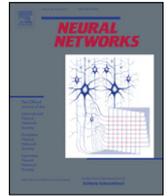




Contents lists available at SciVerse ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

2012 Special Issue

Using a million cell simulation of the cerebellum: Network scaling and task generality

Wen-Ke Li^{a,*}, Matthew J. Hausknecht^b, Peter Stone^b, Michael D. Mauk^a^a Center for learning and memory, Institute for Neuroscience, The University of Texas at Austin, 1 University Station C7000 Austin, TX 78712, USA^b Department of Computer Science, The University of Texas at Austin, 1616 Guadalupe, Suite 2.408, Austin, TX 78701, USA

ARTICLE INFO

Keywords:
Cerebellum
Eyelid conditioning
Cart–pole task

ABSTRACT

Several factors combine to make it feasible to build computer simulations of the cerebellum and to test them in biologically realistic ways. These simulations can be used to help understand the computational contributions of various cerebellar components, including the relevance of the enormous number of neurons in the granule cell layer. In previous work we have used a simulation containing 12000 granule cells to develop new predictions and to account for various aspects of eyelid conditioning, a form of motor learning mediated by the cerebellum. Here we demonstrate the feasibility of scaling up this simulation to over one million granule cells using parallel graphics processing unit (GPU) technology. We observe that this increase in number of granule cells requires only twice the execution time of the smaller simulation on the GPU. We demonstrate that this simulation, like its smaller predecessor, can emulate certain basic features of conditioned eyelid responses, with a slight improvement in performance in one measure. We also use this simulation to examine the generality of the computation properties that we have derived from studying eyelid conditioning. We demonstrate that this scaled up simulation can learn a high level of performance in a classic machine learning task, the cart–pole balancing task. These results suggest that this parallel GPU technology can be used to build very large-scale simulations whose connectivity ratios match those of the real cerebellum and that these simulations can be used guide future studies on cerebellar mediated tasks and on machine learning problems.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

The cerebellum is remarkable owing in part to the large number of granule cells that it contains. Estimates indicate that approximately half of the neurons in the human brain are cerebellar granule cells. A satisfying understanding of the cerebellum must therefore include a clear picture of the computational significance of this vast cell layer. Several factors combine to enhance the feasibility of analyzing the computational properties of the cerebellum and its cellular components: (i) the synaptic organization and synaptic physiology of the cerebellum are especially well characterized (Eccles, Ito, & Szentágothai, 1967; Ito, 1984) and (ii) there are several experimentally tractable behaviors that engage the cerebellum relatively directly, such as eyelid conditioning and adaptation of the vestibule-ocular reflex.

The relationship between eyelid conditioning and the cerebellum in particular has enabled both the construction and biologically relevant evaluation of computer simulations of the

cerebellum. Eyelid conditioning involves the paired presentation of a neutral conditioned stimulus (CS) such as a tone and a reinforcing unconditioned stimulus (US), typically an air puff directed at the eye or peri-orbital electrical stimulation. With many CS+US pairings the CS acquires the ability to elicit a conditioned response—the eyelids close in response to the tone. Previous work has revealed that the CS is conveyed to the cerebellum via mossy fiber inputs (Lewis, LoTurco, & Solomon, 1987) and the US by activation of climbing fiber inputs (Mauk, Steinmetz, & Thompson, 1986), and that output from the cerebellum via the anterior interpositus nucleus drives the expression of the learned responses (McCormick & Thompson, 1984). Eyelid conditioning can thus be used to evaluate a computer simulation of the cerebellum by providing the simulation with CS-like and US-like inputs over mossy fibers and climbing fibers respectively. The rich repertoire of well-characterized behavioral properties of eyelid conditioned then serves as a stringent test of the performance of a simulation.

We have previously shown that a biologically constrained simulation of the cerebellum containing 12,000 granule cells can replicate many (but not all) behavioral properties of eyelid conditioning (Medina, Garcia, Nores, Taylor, & Mauk, 2000). Although this simulation has been the source of new predictions that were born out with experimental tests it remains an open question

* Corresponding author. Tel.: +1 832 594 9904.

E-mail addresses: wenke.li@utexas.edu (W.-K. Li), mhauskn@cs.utexas.edu (M.J. Hausknecht), pstone@cs.utexas.edu (P. Stone), mmauk@mail.clm.utexas.edu (M.D. Mauk).

which (if any) limitations of the simulation are attributable to the relatively small number of granule cells. Toward the ability to address such issues, we report the initial development of a much larger simulation that contains over one million granule cells. This expansion is made feasible by parallel implementation on modern multi-processor graphics processing units (GPUs, e.g. Nvidia GTX580). Here, we compare the basic performance of the simulation to its smaller predecessor and use the larger simulation to begin to address the issue of task generality. While previous simulations have only been tested against the behavioral properties of eyelid conditioning, we have tested the million-cell simulation with eyelid conditioning and with a classic machine learning task: balancing a pole by moving a cart (Cart-pole task). We show that the larger simulation is able to replicate eyelid conditioning and show that it readily learns robust performance in the cart-pole task.

2. Methods

The new simulation is based on the original simulation of Buonomano and Mauk (1994) as modified later by Kalmbach, Voicu, Ohyama, and Mauk (2011), Medina et al. (2000) and Medina and Mauk (1999). The principle change is the nearly 100 fold increase in the number of granule cells, from 12,000 to 1,048,567. Consequently the divergence/convergence ratios of granule cell connectivity could be modified to more closely approximate the ratios observed in the cerebellum (Eccles et al., 1967; Ito, 1984).

2.1. Simulation connectivity

Fig. 1A shows the synaptic relationship among the cells in the cerebellum as implemented in the simulation. As with the previous simulation, the connectivity of the present simulation attempts to capture not only the numerical, divergence and convergence ratios, but also the known spatial relationships between the cell types. The algorithm that converts these constraints to the actual cell-by-cell connectivity of the network is identical to that in the previous simulation (Buonomano & Mauk, 1994; Medina et al., 2000). The only difference in connectivity between the 2 networks is the different connectivity ratios provided to the algorithm, which is discussed below.

