What the cerebellum computes

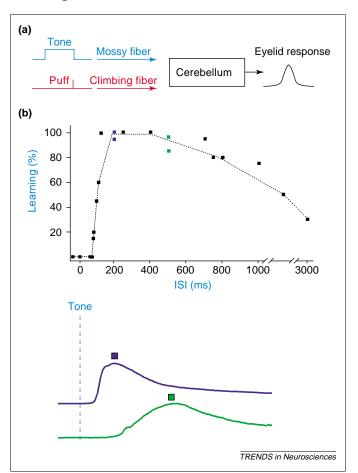## Tatsuya Ohyama, William L. Nores, Matthew Murphy and Michael D. Mauk

Department of Neurobiology and Anatomy, and W.M. Keck Center for the Neurobiology of Learning and Memory, University of Texas Medical School, Houston 6431 Fannin, Houston, TX 77030, USA

The brain is an organ that processes information. Brain systems such as the cerebellum receive inputs from other systems and generate outputs according to their internal rules of information processing. Thus, our understanding of the cerebellum is ultimately best expressed in terms of the information processing it accomplishes and how cerebellar neurons and synapses produce this processing. We review evidence that indicates how Pavlovian eyelid conditioning reveals cerebellar processing to be an example of feedforward control. Eyelid conditioning demonstrates a capacity for learning in the cerebellum that is error driven, associative and temporally specific – as is required for feedforward control. This computation-centered view is consistent with a variety of proposed functions of the cerebellum, including sensory–motor integration, motor coordination, motor learning and timing. Moreover, feedforward processing could be the common link between motor and non-motor functions of the cerebellum.

Many ideas have been proposed regarding the function of the cerebellum, including roles in sensory-motor integration, motor coordination, motor learning, timing, solving the inverse kinematic problem, sensory acquisition, etc. Such statements of function are lacking, in that they are generally descriptions of what is missing when the cerebellum is damaged. Our main thesis is that Pavlovian eyelid conditioning makes it possible to understand the cerebellum in terms of what it computes. We review evidence to suggest that eyelid conditioning provides a relatively direct approximation of the input–output transformations of the cerebellum. The emerging view portrays cerebellar computation as a straightforward example of feedforward processing of inputs in order to improve movement accuracy. Specifically, we show how eyelid conditioning reveals a temporal specificity to cerebellar learning that helps ensure that the faulty motor commands are altered to improve an errant movement. This hypothesis is consistent with ideas that the cerebellum is involved in sensory-motor integration, motor coordination, motor learning and timing.

From a model system to modeling a system

Although initially viewed as a model of associative learning, it is increasingly clear that Pavlovian eyelid conditioning is an especially useful means of studying cerebellar computation and its underlying mechanisms. This transition has occurred with the steady accumulation of evidence to indicate that eyelid conditioning engages the cerebellum in a straightforward and conceptually useful manner. In a typical experiment, the paired presentation of a cue stimulus such as a tone and a reinforcing stimulus such as a puff of air directed at the eye leads to the acquisition of a conditioned response: the eyelid closes in response to the tone (Fig. 1a). Recording, lesion and stimulation studies have shown that

Fig. 1. Eyelid conditioning and the cerebellum. (a) In rabbit eyelid conditioning, a tone serves as the conditioned stimulus (CS) and a puff of air or mild electrical stimulation serves as the unconditioned stimulus (US). The cerebellum is crucial for both learning and performance of conditioned eyelid responses (CRs). Data points are taken from available literature. Learning is measured as CR probability in asymptotically trained animals. (b) Learning depends strongly on the interstimulus interval (ISI): the time between the onset of the CS and that of the US (top). CR probability in asymptotically trained rabbits is maximal for an ISI between 200 ms and 500 ms, decreasing sharply for shorter intervals and more gradually for longer ones. In addition, the CR is timed to the ISI, such that rabbits trained at ISIs of 200 ms (blue) and 500 ms (green) will produce responses that peak near 200 and 500 ms after tone onset (vertical broken line), respectively (bottom). Traces are from individual rabbits trained at the ISI indicated by the data points with corresponding colors in the top panel. The markers above each trace in the bottom panel indicate when the puff is presented relative to the tone on training trials.

