

Computer simulation of cerebellar information processing

Javier F. Medina and Michael D. Mauk

Department of Neurobiology and Anatomy, University of Texas Medical School, 6431 Fannin, Houston, Texas 77030, USA

Correspondence should be addressed to M.D.M. (mmauk@nba19.med.uth.tmc.edu)

Although many functions have been ascribed to the cerebellum, the uniformity of its synaptic organization suggests that a single, characteristic computation may be common to all. Computer simulations are useful in examining this cerebellar computation, as they inherently address function at the level of information processing. Progress is facilitated by factors that make the cerebellum particularly amenable to such analysis. We review progress from two contrasting approaches. Top-down simulations begin with hypotheses about computational mechanisms and then ask how such mechanisms might operate within the cerebellum. Bottom-up simulations attempt to build a representation of the cerebellum that reflects known cellular and synaptic components as accurately as possible. We describe recent advances from these two approaches that are leading to an understanding of what information the cerebellum processes and how its neurons and synapses accomplish this task.

The cerebellum is said to do many things. Recording, imaging and lesion studies indicate that the cerebellum is important for coordination, balance, motor learning and even non-motor tasks¹⁻⁷. Yet its synaptic organization is remarkably uniform⁸⁻¹⁰, suggesting that a single characteristic computation may be accomplished in each case⁴. In this sense, it may be more accurate to say that the cerebellum does one thing: it receives inputs via well-defined afferents and generates outputs according to its internal rules for information processing. Diversity of function then stems from differences in what inputs encode and differences in the downstream targets of cerebellar outputs. Thus, an emphasis on what the cerebellum computes may greatly facilitate our understanding of how different cerebellar regions contribute to various motor and cognitive functions. Although we lack the data required to build useful simulations of many brain systems, several factors combine to make meaningful simulations of the cerebellum feasible. These simulations help determine whether current knowledge is sufficient to explain cerebellar input/output transformations, what important information is missing, and what key experiments can discriminate between competing hypotheses.

Factors that facilitate cerebellar simulations

Simulating a brain system is frequently hindered by incomplete knowledge about synaptic organization and physiology, how inputs are engaged by stimuli and what they encode, and how outputs influence behavior. Owing to the efforts of Eccles, Ito, Szentagothai, Llinas and others, we largely understand the synaptic organization and physiology of the cerebellum⁸⁻¹¹. Similarly, analysis of cerebellar-dependent forms of learning such as Pavlovian eyelid conditioning and adaptation of the vestibulo-ocular reflex has revealed the types of information carried by cerebellar inputs and details of input/output transformations^{12,13}. Together, these advantages are fundamental to the construction and useful application of cerebellar simulations.

The synaptic organization of the cerebellum is well known and lends itself well to simulation (Box 1)⁸⁻¹⁰. There are an enormous number of neurons, but a limited number of neuron types. A rela-

tively simple circuit of these neurons is repeated throughout the entire cerebellum. Inputs arrive via distinct afferents, the mossy fibers and climbing fibers, and output is provided by the cerebellar deep nuclei. Inputs modulate the output activity of deep nuclei cells through two pathways: direct excitatory connections from mossy fibers onto deep nuclei cells and an indirect pathway in the cerebellar cortex involving excitation of Purkinje cells by granule cells and ultimately resulting in inhibition of deep nuclei cells by Purkinje cells. The synaptic physiology, numerical ratios, convergence and divergence ratios, and geometry of projections for most of these cells are known⁸⁻¹¹.

There is also increasingly detailed information available regarding plasticity at the granule cell to Purkinje (*gr*→*Pkj*) synapses in the cerebellar cortex and plasticity in the deep nuclei, possibly at the mossy fiber synapses onto the deep nucleus output cells (*mf*→*nuc*)^{12,14-17}. Together, these findings provide the information and constraints required to build computer simulations composed of detailed representations of cerebellar cells and synapses.

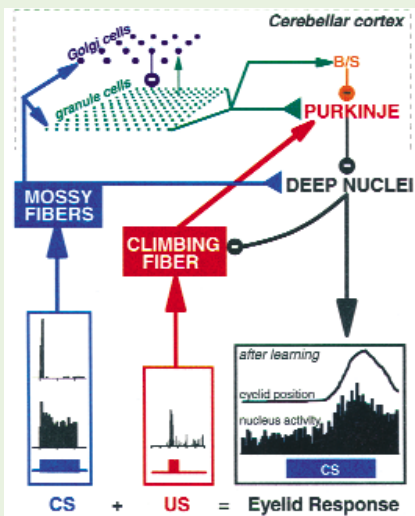
Simulations of the cerebellum also benefit greatly from detailed information about how cerebellar inputs and outputs act under different conditions. Much of this knowledge comes from studies of eyelid conditioning, as its behavioral properties reflect relatively directly the input/output properties of the cerebellum^{12,13}. Eyelid conditioning involves paired presentations of a neutral stimulus such as a tone, and a reinforcing stimulus such as an air puff directed at the eye. Initially, there is only a reflex response to the air puff, but after approximately 100–200 training trials, the tone acquires the ability to elicit conditioned eyelid closure. This learning is bidirectional. In 'extinction', repeated presentation of the tone without the air puff makes the conditioned response gradually disappear¹⁸. Recording, lesion and stimulation studies show relatively direct relationships between eyelid conditioning and cerebellar inputs and outputs. The tone and air puff are conveyed to the cerebellum via mossy fibers¹⁹⁻²² and climbing fibers²³⁻²⁵ respectively, and output from one of the cerebellar nuclei drives the expression of learned responses^{26,27} (Box 1).

These findings are important in three ways. First, they provide

Box 1. Synaptic organization of the cerebellum and its relationship to eyelid conditioning.

