An Adaptive Sensorimotor Network Inspired by the Anatomy and Physiology of the Cerebellum

James. C. Houk
Department of Physiology
Northwestern University, Chicago, IL 60611
Satinder P. Singh, Charles Fisher, Andrew G. Barto
Department of Computer and Information Science
Lederle Graduate Research Center
University of Massachusetts, Amherst, MA 01003

COINS Technical Report 89-108

An Adaptive Sensorimotor Network Inspired by the Anatomy and Physiology of the Cerebellum†

James. C. Houk
Department of Physiology
Northwestern University, Chicago, IL 60611

Satinder P. Singh, Charles Fisher, Andrew G. Barto
Department of Computer and Information Science
Lederle Graduate Research Center
University of Massachusetts, Amherst, MA 01003

to be published as Chapter 14 in *Neural Networks for Control*
edited by W.T. Miller, R.S. Sutton, P.J. Werbos
MIT Press, 1990

COINS Technical Report 89-108
November 1989

Abstract—In this chapter we review the anatomy and physiology of the cerebellum, stressing new knowledge about information processing in cerebellar circuits, novel biophysical properties of Purkinje neurons and cellular mechanisms for adjusting synaptic weights. We then explore the impact of these ideas on designs for adaptive sensorimotor networks. A network is proposed that is comprised of an array of adjustable pattern generators. Each pattern generator in the array produces an element of a composite motor program. Motor programs can be stored, retrieved and executed using adjustable pattern generator modules.

---

† The authors wish to thank Steven Epstein for many stimulating ideas and helpful discussions in the course of this work. We also thank Tom Buchanan and Joyce Keifer for helpful comments on the manuscript and Michael Clarke for helping with the illustrations. This work was supported by ONR contract N00014-88-K-0339.

TABLE OF CONTENTS

1. INTRODUCTION

In spite of much recent progress in the field of robotics, it is widely acknowledged that the control of robot limbs is crude in comparison with the facile manner in which primates use their limbs. While primates may have more efficient effectors (muscles) and better sensors (sensory receptors), we believe that the most important reason for superior performance is because primates have more effective mechanisms for the adaptive control of movement.

The cerebellum is one of the main regions of the brain implicated in the adaptive control of movement (cf. Ito, 1984). This brain structure has long been celebrated as a beautifully organized network of massively interconnected neurons (Eccles, Ito & Szentagothai, 1967). The elegant structure of the cerebellar cortex has inspired the construction of perceptron-like models that are capable of learning complex functional maps between inputs and outputs (Marr, 1969; Albus, 1971), and a variety of adaptive controllers and filters based on these network models have been designed by subsequent authors (cf. Discussion).

These models were inspired mainly by the knowledge of the synaptic connectivity of the cerebellum that was obtained in the late 60's with electrophysiological and neuroanatomical methods (Eccles, Ito & Szentagothai, 1967). Knowledge regarding information flow through cerebellar pathways and cellular properties of neurons was scanty at that time, so major assumptions had to be made. We now know considerably more about these latter topics, although admittedly much is yet to be discovered.

While recent information about synaptic plasticity supports one of the key assumptions of these models, namely that synaptic inputs to cerebellar Purkinje cells are subject to adaptive modification, new data regarding information processing, cerebellar circuits, and biophysical properties of neural membranes were unanticipated by the earlier models. These new data suggest novel approaches that might appreciably impact our concepts of sensorimotor processing in adaptive neural networks. One purpose of this chapter is to provide a brief summary of cerebellar anatomy and physiology that incorporates these recent findings. However, the main purpose is to explore a novel approach to the design of adaptive sensorimotor networks, an approach that is inspired by our new knowledge of cerebellar function.

2. NETWORK ARCHITECTURE

Anatomical and physiological studies have revealed a large number of cell types in the cerebellar network and a formidable number of interconnections between these cell types. In addition to revealing complexity, these studies have also unveiled a high degree of structural order that is expressed at both microscopic and macroscopic levels (Brodal, 1981). On the microscopic level, each small region of the cerebellum contains approximately the same densities and ratios of neuron types, and local connections between the different types show approximately the same patterns from one region to another. On the macroscopic level, different categories of input project to different global regions of the cerebellum, and each of these regions in turn projects differentially to specific targets in the brainstem and diencephalon. In this section we discuss first the notion of modularity that has emerged from the macroscopic order seen in the cerebellar network. We then review the microscopic structure of the modules and lay the groundwork for treating these modules as adjustable pattern generators.

2.1 Modular organization of the cerebellum

The microscopic similarity of local circuitry coupled with macroscopic variety in input-output connectivity has suggested to a number of authors that the cerebellum is organized in a modular fashion (Eccles, Ito & Szentagothai, 1967; Voogd & Bigaré, 1980; Andersson, Ekerot, Oscarsson & Schouenborg, 1987; Ito, 1984). Module boundaries have differed from one author to another, and the module proposed here is yet another design. However, the basic notion has remained quite similar. Each module is presumed to operate in the same manner on whatever inputs come to it. Modules in different regions of the cerebellum are able to perform different functions in a behavioral sense by virtue of their specific input-output connections.

As a first approximation, one can consider the cerebellum as consisting of three medial to lateral regions (cf. Ito, 1984; Stein, 1986). (1) Much of the medial region connects with vestibulospinal and reticulospinal pathways, which fits with its specialization for the control of whole body movements and posture. (2) The intermediate region connects with the rubrospinal and corticospinal pathways, which fits with its specialization for the control of goal-directed limb movements. (3) The lateral region connects prominently with the premotor frontal cortex, and it appears specialized for target prediction, motor planning and other high-level modes of processing. Each of these global regions contains a large number of cerebellar modules.

The hypothesis developed here, that the cerebellum functions as an array of adjustable pattern generators, takes full advantage of the modular concept of cerebellar architecture. However, unlike most previously proposed cerebellar modules, an adjustable pattern generator is a module that includes feedback loops between the cerebellum and other brain structures, in addition to strictly cerebellar circuitry. Since each region of the cerebellum connects with different brainstem or cerebral structures, modular designs for adjustable pattern generators must accommodate to these regional specializations (Houk, 1989). Rather than attempting to discuss several variations of pattern generator modules, our approach here will be to focus on the intermediate cerebellum, and then upon one particular pathway linking the cerebellum to the spinal cord. Thus, we will emphasize the cerebellorubrospinal pathway which has been studied extensively in Dr. Houk's laboratory (cf. Houk & Gibson, 1987). This pathway is known to play a particular role in the control of limb movements (Kuypers, 1981). The reader may wish to consult Ito's (1984) excellent monograph for a discussion of the anatomy and physiology of other cerebellar systems.