The increase of granule cell population from 12,000 to 1,048,567 in the simulation enabled the observed numeric ratios of granule cell connectivity to be more closely approximated than was possible in the previous simulation. Fig. 1B compares the convergence–divergence ratios of connectivity between the previous simulation, the expanded simulation, and the observed ratios in cerebellum (Palkovits, Magyar, & Szentágothai, 1971a, 1971b). Most notably, the increase in granule cell numbers enabled much closer convergence ratio of granule cell to Purkinje cell synapses. Whereas the previous simulation can only achieve 1/160 of the observed ratio, the expanded simulation achieves 1/3 of the observed ratio. The same is true for convergence ratios of granule cell to basket cell connectivity. In addition, the expanded granule cell population allowed for much closer approximation of connectivity ratios between granule cells, Golgi cells, and mossy fibers. Most notably, the expanded simulation was able to achieve 1/2 of the divergence ratio of Golgi cell output to granule cells. The previous simulation suggested that the connectivity among these 3 types of cells is necessary to produce behaviors that require the cerebellum, thus it is crucial that we are able to closely approximate these ratios observed in the cerebellum.

The expansion of number of granule cells allowed us to more closely approximate the convergence ratio of granule cell–Purkinje cells. The simulation modeled a single strip of 32 Purkinje cells without overlapping dendrites, thus each Purkinje cell received a

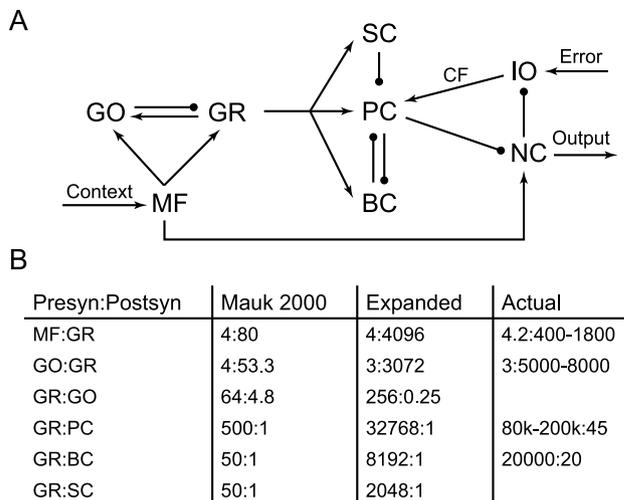


Fig. 1. Connectivity of cells in the cerebellum. A. Synaptic connections among cells. The mossy fibers are thought to carry information about the state of the world and climbing fibers are thought to carry teaching signals. MF: mossy fibers, GO: Golgi cells, GR: granule cells, SC: stellate cells, BC: basket cells, PC: Purkinje cells, IO: inferior olivary cells, NC: nucleus cells, CF: climbing fibers. Arrows indicate excitatory connections and round ends indicate inhibitory connections. B. Ratios of granule cell connectivity. Presyn:postsyn: presynaptic cell to post synaptic cell connectivity. Mauk 2000: previous smaller simulation. Expanded: expanded simulation discussed here. Actual: connectivity ratios observed in the cerebellum, with blank fields indicating unknown. The ratios are listed as convergent: divergent. For example, for mossy fiber output to granule cells observed in the cerebellum, each granule cell receives 4.2 mossy fiber inputs on average and each mossy fiber connects to 400–1800 granule cells.

unique set of 32,768 granule cell inputs. We decided to only model a single layer of Purkinje cells because it captured inputs from all granule cells and was more computationally efficient. We also modeled a strip of 128 basket cells and a strip of 512 stellate cells in the same manner as the Purkinje cells. Thus, each granule cell in our simulation output to exactly one Purkinje cell, one basket cell, and one stellate cell.

2.2. Representation of neurons

The cells in the simulation are implemented identically to that in the previous simulation, with the exception that instances where there are now a greater number of synaptic inputs required rescaling the synaptic inputs (decreasing the maximum synaptic conductance). Briefly, the cells are implemented using a single compartment leaky integrate and fire representation (Buonomano & Mauk, 1994; Medina et al., 2000). In this representation, membrane potential is calculated from synaptic conductances, leak conductances and membrane capacitance. These individual conductances are modeled based on known physiological data for each cell types. The leaky integrate-and-fire representation gains a great deal of computational efficiency by omitting explicit calculation of active (voltage-dependent) conductances. Instead, the influence of these conductances is approximated by (1) action potentials occur when the calculated membrane potential exceeds threshold—these spikes are broadcast as output to the appropriate follower neurons, and (2) threshold increases when an action potential occurs to emulate the absolute and relative refractory periods. After these spike-initiated increases, the threshold decays exponentially back to its normal level. In addition, the synaptic delay is modeled at one millisecond in the simulation. This representation can give rise to a phenomenological model that can be fine-tuned to match published physiological properties of each neuron type. This procedure yields representations that are highly computationally efficient and are suitable for studying the emerging network behavior.

2.3. Mossy fiber–granule–Golgi input network

Mossy fibers provide one of the two major inputs into the cerebellum. The mossy fibers are thought to carry information about the internal and external state of the world, such as limb positions, commands from motor and premotor cortices, tone stimuli, etc.. The mossy fibers diverge extensively onto the granule cell population in the cerebellar cortex. The granule cells then connect to Purkinje cells. Thus, signals from mossy fibers indirectly affect Purkinje cell activity. In addition, the granule cells excite Golgi cells and receive inhibitory inputs from Golgi cells, thus forming a recurrent feedback loop. It is thought that this mossy fiber–granule–Golgi input network performs input pattern separation and timing (Bullock, Fiala, & Grossberg, 1994; Fujita, 1982; Maex & De Schutter, 1998; Marr, 1969; Medina & Mauk, 2000; Moore & Choi, 1997). As discussed previously, the increase in the number of granule cells in the simulation enabled a much closer approximation of the convergence–divergence ratios observed among the three types of cells.

2.4. Climbing fiber inputs

The other major input to the cerebellar cortex is climbing fibers from inferior olivary cells. The climbing fibers make extensive synaptic connection to the Purkinje cell dendrites. Compared to the massive convergence (80,000:1) ratio of granule cell to Purkinje cell synapses, each Purkinje cell only receives input from one climbing fiber. The climbing fiber activity has been shown to be the signal that induces plasticity in the granule to Purkinje cell synapses (Ito, 1989; Ito & Kano, 1982; Lev-Ram, Mehta, Kleinfeld, & Tsien, 2003; Medina, Nores, & Mauk, 2002; Simpson, Wylie, & De Zeeuw, 2011).