Purkinje cells of the cerebellar cortex (dashed box) project to the deep nuclei (black), which transmit cerebellar output to other brain regions. Input reaches the cerebellum via two excitatory pathways: mossy fibers (blue) and climbing fibers (red). Mossy fibers originate from various brain stem nuclei that receive information from cerebral cortex and spinal cord. They influence cerebellar output through direct excitatory connections onto the deep nuclei cells ($mf \rightarrow nuc$ synapses) and through a more complex pathway in the cerebellar cortex that culminates in Purkinje cell inhibition of deep nuclei neurons. This pathway begins with mossy fiber excitation of both granule and Golgi cells, includes three types of inhibitory interneurons (Golgi, basket and stellate cells), and involves the abundant excitatory synapses that granule cells make onto Purkinje cells ($gr \rightarrow Pkj$ synapses; there are about 200,000 for each Purkinje cell, and given that Purkinje cells converge onto deep nucleus cells, ~ 100 million $gr \rightarrow Pkj$ synapses contribute to the output of a single nucleus cell). The other input pathway, the climbing fibers, is strikingly different. Climbing fibers originate in the inferior olive and make direct excitatory connections onto Purkinje cells. Individual climbing fibers contact about 10 Purkinje neurons, and in sharp contrast to the granule cell input, each Purkinje cell receives synaptic connections from a single climbing fiber. This synapse is exceptional for its extensive spatial distribution over the Purkinje cell dendrite, which makes the climbing fiber response in Purkinje cells all-or-none.

During eyelid conditioning, these pathways lead to learning through the modification of synaptic strength at two different cerebellar sites. Eyelid conditioning involves paired presentation of a conditioned stimulus (CS; for example, a tone) and an unconditioned stimulus (US; for example, a puff of air to the eye). Whereas the CS does not initially cause a response, after repeated CS+US presentations, the animal learns a conditioned response: the eyelid closes in response to the CS. Lesion, stimulating and recording evidence from many laboratories suggests that, first, the CS is conveyed to the cerebellum via activation of the mossy fiber input¹⁹ (histograms in blue box; note the minimal temporal variation because mossy fibers are activated either throughout tone presentation or just at its onset). Second, the US is conveyed by activating the climbing fiber input²³ (histogram in red box). Third, increases in the activity of cerebellar output cells in the anterior interpositus nucleus (histogram in black box) drive the expression of the learned eyelid response²⁷ (line above histogram). Learning this conditioned response involves two sites of plasticity: $gr \rightarrow Pkj$ synapses in the cerebellar cortex, and in the cerebellar nucleus, possibly at $mf \rightarrow nuc$ synapses.



some clues about what the cerebellum computes. Essentially, eyelid conditioning shows that repeated presentation of certain patterns of mossy fiber and climbing fiber inputs induces changes in the cerebellar output elicited by the mossy fiber inputs. The uniformity of cerebellar circuitry suggests that these input/output properties may apply generally to all cerebellar tasks. Thus, analysis and simulation of eyelid conditioning can be used as tools to understand how the cerebellum processes information. Second, actual tone-activated mossy fiber¹⁹ and air puff-activated climbing fiber²³ responses can provide meaningful inputs to simulations. Third, comparing the output of a simulation with the behavioral properties of eyelid responses under different conditions provides a simple measure of how well the simulation captures cerebellar input/output transformations.

Temporally specific learning

Eyelid conditioning studies show that the fundamental computation carried out by the cerebellum involves not only changing its output when there is a motor error (signaled by the climbing fiber pathway), but delaying such changes in output to the time when the error is expected. The key behavioral feature in this regard is the extraordinary sensitivity of eyelid conditioning to the interval of time between the onsets of the tone and air puff (the interstimulus interval or ISI; Fig. 1). There is no learning with ISIs less than 100 ms; learning is best for ISIs of 150–500 ms, and gradually worsens for longer ISIs (Fig. 1a)^{28,29}. Moreover, conditioned responses are precisely timed; they begin before, and always peak at, the time when the air puff is expected (Fig. 1b)^{30,31}.

Eyelid conditioning thus reveals that a mossy fiber input that repeatedly predicts a climbing fiber input causes the cerebellar output to increase. This increase is delayed with respect to the mossy

fiber input so that it peaks about when the climbing fiber input is expected (Fig. 1c). This is when cerebellar output was faulty, leading to a motor error and generation of the climbing fiber input, and so the cerebellum has solved a temporal version of the credit assignment problem: that is, when there is a motor error, how should cerebellar output change? The answer is to time the output change to influence the motor commands responsible for the error encoded by the climbing fiber input, so that subsequent motor performance will be improved.

Although the varied behavioral properties of eyelid conditioning reveal other interesting components of cerebellar processing, this capacity for temporally specific learning seems to be the fundamental aspect of cerebellar information processing. Thus, we focus on computational attempts to understand the ISI effect and the learned timing of conditioned eyelid responses.

Top-down simulations of timing and ISI effects

The goals of attempts to understand temporally specific learning using top-down simulations of eyelid conditioning have been stated clearly³². First, devise real-time computational models that describe as much of the known behavioral and physiological evidence as possible. Second, devise an implementation scheme that aligns features of the model with the neural circuits involved. Third, test implications of the model and its implementation in experiments. Although different in their details, most top-down simulations of eyelid conditioning accomplish temporally specific learning with three common features.