2.2 Circuit diagram of a pattern generator module

Figure 1 summarizes some of the main connections in the cerebellorubrospinal pathway, and it also serves as a circuit diagram of the postulated module, an individual adjustable pattern generator. The output of this circuit is a rubrospinal fiber (for simplicity, let us consider just one). Each rubrospinal fiber carries a motor command from the red nucleus (R in Fig. 1) to the spinal cord where it acts upon motor neurons and interneurons, thus contributing to the control of limb movements. Rubrospinal fibers also send collaterals to the lateral reticular nucleus (L) which projects back into the cerebellum, both to the deep cerebellar nucleus (N) and to Purkinje cells (P) in the cerebellar cortex by way

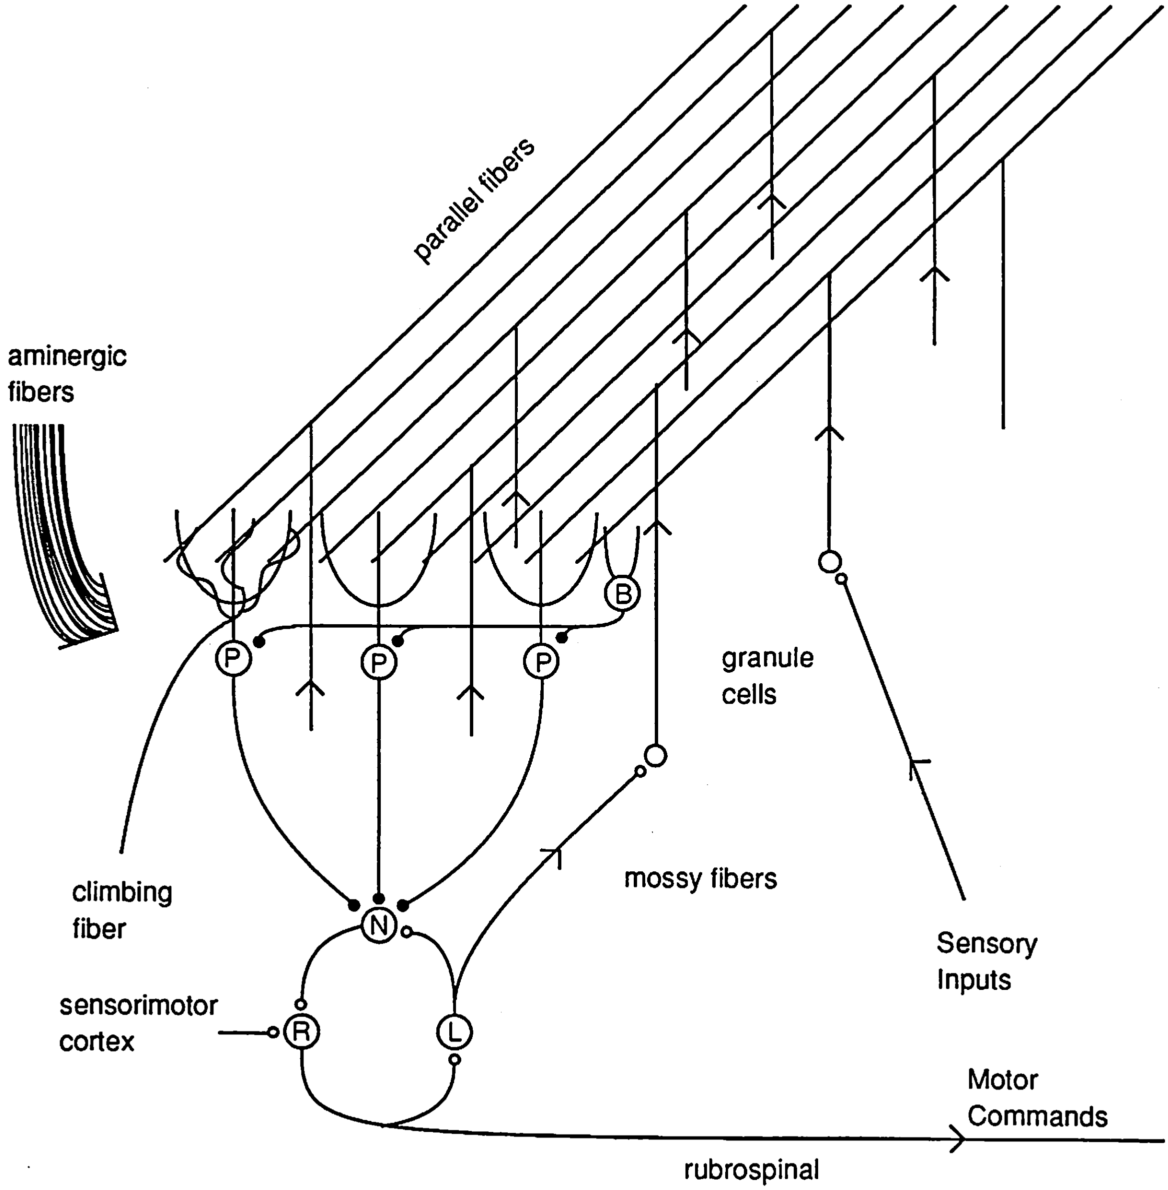

Figure 1: Circuit diagram of an adjustable pattern generator. This diagram applies specifically to cerebellorubrospinal pattern generators (but see Houk, 1989, for extensions to motor cortical loops). Note the positive feedback loop from cerebellar nuclear (N) cells, to red (R) nucleus, to the lateral (L) reticular nucleus, back to nuclear cells. Inhibitory input from a band of Purkinje (P) cells (only 3 shown) located in the cerebellar cortex regulates the intensity of loop activity. Basket (B) cells inhibit bands of Purkinje cells.

of the mossy fiber - parallel fiber pathway that is discussed later. The main input to the red nucleus comes from the cerebellar nucleus, but there are weaker inputs coming from the sensorimotor cortex and from sensory relays in the spinal cord.

The pathway from the cerebellar nucleus to red nucleus and back to the cerebellar nucleus via the lateral reticular nucleus is a positive feedback loop. Inhibitory input to the cerebellar nucleus from Purkinje cells normally restrains this positive feedback, but when the inhibition is blocked, sustained high-frequency discharge is observed. The adjustable pattern generator model treats this tendency for sustained loop activity as the fundamental driving force for the readout of a motor program (Houk, 1987). The detailed, time-intensity profile comprising a motor program is thought to be sculpted out of this tendency for sustained loop activity, the sculptor being the powerful inhibitory input to the loop from cerebellar Purkinje cells.

Purkinje cells are located in the cerebellar cortex where they are distributed in a monolayer. Each cerebellar nuclear cell receives inhibition from several hundred Purkinje cells (cf. Ito, 1984), only 3 of which are shown in Figure 1. The Purkinje cells innervating a given nuclear cell are situated in a longitudinal band that is aligned in an anterior-posterior, or parasagittal, plane. This arrangement is illustrated schematically in Figure 2 by a band of 12 Purkinje cells viewed from the surface of the cerebellar cortex. The box surrounding these cells is meant to outline the domain of cerebellar cortex that is associated with one adjustable pattern generator.

The dendritic trees of Purkinje cells (stippled in Figure 2) are fan-shaped and longitudinally aligned. Parallel fibers running perpendicular to the dendritic fans make synapses at points of crossing to provide a large number of excitatory inputs (up to 200,000) to each Purkinje cell. Parallel fibers are the axons of the numerous small granule cells (2 shown in Figure 1) contained in the granular layer of the cerebellar cortex. Granule cells receive their inputs from mossy fibers which transmit a variety of sensory and efference copy signals that will be described in the next section.