2.5. Relationship between eyelid conditioning and cerebellum

The cerebellum is necessary for several well-characterized types of motor learning such as eyelid conditioning (Garcia, Steele, & Mauk, 1999; Mauk & Thompson, 1987; Perrett, Ruiz, & Mauk, 1993; Raymond, Lisberger, & Mauk, 1996), adaptation of the vestibular ocular reflex (DuLac, Raymond, Sejnowski, & Lisberger, 1995; Lisberger, 1988), and learning smooth pursuit eye movements (Li & Lisberger, 2011; Lisberger, 2010; Lisberger, Morris, & Tychsen, 1987). In eyelid conditioning, the animal is presented with a conditioning stimulus (CS, e.g., tone) for a fixed duration, and at the end of the tone a reinforcing unconditioned stimulus (US) such as a puff of air into eye or peri-orbital electrical stimulation is presented. After repeated presentation of the CS and US, the animal learns to close its eyelid prior to the onset of the US. What makes this learning useful for testing the performance of cerebellar simulations is the relationship between these stimuli and the inputs to the cerebellum. The presentation of the tone CS is conveyed to the cerebellum via the mossy fiber inputs—that is, mossy fiber inputs are necessary and sufficient to convey the CS (Lewis et al., 1987; Steinmetz, Lavond, & Thompson, 1989). Similarly, activation of climbing fiber inputs to the cerebellum is necessary and sufficient to convey the US (Mauk et al., 1986; Türker & Miles, 1986). In addition, recording studies have revealed how mossy fibers and climbing fibers respond to the CS and US, respectively (Aitkin & Boyd, 1978; Sears & Steinmetz, 1991). On the output side, previous work demonstrates that cerebellar output via activity of neurons in the anterior interpositus nucleus is necessary and sufficient to drive the expression of the learned eyelid responses (McCormick & Thompson, 1984). Combined, these factors reveal that cerebellar simulations can be rigorously tested with eyelid-conditioning-like inputs over the mossy fibers and climbing fibers. The well-characterized behavioral properties of eyelid conditioning then serve as the test bed for simulation performance.

2.6. Cerebellar plasticity involved in eyelid conditioning

Two sites of plasticity in the cerebellum are known to be involved in eyelid conditioning: climbing fibers control the induction of plasticity at granule cell to Purkinje cell synapses (Gilbert & Thach, 1977; Ito, 1989; Ito & Kano, 1982; Wang, Denk, & Häusser, 2000), and Purkinje cells appear to control the induction of plasticity at mossy fiber to nucleus synapses (Garcia & Mauk, 1998; Garcia et al., 1999; Kalmbach et al., 2010; Medina, Garcia, & Mauk, 2001; Medina & Mauk, 1999; Ohyama, Nores, Medina, Riusech, & Mauk, 2006; Perrett & Mauk, 1995; Pugh & Raman, 2006, 2008). The implementation of these rules for plasticity for the expanded simulation is identical to the previous simulation.

2.7. Parallel implementation

Owing to the expanded granule cell numbers, we found that a traditional single threaded implementation took around 600 s to process five seconds of simulated time with one-millisecond time steps, which would limit the simulation's usefulness. In order to exploit the modern multi-core processors, we switched our implementation to C++. This allowed us to use multithreading with OpenMP. We tested this implementation on an eight core Intel Xeon workstation. However, we could only achieve a 2× speed up, instead of 6–8× we were hoping for. By profiling the simulation to determine the performance-limiting factor we realized that memory bandwidth is a significant issue. We tallied the amount of data for granule cells, and found that each granule cell required 128 bytes of data, which meant 128 MB of data for a million granule cells. During each time step, all 128 MB of data have to be either read and/or written to. Thus, the memory bandwidth required for the simulation to compute in real time would be 128 GB/s.

On the other hand, we realized that calculating the granule cell activities mostly involved applying identical instructions to large arrays of data. This computation pattern matches the single instruction multiple data (SIMD) pattern particularly well. Modern vector processors such as the Nvidia graphics processing units (GPU) being developed for general computation purposes should excel at this computation. The GPU we used at the time (GTX275) had more than 150 GB/s of memory bandwidth and 240 cores. Utilizing the CUDA C programming extensions we were able to accelerate the simulation to 30 s for five seconds of simulation time, which is in the realm of the runtime we were aiming for. The final challenge was updating the activity of the cells according to their connectivity patterns. This is especially problematic because the connectivity among mossy fibers, Golgi cells, and granule cells is highly random. Consequently, the memory access patterns for these interactions are also very random, and memory latency becomes the primary limiting factor. Fortunately, there is no communication among granule cells, only the large divergence and convergence from mossy fibers and Golgi cells. In addition, there are only 1024 mossy fibers and Golgi cells, so that their action potentials can be stored entirely in the L1 cache of the GPU that has the low latency we needed. Currently, using an Nvidia Fermi GTX580 GPU with the above optimizations, our simulation takes 9 s to run five seconds of simulation time. We also tested how well the simulation scales across multiple GPUs by comparing the performance of the simulation on 1 GTX470 and 2 GTX470 GPUs. We found that 2 GTX470 GPUs were able to accelerate the simulation run time from 12 to 6 s for five seconds of simulation time, which was almost scaling linearly. Finally, we tested how well the performance of the algorithm would scale with fewer cells. We found that, on the GTX580, a 16 thousand cells simulation takes 0.5 s to perform five seconds of simulation time. However, if the scaling were linear, the 16 thousand cells simulation should only take 0.14 s. The sublinear scaling suggests that the performance is CPU bound for small number of cells.

