Input minimization: a model of cerebellar learning without climbing fiber error signals

Thomas J. Anastasio

Department of Molecular and Integrative Physiology; Beckman Institute,

University of Illinois at Urbana/Champaign, 405 North Mathews Ave, Urbana, IL 61801, USA

Received 16 August 2001; accepted 6 September 2001

The cerebellum is critical for motor learning. Current cerebellar learning models follow the Marr/Albus paradigm, in which climbing fibers provide error signals that shape plastic synapses between parallel fibers and Purkinje cells. However, climbing fibers have slow and largely random discharge, and seem unlikely to provide error signals with resolution sufficient to guide cerebellar learning. Parallel fibers carry error signals and could direct the plasticity of their own synapses, but the error signals are carried along with other signals. This report presents the new input minimization (InMin) model, in which Purkinje cells reduce error by minimizing their overall parallel fiber input. The slowly, randomly firing climbing fiber provides only synchronization pulses. InMin offers an alternative that can unify cerebellar findings. *NeuroReport* 12:3825–3831 © 2001 Lippincott Williams & Wilkins.

Key words: Cerebellum; Computational modeling; Learning algorithm; Motor learning; Plasticity; Reinforcement learning; Unsupervised learning

INTRODUCTION

The cerebellum plays a central role in motor learning. A great deal is known about the cerebellum [1–3], but its learning mechanism remains obscure. This report describes a new model of the computation that underlies cerebellar learning. The new model is used to simulate adaptation of the vestibulo-ocular reflex (VOR).

Purkinje cells are the principal cells of the cerebellum. They are excited by the parallel fiber projections of granule cells, and by climbing fiber projections from the inferior olive. Each Purkinje cell receives input from $>100\,000$ parallel fibers but from only one climbing fiber. Parallel fibers produce conventional, simple spike discharges from Purkinje cells, whereas climbing fibers produce prolonged discharges known as complex spikes. Neither simple nor complex spikes are modeled explicitly here. The focus is on the cerebellar learning mechanism.

The VOR stabilizes the retinal image by making eye rotations that accurately counterbalance head rotations [4]. If the VOR is inaccurate, then images slip over the retina during head rotation. The flocculus, which is part of the vestibulo-cerebellum [3], maintains VOR accuracy and mediates VOR adaptation [4]. Granule cells in the flocculus are excited by mossy fibers that carry signals related to VOR performance. Some mossy fibers carry a signal proportional to retinal slip error, with a delay of 0.1 s [5]. Many other mossy fibers carry vestibular signals to the flocculus that are notable for their phase diversity [6]. These signals are relayed by parallel fibers to floccular Purkinje cells, which respond and then modulate the VOR by inhibiting vestibular nucleus neurons [4]. The flocculus can adapt the VOR up or down, and VOR adaptation is associated with adjustments in the responses of floccular Purkinje cells [7–9]. Purkinje cell response adjustment may be brought about through modification of plastic synapses from parallel fibers and inhibitory interneurons. Stellate interneurons receive parallel fiber input and inhibit Purkinje cells. Basket and Golgi cerebellar interneurons are present in some but not all vertebrates [3] and therefore represent specializations that may be ignored in models of basic cerebellar function.

Most models of cerebellar learning follow the Marr/Albus paradigm [10–12], which assumes that plastic synapses are modified through a supervised mechanism using error signals provided by climbing fibers. Support for this assumption is unconvincing. Climbing fiber background discharge can be modulated during sensorimotor behaviors such as the VOR [13]. This modulation is thought [14] to represent a retinal slip error signal that drives VOR adaptation. However, climbing fiber discharge rates are low (about 1 Hz) and have a large random component [15]. Recent analysis [16] reveals that complex spike modulation is weak (2% of simple spike modulation) and discernable only after signal averaging so extensive that it is unlikely to occur in the real cerebellum. The same analysis reveals that complex spike modulation is better correlated with eye velocity than with retinal slip [16]. These results deepen the concern that climbing fibers lack error signals of sufficient precision and temporal resolution to guide VOR adaptation [6]. Numerous lines of evidence suggest more generally that climbing fiber spikes signal events such as the start of movement or unexpected input

(see [17] for review). These disparate views of climbing fiber function, set against the random climbing fiber background discharge, need to be reconciled.