Considering mossy fiber signals to be elements of input vectors, the geometrical arrangement of the mossy fiber - parallel fiber pathway seems optimal for presenting a very large input vector to each Purkinje cell (Figure 2). Each signal is first distributed longitudinally by the branching of a mossy fiber to many granule cells (filled circles along the fiber). Then the axons of the granule cells, the parallel fibers, distribute the signals in the mediolateral direction to intersect the dendritic fans of Purkinje cells. In this manner an input vector with many similar elements can be presented to a large group of Purkinje cells. This group would comprise the Purkinje domains of many adjustable pattern generators.

A second major category of input to Purkinje cells comes from the inferior olivary nucleus. The axons of olivary neurons are called climbing fibers because of the distinctive manner in which they wrap themselves around Purkinje cells (Figure 1), and these fibers also send collaterals to the cerebellar nuclei (not shown for simplicity). Each Purkinje cell receives input from only a single climbing fiber, in marked contrast to the massive convergence of parallel fiber input onto individual Purkinje cells. Climbing fibers produce large excitatory potentials in Purkinje cells that are followed by periods of depression. In addition, and perhaps more importantly, they appear to transmit training signals that adjust the synaptic weights of the parallel fiber inputs onto Purkinje cells, as discussed more fully in section 6.

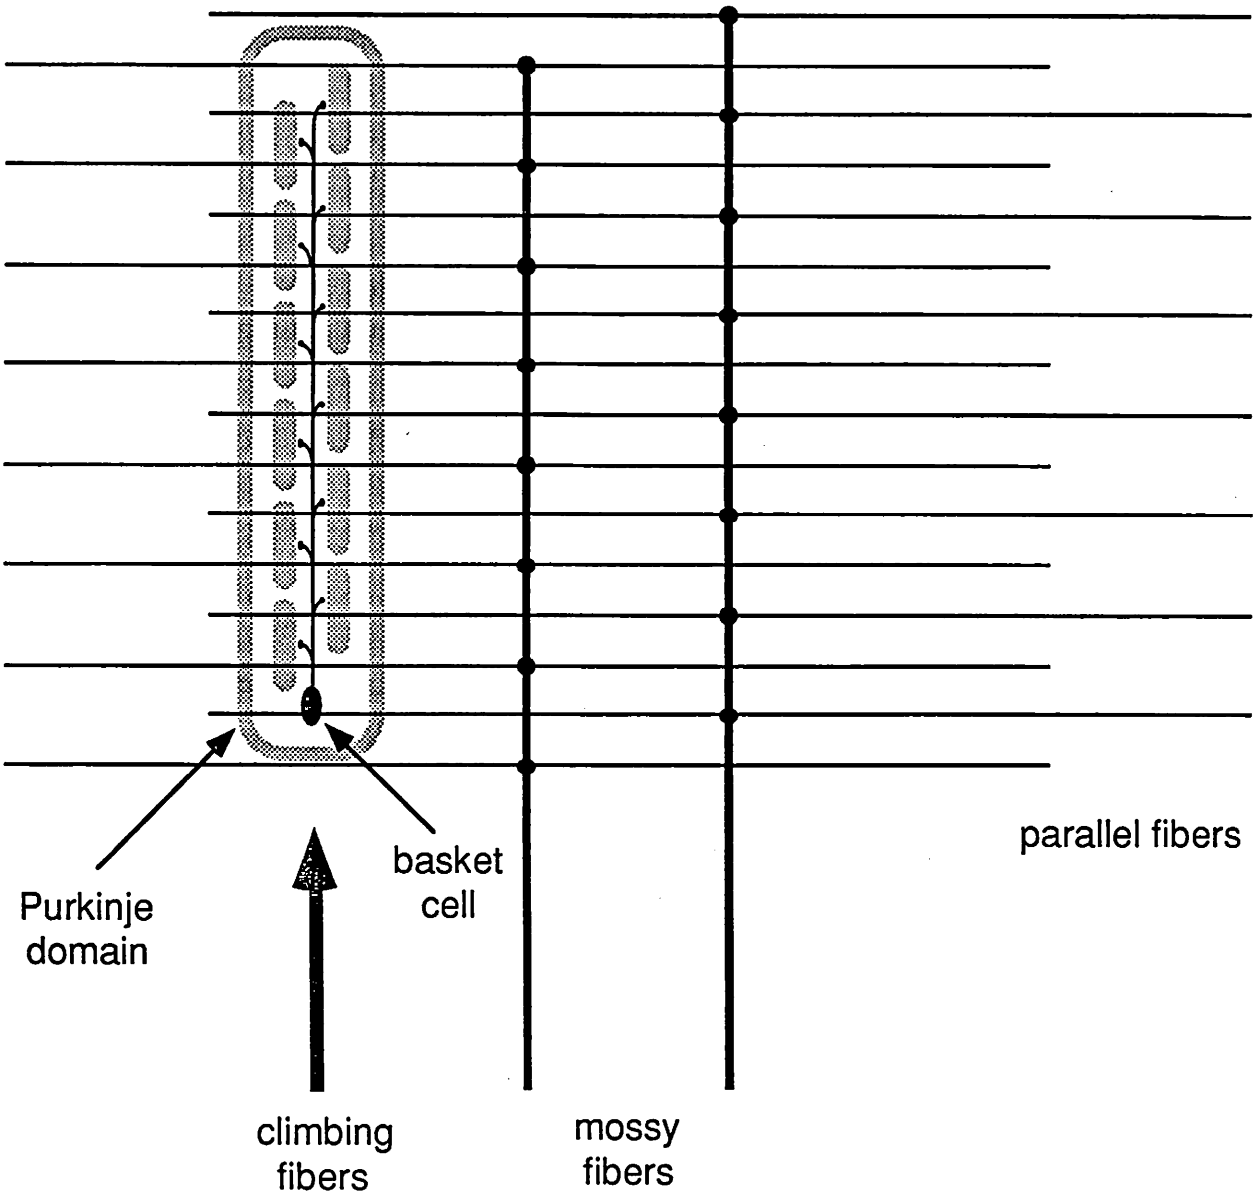

Figure 2: Organization of a Purkinje domain. The view is looking down onto the surface of the cerebellar cortex. Mossy fibers distribute input to a series of granule cells along the longitudinal (shown vertical) direction of the body, within the parasagittal plane of the cerebellum. The diagram shows two mossy fibers as heavy vertical lines punctuated by filled circles representing the granule cells they innervate. Each granule cell axon bifurcates to form a parallel fiber that travels medially and laterally (thin lines running left and right) to innervate large numbers of Purkinje cells (only 12 are shown for simplicity). The dendritic trees of Purkinje cells are fan-shaped and are portrayed as oblong stippled zones as viewed from above. A Purkinje domain is defined as a group of Purkinje cells that innervates a particular nuclear cell. Purkinje cells in a Purkinje domain are innervated by a coherent set of climbing fibers and receive inhibitory input from a common basket cell.

On a more global scale, the projections from the inferior olive to the cerebellum are exquisitely organized into longitudinal bands that are oriented perpendicular to the parallel fibers. Recalling that the Purkinje domains of adjustable pattern generators are also longitudinally organized, it is reasonable to postulate that the entire climbing fiber innervation of an adjustable pattern generator (large arrow in Figure 2) may derive from a small cluster of cells in the inferior olive. Since adjacent olivary neurons typically have similar response properties, the climbing fiber inputs to each adjustable pattern generator could be a fairly coherent training signal.