Corresponding authors: Tatsuya Ohyama (tatsuya.ohyama@uth.tmc.edu), Michael D. Mauk (m.mauk@uth.tmc.edu).
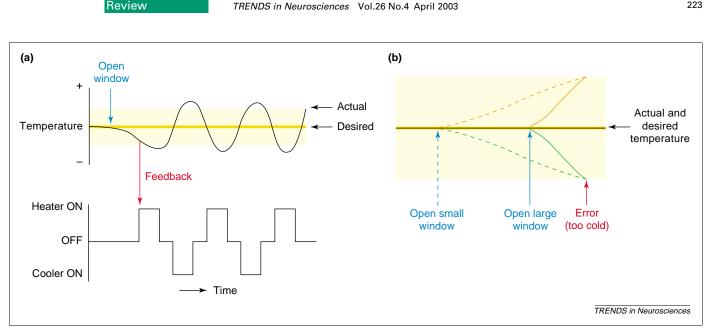
Fig. 2. Feedback and feedforward control. (a) An example of a feedback-regulated thermostat. Though simple and accurate, feedback regulation tends to be slow and attempts to speed it up cause oscillations. In this example, an attempted correction by a powerful heating system (in response to a decrease in temperature caused by opening a window) overshoots the mark, forcing another attempted adjustment by a powerful cooling system that also overshoots the mark, and so on. The upper and lower boundaries of the yellow area indicate the thresholds for detecting temperature change. (b) An example of a feedforward thermostat, which counteracts (orange lines) the changes in temperature predicted by the opening of windows (green lines). Though much quicker than feedback control, feedforward control is more complicated, requiring the ability to detect error-predicting stimuli (windows opening) and to delay anticipatory responses appropriately. Here, the opening of a large window (unbroken green line) predicts a faster decrease in temperature than does the opening of a small window (broken green line). The anticipatory response (increasing room temperature; broken or unbroken orange lines) must therefore be appropriately delayed in response to the predictive stimulus (opening a large or small window). This type of calibration requires associative learning that is temporally specific.

tone and puff presentations are conveyed to the cerebellum, respectively, via activation of its two major types of afferents – the mossy fibers and climbing fibers [1,2]. Furthermore, output of the cerebellum via one of its deep nuclei is responsible for the expression of the learned responses [3,4] (Fig. 1a). This relatively direct mapping of stimuli to inputs and of outputs to responses makes the behavioral properties of eyelid conditioning a good first approximation of the input–output transformations of the cerebellum [5,6]. For example, the various ways that tones and puffs can promote or fail to promote learning suggest the rules that govern how patterns of mossy fiber and climbing fiber inputs engage cerebellar learning [7,8].

Two ways in which eyelid conditioning is influenced by the delay between the onsets of the tone and the puff (the interstimulus interval or ISI) are particularly telling (Fig. 1b). Learning occurs for only a narrow range of ISIs. In rabbits, acquisition of conditioned responses fails to occur with ISIs < 100 ms, is robust with ISIs between 200 and 500 ms (tone before puff), and becomes increasingly poor as the ISI increases beyond the optimal range [9]. The timing of conditioned responses is also influenced by the ISI. Latencies to onset and rise times are gauged so that conditioned responses peak at about the time the puff is normally presented [10]. Behavioral studies indicate that this adaptive timing is learned and is not attributable to the strength of learning where response latencies are short for strong learning and longer for weak learning [11].

Our goal is to relate these time-dependent properties of eyelid conditioning to concrete statements about the computation performed by the cerebellum. In order to do so, it will be useful to begin by reviewing the concept of feedforward control.