The error signals that guide cerebellar learning could be carried by parallel fibers themselves. Due to the incredible abundance of parallel fibers in the cerebellum [1–3], the resolution at which they could encode signals, including error signals, is potentially very high. Because parallel fibers carry error signals along with other signals, overall parallel fiber activity should decrease as error decreases. This report presents a new cerebellar learning algorithm, the input minimization (InMin) algorithm, which trains a network of model Purkinje cells to reduce error by minimizing their input from parallel fibers. Purkinje cell response adjustments due to VOR adaptation are similar in the model and in the real cerebellum. Climbing fiber spikes in InMin provide synchronization pulses that can be interpreted as 'learn now' signals, and are consistent with various views of climbing fiber function. InMin offers an alternative to the Marr/Albus paradigm that can unify findings on cerebellar neurophysiology.

MATERIALS AND METHODS

Model: Briefly, InMin trains individual Purkinje cells to develop responses that are specialized for specific temporal segments of the input, and to adjust the amplitudes of those responses to produce a desired combined output. Climbing fiber spikes provide 'learn now' signals that synchronize training of Purkinje cell response time and amplitude. Adjustments in Purkinje cell responses involve modification of excitatory and inhibitory synapses onto Purkinje cells from parallel fibers and stellate cells, respectively. The modifications are made so as to decrease the overall number of active parallel fibers.

The structure of the InMin model is based on cerebellar anatomy [1–3]. Figure 1 depicts a floccular microzone that contributes to VOR adaptation. A microzone is composed of one climbing fiber and the set of Purkinje cells it excites [2]. The model employs four microzones; it works well with one to a dozen. It receives input $x(t)$ and learns to make actual output $y(t)$ match desired output $z(t)$. Input and desired output signals are zero offset sinusoids representing a vestibular head-rotational velocity signal and an eye-rotational velocity command in a unilateral simplification of the VOR. The input has amplitude one, and the model learns to produce desired outputs having amplitude one (normal VOR), zero (down adaptation), or two (up adaptation). Error $e(t)$ is the difference between desired and actual output: $e(t)=z(t)-y(t)$. The error corresponds to retinal slip, and is available only in the light (dark/light switch closed).

The VOR is simulated at frequency $f=5\,\mathrm{Hz}$ (the 0.2 s period of the VOR cycle is divided into 100 time steps t). Visual following mechanisms are ineffective at 5 Hz [13]. Therefore, VOR performance is practically the same in darkness or light at this frequency. InMin works equally well at high and low frequencies.

To economize on elements, mossy fiber signals take both positive and negative values. The model receives input from mossy fibers m_i ($i=1,...,9$). Mossy fiber one carries a time delayed version of the error: $m_1=e(t-d)$, where $d=0.1\,\mathrm{s}$. Mossy fibers two through nine carry phase shifted

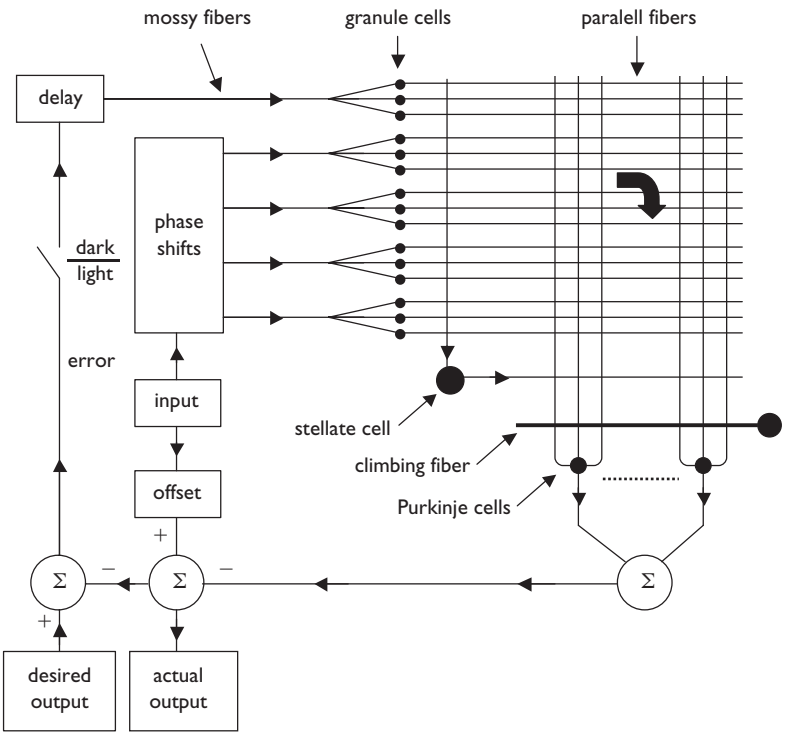

Fig. 1. One microzone in the InMin model of cerebellar learning. A microzone is composed of one climbing fiber and the set of Purkinje cells it excites. The number of microzones, and the number of Purkinje cells in each microzone, can be varied (see text). The responses of Purkinje cells from all microzones are combined to form the output, and all microzones receive the same parallel fiber input.

versions of vestibular input $x(t)$, where ϕ_i varies in steps of 22.5° over a realistic range from 67.5° lead to 90° lag (in phase with head position). The signals carried by mossy fibers m_2 through m_9 at time t are: $m_i(t) = \sin(2\pi ft + \phi_i)$. Eight phase shifts give relatively smooth Purkinje cell responses. Phase shift is continuous in the brain [6].