Considerably less is known about a third type of input to the cerebellum, the innervation supplied by aminergic fibers (Figure 1). The neurotransmitters used by these fibers are two amines, norepinephrine and serotonin, which are usually assumed to function as global modulators of neuronal activity. This would fit with the diffuse nature of the input seen in anatomical studies. These systems seem well suited to broadcast general messages to the whole network.

Several types of inhibitory interneuron are found in the cerebellar cortex, only one of which will be mentioned here. We have chosen to highlight the basket cell (B in Figure 1) because of its potential for coordinating groups of Purkinje cells subserving individual pattern generators. Basket cells receive the same types of input as Purkinje cells, although the innervation is less elaborate. They project asymmetrically to a longitudinal band of about 30 Purkinje cells in the manner illustrated in Figure 2. This pattern seems ideal for sending a common message to the cells in an individual Purkinje domain.

Having described the circuit diagram of an adjustable pattern generator in the cerebellorubrospinal pathway, we now need to discuss the manner in which these modules are likely to process information.

3. INFORMATION PROCESSING

The most direct approach to an assessment of information processing in cerebellar modules is to monitor the signals that are transmitted along the pathways to and from the cerebellum. Comparison of inputs and outputs should suggest the manner in which information is processed as it passes through cerebellar modules. The account provided here is based heavily on studies which have used the cerebellorubrospinal pathway as a model system for exploring sensorimotor processing through the cerebellum (Houk & Gibson, 1987), though it draws upon numerous additional sources, many of which are summarized by Ito (1984).

3.1 Inputs and outputs of cerebellar modules

Mossy fibers originate from several brain sites, transmitting a diversity of information about the internal state of the body and the state of the external environment. Some of these inputs are clearly sensory. They come quite directly from cutaneous, muscle or vestibular receptors. Other mossy fiber inputs are routed more indirectly via the cerebral cortex. They transmit highly processed visual, auditory or somatosensory information. Yet another category of mossy fiber transmits information about central motor commands. These signals are designated efference copy since they transmit copies of the efferent signals that are sent to the spinal cord and onward to muscles. Figure 1 shows an example of an efference copy signal that arises from collaterals of the rubrospinal tract and transmits its information through the lateral reticular nucleus.

In general, the discharge rates of mossy fibers are modulated over a wide dynamic range which permits them to transmit detailed parametric information. An example particularly relevant to the operation of adjustable pattern generators concerns sources of information about the progress of a movement. Signals derived from muscle proprioceptors provide sensory information about actual limb position whereas signals derived from efference copy provide predictions of the position the limb will soon assume if the motor command is successfully completed. Both types of mossy fiber have discharge rates that code limb position, and velocity, over a broad range.

The sole source of climbing fibers is from cells located in the inferior olivary nucleus. These neurons respond to various types of somatosensory, vestibular or visual stimuli. Different cells have different receptive fields, so information about the position of a stimulus or its direction of motion is well preserved. In contrast, olivary neurons appear to ignore information about the intensity or duration of a stimulus. This is because the cells have unusually long refractory periods which typically limits their responses to a single action potential. As a consequence, olivary neurons become selectively responsive to the occurrences of stimuli and not to their parameters. This is why climbing fibers have been characterized as somatic event detectors.

There are also efference copy inputs to the climbing fiber pathway, which appear to be inhibitory. These inputs gate off responsiveness to self-induced (or expected) stimuli, thus converting olivary neurons into detectors of unexpected sensory events. For example, one category of climbing fiber fires when the limb bumps into an object in the course of a movement but not when the movement is gracefully completed. Another category fires when the limb is disturbed at the termination of a movement. Later we speculate on how information of this type might detect dysmetria and then be used as a training signal for the adjustment of the synaptic weights of parallel fiber inputs onto Purkinje cells.

Less is known about the signals transmitted by aminergic inputs. The norepinephrine input comes from cells in the locus coeruleus that are believed to fire in relation to novel stimuli. The serotonergic input comes from cells in the raphe nucleus that discharge at low, steady rates which correlate with the level of arousal.

The signals on the output side of cerebellar modules appear to represent motor commands. The different categories of output cell (Purkinje, cerebellar nuclear and red nucleus) are similar to each other in several important respects. Their signals are closely related to movement performance and poorly related to sensory input. This is why they are considered to be motor, and not sensory, signals. The responses typically lead movement by about 100 msec and may correlate closely with motor parameters such as movement onset, velocity of motion or rate of force development.

The example in Figure 3 shows signals recorded from a red nucleus cell while a monkey subject tracked visual targets (Gibson, Kohlerman & Houk, 1985). The responses correlate closely with movement velocity and appear to represent velocity commands. When the target was a step (panel A), the cell discharged at a high rate for a short-duration, corresponding to the high-velocity, short-duration movement that followed. When the target was a slow ramp (panel B), the tracking movement showed several discrete phases, each of which was preceded by an appropriate velocity command. In contrast to the strong movement-related response, there was no noticeable sensory response to the visual target (target motion is shown by the dashed line in Figure 3).

Output neurons also show little or no response to sensory stimulation produced by passive joint rotation. This finding is particularly remarkable since the mossy fiber input to the cerebellum from muscle proprioceptors is very prominent, as discussed earlier. However, proprioceptive feedback during movement might still be important, so this was tested by interrupting ongoing movements (Harvey, Porter & Rawson, 1979; Gibson, Kohlerman & Houk, 1985). If continuous feedback from the limb were important, this maneuver would be expected to produce prompt and conspicuous changes in discharge rate. In contrast, the interruption of movement produced little or no change in discharge rate, indicating an absence of continuous feedback. However, the responses of red nucleus neurons did show evidence for an unusual kind of feedback (cf. Houk, 1987). There was a prolongation of the velocity command beyond the expected time of its termination in many of the trials in which the movement was interrupted. Later we propose a special type of feedback that accounts for these observations.

