2000-09

A Scalable Model of Cerebellar Adaptive Timing and Sequencing: The Recurrent Slide and Latch (RSL) Model

A scalable model of cerebellar adaptive timing and Sequencing: The recurrent slide and latch (RSL) model

Bradley Rhodes and Daniel Bullock

September, 2000

Technical Report CAS/CNS-2000-021

Boston University Center for Adaptive Systems and
Department of Cognitive and Neural Systems
677 Beacon Street
Boston, MA 02215

A Scalable Model of Cerebellar Adaptive Timing and Sequencing: The Recurrent Slide and Latch (RSL) Model.

Bradley J. Rhodes[1,2] & Daniel Bullock[1]*
brhodes@cns.bu.edu, danb@cns.bu.edu

[1]Department of Cognitive & Neural Systems and Center for Adaptive Systems
Boston University
677 Beacon St., Boston, MA 02215
Phone: 617-353-7858 or –7857
FAX: 617-353-7755

[2]School of Human Movement & Sport Sciences
University of Ballarat
Mt. Helen, VIC 3353, Australia

In press, *Applied Intelligence*

December 15, 2000

Acknowledgements

B.R. was supported in part by the Defense Advanced Research Projects Agency and the Office of Naval Research (ONR N00014-92-J-1309, ONR N00014-93-1-1364, and ONR N00014-95-1-0409).

D.B. was supported in part by the Defense Advanced Research Projects Agency and the Office of Naval Research (ONR N00014-95-1-0409 and ONR N00014-92-J-1309).

*Corresponding author

A Scalable Model of Cerebellar Adaptive Timing and Sequencing: The Recurrent Slide and Latch (RSL) Model.

Bradley J. Rhodes[1,2] & Daniel Bullock[1]
brhodes@cns.bu.edu, danb@cns.bu.edu

[1]Department of Cognitive & Neural Systems, Boston University, Boston, MA 02215
[2]School of Human Movement & Sport Sciences, University of Ballarat,
Mt. Helen, VIC 3353, Australia

Abstract

From the dawn of modern neural network theory, the mammalian cerebellum has been a favored object of mathematical modeling studies. Early studies focused on the fan-out, convergence, thresholding, and learned weighting of perceptual-motor signals within the cerebellar cortex. This led in the proposals of Albus (1971; 1975) and Marr (1969) to the still viable idea that the granule cell stage in the cerebellar cortex performs a sparse expansive recoding of the time-varying input vector. This recoding reveals and emphasizes combinations (of input state variables) in a distributed representation that serves as a basis for the learned, state-dependent control actions engendered by cerebellar outputs to movement related centers. Although well-grounded as such, this perspective seriously underestimates the intelligence of the cerebellar cortex. Context and state information arises asynchronously due to the heterogeneity of sources that contribute signals to compose the cerebellar input vector. These sources include radically different sensory systems – vision, kinesthesia, touch, balance and audition – as well as many stages of the motor output channel. To make optimal use of available signals, the cerebellum must be able to sift the evolving state representation for the most reliable predictors of the need for control actions, and to use those predictors even if they appear only transiently and well in advance of the optimal time for initiating the control action. Such a *cerebellar adaptive timing competence* has recently been experimentally verified (Perrett, Ruiz, & Mauk, 1993). This paper proposes a modification to prior, population, models for cerebellar adaptive timing and sequencing. Since it replaces a population with a single element, the proposed Recurrent Slide and Latch (RSL) model is in one sense maximally efficient, and therefore optimal from the perspective of scalability.

Keywords: recurrent network, sequence learning, adaptive timing, cerebellum

Introduction: Representation of state space and time in the cerebellar cortex

The cerebellum is an ancient neural module implicated in sensory-motor control. Damage to appropriate regions of the cerebellum in human adults causes a lasting breakdown in coordinated action, especially in high speed actions that require synchronous motion by several contributing joints (Bastian, Martin, Keating, & Thach, 1996; Thach, Goodkin, & Keating, 1992). Yet the core ability to generate multi-joint movements, with voluntary selection of spatial targets and continuous voluntary control of movement speed, is not compromised by cerebellar damage. This dissociation between a core generative ability and the ability to realize planned movement in a highly coordinated, state-dependent, manner is

also seen during development. A newborn rat pup soon exhibits the rudiments of its species' basic locomotor patterns. But the rat pup remains a very poorly coordinated locomotor until maturation of the cerebellum (Gramsbergen, 1993). Within a few days thereafter, the rat pup becomes a skilled machine that anticipates and compensates the complex mechanical consequences engendered during episodes of locomotion.

Such facts, taken together with abundant additional information about the structure of cerebellar circuits (Figures 1 and 2), and about extra-cerebellar (e.g., spinal) pattern generating circuits, have begun to build a consensus regarding the role of the cerebellum in most mammals (e.g., Contreras-Vidal, Grossberg, & Bullock, 1997; Ghez & Thach, 2000; Peters & van der Smagt, in press; Spoelstra, Schweighofer, & Arbib, 2000; Wolpert, Miall, & Kawato, 1998). Perceptually and volitionally responsive central pattern generators, which are capable of generating desired trajectories for multi-joint systems, exist outside the cerebellum. Although each joint has an associated negative feedback controller, the local feedback controllers associated with the joints have severe limitations because relatively long feedback lags force them to operate at low gain in order to avoid instability. In multi-joint movements, during which each joint motion perturbs the trajectories of mechanically coupled joints, such low gain feedback control, by itself, implies a very low performance ceiling. However, the effects of those (and any other) perturbations – namely trajectory errors – can be *preempted* if some supra-joint-level controller "sees them coming" and correctly schedules control actions that effectively cancel the expected perturbations. On the basis of its informational inputs, its tightly timed outputs, and the rules by which it adaptively adjusts those outputs, the cerebellum can be treated as such a predictive, supra-joint, controller.

The cerebellum appears to be very well designed for such adaptive error preemption. As schematized in Figure 1b, its outputs via the deep cerebellar nuclear (***DCN***) cell stage are organized into discrete motor control channels that map one-to-one onto feedback error channels. Thus for every feedback control channel, there is a corresponding cerebellar channel whose output excites the same motor command center excited by the feedback error signals carried by that channel. Sitting above each *DCN* channel (Figure 1b) is an entire population of giant Purkinje cells that normally inhibit the *DCN* channel. This Purkinje cell (**PC**) population must learn in *what* predictive contexts, exactly *when* after the detection of those contexts, and by *how much* they should briefly reduce their inhibitory outputs in order to allow transient activation of the *DCN* channel they control.

Several constraints must be simultaneously satisfied if the PC population is to perform the "what", "when" and "how much" functions well.

What predictive context? First, the population of Purkinje cells should be sensitive to a very large set of potentially predictive contexts. The size of this set can be maximized by ensuring that there is not too much overlap between the contexts visible to the different Purkinje cells in the population. As shown in Figure 1b, this is guaranteed by the fact that those PCs arrayed along the longer axis of the elongated band of PCs that affect a given *DCN* channel receive inputs from distinct subsets of parallel fibers (***PF***s). Moreover, Figure 1a shows that each *PF* is a bifurcated axon arising from a granule cell that is most likely to send an output when excited simultaneously by inputs from each of several mossy fiber "rosettes" (complex terminals) carrying context and state signals. Moreover, the separation

between successive rosettes arising from any one mossy fiber (***MF***) is larger (see Figure 1a) than the receptive (dendritic tree) volume for any single granule cell, so each granule cell is contacted by only one rosette from a given *MF*. This means that each granule cell and its associated, bifurcated, parallel fiber tends to signal the presence of a unique combination of context/state variables. But because of *MF* branching along the long axis of the cellular band associated with a given *DCN* channel, the particular variable it codes will be remixed into many (order 1000) distinct combinatorial representations. Finally, there are huge numbers of *MF*s that pass on context and state information, and this information arises asynchronously due to the heterogeneity of sources that contribute signals to the *MF* input vector. These sources include radically different sensory systems – vision, kinesthesia, touch, balance and audition – as well as many stages of the motor output channels.