Feedforward control and learning

Making accurate movements requires sensory input. In principle, this input can be applied to motor commands in two general ways: feedback and feedforward. This distinction can be illustrated using control of room temperature. A standard thermostat controls room temperature via feedback. It activates the heating–cooling system when the thermometer (sensory input) signals that room temperature differs from the target setting. Thus, feedback about current performance is used to generate output commands. Although feedback control does not require learning, only an accurate comparison between actual and target temperature, it is inherently slow. A thermostat can only react, after processing delays, to existing errors in temperature, and cannot predict such errors. Moreover, attempting to speed up a feedback-based system causes oscillations [12] (Fig. 2a).

By contrast, feedforward control combines sensory inputs and previous experience to predict the appropriate output (Fig. 2b). This allows quick reactions because errors are anticipated rather than detected. A hypothetical feedforward thermostat would use input from a variety of sensors to predict heater–cooler output required to maintain desired temperature. For example, such a system would react to the opening of a door on a cold day (the sensory signal) with activation of the heating system before the room temperature dropped. Although feedforward control can be fast, accuracy requires learning from previous experience. This learning involves the use of simple directional error signals (e.g. room too cold) to modify the subsequent responses of the system to similar inputs. If the room becomes too cold after a window is opened, the internal settings of the thermostat would be
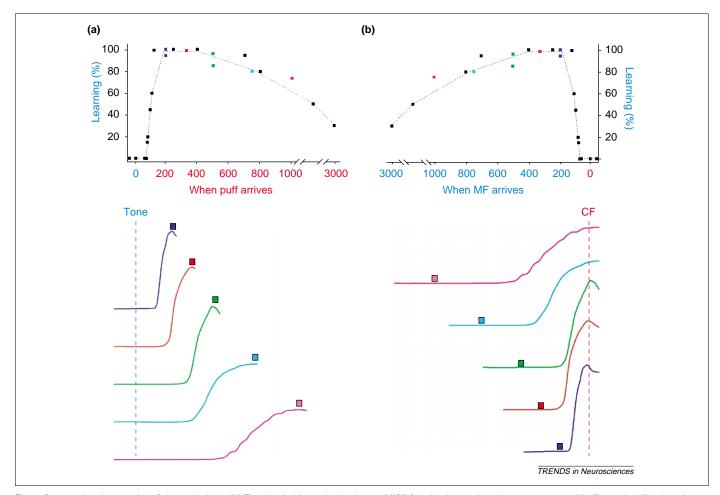
Fig. 3. Computational properties of the cerebellum. (a) The classical interstimulus interval (ISI) function in learning theory, as presented in Fig. 1, describes learning as dependent on the interval by which the onset of the puff trails after that of the tone (top). In addition, the ISI determines the timing of the conditioned response (CR), with longer ISIs producing increasingly delayed CRs (bottom). (b) The ISI function and ISI-specific timing, from the perspective of cerebellar computation. In this view, it is more useful to think of learning as depending on when mossy fiber input (MF) occurs before climbing fiber input (CF) (top). Only a small window of intervals within a range appropriate for fast movements supports learning. In addition, cerebellar output is adjusted so that it occurs just before and peaks near the puff-evoked error (bottom). The accurately timed eyelid closure minimizes the movement error. The markers above each trace in the bottom panel indicate when the onset of MF activity occurs in relation to the CF input on training trials.

altered so that the heater is activated longer the next time that window is opened.

If the cerebellum operates as a feedforward control system, as has been previously suggested [12,13], then it should display this type of associative learning in which output is adjusted for sensory inputs that reliably predict errors. This is precisely the capacity for learning in the cerebellum that is revealed by eyelid conditioning. When specific mossy fibers inputs (tone) reliably predict an error signaled by a climbing fiber input (puff), cerebellar output for subsequent tones is adjusted.

Temporally specific learning
The demands on this learning are made more complex by the likelihood that not all sensory signals will predict errors with the same time delay. The time-dependent properties of eyelid conditioning show that cerebellar learning is well suited to solve this problem.

