbation patterns. Left column: null followed by a brief period of adaptation to field B (termed nb), vs. training that starts with B, then washout, then brief re-adaptation to B (termed Bnb). Middle column: Ab vs. BAb. Right column: Ab vs. gradual BAb. **B**) Performance during the 20 trials of B (re-learning). There is a trend toward faster re-learning for groups that had previously been exposed to B. However, this improved performance is marginal. **C**) Following 20 trials of B, all movements were in error-clamp trials. In all cases, the prior training in B is reflected in the more negative forces that subjects produce in the error-clamp trials. During error-clamp trials, motor output is a mixture of A and B, with the strongest evidence for B exhibited in the gradual condition.

Figure 8.10. Circuitry of the cerebellum and its relationship to eyelid classical conditioning. During the task, mossy fibers convey information from the CS. This input affects grandule cells, and ultimately the Purkinje cells, which inhibit the cells in the deep cerebellar nucleus. Climbing fibers convey information about the US. Learning is driven by changes in two locations: the synapse between granule cell and Purkinje cell, and the synapse between the mossy fiber and deep cerebellar nucleus cell. (From (Medina et al., 2001)).

Figure 8.11. A schematic representation of the two timescales of learning in the cerebellum. **A**) In the naïve animal, Purkinje cells activity is not affected by the stimulus (CS). **B**) In phase 1, a few trials of long-duration CS is presented with US. There is no behavioral response. Yet, the Purkinje cells show a change in their response (reduced activity near offset of CS). **C**) In phase 2, many trials of short-duration CS is presented with US. There is robust behavioral response. This response is due to plasticity in the granule cell-Purkinje cells synapse and mossy fiber-nucleus cell synapse. **D**) In the CS probe trials, a long duration CS is presented without US. The behavioral response is a double blink. (From (Ohyama and Mauk, 2001)).

Figure 8.12. Eyelid traces throughout acquisition and extinction for one animal. A0: before start of training. A5: 5^{th} day of CS-US training. E15: 15^{th} day of extinction training (CS only). The location of the black arrow indicates the time during the session when picrotoxin (PTX) was infused into the cerebellar nucleus. The black portion of each trace represents the time when the CS was present. (From (Medina et al., 2001)).

Reference List

Caithness G, Osu R, Bays P, Chase H, Klassen J, Kawato M, Wolpert DM, Flanagan JR (2004) Failure to consolidate the consolidation theory of learning for sensorimotor adaptation tasks. J Neurosci 24:8662-8671.

Criscimagna-Hemminger SE, Bastian AJ, Shadmehr R (2010) Size of Error Affects Cerebellar Contributions to Motor Learning. J Neurophysiol 103:2275-2284.

Criscimagna-Hemminger SE, Shadmehr R (2008) Consolidation patterns of human motor memory. J Neurosci 28:9610-9618.

Ethier V, Zee DS, Shadmehr R (2008) Spontaneous recovery of motor memory during saccade adaptation. J Neurophysiol 99:2577-2583.

Golla H, Tziridis K, Haarmeier T, Catz N, Barash S, Thier P (2008) Reduced saccadic resilience and impaired saccadic adaptation due to cerebellar disease. Eur J Neurosci 27:132-144.

Kojima Y, Iwamoto Y, Yoshida K (2004) Memory of learning facilitates saccadic adaptation in the monkey. J Neurosci 24:7531-7539.

Kording KP, Tenenbaum JB, Shadmehr R (2007) The dynamics of memory as a consequence of optimal adaptation to a changing body. Nat Neurosci 10:779-786.

Krakauer JW, Ghez C, Ghilardi MF (2005) Adaptation to visuomotor transformations: consolidation, interference, and forgetting. J Neurosci 25:473-478.

Lee JY, Schweighofer N (2009) Dual adaptation supports a parallel architecture of motor memory. J Neurosci 29:10396-10404.

McLaughlin S (1967) Parametric adjustment in saccadic eye movements. Percept Psychophys 2:359-362.

Medina JF, Garcia KS, Mauk MD (2001) A mechanism for savings in the cerebellum. J Neurosci 21.

Ohyama T, Mauk MD (2001) Latent acquisition of timed responses in cerebellar cortex. J Neurosci 21:682-690.

Rabe K, Livne O, Gizewski ER, Aurich V, Beck A, Timmann D, Donchin O (2009) Adaptation to visuomotor rotation and force field perturbation is correlated to different brain areas in patients with cerebellar degeneration. J Neurophysiol 101:1961-1971.

Robinson FR, Soetedjo R, Noto C (2006) Distinct short-term and long-term adaptation to reduce saccade size in monkey. J Neurophysiol 96:1030-1041.

Scheidt RA, Reinkensmeyer DJ, Conditt MA, Rymer WZ, Mussa-Ivaldi FA (2000) Persistence of motor adaptation during constrained, multi-joint, arm movements. J Neurophysiol 84:853-862.

Shadmehr R, Brashers-Krug T (1997) Functional stages in the formation of human long-term motor memory. J Neurosci 17:409-419.

Smith M, Shadmehr R (2004) Modulation of the rate of error-dependent learning by the statistical properties of the task. Advances in Computational Motor Control 3.

Smith MA, Ghazizadeh A, Shadmehr R (2006) Interacting adaptive processes with different timescales underlie short-term motor learning. PLoS Biol 4:e179.

Stollhoff N, Menzel R, Eisenhardt D (2005) Spontaneous recovery from extinction depends on the reconsolidation of the acquisition memory in an appetitive learning paradigm in the honeybee (Apis mellifera). J Neurosci 25:4485-4492.

Straube A, Deubel H, Ditterich J, Eggert T (2001) Cerebellar lesions impair rapid saccade amplitude adaptation. Neurology 57:2105-2108.

Zarahn E, Weston GD, Liang J, Mazzoni P, Krakauer JW (2008) Explaining savings for visuomotor adaptation: linear time-invariant state-space models are not sufficient. J Neurophysiol 100:2537-2548.