Thus the cerebellum has a very rich basis for beginning its search for an answer to the question: *What predictive context?* Local inhibition applied by Golgi cells (see Figure 2) to the granule cells helps restrict the search to the signal combinations coded by strongly excited granule cells. Within the set of contexts represented by such granule cells, search for predictive contexts occurs through a learning process that transpires over a set of performance trials. In some reported cases, 20 trials were sufficient to produce learning that significantly affected performance. Learning is guided by climbing fiber (***CF***) inputs. As shown in Figure 1b, *CF*s originate in a ganglion called the inferior olivary nucleus (***OLN***). Unlike *MF*s, each *CF* distributes to only about 10 PCs, all of which are associated with the same control channel. Moreover, each PC receives only one *CF*. In the interpretation favored here, elevated signal traffic on a PC's climbing fiber input line indicates a performance error, and thus the need for a PC to search for a recent context (as represented by recently active granule cell and *PF*s) upon which that PC can conditionalize generation of an error-preemptive control action. Although the primate cerebellum has billions of granule cells, such high numbers are needed only because of the large number of states that are pertinent for controlling the extraordinarily rich repertoires of behavior found in primates. The cerebellar strategy can be followed with far fewer state-detection elements (or far more: some animals that use a cerebellar strategy adapted for high resolution sensory prediction devote a much larger proportion of their brains to cerebellar tissue.)

When to take action after the detection of a predictive context? Even if a PC can detect a strong predictive relationship between a *PF* input and a *CF* input, it would be far from optimal to be required to take the control action immediately upon detection of the predictive context. There are two reasons. First, such a restriction could as easily amplify error as preempt it. Second, there is high variability in the preprocessing delays exhibited by the many sources of *MF* inputs and in the natural distribution of lags between "leading indicators" and the events they predict. Therefore, however far the predictive context signal on the *PF* input line may lead the time the control action is needed, the control action itself should occur with only a modest lead relative to the receipt of the error signal on the *CF* line. With the mechanism described below, a PC's learned control action, while conditional upon predictive context detection, always occurs with a slight lead relative to the (expected) time of occurrence of the *CF* input signal. This implies an intra-cellular ability to use the detection of an input signal to initiate a clocking process whose "timing out" causes a cellular response.

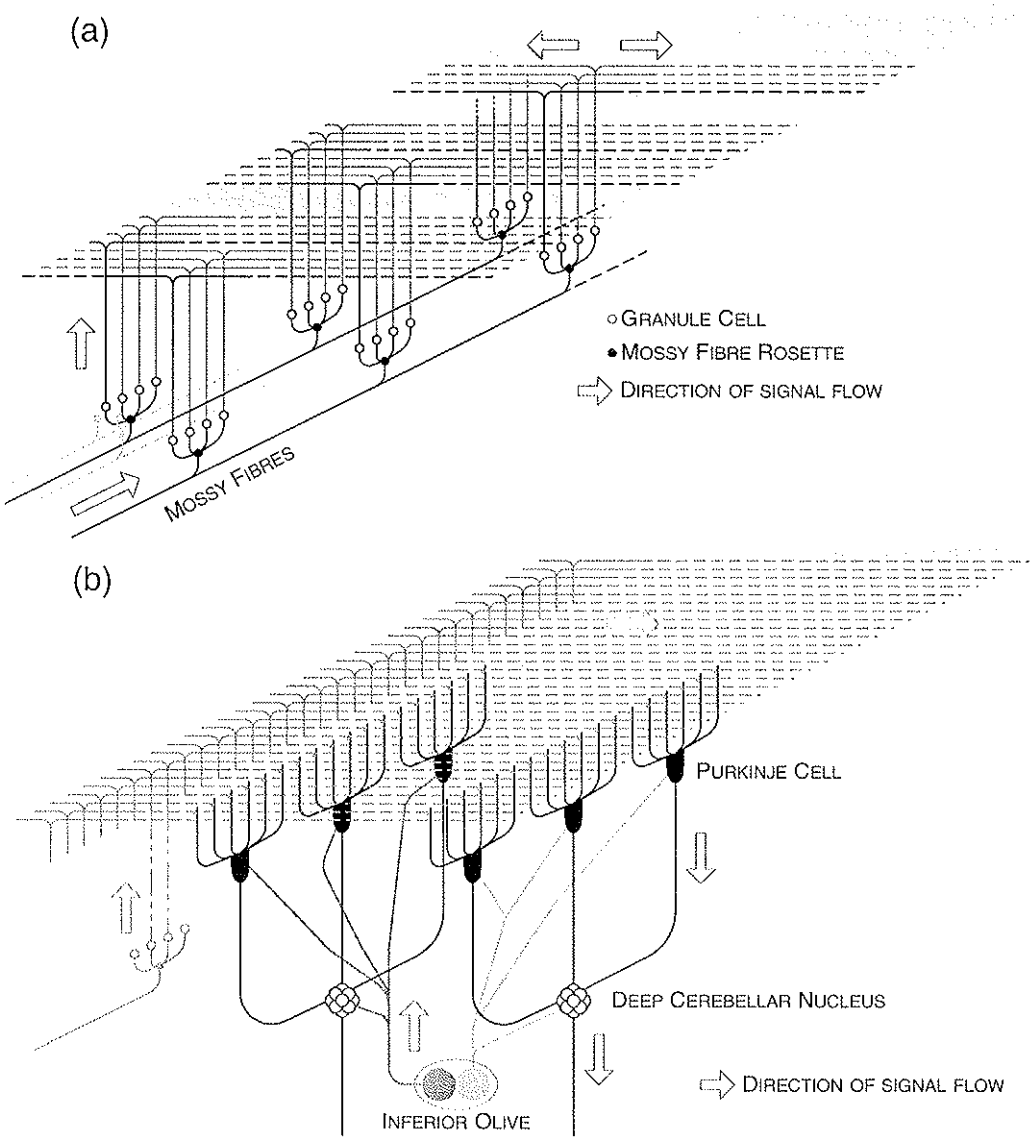

Figure 1. Illustration of the (a) divergent and (b) convergent geometry within the circuitry of the cerebellum. The grey panel represents a parasagittal strip of cerebellar cortex. (a) As described more completely in the text, inputs to the cerebellum, via mossy fibers, undergo expansion recoding at the divergent mossy fiber-granule cell interface, whereupon the resulting state representation is distributed within cerebellar cortex along parallel fibers. (b) Parallel fiber traffic is sifted by giant Purkinje cells, which sample from a huge number of parallel fibers along their parasagittal plane. Numerous Purkinje cells converge on a population of deep nuclear cells representing a given cerebellar output channel. Climbing fiber input from the inferior olivary nucleus, dedicated to this particular channel, innervates this set of Purkinje and nuclear cells. It should be noted that for purposes of clarity, no indication of the sign of the various synapses is included here (but is clearly illustrated in Figure 2).

How much of a cellular response should be generated? After PCs are stimulated by *PF* inputs, they can act by increasing or decreasing their inhibitory output to *DCN* cells. Evidence indicates that a strong *PF-CF* predictive relationship causes a *PF* dependent, lagged reduction of PC inhibition of *DCN* cells, which allows the latter to generate a transient burst of activity. The size of this burst can be adaptively adjusted in several ways. In most proposals, the learning process stops enhancing the size of the *DCN* burst whenever the control action succeeds in canceling the error. This will happen automatically provided that reduction of performance error amplitude in the controlled channel monotonically erodes excitatory input to the *OLN* cells that give rise to the channel's *CF*s. In addition to this implicit negative feedback, there is an explicit inhibitory feedback from *DCN* cells to associated *OLN* cells (Figure 2). This helps guarantee intermittency in the *OLN* output and thereby ensures that the *OLN* signal serves as a teaching signal but *not* as a continuous error signal. This is very important because the *CF* reaches a cluster of *DCN* cells (see Figures 1b & 2) as well as a band of overlying PCs. If it were allowed, a continous *OLN* output could act as part of a long-delay