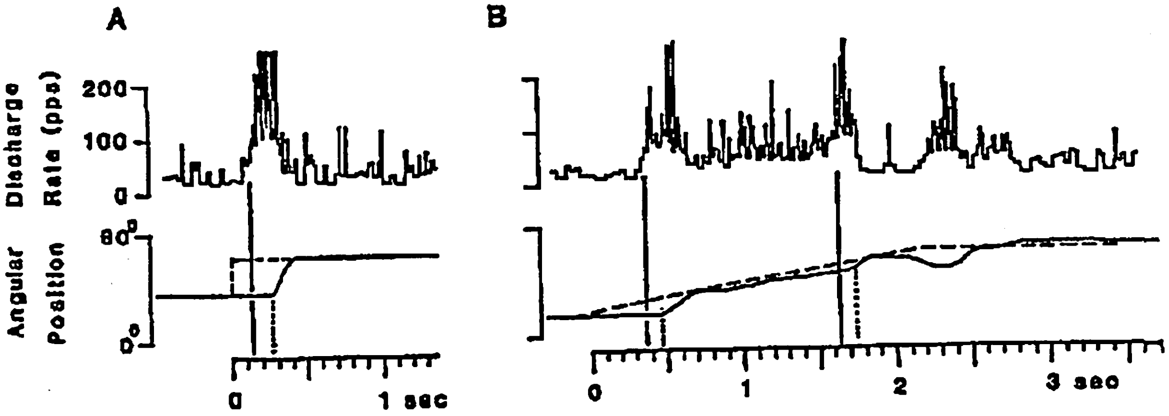

Figure 3: Examples of central motor commands. Panels A and B show responses recorded from a red nucleus cell while the monkey subject was tracking step and ramp targets respectively. The lower traces show the targets (dashed lines) and actual movements (solid lines) while the upper traces show the discharge rate of the cell. Discrete bursts of discharge precede each movement, and each movement segment. The solid vertical lines show some of the burst onsets, and the dotted vertical lines show the corresponding movements. Note that the rate of discharge within each burst corresponds to the velocity of movement, and the duration of each burst corresponds to the duration of the corresponding movement segment. In contrast, discharge rate correlates poorly with the time course of target motion.

In summary, comparison of the signals that enter and leave the cerebellum reveals some prominent differences. Mossy fibers appear to transmit a rich variety of parametric information to the cerebellum. Sensory inputs signal external state, and efference copy inputs furnish information about internal state. In marked contrast, climbing fibers transmit much less parametric information to the cerebellum. Instead they appear to be specialized for the detection of unexpected somatic events that may signify error conditions. The signals on climbing fibers might be suitable for training the cerebellar network. Finally, the signals present at the output stages of cerebellar modules are surprisingly insensitive to sensory inputs and instead appear to represent preprogramed motor commands. The finding that much sensory information enters the cerebellum whereas the outputs follow time courses that are relatively independent of sensory events is particularly striking. This contrast between sensory input and motor output places a major constraint on the type of information processing that must occur in cerebellar modules.

3.2 Processing mode of the cerebellum

The finding of a marked transformation in cerebellar modules between sensory-driven input and movement-related output suggests that motor commands are not created by simply combining inputs to the cerebellum. Instead, it implies that these commands are generated by mechanisms intrinsic to cerebellar modules that are capable of self-sustained activity, independent of the time course of sensory input. This is one of the main reasons for postulating pattern generation as the basic operation performed by cerebellar modules and for referring to these patterned outputs as motor programs.

There are two ways in which self-sustained activity might be produced in a neural network. One is by virtue of recurrent circuits such as the positive feedback loop between the cerebellum and red nucleus that was discussed earlier (Figure 1). The other mechanism is for the neurons themselves to have intrinsic dynamics that makes them capable of self-sustained activity. The model of an adjustable pattern generator developed in the next section incorporates both recurrent loops and intrinsic dynamics as mechanisms for the production of self-sustained activity.

If the extensive sensory input known to enter the cerebellum is not used to shape the time course of motor signals, what might its function be? We would postulate that sensory information has several alternative uses: (1) to preselect motor programs, (2) to trigger their initiation, (3) to influence when they terminate, and (4) to evaluate the success or failure of the patterns in controlling a motor behavior. As an analogy, consider how we control an electronic function generator to produce trains of electrical pulses. First we adjust all the dials on the stimulator to preselect pulse duration and amplitude, train duration and repetition rate. A transient input then triggers the entire preselected sequence. If we are not happy with the way the output looks, we readjust the dials on the function generator before trying it out again.

The mode of operation outlined in the previous paragraph represents a type of adaptive feedforward control. Elsewhere it has been suggested that most physiological control systems use adaptive feedforward mechanisms to generate command or reference signals and then use local feedback systems to translate these commands into actual performance (Houk, 1988a). Confining feedback to lower-level processes with relatively short time delays minimizes the stability problems inherent in closed-loop control systems. Also, one can capitalize on predictive capabilities and other special advantages of feedforward systems. What we actually end up postulating in the next section is a cerebellar module that in most respects operates as an adaptive feedforward processor. However, the postulated module does uses feedback, though in a limited manner that is less subject to the stability problems encountered in conventional feedback systems. Correspondingly, we will refer to this mechanism as a quasi-feedforward process.

4. PATTERN GENERATOR MODULES

In section 2 we described the circuit diagram of a cerebellar module that included reciprocal connections between the cerebellum and the brainstem. While some arguments were given for postulating that these modules function as adjustable pattern generators, we did not explain how they might actually operate. In section 3 we reviewed the signals that enter and leave the cerebellum and concluded that cerebellar modules are repositories of motor programs. They appear to use sensory inputs to adjust parameters and to start a motor program, but then operate in a quasi-feedforward manner until the program runs to completion. In the present section we describe how these features might result if we were to introduce elements capable of undergoing state transitions, inspired by the physiology of Purkinje cells and by the anatomy and physiology of the cerebellorubral loop.

4.1 Cerebellar elements capable of state transitions

Some neurons are reasonably well characterized by quasi-linear static properties which yield graded outputs (firing rates) over a range between threshold and saturation. Correspondingly, in most network models the neurons have outputs that are proportional to weighted sums, or some possibly nonlinear function, of the synaptic inputs. This type of static model may provide a reasonable approximation to the characteristics of some cells in the nervous system.

In contrast, Purkinje cells in the cerebellar cortex have unusual dynamic properties that warrant special consideration, since they might give rise to interesting computational capabilities. Their unusual properties arise because these cells and their dendritic membranes contain a high density and a rich variety of specialized ion channels. Particularly intriguing are the non-inactivating sodium and calcium channels that produce prolonged plateau potentials in response to brief depolarizing currents (Andersson, et al 1984; Llinás & Sugimori, 1980). Plateau potentials give rise to a bistability of membrane potential. A brief positive input leaves the membrane in a depolarized state that persists, and a brief negative input will reset membrane potential back to its prior state. These properties resemble hysteresis, except that the plateau potentials tend to reset automatically after a period of a few hundred msec.

Since plateau potentials are particularly prominent far out on the distal dendrites, each dendritic branch of a Purkinje cell may be capable of independent bistable behavior. As a consequence, the overall firing characteristics of the cell might exhibit multistable properties. While this is an interesting possibility, the lack of experimental evidence for multistability at the present time led us to choose a simpler, more abstract model of Purkinje cells.

The model utilized here is intended to capture the capacity for state transitions in its most elementary form. Purkinje cells are represented as simple bistable devices, having the characteristics graphed in Figure 4A. When the summed synaptic input exceeds an on-threshold value, the Purkinje cell is set to an on-state, and it is then assumed to fire at a constant rate. It stays on until the synaptic input falls below an off-threshold value at which point the cell is set to an off-state, and firing rate falls to zero.

Our model of the recurrent pathway between the cerebellum and red nucleus (Figure 1) also shows a capacity for state transitions, although the postulated mechanisms and characteristics are appreciably different from the simpler bistability of Purkinje cells. The postulated characteristics of the cerebellorubral loop are graphed in Figure 4B. This diagram shows: (1) an off-state characterized by no activity in loop neurons, (2) an on-state characterized by a whole range of intensities of loop discharge rate, and (3) two examples of transitions from the off- to the on-state.