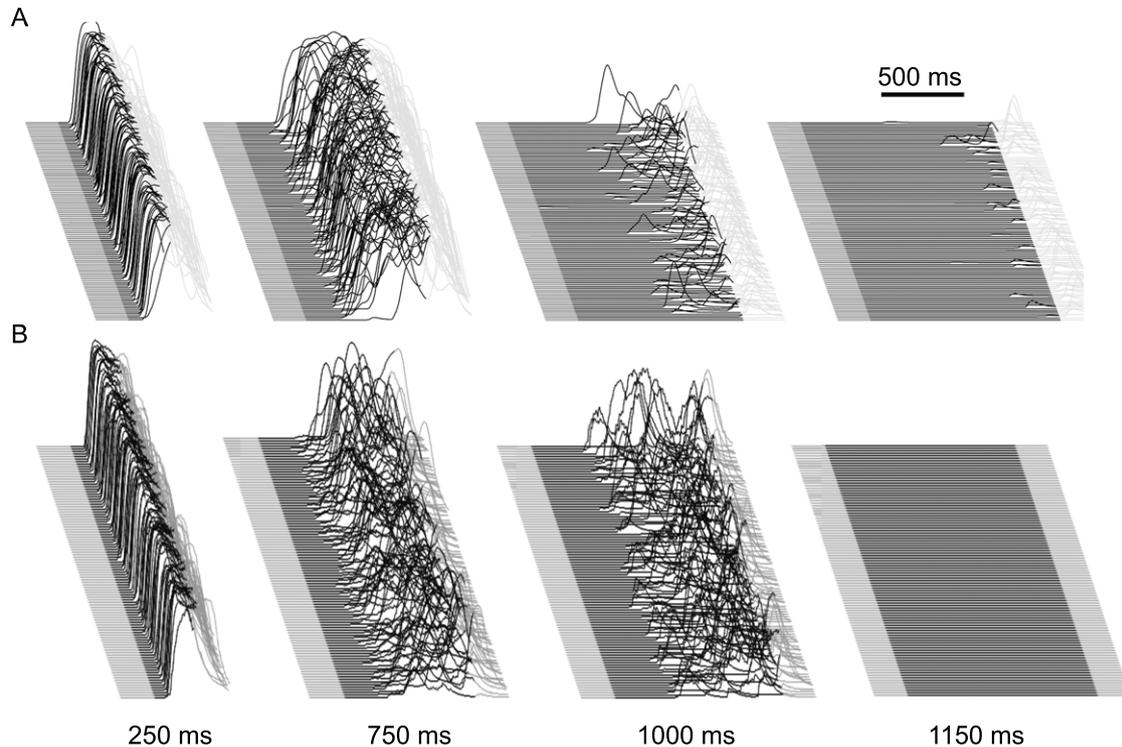


Fig. 2. A comparison of eyelid conditioning performance of the smaller (12,000 granule cell) simulation and the larger (1,048,576 granule cell) simulation. Each sweep is the eyelid response predicted by the output of the simulation, as conveyed by its deep nucleus neurons. Each panel shows 100 trials from well-trained simulations, where each upward deflection is the predicted (learned) closure of the eyelid. The black portion of each sweep indicates the time over which the mossy fiber inputs were active to mimic a conditioned stimulus and thus the upward deflection of the traces in the black portions shows a learned response by the simulation. A. Performance of the smaller simulation trained using inter-stimulus intervals ranging from 250 to 1150 ms. Robust conditioned responding is predicted by the simulation for the 250 and 750 ms intervals. Very poor responding is seen at the 1000 ms interval and essential no learning is evident at 1150 ms. B. By comparison, performance of the larger simulation over a similar range of inter-stimulus intervals. Like eyelid conditioning results from rabbits and the smaller simulation (data not shown) the larger simulation does not learn with an interval of 100 ms. Like rabbits and the smaller simulation, the larger simulation shows robust and well timed responses for 250 and 750 ms intervals. The larger simulation shows more robust learning at the 1000 ms interval than does its smaller predecessor, but falls short of rabbit performance by showing no learning at all at an interval of 1150 ms.

3. Results

3.1. Eyelid conditioning

To examine simulation performance, we tested its ability to emulate proper learning and response timing for eyelid conditioning. We compared the performance of the large simulation to the small simulation to examine the timing performance (Fig. 2). To mimic presentation of a tone CS a small subset of mossy fibers was made to fire in a way that is consistent with published peri-stimulus histograms of mossy fiber responses to auditory stimuli (Aitkin & Boyd, 1978). Briefly 20 of the 1024 mossy fibers had a tonic increase in activity in the presence of the tone CS and 30 of the 1024 mossy fibers had a phasic increase in activity during the “tone” CS. To implement the US input, each climbing fiber underwent a small depolarization sufficient to elicit a spike from the inferior olivary cells on most presentations. The summed output of the 8 deep nucleus cells, integrated over a time span of five ms was used at the “eyelid response” output of the simulation.

As a preliminary test, we first examined whether the simulation is capable of acquiring eyelid responses. Fig. 2B (left panel) shows that the large simulation was able to respond robustly to a CS-US interval of 250 ms. To test the ability of the larger simulation to mimic the learned timing of the responses, we next trained it using CS-US intervals of either 750 ms or 1000 ms. The simulation’s response after learning a 750 ms interval was delayed compared to the 250 ms interval response, but less delayed than the 1000 ms interval response (Fig. 2B). This data generally captures the timing behavior of animals, where the onset of the learned response is delayed depending on the CS-US interval.

Finally, we tested the large simulation on two long intervals at 1150 (Fig. 2B) and 1500 ms (data not shown). The simulation could not learn to either of the two intervals, which is not consistent with the rabbit data. These results are generally consistent with the results from the previous smaller simulation, with one exception: the smaller simulation shows learning with the 750 ms interval and unreliable responses with the 1000 ms interval whereas the larger simulation is capable of more robust responses with the 1000 ms interval. Thus, the expanded simulation shows similar behavior to the previous simulation, with an apparent improvement in its ability to mimic the rabbit data in terms of longer CS-US intervals. Subsequent work will focus on the differences in the simulations that make this improvement possible.

3.2. Cart-pole balancing

In order to examine the generalization of our simulation to other tasks, we choose to apply the simulation to a classic inverted pendulum balancing task (Anderson, 1989). The inverted pendulum rests on a cart that can move on a one dimensional track. The objective is to balance the inverted pendulum for as long as possible by applying forces that move the cart left or right. This task is analogous to balancing tasks that require the cerebellum (Morton & Bastian, 2004). The task involves the coordination of multiple forces to achieve a single task, which is very distinct from eyelid conditioning where there is only a single force (closure of the eyelid) to control. Thus, this was an excellent problem to test the generality of our simulation.

We have connected the simulation to the cart-pole domain as illustrated by Fig. 3A. The state of the cart-pole world, such

as pole angle (Fig. 3B) and pole velocity is transmitted by mossy fibers to the simulation. The parallel fibers from the input network (granule cell axons) then connect to two independent output networks. Each network is identical to the network used in eyelid conditioning, containing a full set of Purkinje, basket, and stellate cells that receive input from the entire granule cell population. Each network also contains the full complement of nucleus cells and inferior olivary cells. Each network has identical parallel fiber–Purkinje cell and mossy fiber–Nucleus cell plasticity rules as that used in eyelid conditioning. Each output network controls pushing the cart in a single direction. The force of the output is simply extracted as the voltage of a leaky integrate model that received all nucleus cell spikes without any delay (as long as the scaling factors are appropriate the actual kinetics of the model had no significant impact). The forces from both output networks are subtracted from each other to yield the net force acting on the cart (Fig. 3C and D, bottom panel). Finally, the error signals about the failure of maintaining pole balance are transmitted back to the inferior olivary cells of each output network without any delay. The output network that is responsible for pushing the cart to the left receives an error when the pole falls to the left, and vice versa. Thus, we assume no explicit communication between the two output networks. With this setup, we explored the various encoding schemes to examine if the simulation can perform this task without tuning the simulation itself.