First, they incorporate the finding that $gr \rightarrow Pkj$ synapses decrease in strength (long-term depression or LTD) when they are active during a climbing fiber input to the Purkinje cell^{15,16}. In this way, the air puff-activated climbing fiber serves as a signal to decrease the

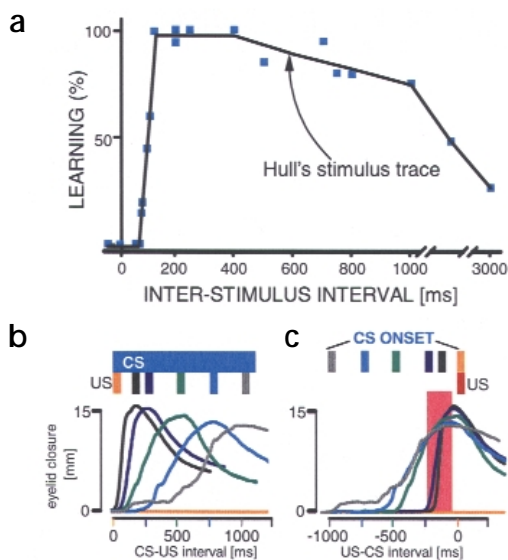


Fig. 1. Time-dependent properties of eyelid conditioning. **(a)** Learning is possible only for a limited range of interstimulus intervals (ISI) between CS and US onset (~ 100 ms to 3 s). As suggested by Hull³⁸, this effect can be explained by considering a hypothetical stimulus trace curve with the time course of the ISI function. **(b)** When an effective ISI is used, learned responses are precisely timed so that maximum eyelid closure occurs near the US onset³⁰. Bottom, sample eyelid responses (upward deflection means eye is closing) illustrate response timing for six different ISIs. In each case, the small rectangle indicates when the US is presented during training, with its associated learned response in the same color. Simultaneous presentation of CS and US (ISI = 0) produces no learning (orange). **(c)** The same eyelid responses as in **(b)** aligned with respect to the climbing fiber input (US). The timing of the learned responses shows how, regardless of the time of the onset of the CS, the learned increases in cerebellar output always peak about 100 ms before the arrival of the climbing fiber input (pink box). Because the errant movement that elicited the climbing fiber input was produced by motor commands near this time, this timing capability ensures that the correct component of the movement is adjusted; note that all responses are timed to peak just before the error.

strength of tone-activated synapses, causing Purkinje cell activity during the tone to decrease. Because Purkinje cells are inhibitory, the increase in cerebellar output associated with the conditioned response is generated by the transient pause in Purkinje cell firing through disinhibition of cerebellar nucleus cells. In support of this hypothesis, Purkinje cell recordings from identified eyelid regions of the cerebellar cortex in conditioned animals demonstrate Purkinje cells that learn to decrease their firing during the tone³³.

Second, the activity of granule cells in top-down models is assumed to vary in time such that different $gr \rightarrow Pkj$ synapses are active at different times during the tone. This feature is crucial for generating appropriately timed conditioned responses because it allows LTD to occur strictly for those $gr \rightarrow Pkj$ synapses that are active when the air puff is presented. Although the activity of individual granule cells *in vivo* remains unknown, top-down simulations propose three different patterns of granule cell activity that could produce temporally specific learning (Fig. 2). Tapped-delay-line models (Fig. 2a) from Moore and colleagues assume that the tone sequentially engages different input elements (for example, mossy fibers)^{32,34}. In contrast, Grossberg's spectrum models assume that granule cells have a variety of time constants so that $gr \rightarrow Pkj$ synapses are activated at different times after the tone (Fig. 2b)³⁵. The most recent variant of this kind of model (Fig. 2b) assumes that an array of differently timed calcium responses occur in Purkinje cells because of variations in the number of metabotropic glutamate receptors³⁶. A third type of model assumes that granule cells activated by the tone oscillate with different frequencies and phases (Fig. 2c)³⁷. Thus, in each model, any given Purkinje cell is receiving a barrage of inputs from many granule cells. However, some of those granule cells fire at a specific time relative to the onset of the tone. Those $gr \rightarrow Pkj$ synapses that become active at the time of the air puff-activated climbing fiber input are weakened by LTD, resulting in a sudden drop in synaptic input (and a transient pause in Purkinje cell activity) specific to the time when the air puff is expected.

Third, the ISI effect in top-down models is produced by a feature, initially derived from Hull's influential stimulus trace idea³⁸, that prevents learning at ISIs shorter than 100 ms or longer than 3 seconds. According to this hypothesis, the strength or some energetic aspect of the tone's neural representation has a time course that parallels the ISI function (Fig. 1a). The air puff is then assumed to promote learning on each trial according to the strength of the tone representation. For example, the tapped delay line model

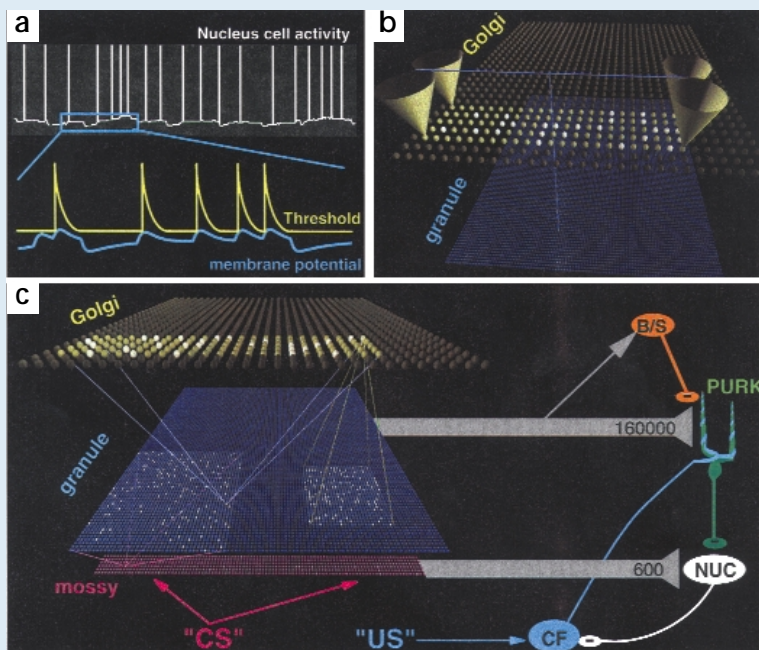
(Fig. 2a) assumes that the tone initiates an eligibility signal with the time course of the ISI function, which determines whether $gr \rightarrow Pkj$ synapses undergo plasticity³⁴. In models of this kind, learning is not possible with certain tone-air puff intervals because although there are active $gr \rightarrow Pkj$ synapses at the time of the air puff-activated climbing fiber input, the value of the hypothetical eligibility signal that is associated with active $gr \rightarrow Pkj$ synapses at those times is too small to cause LTD induction. In spectrum models (Fig. 2b), the differently timed granule cells responses accomplish the same task^{35,36}. In these models, ineffective ISIs cannot support learning simply because there are no $gr \rightarrow Pkj$ synapses active at those times.