The off-state is a condition in which most loop neurons are below threshold, and the output of the loop, transmitted by a rubrospinal fiber, remains constant at zero. The on-state is a condition in which positive feedback sustains continuing loop activity. When the loop is in its on-state, the intensity of activity is assumed to be regulated by the strength of the inhibitory input provided by Purkinje cells (Figure 1). Thus, if a larger number of cells in the Purkinje domain (Figure 2) of this pattern generator are in their on-state, the loop will receive more inhibition and be less active. Alternatively stated, there is a direct proportionality between the intensity of sustained loop activity (expressed as discharge rate in Figure 4B) and the fraction of Purkinje cells that are in their off-state. While the loop's capacity for state transitions is attributed mainly to the coupled effects of positive feedback and thresholds in loop neurons, some of these neurons may have non-inactivating ion channels that could also contribute to the expression of persistent states and state transitions.

Transitions of loop activity from the off- to the on-state are postulated to be triggered by extrinsic inputs to the loop, such as the input from sensorimotor cortex shown in Figure 1. A transition would occur when an input brings a sufficient number of loop cells above threshold, at which point positive feedback would sustain the elevated state. A transition back to the off-state is postulated to occur when all (or most) of the Purkinje cells return to their on-state, a condition of maximum inhibition. These state transitions in turn determine the onsets and offsets of motor programs, as discussed in the next section.

4.2 Execution of a motor program

The execution of a motor program begins with the arrival of a trigger signal, which serves as its start pulse. We assume that the positive feedback loop is initially in its quiescent state. While Figure 1 shows the trigger being sent to the red nucleus, any sufficiently strong transient excitation to any of the cells in the loop would initiate a transition to a state of sustained activity. While we do not pursue it here, a transition might also be initiated by a highly synchronous removal of Purkinje inhibition. Thus, there are several ways to start the execution of a motor program.

Once started, a program is regulated by the state of cells in the Purkinje domain of the pattern generator. If the fraction of cells in the off-state is large, there will be a transition to a state of intense loop activity as shown by transition 1 in Figure 4B. The correspondingly intense discharge of the red nucleus cell that participates in this loop will serve to specify a high movement velocity, and this command will be transmitted to the spinal cord in the rubrospinal fiber (Figure 1). In contrast, if many of the cells in the Purkinje domain are in the on-state, inhibition will restrain positive feedback, the transition will be to a state of modest loop activity (transition 2 in Figure 4B), and the pattern generator will specify a low velocity of movement. If nearly all of the Purkinje cells are in the on-state, the loop will be unresponsive to the trigger, and no motor command will be

A. Purkinje cell

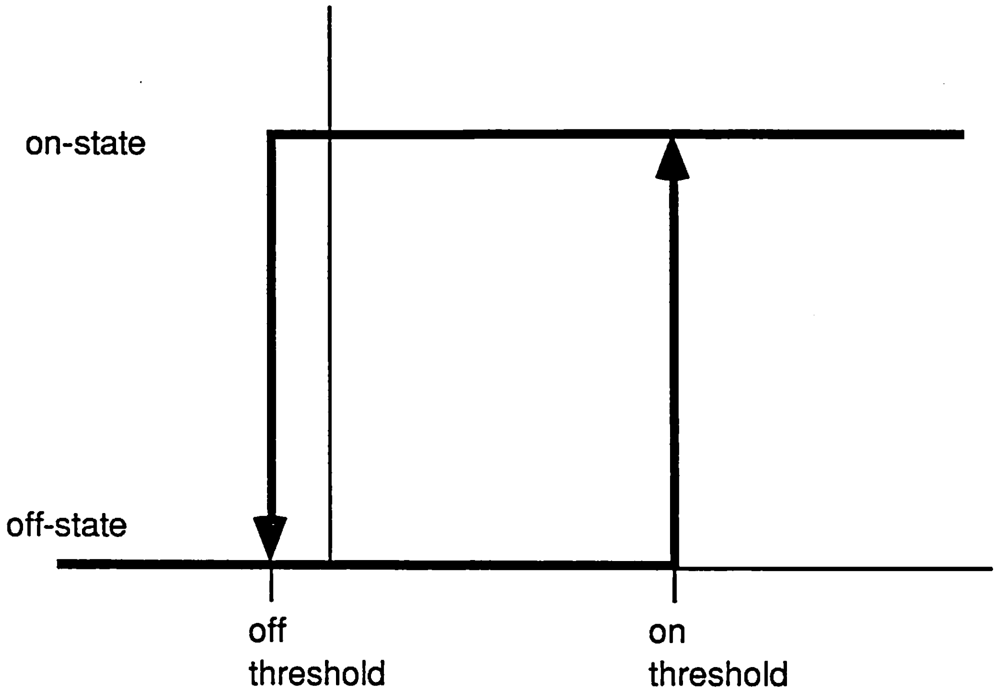

B. Loop

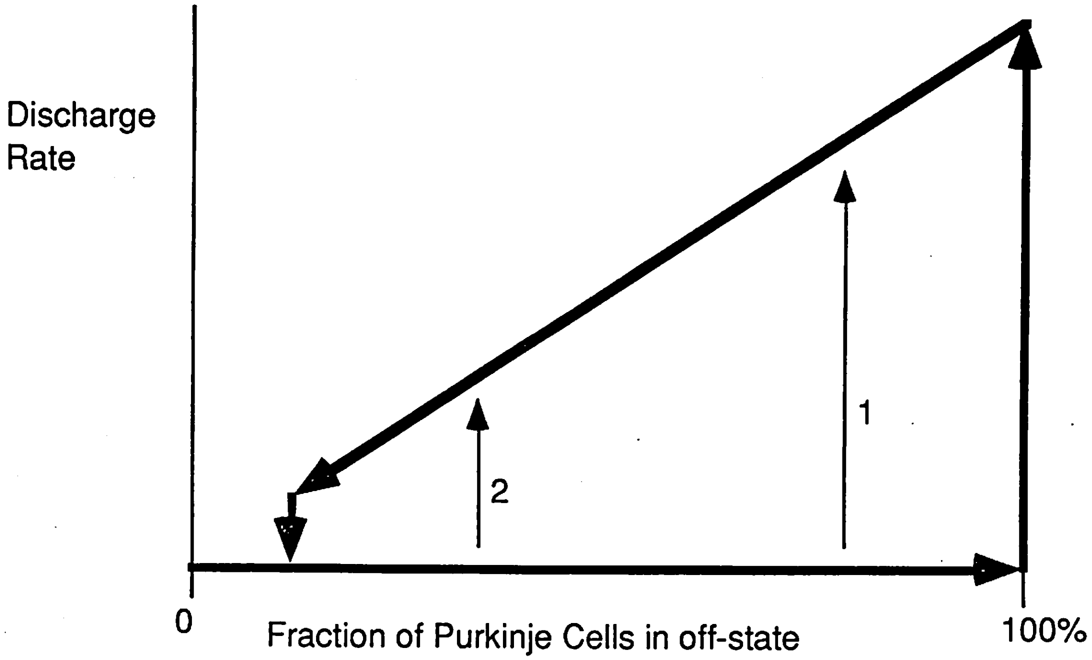

Figure 4: Operating characteristics assumed for Purkinje cells and for the cerebellorubral loop. A. Purkinje cells are assumed to have bistable properties. As a consequence, they are unresponsive to inputs except at two points, the on-threshold where transitions from the off- to the on-state occur and the off-threshold where transitions from the on- to the off-state occur. B. The cerebellorubral loop is also assumed to have off- and on-states, but with more complex features. When the loop is in its off-state, it has a zero output independent of its Purkinje input. Transitions to the on-state are caused by trigger signals, and, after a transition, output (discharge rate) depends on the fraction of cells in the Purkinje domain that are in their off-state.

generated. In this manner, the velocity commanded by an adjustable pattern generator can be preset to values anywhere between zero (no movement) and a maximum level.