We first assumed a very simple binary firing rate encoding scheme for the cart–pole world to explore if the simulation can learn appropriately by tweaking the error signals. We chose to encode 3 variables: pole angular position, angular velocity, and angular acceleration. For simplicity we did not place any physical limits on the velocity or the position of the cart, thus we did not encode these variables. We randomly chose 30 mossy fibers to encode each of the pole variables, out of 1024 mossy fibers. For example, the 30 mossy fibers for pole position are divided into 3 groups of 10. Each of the 3 groups has a preferred pole angle range from the upward midline. Each group has two firing rates: a baseline rate and a response rate that is higher than baseline. When the pole is in the preferred angle range of that group the mossy fibers fire at the response rate. Otherwise the mossy fibers in the group fire at baseline rate. The preferred angle ranges are divided into 3 parts: left, right, and middle, corresponding to when the pole angle is less than $-.0025$ radians, between $-.0025$ and $.0025$ rad, and greater than $.0025$ radians. The pole angular velocity and angular acceleration are encoded similarly. Using this scheme, we next explored the timing of the error signal onset.

Considering that there is a finite limit on the forces that can be generated to push the cart, the physics of the system defines a certain pole angle (symmetrical to the upward midline on each side) where the pole is no longer recoverable. We examined the timing of the error signal relative to this point of no return. We discovered that when the error signal is given after the point of no return, the simulation was able to learn initially to balance the pole. As shown in Fig. 4A, the simulation was able to learn to balance the pole in the 3rd trial, where the pole stayed close to vertical throughout the trial. However, after that successful trial, the simulation fails to retain its performance. After a few more trials, the simulation is able to learn again to balance the pole, but again fails to retain its performance. This appeared to be similar to extinction we observed in eyelid conditioning. When we examined the output network again, we realized that after learning, the nucleus cells increase their firing rate during a successful balancing trial. However, this increase in nucleus cell activity inhibits the inferior olivary cells, and disrupts their equilibrium firing rates, which is a signal to extinguish the responses (Medina et al., 2002). The failure of this system suggested that the inferior olivary cells must maintain a certain level activity even when the nucleus cells are responding appropriately.

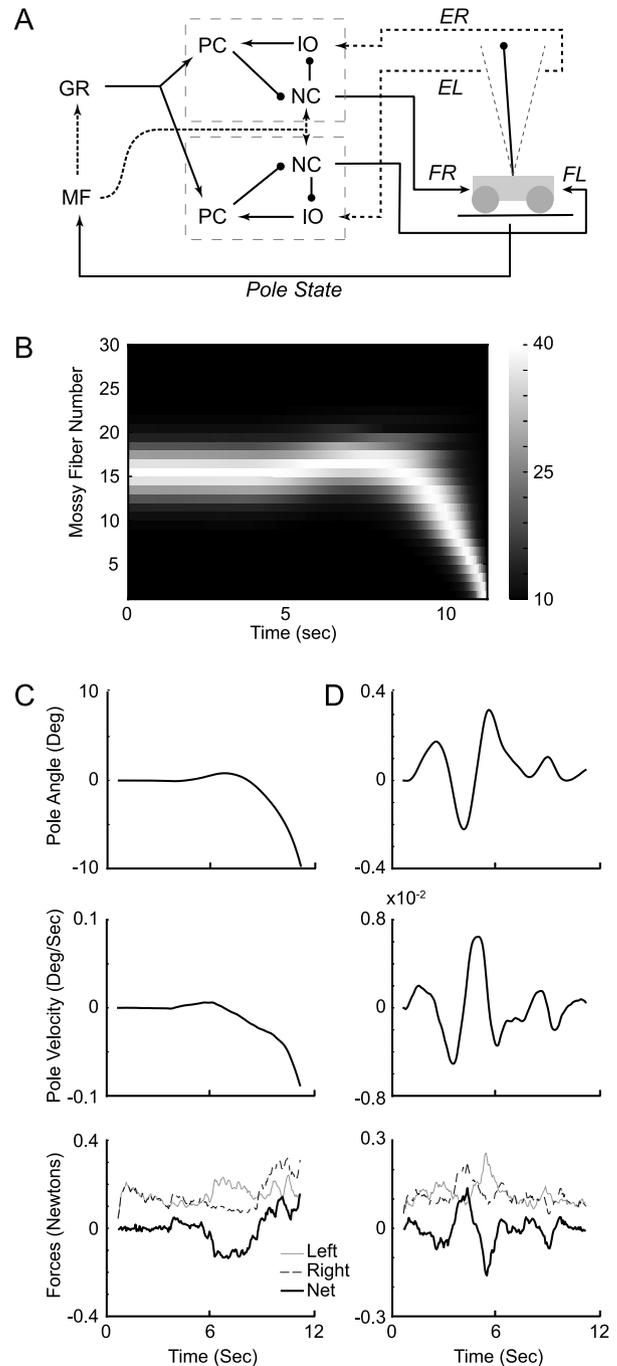


Fig. 3. The simulation setup as applied to the cart–pole balancing problem. A. Schematic of the simulation setup for cart–pole balancing. Information about the pole is encoded as mossy fiber inputs (MF). The mossy fibers indirectly connect to Purkinje cells (PC) through granule cells (GR). Unlike in eyelid conditioning, there are two sets of Purkinje cells, nucleus cells (NC), and inferior olivary cells (IO). The output of nucleus cells in each set is responsible for pushing the cart in one direction (FR: force right, FL: force left). When the pole exceeds a certain position threshold (gray lines on either side of the pole), error signals (ER, error right, EL, error left) are sent to the associated inferior olivary cells. B. Pole angle population mossy fiber firing rate evolution over time for pole position of a failed train as in (C). Scale bar is in Hz. The mossy fibers shown are using the Gaussian encoding (see Fig. 4 and results). C. Evolution of pole states and simulation output over time for a single trial. Top: angle difference in degrees between the pole and upright. Positive is toward the left and negative is toward the right. Middle: $d(\text{angle})/d(t)$ in degrees/second. Same signs as pole angle. Bottom: the output forces on the cart from the two output networks from the simulation. Left: force pushing the cart to the left, Right: force pushing the cart to the right, Net: the left and right forces are subtracted from each other to arrive at the actual net force acting on the cart. D. Same as in (C), for a successful trial. Note the difference in scales.