Each mossy fiber diverges to contact a separate, non-overlapping subset of granule cells g_j ($j = 1,...,1200$). Granule cells 1 through 400 encode delayed error, while granule cells 401 through 1200 encode the eight phase shifted vestibular inputs. Mossy fiber m_1 contacts 400 granule cells, and mossy fibers m_2 through m_9 each contact 100 granule cells. Granule cells are modeled as binary, threshold elements for simplicity. Half the granule cells have positive, and the other half negative, thresholds. For granule cell g_j with positive threshold θ^+ receiving input from mossy fiber m_i: $g_j(t) = \{1$ if $m_i(t) > \theta^+$; 0 otherwise$\}$. Similarly, for granule cell g_j with negative threshold θ^-: $g_j(t) = \{1$ if $m_i(t) < \theta^-$; 0 otherwise$\}$. Thresholds are evenly spaced between ± 2 for the 400 granule cells encoding error, and between and ± 1 for the subsets of 100 granule cells encoding phase shifted vestibular inputs. Together, each subset of granule cells encodes the magnitude of its mossy fiber input at any time according to the number of active granule cells in the subset.

Each granule cell relays its binary value through a parallel fiber, and every parallel fiber projects to all Purkinje cells p_k ($k = 1,...,24$). Each microzone contains six Purkinje cells. InMin works well with up to 20 Purkinje cells per microzone. The parallel-Purkinje weights contained in matrix $\boldsymbol{W}$ are exclusively excitatory, as in the actual cerebellum [1–3]. The vector of Purkinje cell responses $\boldsymbol{p}(t)$ evoked by parallel fiber input is computed as the product of the granule cell (parallel fiber) activation vector $\boldsymbol{g}(t)$ and the parallel-Purkinje weight matrix $\boldsymbol{W}$: $\boldsymbol{p}(t) = \boldsymbol{W}\boldsymbol{g}(t)$. The response $c_k(t)$ of Purkinje cell k to its combined input from granule and stellate cells is the product of the Purkinje cell response $p_k(t)$ to parallel fiber input and the stellate-Purkinje weight s_k: $c_k(t) = s_k\, p_k(t)$. The stellate cell population is represented by a single element that receives non-specific parallel fiber input and maintains a tonic activity of one. Stellate-Purkinje weights s_k are bounded between zero and one, and model a shunting inhibition. The combined parallel fiber and stellate response $c_k(t)$ represents the simple spike activity of Purkinje cell k. The output of the model, corresponding to a VOR eye velocity command, is computed by subtracting the summed response of the Purkinje cells from an offset ($o = 2$) version of the input that represents a vestibular nucleus signal. Purkinje cell responses are scaled by output weight $v = 0.025$. The actual output $y(t)$ of the model is: $y(t) = x(t) + o - v\Sigma_k c_k(t)$.

Climbing fibers in the model produce spikes at a low, random rate. Each microzone has its own climbing fiber. To model the low climbing fiber discharge rate, a climbing fiber fires at most once during a cycle of 5 Hz VOR. The occurrence time t_c of a climbing fiber spike is random, and climbing fiber discharge in one microzone is independent of discharges in other microzones. Updating of both parallel-Purkinje ($\boldsymbol{W}$) and stellate-Purkinje (s_k) weights is initiated by climbing fiber discharge.

An unsupervised, competitive (Kohonen) learning rule [18] updates the parallel-Purkinje weights $\boldsymbol{W}$. These weights initially take random values, uniformly distributed between zero and one. Climbing fiber discharge at time t_c engages a competition among the Purkinje cells in a microzone to find the index r of the cell with the largest response to the parallel fiber activity pattern at that time: $r = \arg\max_k(p_k(t_c))$. The winning Purkinje cell and its two nearest neighbors ($q = (r-1, r, r+1)$, with circular boundary conditions) are trained to become more specific for the current parallel fiber activity pattern by increasing and decreasing their parallel-Purkinje weights from active and inactive parallel fibers, respectively. This is done by adding a scaled version of the parallel fiber activity vector to each parallel-Purkinje weight vector $\boldsymbol{w}_q$, and then normalizing: $\boldsymbol{w}_q(u) = (\boldsymbol{w}_q(u-1) + \alpha \boldsymbol{g}(t_c)^{\mathrm{T}})/||\boldsymbol{w}_q(u-1) + \alpha \boldsymbol{g}(t_c)^{\mathrm{T}}||$. The learning rate α is 0.001, $||\cdot||$ represents the vector norm, and T signifies the transpose operation. Weight updates are indexed by u.