Fig. 2. Top-down mechanisms for timing. **(a)** Tapped-delay-line models. The simple delay-line generates time-varying activity by arranging neurons in a chain so that each synapse (red triangles) adds a discrete delay to the signal originated by the CS. Right, hypothetical raster plots for five mossy fibers arranged top to bottom according to their position in the chain. Appropriate response timing is generated by selectively modifying only those cells that are active at the time of reinforcement by the US. **(b)** Spectrum models assume that neurons can respond to the CS with a broad range of time constants. Thus, if the US is presented shortly after the CS, learning is restricted to neurons with fast dynamics (top green cells). Similarly, the activity of neurons with slower membrane dynamics (bottom green cells) is delayed and thus contributes to response timing for longer CS-US intervals. **(c)** Oscillation models. A pool of pacemaker neurons oscillating at different frequencies can also be used to encode time. Different times during presentation of a CS can be distinguished from each other because the combination of neurons active at any point is unique to that particular time. Thus, response timing arises because a particular combination of cells is weakened through LTD by the US-activated climbing fiber.

review

Box 2. Building a bottom-up computer simulation of cerebellar cortex.

(a) Integrate-and-fire representations of neurons. Activity of a nucleus cell (top) and an expanded view of its membrane potential (bottom) during a simulation. The integrate-and-fire model does not, as commonly assumed, sum inputs until a threshold is reached. Instead, it uses the Goldman-Katz equation to solve for membrane potential (blue line) given leak and synaptic conductances (note EPSPs and IPSPs). The main simplification is that because there are no active conductances, action potentials are assumed to occur when membrane potential reaches threshold (yellow line). Even without active conductances, absolute and relative refractory periods can be implemented via spike-activated changes in threshold. For example, when a spike occurs, threshold is increased and then decays back to its normal level. Similarly, afterhyperpolarizations can be implemented by having each spike transiently increase a potassium conductance. (b, c) Wiring the network. The goal of building a bottom-up simulation is to duplicate the numerical, divergence and convergence ratios of cells and the geometry of their projections. This requires a simulation with a very large number of granule cells (here 10,000, blue layer). There are 900 Golgi cells (yellow), 600 mossy fibers (purple), 20 Purkinje cells (green, only one shown) and 6 nucleus cells (white, one shown). Because there are over 300,000 synapses in this network, it is not possible to specify each by hand; instead, the network is built by specifying appropriate rules of connectivity. (b) An example of how the connections of one granule cell onto its target Golgi cells are determined. The geometry of projections is determined by limiting possible connections to a defined region (bright yellow Golgi cells), with its shape and size determined by the length of the granule cell axon (parallel fiber, light blue) and the diameter of the Golgi cell dendritic trees (yellow cones). The Golgi cells within this region that actually receive a synapse (white Golgi cells) are determined by specifying divergence (how many Golgi cells a single granule cell contacts) and convergence ratios (how many granule cell synapses each Golgi cell receives). (c) This process is then repeated for all cells and all cell types in the network. This includes a large number of *gr*→*Plkj* synapses (gray), *mf*→*nuc* synapses (also gray), Purkinje cell inhibitory synapses onto nucleus cells and nucleus cell synapses onto the climbing fiber.



How these models accomplish response timing and the ISI effect is instructive, but the built-in nature of the mechanisms means that the utility of these models arises mostly from the specificity of their predictions, which can be used to guide research. In this context, tapped-delay lines, neurons with a broad range of membrane dynamics, and oscillatory cells all are possible candidates capable of generating and making use of temporal information. Because granule cells are too small and densely packed to be recorded individually *in vivo* by conventional physiological techniques, whether or not the cerebellum uses any of these mechanisms to encode time remains an intriguing question whose resolution awaits further technological advances.

Bottom-up simulations

Following the initial success of top-down models in proposing hypothetical mechanisms and identifying key features of temporally specific learning by the cerebellum, recent studies using a different modeling approach have provided further insights into how the cerebellum processes information. In contrast to the fundamentally hypothetical nature of top-down models, bottom-up approaches begin by building computer simulations that are constrained by cerebellar anatomy and physiology. The point is to avoid building in assumptions about how certain components function. For example, rather than assuming that all cells of a particular type are similar and so only a few will suffice, a bottom-up approach tries to adhere to numerical cell ratios as well as to divergence/convergence ratios and connection geometry, making bottom-up sim-

ulations large and computationally intensive. They also require extremely large numbers of synapses (more than 250,000 in our smallest simulations), precluding the possibility of specifying connections by hand. Instead, in a way that somewhat parallels development, simulations are built by using empirically determined constraints to guide formation of connections (Box 2).