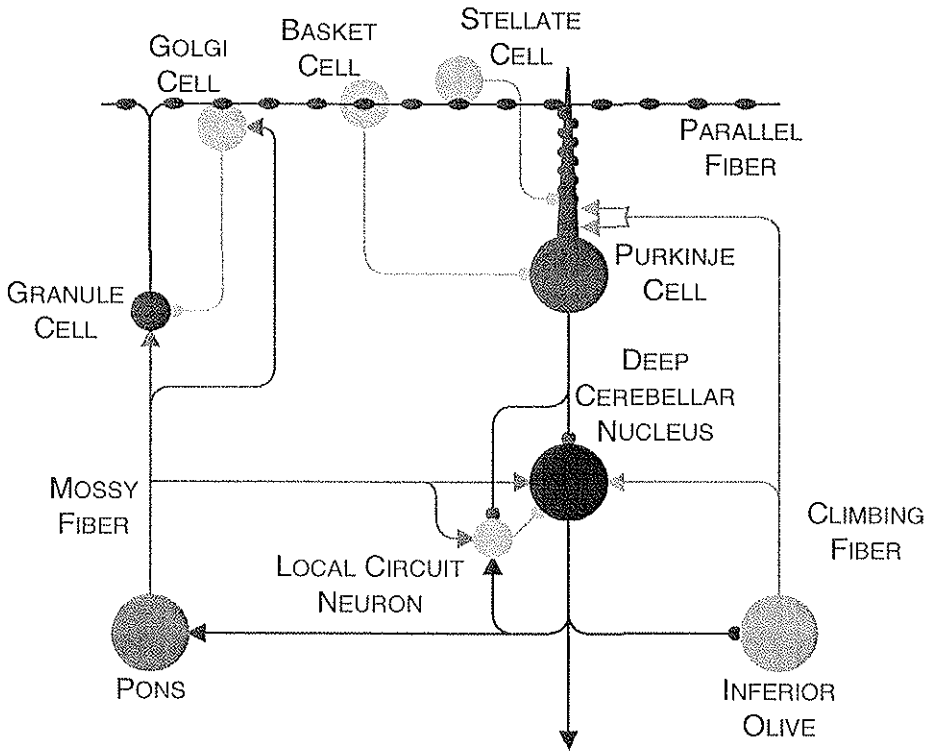

Figure 2. Circuit layout of the RSL model. The figure illustrates major cerebellar cell types and their connectivity. As noted in the text, the basket and stellate inhibitory interneurons were not explicitly modeled, but are included here for completeness. Excitatory connections are represented with arrows and inhibitory connections with small circles such that the flow of signals around the circuit follows that indicated in Figure 1. In addition to the *inter*cellular pathways shown, the model defines signaling processes within a 'second messenger' pathway inside the Purkinje cells. Through such pathways, mediated by an *intra*cellular metabolic cascade, an input to a cell can have a delayed activating/deactivating effect on the cell that is different than the immediate effect of the input.

feedback controller that could easily be destabilizing. Instead, the intermittent, burst *OLN* signal acts as a teaching signal whose only significant role is to modify how the cerebellum will behave in response to highly similar future contexts.

Some Limitations of Prior Models of Cerebellar Adaptive Timing

Numerous proposals have been made regarding cerebellar mechanisms that might create the ability to adaptively vary the interval between onset of a "leading indicator" context and the appropriate time for a control action, namely a short time before the expected onset time of the *CF* input. Two prominent classes of such models are *spatial pattern timing* and *spectral timing.*

Spatial Pattern Timing

In this type of cerebellar timing model, proposed by Buonomano & Mauk (Buonomano & Mauk, 1994; see also prior work of Church & Broadbent, 1990), a contextual input defined at the *MF* stage works with cerebellar cortical recurrence to induce a dissipating sequence of partially overlapping granule stage activations. Each pattern in the sequence becomes briefly eligible to serve as a basis for generating a learned PC response, and it begins to serve that role if its brief interval of eligibility coincides with a *CF* input. This proposed mechanism inherently requires a large population of granule cells to achieve timing competency. Also, by using granule cell activations to code time delays from an earlier event, it implicitly recasts them as time delay elements rather than as combinatorial state detectors.

Spectral Timing

In this type of model, each context signal induces an entire spectrum of transient delayed responses. Whichever transient element of the spectrum coincides with *CF* activation begins to be used to control timing of the PC's response to the onset of the context. Bullock, Fiala & Grossberg (1994) proposed that this spectrum emerged at the granule cell stage, so with respect to the synapse between *PF* and PC, the spectrum generation could be described as "presynaptic". Because the presynaptic spectrum emerged at the granule cell body, as opposed to the *PF* terminals, this proposed mechanism required one granule cell for each spectral component. Thus it was also inherently a population model, and it relied on the idea that granule cells multiplex context detection and time delays.

While multiplexing of context detection with timing variability may be possible as in the two models just discussed, it is not optimal. Ideally, any state should be able to control a full range of delayed actions. This type of limitation was overcome in the spectral timing model of Fiala, Grossberg & Bullock (1996), through serialization of state/context detection and timing control. On the basis of accumulating data regarding slow post-synaptic signal processing by a "second messenger" cascade in PCs, Fiala et al. proposed that a spectrum arose postsynaptically – i.e., within the PCs themselves in the dendritic spines associated 1-to-1 with all *PF* to PC synapses. Although this proposal reduced the required number of granule cells, and made each able to engender a full spectrum of delayed actions, the mechanism required a population of post-synaptic sites to generate the spectrum. Because each *PF* synapses only once on a given PC, the required population was constructed from dendritic spines of distinct PCs, whereon a predetermined distribution of receptor densities

was assumed to act as a distribution of gains sufficient to create a spectrum of delayed signals subsequent to granule activation.

Current Proposal: The Recurrent Slide and Latch (RSL) Network

The Fiala et al. (1996) spectral timing solution scales better than it first appears, because each PC receives single synapses from many thousands of passing *PF*s, so PCs can be reused to mediate timed responses to an equally large number of contexts. Nevertheless, it would be better – and in one sense optimal – if *any* granule-detected context could be responded to after *any* learned interval by *any* PC to which the granule cell projects via its *PF*. For applications, such a more efficient, scalable, design may be critical. Moreover, the lack of direct evidence of preformed delay spectra leaves open the possibility that a more efficient network may also be a more biologically accurate model.

Although Fiala et al. realized that the model could be made more efficient by assuming that the receptor density at each *PF* to PC synapse was itself modifiable, they did not see a way that such a model could explain a well established property of timed response learning in animals: the multi-peak phenomenon. Essentially, animals can be trained to exhibit multiple timed responses to a single contextual cue or state if the onset of this cue is followed by *CF* activity at one interval (say after 250 ms) on 50%of training trials and at a different interval (e.g., 750 ms) on the remaining 50% of trials. While readily explained in a spectral model, this effect is not immediately explicable in a model that replaces the spectrum with an adjustable receptor density that can be used to learn any single desired delay.

Recently, Steuber and Willshaw (1997) developed a variant of the Fiala et al. model in which they showed that "timing can be learnt by a single Purkinje cell" if the receptor density variable in the latter model is brought under adaptive control. In particular, they assumed that the initial density value is set large, which implies a relatively short delay. If needed, a learning process based on *CF* teaching signals causes a reduction of the post-synaptic receptor density and a consequent lengthening of the delay between context onset and PC response. The Steuber and Willshaw proposal did not address the phenomenon of multiple timed responses to a single context. Moreover, the model was reported to span only intervals between 200 and 1000 ms, which falls well short of the 90 ms to 4000 ms timing range that appears to be within the competence of the real cerebellum.

The current work began from a slightly different perspective, that of learned behavioral sequences (Rhodes, 1999). We sought a model that preserved the serialization of state and timing control as well as the core biochemical basis of the Fiala et al. model. Indeed, support for the biochemical and post-synaptic timing assumptions of that model has been striking (Finch & Augustine, 1998; Raymond & Lisberger, 1998; Takechi, Eilers, & Konnerth, 1998). We also sought greater efficiency, as in the learnable delay variant developed by Steuber and Willshaw (1997). However, because of our interest in learned sequences, we required a natural solution that could explain multiple timed responses to a single contextual cue – preferably without introducing any new circuitry. The result is the **Recurrent Slide and Latch (RSL)** network model. Like the earlier models of Buonomano & Mauk (1994) and Bullock et al. (1994), it incorporates recurrence into cerebellar timing control, but it emphasizes between-response (output to input) rather than within-response recurrence. In that sense, its use of recurrence is more closely related to (non-cerebellar)

models by Jordan (e.g., Jordan, 1985) and by Dominey and colleagues (e.g., Dominey, Arbib, & Joseph, 1995).