In the thermostat analogy, some sensors might be activated when a small window is opened, and others activated by opening a larger window. On a cold day, these signals would predict the same error (room too cold), but with different delays (Fig. 2b). Similarly, the appropriate

responses to learn are differently timed: turn on the heater rapidly after receiving an input connected to the large window, but slowly activate the heater in response to a signal connected to the small window (Fig. 2b). This example illustrates that learning for a feedforward system should adjust outputs so that they are time-locked to the occurrence of the errors, not to the sensory inputs.

These basic ideas can be applied to produce more specific and concrete expectations about the temporal properties that cerebellar learning should display. When an error signal arrives via a climbing fiber input, the faulty cerebellar output that should be altered by learning has already occurred ~100–200 ms ago (Fig. 3a, top). This reflects the time required for cerebellar output to be translated into movement, as well as for the movement error to occur, be detected and conveyed to the cerebellum. Thus, no learning should occur for mossy fiber inputs that arrive <100 ms before the climbing fiber error, because they occur after the offending cerebellar output. Although learning should occur for mossy fiber inputs that arrive over a range of delays, as illustrated above, this range cannot be infinite. Otherwise, each climbing fiber input would induce learning for all prior mossy fiber inputs and

there would be no specificity to the learning. These constraints define the period over which mossy fiber inputs could be useful predictors of the climbing fiber error signal (Fig. 3b, top). Learning should occur for mossy fiber inputs that arrive within a limited temporal window, defined by a relatively sharp boundary around 100 ms and falling off gradually at the other boundary. Figure 3 illustrates that this is the temporal window of learning displayed by eyelid conditioning (compare Fig. 3a with Fig. 3b, top).

Although cerebellar learning can be engaged by mossy fiber signals that precede climbing fiber error signals over a range of time delays, the appropriate change in cerebellar output should occur just before the error signal, when the faulty motor commands were issued. This is the same temporal specificity displayed by conditioned eyelid responses, which are timed to occur just before the error signal (Fig. 3b, bottom). Thus, eyelid conditioning reveals that when a mossy fiber input consistently predicts a climbing fiber error signal with a specific time delay, cerebellar learning will adjust output at the appropriate time so that the errant motor commands are altered. In summary, eyelid conditioning reveals that the temporal properties of cerebellar learning are appropriate for predictive feedforward control.

Mechanisms of temporally specific learning

Analysis of cerebellum-dependent behaviors such as eyelid conditioning and adaptation of the vestibulo–ocular reflex (VOR) has revealed a great deal about how cerebellar neurons and synapses accomplish this computation [8,14]. Evidence from such studies indicates that plasticity in both the cerebellar cortex and cerebellar nuclei is involved [15,16]. For learning that increases cerebellar output, climbing fiber inputs drive the induction of plasticity in the cortex. Recent evidence indicates that spontaneous climbing fiber activity is required to maintain equilibrium in the cerebellar cortex, and that inhibition of climbing fibers below the spontaneous level is required for learning that decreases cerebellar output [17,18]. This is consistent with the behavior of climbing fibers during increases and decreases in the gain of the VOR [19].

In the cerebellar nuclei, recent studies using recording and electron microscope analyses demonstrate that eyelid conditioning is associated with an increase in excitatory synapses [20,21]. The hypothesis that induction of deep nucleus plasticity is driven by input from the cerebellar cortex [6] is supported by theoretical analyses of VOR adaptation [22] and eyelid conditioning [23], and by data suggesting that eyelid conditioning is blocked by permanent or reversible lesions of the cerebellar cortex [24,25]. A non-synaptic form of plasticity induced *in vitro* also requires specific patterns of Purkinje cell activity [26].

Evidence indicates that the computations associated with the temporal specificity of cerebellar learning occur in the cerebellar cortex [27]. Lesions of the cerebellar cortex abolish the learned timing of conditioned eyelid responses [28]. After permanent lesions or reversible inactivation of the cerebellar cortex, conditioned eyelid responses are no longer properly timed and instead default to a relatively fixed, short latency [15,24,25,27–30].