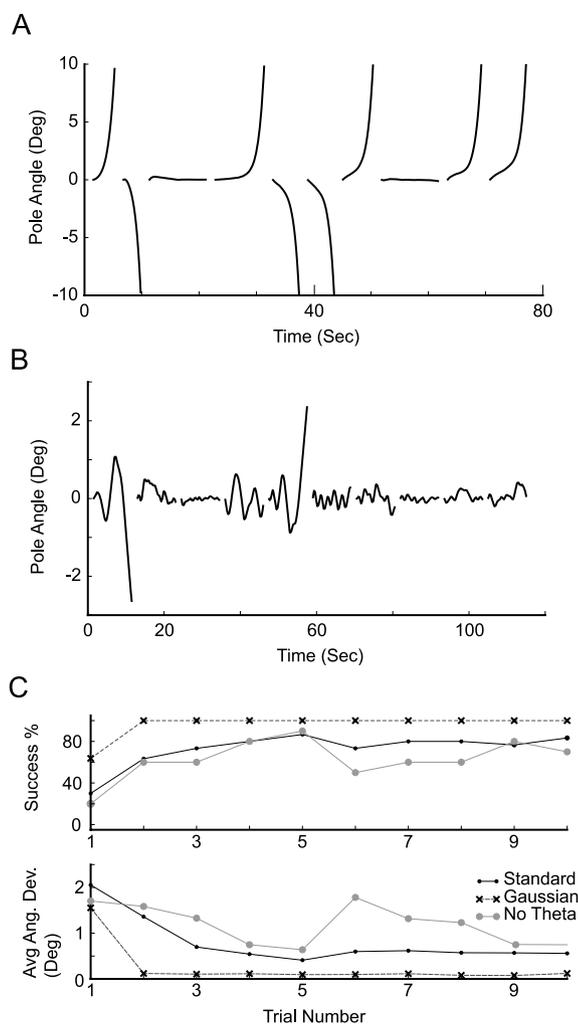


Fig. 4. The simulation is generalizable from eyelid conditioning to cart–pole balancing. A. Pole balance performance where the error is only sent after point of no return. Pole position relative to upright is plotted against time. Each line represents a trial, which terminates at the point of failure, or exceeding 10 s. B. Pole balance performance where the error rate is proportional to pole position deviation from upright. C. Pole balance performance for 3 mossy fiber encoding schemes. Top: success percentage as a function of trial number for 10 repetitions. Success is defined as keeping the pole balanced for 10 s. Bottom: average angular deviation of pole from upright for each trial. Standard: binary encoding of pole position, pole velocity, and pole acceleration. Gaussian: Gaussian encoding of the same variables. No theta: Gaussian encoding of pole velocity, acceleration, and a linear combination of pole velocity and acceleration, see Section 3.2 for more details.

Thus, we tried to encode the error signal as proportional to the pole angle. Specifically, the probability that the inferior olivary cells are stimulated is proportional to the pole angle relative to midline. This ensured that even during successful trials where the pole is balanced by the nucleus cell output, inferior olivary cells could still be active to prevent extinction. Fig. 4B illustrates that this encoding scheme was able to retain the ability to balance the pole.

We next explored the encoding schemes for the world state variables. We first explored Gaussian encoding of the pole angular position (Fig. 3B), angular velocity, and angular acceleration. Again, we use pole angular position as an example to illustrate the encoding scheme. Each of the 30 mossy fibers for angular position was assigned a preferred angle where the firing rate is maximal. The actual firing rate of the mossy fiber is then dependent on how far the angular position of the pole is from the preferred angle, transformed with the Gaussian function. We find that the

simulation could learn robustly in the presence of both binary rate encoding and Gaussian rate encoding (Fig. 4C). We also explored an encoding scheme similar to that observed in VOR (Lisberger & Pavelko, 1986), where pole angular velocity, angular acceleration, and an angular pulse velocity (linear combination of velocity and acceleration) was encoded. The pole angular position was not encoded in this scheme. As shown in Fig. 3D (no theta), the simulation could not learn very robustly with this encoding scheme. This would indicate that our knowledge about how the cerebellum achieves coordination between multiple output networks is still incomplete.

4. Discussion

We have demonstrated the ability to increase the scale of a well-characterized computer simulation of the cerebellum. Through the application of GPU parallel processing the number of granule cells in this simulation can be increased from 12,000 to over one million. In doing so, the execution speed has been maintained at a level that permits sensible analysis and progress. On a high performance workstation the smaller simulation runs at real time or slightly better—four seconds of execution to simulate five seconds. Although the larger simulation that we characterized here implements almost 100-fold more granule cells, it requires approximately twice the execution time using the Nvidia Fermi GTX580 GPU. Moreover, in the current work we find that execution time decreases proportionally with 2 GPU boards. With this scaling factor we estimate that a simulation that implements 10 million granule cells can run at 0.5× to 1× real time on a standard workstation computer with eight next generation GPUs, if we are to assume that the scaling can be maintained for 8 GPUs, that the next generation GPUs are faster than the current models, and there are no additional overhead when dealing with 10 million granule cells. If we are to take that a para-sagittal module of the cerebellum involves around 150 million granule cells, it is imaginable that we will be able to build such a simulation within the next 2–4 generations of computer hardware. These advances highlight that it is now possible to address with computer simulations the question of the computational value of the cerebellum's enormous layer of granule cells.

The ability to expand the number of granule cells by 100 fold over the original simulation has the important advantage of allowing a closer approximation of the connectivity ratios observed in the real cerebellum. The motivation for better approximating connectivity ratios is to begin exploring the question of the computational properties of the unique connectivity observed with cerebellar granule cells. Given that the cerebellar granule cells account for over 50% of all neurons in the entire human central nervous system, and that this characteristic of very large numbers is remarkably well conserved in evolution, it is conceivable that there are underlying characteristic computational functions behind such connectivity. In fact, Marr (1969) in the first coherent theory about the cerebellar cortical computation, hypothesized about the role of this connectivity in mossy fiber input pattern separation. Understanding the computational functions of this connectivity will be a significant step forward in our understanding of how the cerebellum functions.