As in Fiala et al. (1996), the substrate of delayed responses is proposed to be a second messenger response within PC dendritic spines contacted by *PF*s. In the current proposal, however, a learning mechanism allows each second messenger response to change its time course. As shown in Figure 3, each delayed second messenger response has a rising and a falling edge, with the peak representing when the maximal cerebellar output will occur. When *PF* input reliably predicts later *CF* input, adaptive weight changes accumulate and cause the time of this peak to **slide** (forward in time) toward the time at which the error/teaching signal (*CF* burst) occurs. When the slide results in slight temporal precedence

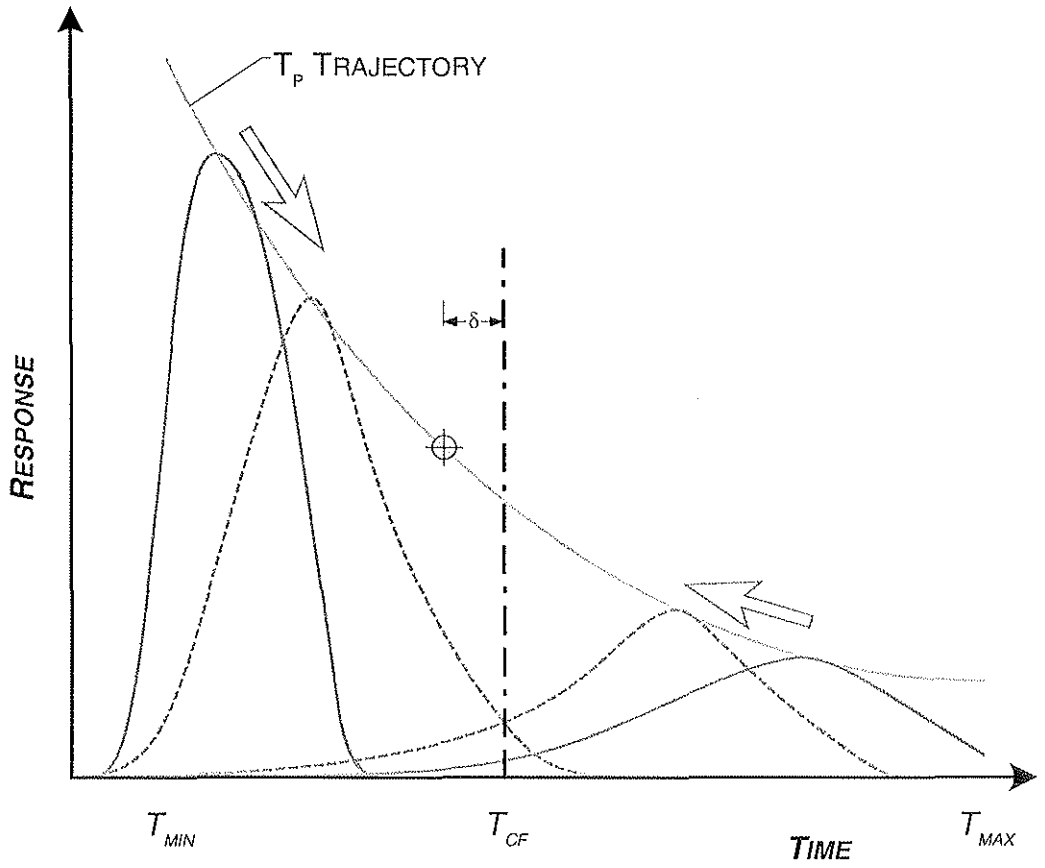

Figure 3. Schematic operation of the adaptive second messenger response that occurs within the Purkinje cells of the RSL model. The time to peak (T_P) of the *SMOut* response can vary along the T_P trajectory indicated – ranging from T_{MAX} to T_{MIN}. At baseline it starts at T_{MAX} and the occurrence of a climbing fiber teaching signal – at T_{CF} – causes T_P to slide forward toward T_{CF} as is illustrated toward the right of the figure where the response indicated by the solid line changes to that indicated by the dashed line. After sufficient trials, T_P will have advanced until it precedes T_{CF} by a small lead (δ) and there become latched (indicated by the cross-hairs). In the absence of a climbing fiber signal, T_P will drift backward in time toward T_{MAX} as illustrated toward the left of the figure.

of the peak relative to the *CF* input (see Figure 3), weight change ceases, and the time of the peak is thereby **latched** at a slight lead relative to the expected time of the error signal. In the absence of an error signal to latch to (or move towards), the peak slides (backward in time) toward some finite maximal time. As shown in Figure 3, the peak's amplitude also decreases markedly, which means it comes to have a negligible effect on Purkinje cell response. As in the Steuber & Willshaw (1997) model, the slide and latch process removes the need to postulate a population to carry a spectrum of preformed lags as a precondition for adaptive timing. Thus the RSL model is not a spectral timing model.

Two major differences and one minor one separate the RSL model from the Steuber & Willshaw variant (1997) of the Fiala et al. post-synaptic timing model. First, the RSL model can cover the full 90-4000 ms range, but this is a minor difference. More important is that in the RSL model the *CF* signals generally induce forward movement of the PC response. Because of the time based discounting achieved by the inverse relationship between time delay and response amplitude, this assures that *PF* signals that are spuriously correlated with *CF* signals over a few trials will not thereby be able to generate significant control actions, which could occur in the Steuber & Willshaw variant. Equally important, the RSL model explains how the cerebellum can generate multi-peaked responses in which secondary peaks can be just as vigorous as (of amplitude equal to) primary responses. As Figure 2 shows, recurrent connections from the *DCN* cells to the pons allow cerebellar outputs to be added to the state representation forwarded to the cerebellar cortex via the *MF* system. This means that in every case that requires multiple timed responses, the occurrence of the first response can serve to help define a new combinatorial context that can initiate a new timing episode.

Formal Specification of Model Circuit and Functional Components

The model is specified as a system of ordinary differential equations. Symbol definitions, parameters, and parameter values appear in Table 1. The terms of the presentation are drawn from the classical (a.k.a. Pavlovian) conditioning paradigm of experimental psychology. In this protocol, the experimenter sets up a predictive relationship between an arbitrary conditioned stimulus (***CS***) and an unconditioned stimulus (***US***) that is already capable of evoking an unconditioned response. In particular, the *CS* onset precedes the *US* onset by a fixed time called the inter-stimulus interval (**ISI**). As training progresses, an animal exposed to the predictive relationship begins to generate a new response that resembles the unconditioned response, but with two notable differences. First, its conditional probability of production, given the *CS*'s occurrence, gradually approaches 1.0. Second, its production begins just before the expected time of arrival of the *US*. Therefore, it is an anticipatory, *adaptively timed conditional response* (**ATCR**). By selecting a corresponding ISI, the experimenter can train an ATCR to appear at any time between 90 and 4000 ms after the *CS* onset. Here the role of the *US* is analogous to any internally generated error signal capable of evoking a reactive compensatory control action, and the role of the *CS* is analogous to any internal or external state/context whose occurrence reliably precedes the error. Any mechanism capable of learning to generate ATCRs to arbitrary *CS*s can therefore solve the adaptive error preemption problem. The classical conditioning paradigm is therefore probing a very general learning ability.

Table 1. List of symbols and parameters values.