Stellate-Purkinje weights, set initially to the midrange value 0.5, are modified using a form of reinforcement learning [19] based on weight perturbation [20]. Climbing fiber discharge at time t_c starts a counter in Purkinje cell r with the largest response to parallel fiber input at that time. Each Purkinje cell makes its largest contribution to the output during the temporal segment for which it is most specific. However, that cell's effect on error will not be reflected by parallel fiber activity until after the retinal slip delay d has elapsed. The counter counts down this delay, and provides a simple implementation of an eligibility trace [19]. To accomplish reinforcement learning, a Purkinje cell transiently perturbs its stellate-Purkinje weight during its preferred temporal segment, and then registers the change in the number of active parallel fibers due to the perturbation when it becomes eligible. The Purkinje cell restores the perturbation if the number of active parallel fibers decreases, signifying a decrease in error.

The reinforcement sequence occurs as follows. At time $t = t_c + 1$, the stellate-Purkinje weight $s_r(t)$ of winning Purkinje cell r is perturbed by $b_r(u) = \beta\eta(u)$, where $\beta = 0.1$ and η is a mean 0 variance 1, normally distributed random deviate. At time $t = t_c + 2$ the perturbation $b_r(u)$ is removed from $s_r(t)$. The overall number of active parallel fibers at time $t = t_c + d + 1$ and at time $t = t_c + d + 2$ are held in variables n_b and n_a, respectively. The change in the number of active parallel fibers due to the effect of the perturbation on error is reflected by the difference between n_b and n_a. The tolerance $h = 4$ is the maximum change in number of active parallel fibers that occurs in the absence of any change in error. If the difference between n_b and n_a is $> h$, then the perturbation has reduced the error. Perturbation $b_r(u)$ is added back to $s_r(t)$ if the perturbation decreases the overall number of active parallel fibers within tolerance h: $s_r(t) = \{s_r(t) + b_r(u)$ if $n_b - n_a > h$; $s_r(t)$ otherwise$\}$, where $t = t_c + d + 2$.

In simulating VOR adaptation with InMin, unsupervised and reinforcement learning occur simultaneously. The model is first exposed to the vestibular input only, during a dark stage in which retinal slip error is absent. This is meant to simulate prenatal and/or early post-natal development (see below). Purkinje cells develop temporal specificity at this stage. Then error is introduced (dark/light switch closed) and the model is trained on the normal

VOR (amplitude one). The VOR is down adapted (amplitude zero) or up adapted (amplitude two) from the normal state. One unsupervised and one reinforcement update occur in each microzone on each training cycle. Training ceases when the mean squared difference between desired and actual outputs is brought within a tolerance of 0.01. The algorithm requires a few hundred training cycles to learn the normal VOR, and a few thousand to down-adapt or up-adapt it.

RESULTS

The actual output of the model closely matches the desired output following normal VOR training, and following down and up adaptation (Fig. 2a). The model produces the actual output by subtracting the weighted sum of the Purkinje cell responses from a version of the input offset by two, which represents a vestibular nucleus signal (bold curve in Fig. 2a). The offset input and the actual and desired outputs all have the same phase. In order to produce the normal VOR (amplitude one), the Purkinje cells as a group must evenly reduce the vestibular nucleus offset by two over the input cycle. The InMin algorithm accomplishes this by adjusting the stellate weight to each Purkinje cell to roughly 0.4. Production of the normal VOR is greatly facilitated by the even spread over the input cycle of Purkinje cell response phases (Fig. 2b).

To down-adapt the VOR, Purkinje cells need to reduce the offset more in-phase and less out-of-phase. The InMin algorithm accomplishes this by decreasing and increasing the stellate inhibition of Purkinje cells with in-phase and out-of-phase responses, respectively. Consequently, the responses of in-phase and out-of-phase Purkinje cells are increased and decreased, respectively, relative to normal (Fig. 2c). The reverse occurs for up adaptation of the VOR (Fig. 2d). Changes from normal in the responses of Purkinje cells following down and up adaptation are shown in Fig. 3. Down adaptation is associated with an increase for in-phase and a decrease for out-of-phase Purkinje cell responses. Conversely, up adaptation is associated with an increase for out-of-phase and a decrease for in-phase Purkinje cell responses. These modeling results accord with experimental observations.