Building bottom-up simulations requires knowledge not only about how the cerebellum is wired, but also about the strengths of different synapse types and how the postsynaptic cell integrates their conductances. Fortunately, many important details of cerebellar connectivity are known from *in vitro* studies³⁹. However, as discussed in another review in this collection⁴⁰, dendritic processing is very complex and includes many nonlinear effects. Given the amount of computer power available today, many of these important details cannot be incorporated in large-scale computer simulations. Thus, bottom-up simulations undoubtedly contain errors of omission. The point is to avoid building in assumptions and thus to minimize errors of commission.

One disadvantage of bottom-up simulations is that they often fail to reproduce some of the behavioral properties of eyelid conditioning (for example, generating poorly timed responses or allowing learning for tone-air puff intervals that are known to be ineffective). Because bottom-up simulations do not address these failures by incorporating hypothetical mechanisms that ensure success, they may initially lack the ability to make testable predictions. Yet, even failures can indicate the need for more information and can guide empirical studies by providing hints about what is

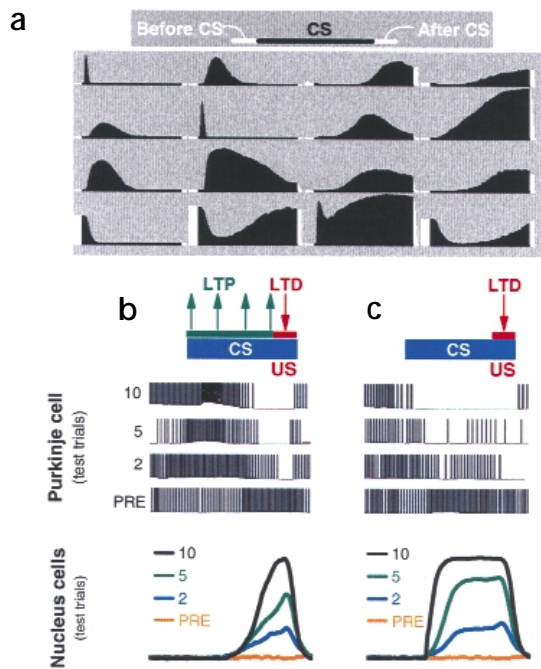


Fig. 3. Factors contributing to response timing in the simulations. (a) Example histograms for 16 of the 10,000 simulated granule cells during 500 presentations of a 1-s CS (black). For reference, the activity of each cell during 100-ms periods before and after the CS is shown (white). Simulated granule cells showed both increased activity (top three rows) and decreased activity (bottom row) at some point during CS presentation. The ability of simulated granule cells to fire at different times during the CS arises naturally as a consequence of the complex patterns of granule cell excitation and inhibition that result from dynamic interactions among mossy fibers, granule cells and Golgi cells. (b) Activity of simulated nucleus and Purkinje cell in test trials (CS presented without US) at different times during acquisition training. Numbers at left, training session from which the response was taken. Simulated Purkinje cell spontaneous activity is approximately 60 Hz, and before training (pre), it is relatively unchanged by CS presentation (blue). As training proceeds and LTD is induced at *gr*→*Pkj* synapses during presentation of the US (red bar), the Purkinje cells learn to decrease their activity late in the CS. In addition, LTP induction during the unreinforced portions of the CS (green bar) results in increased Purkinje cell activity during the early parts of the CS. Bottom, both these changes in Purkinje cell activity contribute to response timing by suppressing nucleus cell responding early in the CS and disinhibiting responding later in the CS. (c) When LTP at the *gr*→*Pkj* synapses is artificially disengaged during the CS, the simulation learns to respond, but the timing of the nucleus cell responses (bottom) is inappropriately broad. This reflects Purkinje cell activity (top) decreasing late in the CS without being able to increase early in the CS.

missing. Then when (if) bottom-up simulations are successful in generating some aspect of eyelid conditioning, it can be relatively straightforward to determine how they do so, generating specific empirically testable predictions.

Our initial simulations illustrate how failures can provide insights⁴¹. Top-down models had suggested that temporally specific learning requires different *gr*→*Pkj* synapses to be active at different times during the tone. Our initial bottom-up simulations revealed that the intrinsic circuitry of the cerebellar cortex naturally produces this time-varying activity during the tone in a subset of granule cells (Fig. 3a, peri-stimulus histograms). However, this same simulation revealed that during realistic tone-like inputs (which include mossy fibers that fire tonically throughout the entire tone; Box 1), the geometry of cerebellar cortical connections caused many granule cells to become activated through large portions of the tone (Fig. 3a). This coarse temporal coding, in contrast to the sharp temporal specificity of granule cells proposed in top-down models, resulted in abnormally broad responses and thus failed to generate temporally specific learning. We sought to address this serious deficiency with a more complete representation of the cerebellum and, as information became available, by adding quantitative details about synaptic conductances and conditions for plasticity.

Perhaps the most surprising result from these more complete simulations was that broad responses could be sharpened by incorporating evidence that in addition to LTD, *gr*→*Pkj* synapses can also undergo LTP when active without a climbing fiber input (Fig. 3b)³⁹. When this evidence for bidirectional plasticity of *gr*→*Pkj* synapses was included, the simulation still acquired responses (through LTD at *gr*→*Pkj* synapses and resulting decreases in Purkinje cell activity) about when the air puff-activated climbing fiber input occurred. Unexpectedly, however, the simulation also learned to suppress responses (through LTP at other *gr*→*Pkj* synapses and resulting increases in