Once started, the program will continue to execute in essentially an open-loop manner until parallel fiber inputs turn the off-state Purkinje cells back on. This occurs when the net synaptic input reaches the on-threshold of the Purkinje cells (Figure 4A). Figure 5 shows how position signals, conveyed by efference copy and proprioception in combination, can be used to decide when to terminate a motor program. In this example, we will, for simplicity, neglect the possibility of using velocity signals to enhance performance. Furthermore, we will not attempt to explain how rubrospinal efference copy signals, which code movement velocity, are converted into position signals. Although the mechanism for conversion is not understood, efference copy position signals have been recorded from mossy fiber terminals in the cerebellum (Van Kan, Houk & Gibson, 1987). We simply make use of these signals as one of the mechanisms for terminating a motor program.

The top trace in Figure 5 shows a motor program that commands velocity. It was initiated by the trigger signal shown at the bottom. The second trace represents the resultant movement, which is delayed by about 100 msec due to transmission time in motor pathways and mechanical lags in the limb. The third and fourth traces represent the two types of position signal, efference copy and proprioceptive, that are returned to Purkinje cells via mossy and parallel fibers. These signals inform the cerebellum respectively about the anticipated and actual progress of a movement. Their combined input (fifth trace) excites Purkinje cells, and the motor program is terminated when the combined input crosses the on-threshold of the cells.

Note that the motor program has to be terminated approximately 100 msec before the movement ends due to the neural and mechanical delays. One way to do this is to rely solely on efference copy information. Since these signals are propagated through rather direct pathways, they suffer little delay in transit back to Purkinje cells. However, they do not represent the actual movement, but only a prediction of the position the limb will achieve 100 msec later if nothing interferes with the movement. Another disadvantage of relying solely on efference copy position signals is that they are unlikely to supply accurate information about absolute limb position. Since efference copy signals are derived from collaterals of fibers that transmit velocity commands, the signals must be integrated to provide position information. Integration nearly always involves drift, so these signals are unlikely to supply reliable information about absolute limb position.

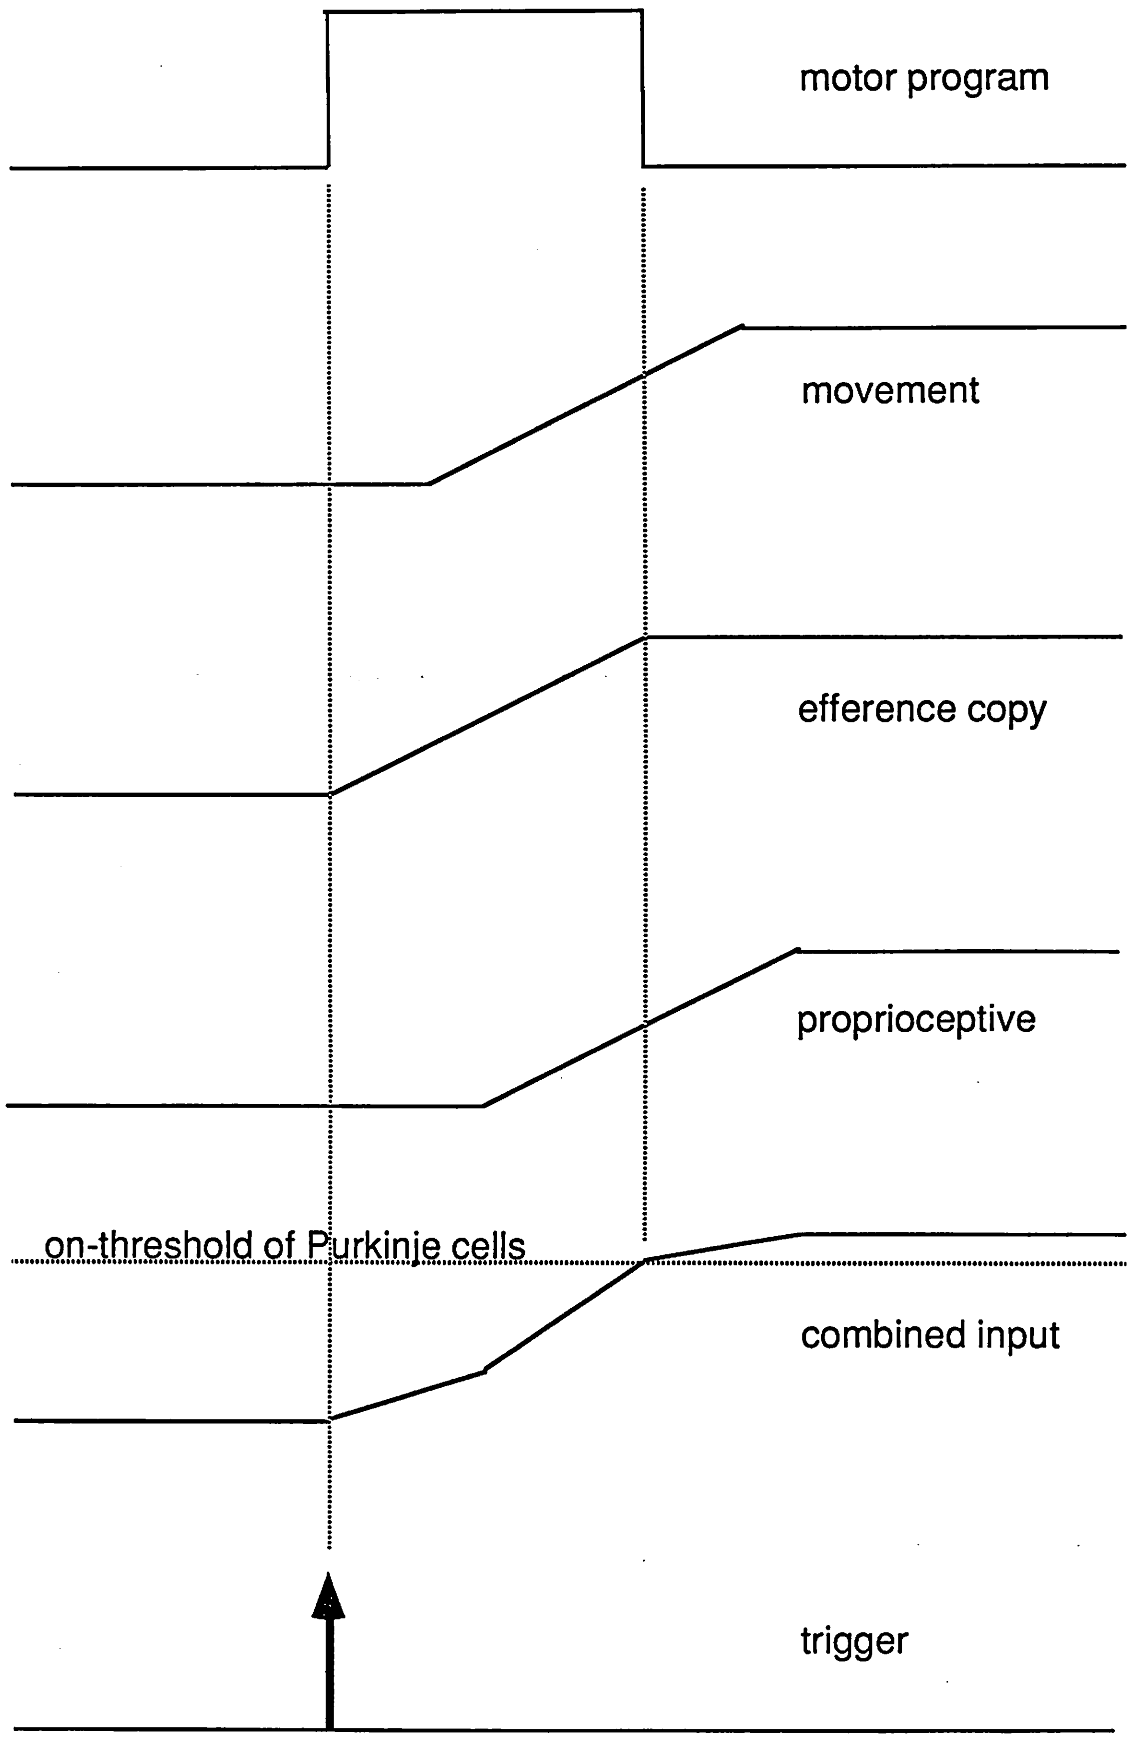
motor program
movement
efference copy
proprioceptive
on-threshold of Purkinje cells
combined input
trigger