Currently, there is little empirical evidence to indicate how the cerebellar cortex learns to delay responses until the appropriate time. Instead, ideas about the particulars of this computation have been pursued mostly through computer simulations. A large-scale simulation of the cerebellum (~12 000 simulated neurons), which incorporates climbing fiber-driven plasticity in the cortex and Purkinje cell-driven plasticity in the nucleus, successfully captures the basic behavioral properties of eyelid conditioning [5]. These include acquisition and extinction of conditioned responses, the appropriate timing of responses and the dependence of learning on the ISI.

Timing requires temporal coding, such as a population of cells where different subsets of cells are active at different times during or after an input. In the simulation, this is accomplished with granule cells, the most abundant neuron type in the cerebellum. Interactions between the populations of simulated granule and Golgi cells make different granule cells active at different times during an input that simulates a tone [31]. The learned responses of the simulation are timed to occur just before the error signal (puff), because the climbing fiber input specifically induces plasticity in synapses active at that time. Extinction processes also operate to suppress responding early in the simulated tone and sharpen response timing. An empirical test produced results consistent with this prediction [29].

Extinction also plays an important role in the simulation for dependence of learning on the ISI [5]. As in eyelid conditioning, learning in the simulation declines as the time interval between mossy fiber and climbing fiber input increases (Fig. 3, top). This decline occurs because, on each training trial, acquisition occurs for synapses that are active towards the end of the simulated tone, and extinction occurs for synapses that are active at other times. As the ISI increases, extinction increasingly dominates acquisition until, with a long enough ISI, there is no net gain each trial.

Such ideas derived from simulations are best thought of as food for thought. In some cases, mechanistic predictions have been made from successful simulations, with positive results. With continued interplay between simulations and experiments, the precise details of how cerebellar neurons implement temporally specific learning for feedforward control will hopefully become clear.

Concluding remarks

Pavlovian eyelid conditioning has emerged as a valuable tool for studying cerebellar computation. Eyelid conditioning engages the cerebellum in a relatively direct way: mossy and climbing fibers, respectively, convey information about the tone and the puff to the cerebellum, and cerebellar output drives the expression of the conditioned response. Because of this straightforward mapping, the behavioral properties of eyelid conditioning are a reasonable first approximation of the input–output transformations accomplished by the cerebellum. The computation that is revealed is remarkably consistent with the requirements for predictive feedforward control of movements: learning is error driven, associative and temporally specific. Through experiment and simulation,

the systems-level mechanisms that underlie this computation are becoming increasingly clear.

The disadvantage of eyelid conditioning is its artificiality, which should raise concerns that this computation might not apply to all cerebellum-dependent processes. There are three factors, however, that support its generality. First, the regularity of connectivity seen across the cerebellar cortex suggests that the computation performed in one region applies to all regions. Second, the computation revealed by eyelid conditioning is independent of precisely what mossy fiber inputs encode. Whether activated by a tone, the position of a limb or some cognitive process, the cerebellum will process mossy fiber inputs as arbitrary signals that do or do not reliably predict events. For climbing fiber inputs, the only constraint is that what is encoded should convey the need for cerebellar output to be greater at that time. Thus, the increasing evidence that lateral regions of the cerebellum in humans are engaged during cognitive rather than motor tasks is not inconsistent with the computation revealed by eyelid conditioning [32].

The third and strongest piece of evidence comes from similarities between cerebellar contributions to eyelid conditioning and to control of eye movements. Strong parallels in patterns of input and demands on outputs between eyelid conditioning and VOR adaptation have been noted previously, particularly in learned timing [14,33]. Studies on saccades also illustrate the contribution of temporally specific learning. Multiple cerebellar outputs control saccades [34], and their inactivation makes saccades grossly inaccurate [35]. We have recently suggested how the computation revealed by eyelid conditioning explains the adaptation of saccades [36]. Finally, smooth-pursuit eye movements require the cerebellum and display predictive properties consistent with feedforward control [37,38]. Moreover, a simulated network based on the cerebellum has been shown to produce accurate and predictive smooth pursuit [39].