Symbol		*Parameter*	*Value*
MF	Mossy fiber	τ_{MF}	0.01
		γ_{MF}	10
		δ_{MF}	0.2
GrC	Granule cell	τ_{GrC}	0.01
		δ_{GrC}	4
		Γ_{GrC}	0.1
PF	Parallel fiber	α	12
		β	4
		Γ_{PF}	0.5
G	Golgi cell	γ_{G}	5
		δ_{G}	0.1
		φ_{G}	0.02
PC	Purkinje cell		
WC	*PF*-PC synapse weight in cerebellar cortex	ω_{WC}	0.5
		λ_{WC}	10
SMOn	'On' component of second messenger pathway	γ_{SMOn}	15
		δ_{SMOn}	4
		Γ_{SMOn}	0.1
		ε_{SMOn}	0.0001
SMOut	'Output' component of second messenger pathway	Γ_{SMOut}	0.01
SMOff	'Off' component of second messenger pathway	γ_{SMOff}	10
		δ_{SMOff}	0.1
DCN	Deep cerebellar nuclear cell	Γ_{DCN}	0.2
		φ	0.05
LCN	Local circuit neuron in cerebellar nucleus	γ_{LCN}	30
		δ_{LCN}	5
		χ_{LCN}	200
WN	*MF-DCN* synapse weight in cerebellar nucleus	ω_{WN}	10
		λ_{WN}	0.001
CF	Climbing fiber	ε_{CF}	0.01
		δ_{CF}	0.4
OLN	Inferior olivary neuron	τ_{OLN}	0.01

Cells of the pons (Figure 2) give rise to *MF*s, which transmit *CS*-induced signals. Activity MF_j of the *j*th mossy fiber changes according to the equation

$$\tau_{MF}\frac{dMF_j}{dt} = \left(1 - MF_j\right)\cdot \gamma_{MF}\cdot CS_j - \delta_{MF}\cdot MF_j . \tag{1}$$

To simplify the presentation in order to focus on adaptive timing, each model *MF* excites only one granule cell, each of which receives only one *MF* input. Thus the *j*th granule cell, GrC_j, is governed by the equation

$$\tau_{GrC}\frac{dGrC_j}{dt} = \left(1 - GrC_j\right)\cdot MF_j - \left(GrC_j + \Gamma_{GrC}\right)\cdot \delta_{GrC}\cdot G_j \tag{2}$$

in which G_j represents the magnitude of inhibitory input from Golgi cell j.

GrC activity produces parallel fiber (*PF*) activity according to the sigmoid function

$$PF_j = \frac{\alpha\cdot\left(\left[GrC_j - \Gamma_{PF}\right]^+\right)^2}{1+\beta\cdot\left(\left[GrC_j - \Gamma_{PF}\right]^+\right)^2}, \tag{3}$$

where $[x]^+ = \begin{cases} x, & x \geq 0 \\ 0, & \text{otherwise} \end{cases}$.

This *PF* activity has two sites of excitatory influence. The first is upon Golgi cells, which are also excited by appropriate *MF*s, such that activity G_j in the *j*th Golgi cell varies according to the equation

$$\frac{dG_j}{dt} = \left(1 - G_j\right)\cdot \gamma_G \cdot \left(PF_j + \varphi_G \cdot MF_j\right) - \delta_G \cdot G_j . \tag{4}$$

Parallel fiber (*PF*) activity also excites each of the Purkinje cells (PC). Activity PC_i of the *i*th Purkinje cell changes according to the equation

$$PC_i = 1 + \sum_j PF_j \cdot WC_{ij} - \sum_j SMOnt_{ij}, \tag{5}$$

so PCs receive unit tonic excitation, summed excitation by *PF* activity filtered through associated *WC* weights, and a timed inhibition produced by the *SMOnt* signal, to be explained soon. Although inhibitory interneurons have not been included here, we hypothesize that in animals the stellate cells play critical roles in gain control and may also play a role in producing the lower bound on the range of ISIs that reliably induce ATCRs. The excitatory, direct *PF* input to PCs is filtered through an adaptive set of weights. The weight from the *j*th *PF* to the *i*th PC cell, WC_{ij}, changes according to the equation

$$\frac{dWC_{ij}}{dt} = \underbrace{\omega_{WC} \cdot (0.9 - WC_{ij}) \cdot PF_j}_{\text{LTP}} - \underbrace{\lambda_{WC} \cdot WC_{ij} \cdot SMOn \cdot r(SMOut) \cdot CF_i}_{\text{LTD}}, \quad (6)$$

$$\text{where } r(x) = \begin{cases} 1, & dx/dt \geq 0 \\ 0, & \text{otherwise} \end{cases},$$

and the climbing fiber (*CF*) signal induced by the *US* is defined below. The *WC* value is constrained to the [0,0.9] interval. The labeling braces indicate that the net weight change depends on the balance between a weight incrementing process, long term potentiation (LTP), and a weight decrementing process, long term depression (LTD).

Equation 6 is substantially different from any used in spectral timing models. Explaining this equation requires prior definition of the *SMOn* and *SMOut* signals, so these will be addressed now, and the *WC* equation revisited later. The *SMOn* signal is modeled to represent activity of an intracellular second messenger (SM) pathway within Purkinje cell (PC) dendritic spines. This is similar in concept to the heart of the Fiala et al. (1996) spectral timing model, although the present formulation abstracts away from the biochemical details of the 1996 model. The activity of the 'on' component of the second messenger pathway in the *j*th dendritic spine of the *i*th PC, $SMOn_{ij}$, varies according to

$$\frac{dSMOn_{ij}}{dt} = \gamma_{SMOn} \cdot (1 - WC_{ij}) \cdot \left[\begin{array}{c} (1 - SMOn_{ij}) \cdot (PF_j + f(SMOn_{ij}, \varepsilon_{SMOn}) + SMOut_{ij}) \\ - (SMOn_{ij} + \Gamma_{SMOn}) \cdot \delta_{SMOn} \cdot SMOff_{ij} \end{array} \right], \quad (7)$$

$$\text{where } f(x, \Gamma) = \begin{cases} x, & x \geq \Gamma \\ 0, & \text{otherwise} \end{cases},$$

$$\text{and } SMOut_{ij} = \frac{\alpha \cdot \left(\left[SMOn_{ij} - \Gamma_{SMOut} \right]^+ \right)^2}{1 + \beta \cdot \left(\left[SMOn_{ij} - \Gamma_{SMOut} \right]^+ \right)^2}.$$

This *SMOn* component receives excitatory input from parallel fiber (*PF*) activity via metabotropic glutamate receptors (mGluRs). When the activity in this pathway exceeds a very low threshold, it receives moderate positive feedback via the linear threshold function, f( ). Thus this activity can sustain itself and grow slowly until it reaches a higher threshold. When this higher threshold is achieved, the *SMOut* component becomes active, with its activity level being determined by a nonlinear sigmoid function. Presence of *SMOut* activity, which itself rises quickly once its threshold has been exceeded, provides a greater degree of positive feedback to *SMOn* activity. The result is a sharper increase in *SMOn* activity than was achieved due to the linear feedback alone. The overall rate of *SMOn* response to the excitatory influences is dependent on the value of *WC* – which is therefore acting as a variable time constant for *SMOn*.

SMOut activity is responsible for the timed inhibition of Purkinje cell (PC) activity. It also activates the 'off' component of the second messenger pathway, which provides inhibitory input to *SMOn* and therefore eventually shuts down its memory of the contextual event that

originally engaged it. Activity in the 'off' component, $SMOff_{ij}$, of the second messenger pathway in the jth dendritic spine of the ith PC changes according to the equation

$$\frac{dSMOff_{ij}}{dt} = \left(1 - SMOff_{ij}\right) \cdot \gamma_{SMOff} \cdot SMOut_{ij} - \delta_{SMOff} \cdot SMOff_{ij}. \tag{8}$$

This *SMOff* component assumes the inhibitory role that was performed by interneurons in some earlier spectral timing models. Once engaged, *SMOff* activity reduces due to slow passive decay. This makes it persistent enough to reduce *SMOn* activity, first to a value below the threshold for *SMOut* activity and then below the linear feedback function threshold.

Reexamining equation (6) governing change in the *WC* reveals that the second messenger system plays a critical role in the induction of long term depression (**LTD**) of these weights. As in many prior cerebellar models, climbing fiber (*CF*) activity is required to induce LTD, but here direct parallel fiber (*PF*) activity is replaced by the intracellular *SMOn* component to make this form of learning conditional upon temporal coincidence between a memory trace of recent contextual input and the *CF*-mediated error signal. A further constraint on LTD has been included to ensure that the weight change ceases once the peak of the second messenger response coincides with the occurrence of the error signal. If the weight is initially maximal, then as trials are conducted with the error signal occurring at a fixed interval after the context signal, the relevant *WC* is decreased via LTD. This reduction changes the time constant governing the response rate of the second messenger pathway, which consequently speeds up. Its peak moves forward in time across trials until it coincides with the error signal itself, whereupon the slide stops and the response becomes latched by the error signal. As the peak moves forward (right to left in Figure 3) it also grows in magnitude, eventually being able to inhibit Purkinje cell (PC) activity to the extent necessary to permit a deep cerebellar nuclear (*DCN*) output signal.