As a first result, the expanded simulation's ability to learn more robustly to longer inter-stimulus intervals could be due increase in reliability of having a larger granule population output to Purkinje cells. On the other hand, the lack of improvement in the timing response to both 750 and 1000 ms intervals also points to the possibility that there are features in the connectivity of the input network that we do not yet fully appreciate.

We have also used this larger cerebellar simulation to begin to explore the issue of task generality. The different areas of the

cerebellum are known to perform different motor functions such as balance and fine movements (Ito, 1984). The network architecture of the cerebellum is remarkably uniform throughout its entire structure (Eccles et al., 1967; Ito, 1984). In addition, the network connectivity is evolutionarily well conserved in mammals, and the principal features of parallel fibers, Purkinje cells, climbing fibers are observed in the cerebellum of other vertebrates. This suggests that this particular network architecture performs a characteristic computation that can be applied to a variety of tasks. However, it is certainly not obvious that we could elucidate this computation, if it exists, by studying eyelid conditioning. Considering that cart–pole balancing is arguably a completely different task from eyelid conditioning, it was not expected that our simulation—that was constructed entirely from understanding eyelid condition—could work for cart–pole balancing. The fact that the simulation, without tuning its intrinsic parameters, succeeded in performing this task suggests that such a characteristic computation for the cerebellar network architecture exists. This common computation could be determining the timing and the amplitude (Kreider & Mauk, 2010) of the commands necessary for correct motor output. In addition, this computation would need to include mechanisms to adapt to new motor tasks as well as fluctuations in the motor output system. The success of the simulation suggests that we can study this computation with eyelid conditioning.

Applying the simulation to cart–pole balancing also let us to begin to address a feature of more complex motor movements that is missing in eyelid conditioning: multiple muscle coordination. It is known that the cerebellum is necessary for smooth coordination between multiple muscles. It is thought that each muscle is driven by a specific part of the cerebellum (Ito, 1984). We examined this coordination issue in the simulation with cart–pole by starting with a naïve assumption, that there is no explicit communication between the two output networks that control the two forces on the cart. The simulation's success at cart–pole suggests that this naïve assumption is sufficient for this particular task. This does not exclude the possibility that such communication might make the simulation perform better in this task, and that such communication might be necessary for more complex coordination tasks. The expanded simulation provides us with the tool to explore further into this question, as well as the computational properties of large granule cell populations and their network connectivity.