DISCUSSION

The InMin model is used to simulate the effects of adaptation on the responses of floccular Purkinje cells that are normally active during the VOR [7–9]. The broad spread of Purkinje cell response phase produced by the InMin model is in agreement with the broad range of phase found for Purkinje cells in the flocculus [7–9]. Down adaptation of the VOR is associated with an increase and decrease relative to normal for in-phase and out-of-phase Purkinje cell responses, respectively, whereas up adaptation is associated with the opposite pattern [8,9]. The InMin model successfully simulates these findings.

The correspondence between simulated and real data supports InMin as a viable model of the computation that underlies cerebellar learning. The essential feature of InMin is that Purkinje cells learn by minimizing their input from parallel fibers. Input minimization is achieved through a combination of unsupervised and reinforcement learning, but the details of implementation are not critical. Other mechanisms may be substituted for Kohonen's rule in the implementation of unsupervised learning. More elaborate

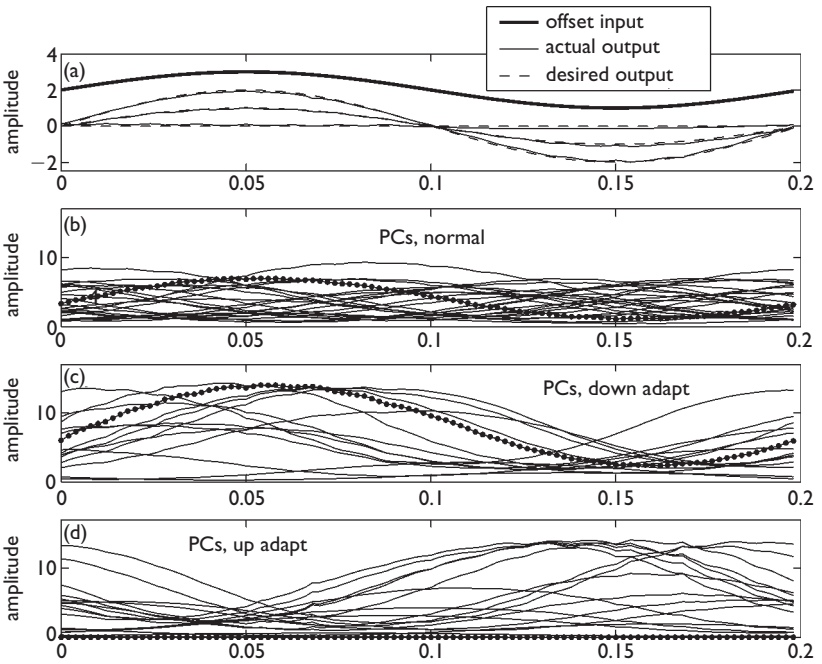

Fig. 2. Output and Purkinje cell responses for the normal VOR, and after down and up adaptation. Actual matches desired output in all cases (**a**). The offset input, representing a vestibular nucleus signal, is shown in bold (**a**). Purkinje cells develop temporally specific responses (**b**). Down adaptation increases and decreases Purkinje cell responses that are respectively in-phase and out-of-phase with the vestibular signal (**c**). The opposite occurs for up adaptation (**d**). Dots mark the response of the same Purkinje cell (**b**,**c**,**d**). PC, Purkinje cell.

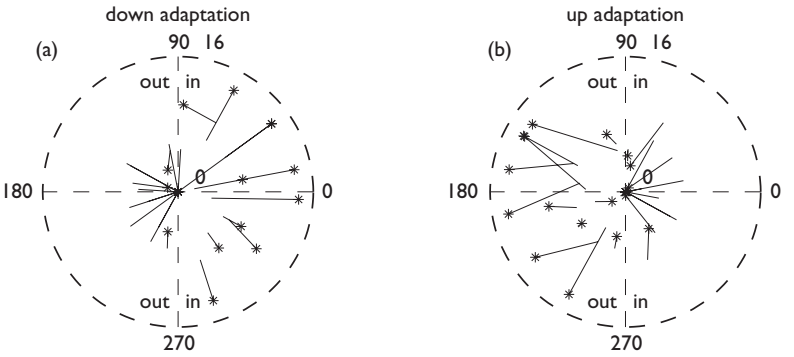

Fig. 3. Polar plots showing the amplitude and phase of Purkinje cell responses before and after VOR adaptation. The base of each line segment marks normal amplitude and phase while the asterisk marks adapted amplitude and phase. Down adaptation increases and decreases the amplitudes of Purkinje cell responses that are, respectively, in-phase and out-of-phase with the vestibular signal (**a**). The opposite occurs for up adaptation (**b**). Several vectors coincide in each plot and cannot be distinguished. Maximal amplitude is 16.