The remaining elements of the RSL model are very similar to elements of prior cerebellar timing models (e.g., Bullock et al., 1994; Fiala et al., 1996) which may be consulted for additional discussion. Activity DCN_i of the ith deep cerebellar nucleus cell changes according to

$$\frac{dDCN_i}{dt} = MF_j \cdot WN_{ji} - \left(DCN_i + PC_i + \frac{\left(\left[LCN_i - \Gamma_{DCN}\right]^+\right)^2}{\varphi^2 + \left(\left[LCN_i - \Gamma_{DCN}\right]^+\right)^2} \right). \tag{9}$$

The local circuit neuron (*LCN*) associated with the ith channel, LCN_i, changes activity according to the equation

$$\frac{dLCN_i}{dt} = \left(1 - LCN_i\right) \cdot \gamma_{LCN} \cdot \left(\sum_j MF_j + DCN_i \right) - \left(\delta_{LCN} \cdot LCN_i + \chi_{LCN} \cdot PC_i\right). \tag{10}$$

The adaptive weight between the jth *MF* and the ith *DCN*, WC_{ji}, changes according to the equation

$$\frac{\mathrm{d}WN_{ji}}{\mathrm{dt}} = MF_i \cdot \left(\omega_{WN} \cdot CF_i \cdot \left(1 - WN_{ji}\right) - \lambda_{WN} \cdot WN_{ji}\right). \tag{11}$$

The climbing fiber activity CF_i that provides a teaching signal to both sets of adaptive weights associated with the ith channel is governed by the equation

$$CF_i = \mathrm{H}(OLN_i, \varepsilon_{CF}) \cdot \left(1 - \delta_{CF} \cdot \mathrm{e}^{-2[OLN_i - \varepsilon_{CF}]^+}\right), \tag{12}$$

where H( ) is a thresholded Heaviside function, $\mathrm{H}(x, \mathrm{T}) = \begin{cases} 1, & x \geq \mathrm{T} \\ 0, & \text{otherwise} \end{cases}$.

Activity OLN_i in the ith OLN unit changes according to the equations

$$\tau_{OLN} \frac{\mathrm{d}OLN_i}{\mathrm{dt}} = -OLN_i^+ + US_i - OLNInh_i, \text{ and} \tag{13}$$

$$\tau_{OLN} \frac{\mathrm{d}OLNInh_i}{\mathrm{dt}} = -OLNInh_i + US_i. \tag{14}$$

Simulations Demonstrating Model Learning and Performance

A number of small scale simulations have been performed to test adaptive timing in the RSL model. Constant parameters were adjusted empirically to yield appropriate behavior, but no attempt was made to optimize them. Similar to those presented in Bullock et al. (1994), the simulations were of classical conditioning trials. In most cases, a unit valued *CS* input was present throughout each trial. After an ISI of 1000 ms., a unit valued *US* stimulus with a duration of 0.01 ms (equivalent to one simulation time step) was presented. Because the system scales linearly with the number of elements, we describe results for only one contextual input channel, one output channel, and one error signal.

In the first simulation, there was no carry over of learning from one simulation run to the next. Instead, the *WC* value was explicitly set for each simulation run in order to illustrate the behavior of the model as *WC* varied systematically across its range. Thus, from its maximal value of 0.9, *WC* was decremented using steps of 0.05 to its minimal value, 0. Figure 4 illustrates the activations within the principal components of the model. The 'input' components – the granule cell (*GrC*) stage, the parallel fibers (*PF*), and the Golgi cells (*G*) – are presented in the upper three panels. Since the conditioned stimulus (*CS*) remains constant across runs, there is no change in the activity of these three components for different simulation runs – hence the appearance of a single line in each panel. The *PF* activation is transient and follows shortly after the *CS* onset. Mossy fiber (*MF*) activity stays moderately high throughout the simulated interval because the *CS* remains present.

The timecourse of activity within the second messenger components is illustrated in the next three panels. Here the activities are different between the various simulation runs. Consider first the *SMOn* responses (in the fourth panel). When *WC* is maximal there is an initial response due to the *PF* input that is sufficient to engage the linear positive feedback

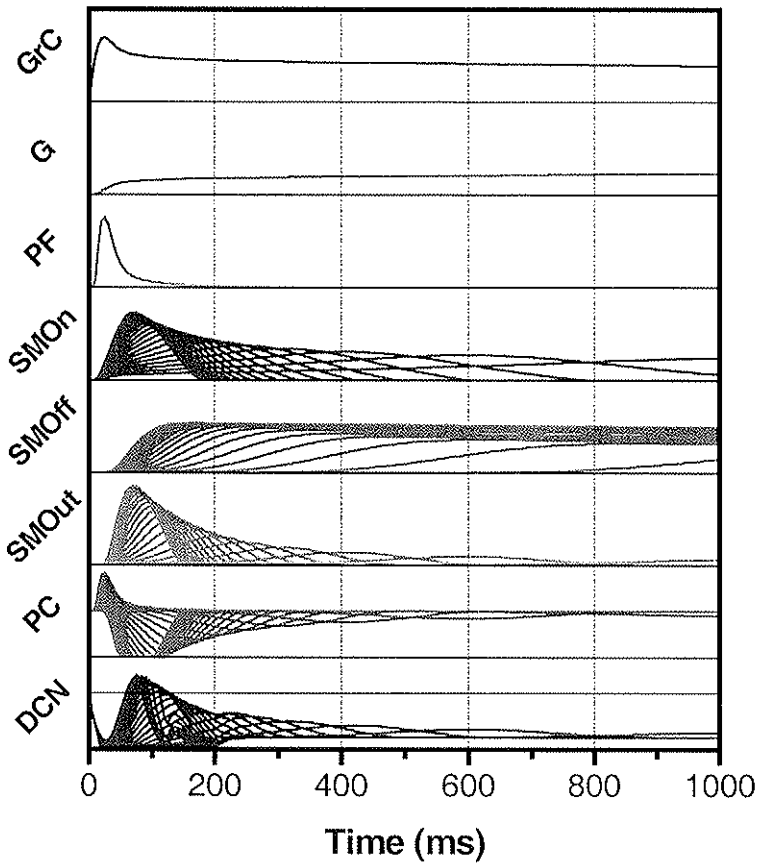

Figure 4. Timecourse of activities within the 'input' components (top three panels), second messenger components within the Purkinje cell (panels 4 through 6), and output components (lower two panels) of the RSL model during the non-learning simulation demonstrating model behavior. Traces arise from a series of simulations where the *WC* value was decremented by a fixed amount from one simulation run to the next to illustrate how the second messenger component kinetics changed in response and the effect these changes have on downstream targets – Purkinje cells and deep cerebellar nucleus neurons. The solid line in the *DCN* panel indicates the threshold level for *DCN* output. *GrC* – granule cell; *G* – Golgi cell; *PF* – parallel fiber; *SMOn* – second messenger onset component; *SMOut* – second messenger output response; *SMOff* – second messenger offset component; P – Purkinje cell; *DCN* – deep cerebellar nucleus.

component and *SMOn* activity slowly rises. Because the time delay is proportional to *WC*, the *SMOn* response is most sluggish in this run. The particular line in the *SMOn* panel being referred to in this case is the one that is lowest throughout the simulated interval. Eventually, however, *SMOn* activity does reach the threshold required to activate the *SMOut* component. This can be seen to occur after approximately 600 ms of the trial has elapsed. The engagement of the *SMOut* component is also sluggish and so a rapid increase in the rate of *SMOn* increase does not occur. Onset of *SMOff* activity is also slow, starting to build up only at the end of the run interval.