References

- Aitkin, L. M., & Boyd, J. (1978). Acoustic input to the lateral pontine nuclei. *Hearing Research*, 1(1), 67–77.
- Anderson, C. W. (1989). Learning to control an inverted pendulum using neural networks. *IEEE Control Systems Magazine*, 9, 31–37.
- Bullock, D., Fiala, J. C., & Grossberg, S. (1994). A neural model of timed response learning in the cerebellum. *Neural Networks*, 7(94), 1101–1114.
- Buonomano, D. V., & Mauk, M. D. (1994). Neural network model of the cerebellum: temporal discrimination and the timing of motor responses. *Neural Computation*, 6(1), 38–55. <http://dx.doi.org/10.1162/neco.1994.6.1.38>
- DuLac, S., Raymond, J. L., Sejnowski, T. J., & Lisberger, S. G. (1995). Learning and memory in the vestibulo-ocular reflex. *Annual Review of Neuroscience*, 18, 409–441.
- Eccles, J. C., Ito, M., & Szentágothai, J. (1967). *The cerebellum as a neuronal machine*. New York: Springer-Verlag (p. 335).
- Fujita, M. (1982). Adaptive filter model of the cerebellum. *Biological Cybernetics*, 206, 195–206.
- Garcia, K. S., & Mauk, M. D. (1998). Pharmacological analysis of cerebellar contributions to the timing and expression of conditioned eyelid responses. *Neuropharmacology*, 37, 471–480.
- Garcia, K. S., Steele, P. M., & Mauk, M. D. (1999). Cerebellar cortex lesions prevent acquisition of conditioned eyelid responses. *Journal of Neuroscience*, 19(24), 10940–10947.
- Gilbert, P. F. C., & Thach, W. T. (1977). Purkinje cell activity during motor learning. *Brain Research*, 128, 309–328.
- Ito, M. (1984). *The cerebellum and neural control*. New York: Raven Press (p. 580).
- Ito, M. (1989). Long-term depression. *Annual Review of Neuroscience*, 12, 85–102.
- Ito, M., & Kano, M. (1982). Long-lasting depression of parallel fiber–Purkinje cell transmission induced by conjunctive stimulation of parallel fibers and climbing fibers in the cerebellar cortex. *Neuroscience Letters*, 33, 253–258.
- Kalmbach, B. E., Davis, T., Ohyama, T., Riusech, F. A., Nores, W. L., & Mauk, M. D. (2010). Cerebellar cortex contributions to the expression and timing of conditioned eyelid responses. *Journal of Neurophysiology*, 103(4), 2039–2049. <http://dx.doi.org/10.1152/jn.00033.2010>.
- Kalmbach, B. E., Voicu, H., Ohyama, T., & Mauk, M. D. (2011). A subtraction mechanism of temporal coding in cerebellar cortex. *Journal of Neuroscience*, 31(6), 2025–2034. <http://dx.doi.org/10.1523/JNEUROSCI.4212-10.2011>.
- Kreider, J. C., & Mauk, M. D. (2010). Eyelid conditioning to a target amplitude: adding how much to whether and when. *Journal of Neuroscience*, 30(42), 14145–14152. <http://dx.doi.org/10.1523/JNEUROSCI.3473-10.2010>.
- Lev-Ram, V., Mehta, S. B., Kleinfeld, D., & Tsien, R. Y. (2003). Reversing cerebellar long-term depression. *Proceedings of the National Academy of Sciences of the United States of America*, 100(26), 15989–15993.
- Lewis, J. L., LoTurco, J. J., & Solomon, P. R. (1987). Lesions of the middle cerebellar peduncle disrupt acquisition and retention of the rabbit's classically conditioned nictitating membrane response. *Behavioral Neuroscience*, 101(2), 151–157. <http://dx.doi.org/10.1037/0735-7044.101.2.151>.
- Li, J. X., & Lisberger, S. G. (2011). Learned timing of motor behavior in the smooth eye movement region of the frontal eye fields. *Neuron*, 69(1), 159–169. <http://dx.doi.org/10.1016/j.neuron.2010.11.043>.
- Lisberger, S. G. (1988). The neural basis for learning of simple motor skills. *Science*, 242(80), 728–735.
- Lisberger, S. G. (2010). Visual guidance of smooth-pursuit eye movements: sensation, action, and what happens in between. *Neuron*, 66(4), 477–491. <http://dx.doi.org/10.1016/j.neuron.2010.03.027>.
- Lisberger, S. G., Morris, E. J., & Tychsen, L. (1987). Visual motion processing and sensory-motor integration for smooth pursuit eye movements. *Annual Review of Neuroscience*, 10, 97–129. <http://dx.doi.org/10.1146/annurev.neuro.10.1.97>.
- Lisberger, S. G., & Pavelko, T. A. (1986). Vestibular signals carried by pathways subserving plasticity of the vestibulo-ocular reflex in monkeys. *Journal of Neuroscience*, 6, 346–354.
- Maex, R., & De Schutter, E. (1998). Synchronization of Golgi and granule cell firing in a detailed network model of the cerebellar granule cell layer. *Journal of Neurophysiology*, 80, 2521–2537.
- Marr, D. (1969). A theory of cerebellar cortex. *Journal of Physiology (London)*, 202(2), 437–470.
- Mauk, M. D., Steinmetz, J. E., & Thompson, R. F. (1986). Classical conditioning using stimulation of the inferior olive as the unconditioned stimulus. *Proceedings of the National Academy of Sciences of the United States of America*, 83, 5349–5353.
- Mauk, M. D., & Thompson, R. F. (1987). Retention of classically conditioned eyelid responses following acute decerebration. *Brain Research*, 403(1), 89–95. [http://dx.doi.org/10.1016/0006-8993\(87\)90126-0](http://dx.doi.org/10.1016/0006-8993(87)90126-0).
- McCormick, D. A., & Thompson, R. F. (1984). Cerebellum: essential involvement in the classically conditioned eyelid response. *Science*, 223(80), 296–299.
- Medina, J. F., Garcia, K. S., & Mauk, M. D. (2001). A mechanism for savings in the cerebellum. *Journal of Neuroscience*, 21(11), 4081–4089.
- Medina, J. F., Garcia, K. S., Nores, W. L., Taylor, N. M., & Mauk, M. D. (2000). Timing mechanisms in the cerebellum: testing predictions of a large-scale computer simulation. *Journal of Neuroscience*, 20(14), 5516–5525.
- Medina, J. F., & Mauk, M. D. (1999). Simulations of cerebellar motor learning: computational analysis of plasticity at the mossy fiber to deep nucleus synapse. *Journal of Neuroscience*, 19(16), 7140–7151.
- Medina, J. F., & Mauk, M. D. (2000). Computer simulation of cerebellar information processing. *Nature Neuroscience*, 3, 1205–1211. <http://dx.doi.org/10.1038/81486>.
- Medina, J. F., Nores, W. L., & Mauk, M. D. (2002). Inhibition of climbing fibres is a signal for the extinction of conditioned eyelid responses. *Nature*, 416, 330–333.
- Moore, J. W., & Choi, J. S. (1997). Conditioned response timing and integration in the cerebellum. *Learning & Memory*, 4, 116–129. <http://dx.doi.org/10.1101/lm.4.1.116>.
- Morton, S. M., & Bastian, A. J. (2004). Cerebellar control of balance and locomotion. *Neuroscientist*, 10(3), 247–259. <http://dx.doi.org/10.1177/1073858404263517>.
- Ohyama, T., Nores, W. L., Medina, J. F., Riusech, F. A., & Mauk, M. D. (2006). Learning-induced plasticity in deep cerebellar nucleus. *Journal of Neuroscience*, 26(49), 12656–12663. <http://dx.doi.org/10.1523/JNEUROSCI.4023-06.2006>.
- Palkovits, M., Magyar, P., & Szentágothai, J. (1971a). Quantitative histological analysis of the cerebellar cortex in the cat, II, cell numbers and densities in the granular layer. *Brain Research*, 32(1), 15–30.
- Palkovits, M., Magyar, P., & Szentágothai, J. (1971b). Quantitative histological analysis of the cerebellar cortex in the cat, III, structural organization of the molecular layer. *Brain Research*, 34, 1–18.
- Perrett, S. P., & Mauk, M. D. (1995). Extinction of conditioned eyelid responses requires the anterior lobe of cerebellar cortex. *Journal of Neuroscience*, 15(3), 2074–2080.
- Perrett, S. P., Ruiz, B. P., & Mauk, M. D. (1993). Cerebellar cortex lesions disrupt learning-dependent timing of conditioned eyelid responses. *Journal of Neuroscience*, 13(4), 1708–1718.
- Pugh, J. R., & Raman, I. M. (2006). Potentiation of mossy fiber EPSCs in the cerebellar nuclei by NMDA receptor activation followed by postinhibitory rebound current. *Neuron*, 51, 113–123.

- Pugh, J. R., & Raman, I. M. (2008). Mechanisms of potentiation of mossy fiber EPSCs in the cerebellar nuclei by coincident synaptic excitation and inhibition. *Journal of Neuroscience*, 28(42), 10549–10560.
- Raymond, J. L., Lisberger, S. G., & Mauk, M. D. (1996). The cerebellum: a neuronal learning machine? *Science*, 272(80), 1126–1131.
- Sears, L. L., & Steinmetz, J. E. (1991). Dorsal accessory inferior olive activity diminishes during acquisition of the rabbit classically conditioned eyelid response. *Brain Research*, 545, 114–122.
[http://dx.doi.org/10.1016/0006-8993\(91\)91276-7](http://dx.doi.org/10.1016/0006-8993(91)91276-7).
- Simpson, J. I., Wylie, D. R. W., & De Zeeuw, C. I. (2011). More on climbing fiber signals and their consequence(s). *Behavioural Brain Research*, 19(03), 496–498.
<http://dx.doi.org/10.1017/S0140525X00081991>.
- Steinmetz, J. E., Lavond, D. G., & Thompson, R. F. (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.
- Türker, K. S., & Miles, T. S. (1986). Climbing fiber lesions disrupt conditioning of the nictitating membrane response in the rabbit. *Brain Research*, 363, 376–378.
- Wang, S. S., Denk, W., & Häusser, M. (2000). Coincidence detection in single dendritic spines mediated by calcium release. *Nature Neuroscience*, 3(12), 1266–1273. <http://dx.doi.org/10.1038/81792>.