forms of reinforcement learning can extend the capability of InMin. Making stellate-Purkinje weight perturbations on alternate climbing fiber spikes, and comparing parallel fiber activity between successive eligibilities, would allow adaptation of ballistic movements like throws, saccades, and eye blinks, where error signals might be available only after the movement. The goal of this initial report is to demonstrate InMin in its simplest form using standard implementations of unsupervised and reinforcement learning [18–20].

While speculation on possible neurobiological implementations of InMin is beyond the focus of this article, it may be noted that available evidence is consistent with InMin and the possibility that both unsupervised and reinforcement learning take place in the cerebellum. It has been suggested [2] that mutually inhibitory connections between Purkinje cells mediate competitive interactions among them. This competition could contribute to the pauses in simple spike activity that immediately follow complex spikes [21]. Complex spikes are not explicitly simulated, but climbing fiber discharge essentially reads out the sensitivity of model Purkinje cells to the current parallel fiber activity pattern, in a manner analogous to that originally proposed by Eccles [1].

For real Purkinje cells, the pauses in simple spike rate induced by climbing fiber spikes are followed by transient changes in responsiveness. The transients are predominantly excitatory but can also be inhibitory [21,22]. They might result from activation by a climbing fiber spike of an increase in excitation from active parallel fibers, and a variable increase or decrease in inhibition from stellate cells. An increase in parallel-Purkinje excitation due to climbing fiber spikes would be consistent with unsupervised learning, as repeated pairing would maintain the strength of parallel-Purkinje synapses that are co-active with the climbing fiber. The decrease in strength of parallel-Purkinje synapses that are not co-active with the climbing fiber may be mediated by the observed decrease in Purkinje cell responsiveness to parallel fiber inputs that occur in the absence of climbing fiber spikes [23,24]. A variable excitatory or inhibitory transient, due to possible perturbation of stellate inhibition, would add to or subtract from the increase in parallel fiber excitation. This is consistent with the finding that all Purkinje cells show both excitatory and inhibitory transients, even cells that show predominantly excitatory transient changes in responsiveness [21]. A transient, variable change in stellate-Purkinje inhibition would be consistent with the perturbations elicited during the InMin reinforcement sequence, particularly as real transients can lead to permanent changes in Purkinje cell responsiveness [21,22]. The stellate-Purkinje synapse is GABAergic [2]. This is consistent with its presumed role in InMin as a shunt and as a modifiable element [25].

The InMin algorithm requires that Purkinje cells implement an eligibility trace, and also represent quantities proportional to the stellate-Purkinje weight perturbation and to numbers of active parallel fiber inputs. In another recent model of cerebellar learning [26], the simulated dynamics of second messenger systems implement the Purkinje cell eligibility trace. Purkinje cells abound in second messenger systems [27]. As for the eligibility trace, it is possible that signaling molecules could represent the quantities required for InMin.

The InMin algorithm can be used to simulate how two forms of learning, unsupervised and reinforcement, may interact in the cerebellum. InMin training is most stable when these two learning processes occur separately, with unsupervised learning first establishing temporally specific responses for each Purkinje cell, and reinforcement learning then adjusting those responses to produce a desired combined response. The realistic strategy employed here involves simultaneous learning with an initial developmental stage lacking error signals, during which unsupervised learning essentially occurs alone, followed by the introduction of error, after which unsupervised and reinforcement learning occur together. Studies on human infants underscore the importance of early vestibular experience on later motor learning [28].

InMin requires that each Purkinje cell in a microzone become specialized for a specific temporal segment of the input. Unsupervised learning fails to produce temporal specificity if the climbing fiber firing rate is too high ($\geqslant 10$ Hz). Lack of Purkinje cell temporal specificity deprives reinforcement learning of its ability to adapt separate temporal segments of the output. The low firing rate of real climbing fibers [1,15] may help ensure that Purkinje cells develop temporal specificity. Purkinje cells in mature

animals receive only one climbing fiber each [1–3], perhaps because innervation by multiple climbing fibers would effectively increase the frequency of their climbing fiber input, and thereby prevent development of temporal specificity. This may explain why mutant mice that maintain multiple climbing fiber innervation of Purkinje cells into adulthood exhibit severe deficits in motor learning [29].