As the value of *WC* progressively decreases, the rate of *SMOn* activity buildup increases and there is earlier and more robust activation of *SMOut* and *SMOff* respectively. Even after the first decrement of 0.05, giving *WC* a value of 0.85, the kinetics of these components are much altered. A definite peak in both *SMOn* and *SMOut* activity is clearly evident. Thus, it does not take a very large weight change to substantially change the behavior of these components. As *WC* is decremented further the kinetics continue to accelerate. So the peak of *SMOn* and *SMOut* activity arrives ever earlier in the simulated interval as the value of *WC* shrinks. There is a diminishing return in the amount of peak shift that is achieved by a constant amount of weight change as the weight approaches its minimum value.

The effects of the changing behavior of *SMOut* (across simulation runs) on the output stages of the RSL model are presented in the lower two panels of Figure 4. The influence of *SMOut* activity on Purkinje cell (PC) activity is apparent. Also apparent is the undelayed excitatory influence of the transient parallel fiber (*PF*) input, which increases PC activity above its tonic baseline. This influence is down modulated by LTD, i.e., by the progressive reduction of *WC* values. *SMOut* inhibition of PC activity releases the deep cerebellar nuclear (*DCN*) cells from inhibition, allowing them to exceed threshold even without the benefit of *WN* learning.

These simulation runs demonstrate that the learned delay proposal being made here does have the types of properties desired as *WC* values change. The next step was to perform an actual conditioning simulation where the weight change produced by each trial is carried forth to subsequent trials. An ISI of 500 ms was used and 19 learning trials were simulated. However, unlike in the simulations of Figure 4, there was no direct manipulation of *WC* throughout the block of trials. Instead, *WC* was adjusted by the learning law, Equation6. The expectation was confirmed that the *WC* would decrease across trials and move the *SMOut* peak forward until a trial on which it slightly preceded *US* arrival, whereupon the peak shift would stop.

Because the contextual input is exactly the same as in the prior simulation, the responses of the granule cell (*GrC*), parallel fiber (*PF*), and Golgi cell (*G*) components are identical to those depicted in Figure 4, and therefore will not be re-illustrated. The responses of the second messenger components are presented in the upper three panels of Figure 5, wherein the *SMOn* and *SMOut* peaks do indeed move forward in time until they reach the time at which the unconditioned stimulus (*US*) has been presented. After arriving at the correct delay, these peaks oscillate slightly around the appropriate point in time as the LTD and LTP components of the *WC* equation compete near the point of equilibrium.

The lower two panels in Figure 5 illustrate the resultant responses of the Purkinje (PC) and deep cerebellar nucleus (*DCN*) cells throughout the simulation. Towards the end of the 19 trial training sequence, when the peak of *SMOut* activity has stopped sliding towards the *US* onset time, the peak PC inhibition level also equilibrates. In addition to changes in *WC*, there are also changes in *WN*. As *WN* increases across practice, the mossy fiber (*MF*) signal, which persists throughout the trial, is able to provide increased drive to the *DCN* cell, so even a small drop in PC activity allows a suprathreshold *DCN* response to be emitted.

To examine the timecourse of *WC* evolution across a set of learning trials, a simulation was conducted that comprised 100 trials at each of four inter-stimulus intervals (ISIs), with all

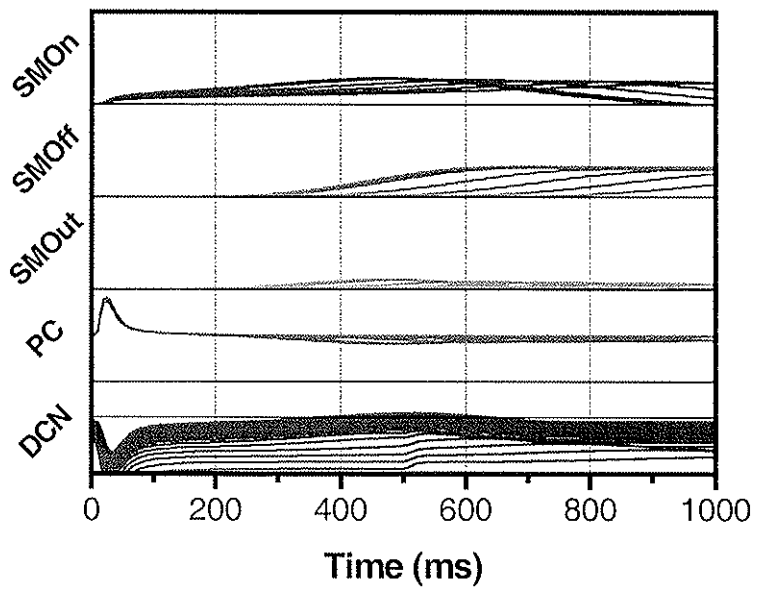

Figure 5. Timecourse of activities of the second messenger components within the Purkinje cell (upper three panels) and 'output' components (lower two panels) of the RSL during a classical conditioning simulation where *US* onset occurred at 500 ms and where *WC* changes made according to the learning law presented in Equation (6) accumulated across trials. The solid line in the *DCN* panel indicates the threshold level for *DCN* output. *SMOn* – second messenger onset component; *SMOut* – second messenger output response; *SMOff* – second messenger offset component; P – Purkinje cell; *DCN* – deep cerebellar nucleus.

other conditions remaining constant from the preceding simulation. A number of features of the *WC* trajectories, shown in Figure 6, are noteworthy. The small oscillations around the asymptotic time of the delayed peak become smaller in amplitude as ISI decreases. Another phenomenon is that the weight decrease is more rapid for longer ISIs. This occurs because *SMOn* activity, which plays a role in determining the amount of LTD that occurs in response to a climbing fiber (*CF*) burst, has more time to build up, due to the linear positive feedback, at longer ISIs. This ISI-learning rate relation is a property of this implementation of the slide and latch model, and is not a defining feature of this class of models.

In Figure 6, the trace for ISI equal to 250 ms requires special comment. This case was used to demonstrate that unlearning ("extinction") of an ATCR is within the scope of this formulation. For the 250 ms ISI case, the *US* was presented only for the first 50 trials. So, in the absence of climbing fiber (*CF*) input, continuation of the same mossy fiber (*MF*) input that was conditioned in the first 50 trials now produces LTP of *WC* values, and the response extinguishes.

Blocks of practice trials thus enable the learning and unlearning of the same sort of ATCRs in the RSL model as were produced by prior spectral timing models, but with no need for a population of elements preset to cover the entire learnable spectrum of ISIs.

Additional advantages from output-to-input recurrence.

In the recurrent implementation recommended by known cerebellar anatomy (Figure 2), the *DCN* output is fed back to form a new component of the *MF* representation of context/state. This new *MF* activity will lead to the formation of new context representations at the granule cell stage to be propagated by distinct *PF*s to PCs. These new *PF* inputs have normal eligibility to participate in ATCR learning guided by further *US* inputs. Therefore any PC cells that both receive the recurrence-excited *PF* inputs and project to the same *DCN* area can be recruited by slide & latch learning to enable a second timed output from the same *DCN* area. The recurrent slide and latch (RSL) model can therefore explain the multiple timed responses observed after classical conditioning in which