Temporally specific learning thus encapsulates many of the functions that have previously been proposed for the cerebellum, including sensory–motor integration, motor learning, timing and motor coordination. This role offers an explanation for the dysmetric and uncoordinated movements seen in individuals with cerebellar pathology. A detailed understanding of this cerebellar computation could represent the foundation for further insights into how the cerebellum contributes to non-motor functions [32].

References

1 Mauk, M.D. *et al.* (1986) Classical conditioning using stimulation of the inferior olive as the unconditioned stimulus. *Proc. Natl. Acad. Sci. U. S. A.* 83, 5349–5353

2 Steinmetz, J.E. *et al.* (1989) Classical conditioning in rabbits using pontine nucleus stimulation as a conditioned stimulus and inferior olive stimulation as an unconditioned stimulus. *Synapse* 3, 225–233

3 McCormick, D.A. and Thompson, R.F. (1984) Cerebellum: essential involvement in the classically conditioned eyelid response. *Science* 223, 296–299

4 Berthier, N.E. and Moore, J.W. (1990) Activity of deep cerebellar nuclear cells during classical conditioning of nictitating membrane extension in rabbits. *Exp. Brain Res.* 83, 44–54

5 Medina, J.F. and Mauk, M.D. (2000) Computer simulation of cerebellar information processing. *Nat. Neurosci.* 3, 1205–1211

6 Mauk, M.D. and Donegan, N.H. (1997) A model of Pavlovian eyelid conditioning based on the synaptic organization of the cerebellum. *Learn. Mem.* 4, 130–158

7 Gormezano, I. *et al.* (1983) *Progress in Psychobiology and Physiological Psychology* (Sprague, J.M., Epstein, A.N., *et al.* eds), pp. 197–275, Academic Press

8 Medina, J.F. *et al.* (2000) Mechanisms of cerebellar learning suggested by eyelid conditioning. *Curr. Opin. Neurobiol.* 10, 717–724

9 Schneiderman, N. and Gormezano, I. (1964) Conditioning of the nictitating membrane of the rabbit as a function of CS–US interval. *J. Comp. Physiol. Psychol.* 57, 188–195

10 Millenson, J.R. *et al.* (1977) Classical conditioning of the rabbit's nictitating membrane response under fixed and mixed CS–US intervals. *Learn. Motiv.* 8, 351–366

11 Mauk, M.D. and Ruiz, B.P. (1992) Learning-dependent timing of Pavlovian eyelid responses: differential conditioning using multiple interstimulus intervals. *Behav. Neurosci.* 106, 666–681

12 Kawato, M. and Gomi, H. (1992) The cerebellum and VOR/OKR learning models. *Trends Neurosci.* 15, 445–453

13 Ito, M. (1970) Neurophysiological aspects of the cerebellar motor control system. *Int. J. Neurol.* 2, 162–176

14 Raymond, J.L. *et al.* (1996) The cerebellum: a neuronal learning machine? *Science* 272, 1126–1131

15 Perrett, S.P. and Mauk, M.D. (1995) Extinction of conditioned eyelid responses requires the anterior lobe of cerebellar cortex. *J. Neurosci.* 15, 2074–2080

16 Lisberger, S.G. *et al.* (1994) Neural basis for motor learning in the vestibuloocular reflex of primates. II. Changes in the responses of horizontal gaze velocity Purkinje cells in the cerebellar flocculus and ventral paraflocculus. *J. Neurophysiol.* 72, 954–973

17 Medina, J.F. *et al.* (2002) Inhibition of climbing fibers is a signal for the extinction of conditioned eyelid responses. *Nature* 416, 330–333

18 Kenyon, G.T. *et al.* (1998) A mathematical model of the cerebellar-olivary system I: self-regulating equilibrium of climbing fiber activity. *J. Comput. Neurosci.* 5, 17–33