Whereas parallel fibers encoding error change their activity pattern during learning, those encoding vestibular inputs do not. The vestibular parallel fibers allow the unsupervised mechanism to maintain Purkinje cell temporal specificity, even as the reinforcement mechanism changes the number of active error parallel fibers. Unsupervised and reinforcement learning interact well when the ratio of vestibular to error parallel fibers is greater than two. Estimates of the ratio of non-error to error cerebellar input elements (mossy fibers, granule cells, or parallel fibers) range from six to 36 [5,6]. A ratio in this range would guarantee stable learning with InMin. The absolute number of parallel fibers is not critical to InMin, but the smoothness and resolution of the output increase as the number of parallel fibers increases.

The output of the cerebellar model, which represents an eye velocity command, can be fed back to the Purkinje cells as an efference copy signal and encoded by another subset of parallel fibers. Efference copy and error both change during learning, but efference copy parallel fibers do not disrupt reinforcement learning if they are few in comparison with error parallel fibers. In that case, adaptive perturbations produce a net decrease in the overall number of active parallel fibers, because the decrease in number active is greater for error than for efference copy parallel fibers. The algorithm works well when there are about four times as many error as efference copy parallel fibers. This result is compatible with findings that efference copy inputs play a minor role in modulating Purkinje cell responses in the flocculus [8,9,30].

The simple threshold mechanism used here to model granule cell responses results in a rather dense parallel fiber activity pattern. Recent work [31] shows that a sparse (low overlap) recoding of mossy fiber signals by granule cells can improve learning in Marr/Albus models. Because sparse parallel fiber activity patterns are more easily distinguished by Purkinje cells, sparse granule cell recoding would enhance the development of Purkinje cell temporal specificity and so improve InMin learning as well.

The InMin algorithm reduces error by minimizing the number of active parallel fibers. For simplicity, parallel fibers in the model are binary, but the InMin algorithm would also work for real valued parallel fibers. In that case, InMin would minimize overall parallel fiber activity, rather than the number of active parallel fibers. Experimentally demonstrating that overall parallel fiber activity decreases with error would provide strong support for the model.

The observed randomness in climbing fiber discharge has caused problems for previous models of cerebellar learning based on the Marr/Albus paradigm [15], but poses no problem for InMin. According to InMin, a climbing fiber spike can occur anytime to synchronize the parallel-Purkinje and stellate-Purkinje weight adjustments of whichever Purkinje cell is the most specialized for the input at that time. Once a task has been learned, additional climbing fiber spikes produce little or no change in either parallel-Purkinje or stellate-Purkinje synapses. The random occurrence of climbing fiber spikes at all input phases is important for learning using the InMin algorithm, because it ensures that all temporal segments of the input will have Purkinje cells that have become specialized for them. Experiments show that climbing fiber discharge can be correlated with seemingly disparate signals and events including retinal slip error, movement onset, or unexpected input [17]. The model suggests that such correlations are unnecessary, but might be advantageous in that they would initiate more cerebellar learning at critical times, as when error is high, at the start of movement, or when unexpected input is encountered. InMin unifies the various views of climbing fiber function by seeing climbing fiber spikes as 'learn now' signals.

A previous model that simulated the effects of habituation on VOR behavior inspired the InMin algorithm. VOR habituation, in which prolonged rotation decreases VOR responsiveness, is associated with frequency-specific and nonlinear behaviors that were simulated using a non-adaptive model of the cerebellum based on pattern correlation [32]. As for InMin, the pattern correlation model contains Purkinje cells specialized for specific temporal segments of the input.

An approach somewhat similar to InMin is described in a preliminary report of a model of cerebellar control of saccades [33], in which parallel fibers carry error signals and climbing fibers provide timing cues. In that model, however, Purkinje cells are pre-set to respond to selected subsets of parallel fibers, and climbing fibers fire only after saccades. That model would fail under the realistic circumstances described for InMin, in which parallel fibers carrying error signals are not distinguished from those carrying other signals, and in which climbing fibers fire randomly. That model also apparently uses some form of supervised learning. On that crucial point it has more in common with the Marr/Albus paradigm than with InMin.