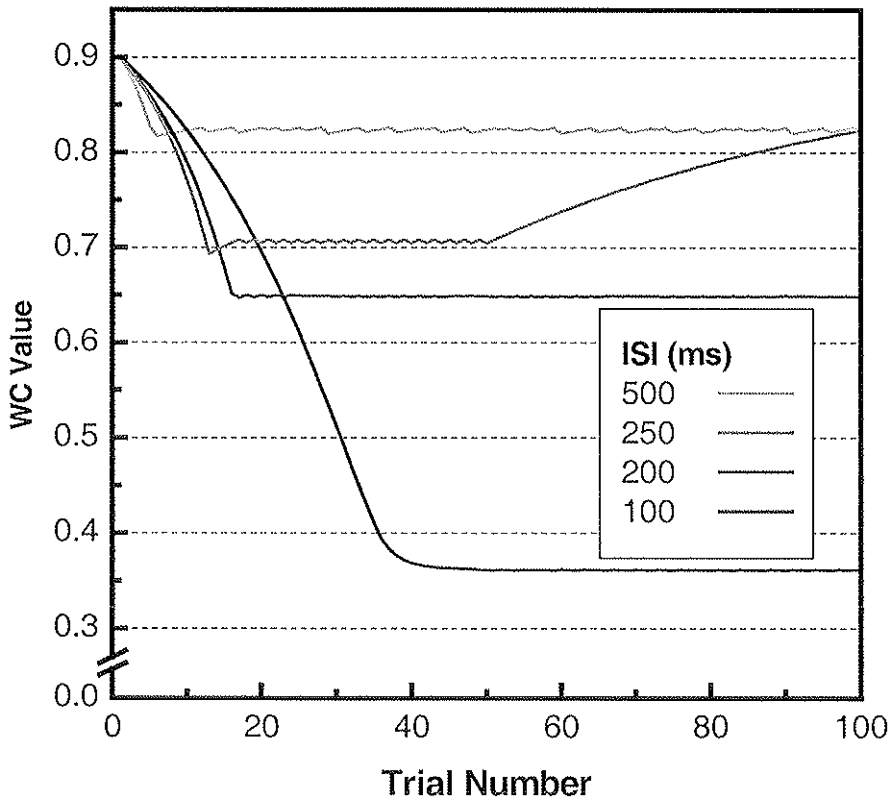

Figure 6. Change in *WC* value as a function of trial number during classical conditioning simulations where *US* onset occurred at various ISIs. In the simulation using a 250 ms ISI, the *US* occurred only in the first 50 trials – thereby inducing learning and providing an illustration of the reversible nature of *WC* values in the absence of teaching signals.

the *US* probabilistically follows a single *CS* by one fixed interval (e.g., 200 ms) on 50% of the trials and by a different fixed interval (e.g., 700 ms) on the other 50% of trials. Although a spectral timing model can easily explain multiple timed peaks without assuming recurrence, anatomical data suggest that recurrence is a robust feature of the cerebellar circuit. The assumption that recurrence is obligatory for stable multi-timed response leads to two unique predictions of RSL relative to a spectral timing model without recurrence. First, multiple timed peaks should not be learnable to a single *CS* if the *US* intervals are separated by less than approximately 90 ms. This is due to the minimal lag through the recurrent pathway, which includes a re-initiated second messenger cascade. Second, the secondary peak's magnitude should be increased relative to the primary peak that would occur after the same post-*CS* interval (at the same ISI) in a 1 *CS*, 1 *US* paradigm. This is because in RSL secondary peaks are actually timed relative to the onset of the primary *DCN* response, not relative to the initial *CS* input. More robust secondary responses are another advantage of the RSL model. By contrast, in a spectral timing model without recurrence, secondary responses are necessarily weaker than primary responses.

Conclusions

The recurrent slide and latch (RSL) model has several computational advantages over prior cerebellar models of ATCRs. First, the RSL model does not require a population to perform adaptive timing. Indeed, every state detector can depress any PC to which it projects after an independently adjustable lag via a single synapse. This makes the RSL model much more scalable than either spectral timing or spatial pattern models. Second, output-to-input recurrence allows non-initial elements of a response sequence to be timed relative to the immediate prior response rather than relative to the initiating stimulus for the sequence. In models of the type examined here, for all of which longer timed delays imply lower amplitude responses, such within-sequence resetting of the clock makes it much easier to maintain vigorous responses through a multi-response sequence.

The RSL mechanism also has some advantages as a biological model. First, it uses output to input recurrence, which was omitted from prior spectral timing and spatial patterns models of ATCR learning by cerebellum. Second, the RSL mechanism can fully incorporate the mGluR cascade model of Fiala et al. (1996), which elucidates an impressive range of physiological and biochemical observations. Third, the RSL model implicitly addresses recent data (e.g., Facchinetti, Hack, & Balazs, 1998) on cross-talk between LTD of AMPA receptors (corresponding to the *WC*s in the RSL model) and the kinetics of processing in the mGluR (second messenger) cascade. However, much more work is needed to fully address the full range of data on what has recently been shown to be an astonishing number of adaptive parameters of the cerebellar system.

References

Albus, J. S. (1971). A theory of cerebellar function. *Mathematical Biosciences, 10*, 25-61.

Albus, J. S. (1975). A new approach to manipulator control. *Journal of Dynamic Systems, Measurement, and Control, 97*, 220-227.

Bastian, A. J., Martin, T. A., Keating, J. G., & Thach, W. T. (1996). Cerebellar ataxia: abnormal control of interaction torques across multiple joints. *Journal of Neurophysiology, 76*(1), 492-509.

Bullock, D., Fiala, J. C., & Grossberg, S. (1994). A neural model of timed response learning in the cerebellum. *Neural Networks, 7*(6-7), 1101-1114.

Buonomano, D. V., & Mauk, M. (1994). Neural network model of the cerebellum: Temporal discrimination and the timing of motor responses. *Neural Computation, 6*(1), 38-55.

Church, R. M., & Broadbent, H. A. (1990). Alternative representations of time, number, and rate. *Cognition, 37*(1-2), 55-81.

Contreras-Vidal, J. L., Grossberg, S., & Bullock, D. (1997). A neural model of cerebellar learning for arm movement control: cortico-spino-cerebellar dynamics. *Learning and Memory, 3*(6), 475-502.

Dominey, P. F., Arbib, M., & Joseph, J. P. (1995). A model of corticostriatal plasticity for learning oculomotor associations and sequences. *Journal of Cognitive Neuroscience, 7*(3), 311-336.

Facchinetti, F., Hack, N. J., & Balazs, R. (1998). Calcium influx via ionotropic glutamate receptors causes long lasting inhibition of metabotropic glutamate receptor-coupled phosphoinositide hydrolysis. *Neurochemistry International, 33*(3), 263-70.

Fiala, J. C., Grossberg, S., & Bullock, D. (1996). Metabotropic glutamate receptor activation in cerebellar Purkinje cells as substrate for adaptive timing of the classically conditioned eye- blink response. *Journal of Neuroscience, 16*(11), 3760-3774.

Finch, E. A., & Augustine, G. J. (1998). Local calcium signalling by inositol-1,4,5-trisphosphate in Purkinje cell dendrites. *Nature, 396*(6713), 753-6.

Ghez, C., & Thach, W. T. (2000). The cerebellum. In E. R. Kandel, J. H. Schwartz, & T. M. Jesell (Eds.), *Principles of Neuroscience* (4 ed., pp. 832-852). New York: McGraw-Hill.

Gramsbergen, A. (1993). Consequences of cerebellar lesions at early and later ages: clinical relevance of animal experiments. *Early Human Development, 34*(1-2), 79-87.

Jordan, M. I. (1985). *The learning of representations for sequential performance.* Unpublished doctoral dissertation, U California, San Diego.

Marr, D. (1969). A theory of cerebellar cortex. *Journal of Physiology (London), 202*(2), 437-70.

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), 1708-18.

Peters, J., & van der Smagt, P. (in press). A scalable approach to cerebellar control. *Applied Intelligence.*

Raymond, J. L., & Lisberger, S. G. (1998). Neural learning rules for the vestibulo-ocular reflex. *Journal of Neuroscience, 18*(21), 9112-29.

Rhodes, B. J. (1999). *Learning-driven changes in the temporal characteristics of serial movement performance: A model based on cortico-cerebellar cooperation.* Unpublished doctoral dissertation, Boston University, Boston.

Spoelstra, J., Schweighofer, N., & Arbib, M. A. (2000). Cerebellar learning of accurate predictive control for fast-reaching movements. *Biological Cybernetics, 82*(4), 321-33.

Steuber, V., & Willshaw, D. J. (1997). How a single Purkinje cell could learn the adaptive timing of the classically conditioned eye-blink response. *Lecture Notes in Computer Science, 1327*, 115-120.

Takechi, H., Eilers, J., & Konnerth, A. (1998). A new class of synaptic response involving calcium release in dendritic spines. *Nature, 396*(6713), 757-60.

Thach, W. T., Goodkin, H. P., & Keating, J. G. (1992). The cerebellum and the adaptive coordination of movement. *Annual Review of Neuroscience, 15*, 403-442.

Wolpert, D. M., Miall, R. C., & Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Sciences, 2*, 338-347.