19 Raymond, J.L. and Lisberger, S.G. (1998) Neural learning rules for the vestibulo–ocular reflex. *J. Neurosci.* 18, 9112–9129

20 Kleim, J.A. *et al.* (2002) Synapse formation is associated with memory storage in the cerebellum. *Proc. Natl. Acad. Sci. U. S. A.* 99, 13228–13231

21 Freeman, J.H. *et al.* (2002) Learning-specific enhancement of mossy fiber synaptic input to the interpositus nucleus. *Soc. Neurosci. Abstr.* 28, 793

22 Miles, F.A. and Lisberger, S.G. (1981) Plasticity in the vestibulo-ocular reflex: a new hypothesis. *Annu. Rev. Neurosci.* 4, 273–299

23 Medina, J.F. and Mauk, M.D. (1999) Simulations of cerebellar motor learning: computational analysis of plasticity at the mossy fiber-to-deep nucleus synapse. *J. Neurosci.* 19, 7140–7151

24 Garcia, K.S. *et al.* (1999) Cerebellar cortex lesions prevent acquisition of conditioned eyelid responses. *J. Neurosci.* 19, 10940–10947

25 Bao, S. *et al.* (2002) Cerebellar cortical inhibition and classical eyeblink conditioning. *Proc. Natl. Acad. Sci. U. S. A.* 99, 1592–1597

26 Aizenman, C.D. and Linden, D.J. (2000) Rapid synaptically driven increases in the intrinsic excitability of cerebellar deep nuclear neurons. *Nat. Neurosci.* 3, 109–111

27 Ohyama, T. and Mauk, M.D. (2001) Latent acquisition of timed responses in cerebellar cortex. *J. Neurosci.* 21, 682–690

28 Perrett, S.P. *et al.* (1993) Cerebellar cortex lesions disrupt learning-dependent timing of conditioned eyelid responses. *J. Neurosci.* 13, 1708–1718

29 Medina, J.F. *et al.* (2000) Timing mechanisms in the cerebellum: testing predictions of a large-scale computer simulation. *J. Neurosci.* 20, 5516–5525

30 Medina, J.F. *et al.* (2001) A mechanism for savings in the cerebellum. *J. Neurosci.* 21, 4081–4089

31 Buonomano, D.V. and Mauk, M.D. (1994) Neural network model of the cerebellum: temporal discrimination and the timing of motor responses. *Neural Comput.* 6, 38–55

32 Leiner, H.C. *et al.* (1993) Cognitive and language functions of the human cerebellum. *Trends Neurosci.* 16, 444–447

33 Carey, M. and Lisberger, S. (2002) Embarrassed, but not depressed: eye opening lessons for cerebellar learning. *Neuron* 35, 223–226

34 Fuchs, A.F. *et al.* (1993) Role of the caudal fastigial nucleus in saccade

generation. I. Neuronal discharge pattern. *J. Neurophysiol.* 70, 1723–1740

35 Robinson, F.R. *et al.* (1993) Role of the caudal fastigial nucleus in saccade generation. II. Effects of muscimol inactivation. *J. Neurophysiol.* 70, 1741–1758

36 Mauk, M.D. *et al.* (2000) Cerebellar function: coordination, learning or timing? *Curr. Biol.* 10, R522–R525

37 Kettner, R.E. *et al.* (1996) Predictive smooth pursuit of complex two-dimensional trajectories in monkey: component interactions. *Exp. Brain Res.* 108, 221–235

38 Leung, H.C. and Kettner, R.E. (1997) Predictive smooth pursuit of complex two-dimensional trajectories demonstrated by perturbation responses in monkeys. *Vis. Res.* 37, 1347–1354

39 Kettner, R.E. *et al.* (1997) Prediction of complex two-dimensional trajectories by a cerebellar model of smooth pursuit eye movement. *J. Neurophysiol.* 77, 2115–2130