Purkinje cell activity) at other times during the same tone that were not reinforced by the climbing fiber input associated with the air puff (Fig. 3b). After learning, the net effect of these synaptic modifications was a temporally sharp Purkinje cell response, with a characteristic pause in activity around the time of the air puff that had been sharpened by a preceding increase in activity (Fig. 3b, Purkinje cell activity at training day 10). This biphasic pattern of Purkinje cell activity had been recorded from the cerebellar cortex of conditioned animals³³ and may be ideally suited to generate precisely timed responses because the deep cerebellar nuclei cells responsible for generating the conditioned response are inhibited by Purkinje cells and respond with a pronounced rebound depolarization and a burst of spikes when released from inhibition⁴². Thus, the simulation suggests that in contrast to the proposals of top-down models, the cerebellum may not require a high degree of temporal specificity in granule cell activity to accomplish temporally specific




Fig. 4. Partial lesions of the cerebellar cortex. Simulations (left) and rabbits (right) display a similar pattern of responding following a partial lesion of the cerebellar cortex. Each trace represents the average output of the simulated nucleus cells (left) or the average eyelid response (right) over 12 CS test trials. Colors indicate how many sessions after the lesion each trace was measured. Both simulations and rabbits were trained before the lesion and the 'pre' sweep shows responses given over the two sessions (240 trials for the simulations) before the lesion. Initially after the lesion (1 post), simulations and rabbits produced responses with a short-latency onset. With additional training, response timing gradually recovered to normal levels (6 post). In the simulations, this recovery was possible because the Purkinje cells that were spared by the lesion underwent additional learning (LTP at *gr*→*Pkj* synapses), acquiring even more robust increases in responding early in the CS and thus suppressing the short-latency component.

review




Fig. 5. Competition between acquisition and extinction contributes to the ISI effect in the simulations. (a) Summary of data from two studies that examined the ISI function in rabbits^{28,29}. Learning was fast and robust with ISIs of 200 ms (black line) and 400 ms (orange line). For the 800 ms ISI (purple line), learning was weaker, and for the 2 s ISI (pink line), there was little evidence of learning. (b) When the simulations are trained with these same ISIs, the rates of learning are comparable to those shown in (a). In the simulations, learning declines as the ISI is lengthened because in each training trial, the portion of the CS not reinforced by the US-climbing fiber (that is, LTP at *gr*→*Pkj* synapses, green bar) grows relative to the portion that is reinforced by the US-climbing fiber (that is, LTD at *gr*→*Pkj* synapses, red bar).

learning. Under conditions of coarse temporal coding by the granule cells, the simulation generated properly timed responses by learning to respond at the appropriate time through the induction of LTD while learning to suppress responses when they were not required through the induction of LTP.

These suggested mechanisms make empirically testable predictions. We recently tested the behavioral effects predicted for partial lesions of the cerebellar cortex³⁹. Initially after a lesion that removed ~50% of the simulated Purkinje cells, the simulation displayed maladaptive responses with a very short latency after tone onset (Fig. 4a). This characteristic disruption of response timing shows that part of what the simulated Purkinje cells had learned to do before the lesion was to suppress responding during the early parts of the tone. With further training in the simulations, the onset of the response recovered to pre-lesion levels, indicating a spared capacity of the remaining Purkinje cells to further suppress responding through the induction of LTP at *gr*→*Pkj* synapses active during the early parts of the tone (Fig. 4a). We observed the same characteristic pattern of responding in rabbits after small, electrolytic lesions of the cerebellar cortex³⁹ (Fig. 4b). These data are consistent with the predicted importance of LTP at *gr*→*Pkj* synapses in generating temporally specific learning by suppressing responses during the unreinforced parts of the tone.

The hypothesized competition between the LTP that is induced during the early parts of the tone and the LTD that occurs during the air puff-activated climbing fiber input also offers a new explanation for the ISI effect (Fig. 5). Essentially, as the duration of the tone–air puff interval increases, opportunities for LTP induction during the tone, before the air puff activates climbing fibers, come to outweigh those for LTD induction. Thus, as the tone–air puff interval lengthens, *gr*→*Pkj* synapses that are active through large portions of the tone undergo more LTP than LTD, which prevents them from contributing to the pause in Purkinje cell activity required for generating responses. This hypothesis makes several predictions that could be tested by manipulating the amount of LTP experienced by *gr*→*Pkj* synapses during each conditioning trial. For example, extending the tone past the air puff while keeping the tone–air puff interval constant may increase the opportunities for LTP relative to LTD. This manipulation is thus expected to retard learning for a particular set of tone–air puff intervals. In contrast, selectively blocking LTP while leaving LTD unaffected is expected to allow normally ineffective tone–air puff intervals to support robust conditioning. As with top-down simulations, the specificity

of these predictions illustrates how bottom-up simulations can also be useful in guiding empirical research.

Summary

Computer simulations are beginning to contribute to our understanding of cerebellar information processing, both in terms of reproducing key behavioral features of cerebellar tasks and in terms of making empirically testable predictions. In particular, simulations of eyelid conditioning have focused empirical studies on key issues while also serving as a rough measure of progress—how well our understanding of different cerebellar components combines to explain how the cerebellum accomplishes temporally specific motor learning. Because the temporally specific learning in eyelid conditioning is a common property shared by other forms of cerebellar tasks¹², these studies are also likely to reveal general features about the input/output transformation of the cerebellum. Thus, given the remarkably uniform synaptic organization of the cerebellum throughout its different regions, it is possible that what the eyelid

region of the cerebellum computes during eyelid conditioning is similar to the computation carried out by other cerebellar regions during other cerebellar tasks. In this context, the development of a computer simulation constrained by the anatomy and physiology of the cerebellum and capable of reproducing key behavioral aspects of the growing list of cerebellar tasks^{1–7} represents an important challenge for future computational studies.

ACKNOWLEDGEMENTS

Supported by MH57051 and MH46904.