CONCLUSION

The Marr/Albus paradigm is based on supervised learning, and requires a continuous and precise error measure. It assumes that climbing fibers provide such an error measure, but real climbing fibers, which fire slowly and randomly, are unlikely to do so. Parallel fibers carry error signals along with other signals, and overall parallel fiber activity could provide a continuous and precise index of the absolute level of error. Thus, the amount of parallel fiber input could serve as a (negative) reinforcement signal. The InMin algorithm combines unsupervised and reinforcement learning, and climbing fiber spikes synchronize both processes. Unsupervised learning causes Purkinje cell responses to become specialized for specific temporal segments of the input. Reinforcement learning adjusts the sizes of those responses to minimize the parallel fiber input to Purkinje cells, and thereby reduce error and achieve a desired output. The InMin algorithm provides an alternative to the Marr/Albus paradigm that is more compatible with cerebellar findings.

REFERENCES

1. Eccles JC, Ito M and Szentàgothai J. *The Cerebellum as a Neuronal Machine*. New York: Springer; 1967.
2. Llinás RR and Walton KD. Cerebellum. In: Shepherd GM, ed. *The Synaptic Organization of the Brain*. New York: Oxford University Press; 1990, pp. 214–245.
3. Butler AB and Hodos W. *Comparative Vertebrate Neuroanatomy: Evolution and Adaptation*. New York: Wiley-Liss; 1996, pp. 180–200.
4. Wilson V and Melvill Jones G. *Mammalian Vestibular Physiology*. New York: Plenum Press; 1979.
5. Noda H. *Ann NY Acad Sci* **374**, 465–475 (1981).
6. Miles FA, Fuller JH, Braitman DJ *et al*. *J Neurophysiol* **43**, 1437–1476 (1980).
7. Dufossé M, Ito M, Jastreboff PJ *et al*. *Brain Res* **150**, 611–616 (1978).
8. Watanabe E. *Brain Res* **297**, 169–174 (1984).
9. Watanabe E. *Neurosci Res* **3**, 20–38 (1985).
10. Houk JC, Buckingham JT and Barto AG. *Behav Brain Sci* **19**, 368–383 (1996).
11. Marr D. *J Physiol Lond* **202**, 437–470 (1969).
12. Albus JS. *Math Biosci* **10**, 25–61 (1971).
13. Raymond JL and Lisberger SG. *Learn Mem* **3**, 503–318 (1997).
14. Ito M. *Annu Rev Neurosci* **5**, 275–296 (1982).
15. Keating JG and Thach WT. *J Neurophysiol* **73**, 1329–1340 (1995).
16. Kobayshi Y, Kawano K, Tekemura A *et al*. *J Neurophysiol* **80**, 832–848 (1998).
17. Simpson JI, Wylie DR and DeZeeuw CI. *Behav Brain Sci* **19**, 384–398 (1996).
18. Kohonen T. *Self-organizing Maps, 2nd edn*. Berlin: Springer–Verlag; 1997.
19. Sutton R and Barto A. *Reinforcement Learning: An Introduction*. Cambridge: MIT Press; 1998.
20. Venkatesh SS. *J Comp Sys Sciences* **46**, 198–217 (1993).
21. Sato Y, Miura A, Fushiki H *et al*. *J Neurophysiol* **68**, 2051–2062 (1992).
22. Ebner TJ, Yu Q and Bloedel JR. *J Neurophysiol* **50**, 205–219 (1983).
23. Hartell NA. *Neuron* **16**, 601–610 (1996).
24. Finch EA and Augustine GJ. *Nature* **396**, 753–756 (1998).
25. Kano M, Rexhausen U, Dreessen J *et al*. *Nature* **356**, 601–604 (1992).
26. Kettner RE, Mahamud S, Leung H-C *et al*. *J Neurophysiol* **77**, 2115–2130 (1997).
27. Linden DJ and Connor JA. *Annu Rev Neurosci* **18**, 319–357 (1995).
28. Clark DL, Cordero L, Goss KC *et al*. *Biol Neonate* **56**, 306–314 (1989).
29. Chen C, Kano A, Abeliovich L *et al*. *Cell* **83**, 1233–1242 (1995).
30. Miyashita Y. *Exp Brain Res* **55**, 81–90 (1984).
31. Schweighofer N, Doya K and Lay F. *Neuroscience* **103**, 35–50 (2001).
32. Anastasio TJ. *Neural Networks* **14**, 1–22 (2001).
33. Krichmar JL, Ascoli GA, Hunter L *et al*. A model of cerebellar saccadic motor learning using qualitative reasoning. In: Mira J, Moreno-Díaz R and Cabestany J, eds. *Biological and Artificial Computation: From Neuroscience to Technology*. Berlin: Springer-Verlag; 1997, pp. 133–145.

Acknowledgements: I thank A. Barto, D. Bullock, A. Fagg, Y. Gad, S. Grossberg, R. Kettner, J. Malpeli, and J. Reichler for comments and suggestions. This research was funded by NIMH grant MH50577.