RECEIVED 15 JUNE; ACCEPTED 26 SEPTEMBER 2000

- Llinas, R. & Welsh, J. P. On the cerebellum and motor learning. *Curr. Opin. Neurobiol.* **3**, 958–965 (1993).
- Courchesne, E. & Allen, G. Prediction and preparation, fundamental functions of the cerebellum. *Learn. Mem.* **4**, 1–35 (1997).
- Bower, J. M. Control of sensory data acquisition. *Int. Rev. Neurobiol.* **41**, 489–513 (1997).
- Bloedel, J. R. Functional heterogeneity with structural homogeneity: How does the cerebellum operate? *Behav. Brain Sci.* **15**, 666–678 (1992).
- Thach, W. T., Goodkin, H. P. & Keating, J. G. The cerebellum and the adaptive coordination of movement. *Annu. Rev. Neurosci.* **15**, 403–442 (1992).
- Ivry, R. B. The representation of temporal information in perception and motor control. *Curr. Opin. Neurobiol.* **6**, 851–857 (1996).
- Schmahmann, J. D. *International Review of Neurobiology. The Cerebellum and Cognition* (Academic, San Diego, California, 1997).
- Voogd, J. & Glickstein, M. The anatomy of the cerebellum. *Trends Neurosci.* **21**, 370–375 (1998).
- Eccles, J. C., Ito, M. & Szentágothai, J. *The Cerebellum as a Neuronal Machine* (Springer, Berlin, New York, 1967).
- Ito, M. *The Cerebellum and Neural Control* (Raven, New York, 1984).
- Llinas, R. in *Handbook of Physiology Vol. II The Nervous System* (ed. Brokks, V. B.) 831–976 (Am. Physiol. Soc., Bethesda, Maryland, 1981).
- Raymond, J. L., Lisberger, S. G. & Mauk, M. D. The cerebellum: a neuronal learning machine? *Science* **272**, 1126–1131 (1996).
- Kim, J. J. & Thompson, R. F. Cerebellar circuits and synaptic mechanisms involved in classical eyeblink conditioning. *Trends Neurosci.* **20**, 177–181 (1997).
- Mauk, M. D. Roles of cerebellar cortex and nuclei in motor learning: contradictions or clues? *Neuron* **18**, 343–346 (1997).
- Sakurai, M. in *The Olivocerebellar System in Motor Control* (ed. Strata, P.) 221–230 (Springer, New York, 1989).
- Hirano, T. Depression and potentiation of the synaptic transmission between a granule cell and a Purkinje cell in rat cerebellar culture. *Neurosci. Lett.* **119**, 141–144 (1990).
- Racine, R. J., Wilson, D. A., Gingell, R. & Sunderland, D. Long-term potentiation in the interpositus and vestibular nuclei in the rat. *Exp. Brain Res.* **63**, 158–162 (1986).
- Schneiderman, N., Fuentes, I. & Gormezano, I. Acquisition and extinction of the classically conditioned eyelid response in the albino rabbit. *Science* **136**, 650–652 (1962).
- Aitkin, L. M. & Boyd, J. Acoustic input to the lateral pontine nuclei. *Hearing Res.* **1**, 67–77 (1978).
- Steinmetz, J. E., Lavond, D. G. & Thompson, R. F. Classical conditioning of the

- rabbit eyelid response with mossy fiber stimulation as the conditioned stimulus. *Bull. Psychonom. Soc.* **28**, 245–248 (1985).
21. Steinmetz, J. E., Logan, C. G. & Thompson, R. F. in *Cellular Mechanisms of Conditioning and Behavioral Plasticity* (eds. Woody, D. L., Alkon, D. L. & McCaugh, J. L.) 143–148 (Plenum, New York, 1988).
 22. Lewis, J. L., LoTurco, J. J. & Solomon, P. R. Lesions of the middle cerebellar peduncle disrupt acquisition and retention of the rabbit's classically conditioned nictitating membrane response. *Behav. Neurosci.* **101**, 151–157 (1987).
 23. Sears, L. L. & Steinmetz, J. E. Dorsal accessory inferior olive activity diminishes during acquisition of the rabbit classically conditioned eyelid response. *Brain Res.* **545**, 114–122 (1991).
 24. Mauk, M. D., Steinmetz, J. E. & Thompson, R. F. Classical conditioning using stimulation of the inferior olive as the unconditioned stimulus. *Proc. Natl. Acad. Sci. USA* **83**, 5349–5353 (1986).
 25. McCormick, D. A., Steinmetz, J. E. & Thompson, R. F. Lesions of the inferior olivary complex cause extinction of the classically conditioned eyeblink response. *Brain Res.* **359**, 120–130 (1985).
 26. McCormick, D. A., Clark, G. A., Lavond, D. G. & Thompson, R. F. Initial localization of the memory trace for a basic form of learning. *Proc. Natl. Acad. Sci. USA* **79**, 2731–2735 (1982).
 27. McCormick, D. A. & Thompson, R. F. Cerebellum: essential involvement in the classically conditioned eyelid response. *Science* **223**, 296–299 (1984).
 28. Schneiderman, N. & Gormezano, I. Conditioning of the nictitating membrane of the rabbit as a function of CS–US interval. *J. Comp. Physiol. Psychol.* **57**, 188–195 (1964).
 29. Smith, M. C., Coleman, S. R. & Gormezano, I. Classical conditioning of the rabbit's nictitating membrane response at backward, simultaneous, and forward CS–US intervals. *J. Comp. Physiol. Psychol.* **69**, 226–231 (1969).
 30. Mauk, M. D. & Ruiz, B. P. Learning-dependent timing of Pavlovian eyelid responses: conditioning using multiple interstimulus intervals. *Behav. Neurosci.* **106**, 666–681 (1992).
 31. Millenson, J. R., Kehoe, E. J. & Gormezano, I. Classical conditioning of the rabbit's nictitating membrane response under fixed and mixed CS–US intervals. *Learn. Motiv.* **8**, 351–366 (1977).
 32. Moore, J. W. & Choi, J. S. Conditioned response timing and integration in the cerebellum. *Learn. Mem.* **4**, 116–129 (1997).
 33. Hesslow, G. & Ivarsson, M. Suppression of cerebellar Purkinje cells during conditioned responses in ferrets. *Neuroreport* **5**, 649–652 (1994).
 34. Moore, J. W., Desmond, J. E. & Berthier, N. E. Adaptively timed conditioned responses and the cerebellum: a neural network approach. *Biol. Cybern.* **62**, 17–28 (1989).
 35. Bullock, D., Fiala, J. C. & Grossberg, S. A neural model of timed response learning in the cerebellum. *Neural Networks* **7**, 1101–1114 (1994).
 36. Fiala, J. C., Grossberg, S. & Bullock, D. Metabotropic glutamate receptor activation in cerebellar Purkinje cells as substrate for adaptive timing of the classically conditioned eye-blink response. *J. Neurosci.* **16**, 3760–3774 (1996).
 37. Gluck, M. A., Reifsnider, E. S. & Thompson, R. F. in *Neuroscience and Connectionist Theory* (eds. Gluck, M. A. & Rumelhart, D. E.) 131–185 (Lawrence Erlbaum, Hillsdale, New Jersey, 1990).
 38. Hull, C. L. *Principles of Behavior, an Introduction to Behavior Theory* (Appleton-Century-Crofts, New York, 1943).
 39. Medina, J. F., Garcia, K. S., Nores, W. L., Taylor, N. M. & Mauk, M. D. Timing mechanisms in the cerebellum: testing predictions of a large-scale computer simulation. *J. Neurosci.* **20**, 5516–5525 (2000).
 40. Koch, C. & Segev, I. Single neurons and their role in information processing. *Nat. Neurosci.* **3**, 1171–1177 (2000).
 41. Buonomano, D. V. & Mauk, M. D. Neural network model of the cerebellum: temporal discrimination and the timing of motor responses. *Neural Comput.* **6**, 38–55 (1994).
 42. Aizenman, C. D. & Linden, D. J. Regulation of the rebound depolarization and spontaneous firing patterns of deep nuclear neurons in slices of rat cerebellum. *J. Neurophysiol.* **82**, 1697–1709 (1999).

Viewpoint • What does 'understanding' mean?

When Ed Lewis in my department won a Nobel Prize a few years ago, our chair organized a party. On my way there, I overheard an illustrious chemist offer, "Hey, at least one smart biologist," making his colleagues chuckle. Nothing new in academia, the land of the high-minded yet curiously parochial primate. Why bring this up? Because science starts with human interactions: if we want theory and experimental neuroscience to strengthen each other, we must hope for people with different cultures, expertise, perspectives and footwear to leave their prejudices at the door and learn to better appreciate each other's strengths. This is not easy to achieve when human nature makes us shun the unfamiliar, when the structure of academic institutions imposes borders between disciplines, and when reductionist approaches alone undeniably produce so much concrete knowledge. So, if reductionism works so well—as it has in the history of neuroscience—why should we care about bringing theory (and theorists) into the kitchen? It all boils down, it seems to me, to a classical philosophical question: what does 'understanding' mean? Upon reflection, it is depressing, if not scandalous, to realize how rarely I ask myself this. As an experimentalist, I would consider most of what my lab does as descriptive; at best, we try to tie one observation to another through some causal link. Most of what we try to explain has a mechanistic underpinning; if not, a manuscript reviewer, editor or grant manager usually reminds us that is what this game is about. And we all go our merry way filling in the blanks. This is, in my view, where theorists most enrich what we do. Theorists, through their training, bring a different view of explanatory power. Causal links established by conventional, reductionist neurobiology are usually pretty short and linear, even when experiments to establish those links are horrendously complex: molecule M phosphorylates molecule N, which causes O; neuron A inhibits neuron B, 'sharpening' its response characteristics. This beautiful simplicity is the strength of reductionism and its weakness. To understand the brain, we will, in the end, have to understand a system of interacting elements of befuddling size and combinatorial complexity. Describing these elements alone, or even these elements and all the links between them, is obviously necessary but, many would say, not satisfyingly explanatory. More precisely, this kind of approach can only explain those phenomena that reductionism is designed to get at. It is the classical case of the lost key and the street lamp; we often forget that the answers to many fundamental questions lie outside of the cone of light shed by pure analysis (in its etymological sense). I am interested in neuronal systems. In most cases, a system's collective behavior is very difficult to deduce from knowledge of its components. Experience with many systems of neurons under varied regimes could, in theory, eventually give me a good intuitive knowledge of their behavior: I could predict how system S should behave under certain conditions. Yet my understanding of it would be minimal, in the sense that I could not convey it to someone else, except by conveying all my past experience. This is one of the many places where theorists can help me. Much of what we need to provide a deeper understanding of these distributed phenomena may already exist in some corner of the theory of dynamical systems, developed by mathematicians, physicists or chemists to understand or describe other features of nature. If it does not, maybe it can be derived. But the first step is to map my biological system onto the existing theoretical landscape. This is where the challenge (and fun) lies—and where sociological forces must be tamed. In brief, neuroscience is, to me, a science of systems in which first-order and local explanatory schemata are needed but not sufficient. Reductionism, by its nature, takes away the distributed interactions that underlie the global properties of systems. Theoretical approaches provide different means to simplify. We must thus learn to understand, rather than avoid complexity: simplicity and complexity often characterize less the object of study than our understanding of it. Maybe one day, neuroscience textbooks will finally start slimming down....

GILLES LAURENT

Division of Biology, California Institute of Technology, Pasadena, California 91125, USA

e-mail: laurentg@caltech.edu