Figure 5: Execution of a motor program. Time plots show relations between (1) a motor program that commands velocity, (2) the resultant movement, (3) a position signal provided by efference copy, (4) a position signal provided by muscle proprioceptors, (5) the combined position input (weighted sum of 3 & 4), crossing the on-threshold of Purkinje cells, and (6) the trigger signal that starts the motor program.

Proprioceptive mossy fibers are capable of signalling absolute limb position, and they naturally detect when disturbances prevent movements from being completed. However, from the standpoint of a pattern generator module, this information is delayed by more than 100 msec. If pattern generators used conventional feedback from proprioceptors, the system would become unstable and oscillate. In contrast, for the proposed model the effect of the delay is negligible when the movements are slow, since position would not change much during the delay period and program termination would occur at the appropriate time. However, proprioceptive signals would show progressively greater errors as movements become more rapid, and this would introduce delays in program termination. Our suggested solution is to use proprioceptive and efference copy signals in combination, which helps to overcome the disadvantages of using one or the other individually.

This type of quasi-feedforward, limited-feedback mechanism offers partial compensation for disturbances, without the danger of instability oscillations. Instability is prevented since the delayed feedback from the limb is used only momentarily to cause a state transition in Purkinje cells; before and after this switching point feedback has no effect on the output of the pattern generator. This limited use of feedback can in fact be useful. If an obstacle in the environment interferes with a movement sufficiently before the time of its expected termination, this will reduce the proprioceptive component of the combined input and thus delay the switching time of Purkinje cells. As a consequence, the duration of the velocity command would be extended in a manner that matches the observations reviewed in section 3. This extension of the motor command might in some cases be sufficient to overcome the obstacle and allow at least partial compensation for the disturbance.

5. ADJUSTING A PATTERN GENERATOR

In the previous section we discussed how a simple motor program specifying the velocity and duration of a movement could be executed in a quasi-feedforward manner by a pattern generator module in the cerebellum. While we indicated factors that determine the magnitude and duration of a motor program, we did not discuss how these parameters might be adjusted in order to control the velocity and duration of a movement. In the present section we discuss some mechanisms that are proposed to take place in a preselection period just before a motor program is executed and others that operate during the preselection period but also continue to function during program execution. These mechanisms allow the commanded velocity and duration to be adjusted to accommodate different initial and final positions and different degrees of urgency.

5.1 Moving to different final positions

The pattern generator module discussed in the previous section is tuned to produce movement to a particular final position by virtue of the synaptic strengths of position signals returned to its Purkinje domain via parallel fibers. The important quantity is the combined position input that derives from efference copy and proprioceptive components of the input vector (Figure 5). In the example given previously, we tacitly assumed that each cell in the Purkinje domain of the pattern generator receives an equivalent strength of this combined input. As a consequence, the off-state cells turn back on nearly synchronously as the combined input reaches a particular endpoint, thus commanding the movement to stop at a particular final position.

Now we wish to consider how adjustments of this endpoint parameter might be used to control movements to different target positions. There are two ways in which the endpoint of a pattern generator might be adjusted. One is to change the synaptic weights of all the efference copy and proprioceptive position signals. We consider this to be a long-term mechanism that is important in motor learning. Synaptic mechanisms for storing motor programs by calibrating targets to endpoints are discussed in the section on training.

For the moment we wish to consider immediate mechanisms that might permit different targets to adjust the endpoint of a pattern generator on a trial-to-trial basis. This requires additional parallel fiber inputs that transmit target information to all the cells in the Purkinje domain. These inputs would add to or subtract from the combined input shown in Figure 5 (subtraction would require inhibitory interneurons). Signals that add would cause the program to terminate sooner, at a shorter endpoint. Signals that subtract would cause termination at a later time, at a more distant endpoint. Note that the parameter that is specifically adjusted is endpoint, rather than program duration.

Mossy fibers bring to the cerebellum a variety of information that might be useful in specifying the end positions of movements. However, this information is coded in a variety of coordinate systems. Signals in the visual system are typically in retinal coordinates, though processed visual information in some cases incorporates eye position information derived from efference copy and/or head position information derived from the vestibular apparatus or from neck proprioceptors. Signals in the auditory system are typically in head coordinates. Some regions of the brain are thought to code information about remembered positions of targets in space. Sequential models have been proposed for combining various information to form maps of invariant target position (cf. Kuperstein, 1988). However, since Purkinje cells are exposed to virtually all of this information by way of mossy and parallel fiber input, it is conceivable that invariant target position information is assembled as a single computational step in the cerebellum.

These issues concerning how target information might be coupled to the execution of motor programs in the cerebellum clearly warrants further investigation. However, for the present purposes we wish to maintain our emphasis on the motor output side of the problem. Correspondingly, we provide a relatively simple example to demonstrate the essence of the proposed computational mechanism.

Consider a set of visual parallel fibers that code the position of a target in retinal coordinates. For didactic simplicity, we assume that each parallel fiber in this set has a small visual receptive field that does not overlap with its neighbors, though the proposed mechanism could easily be extended to fibers with large, overlapping receptive fields. We further assume that one of the mossy fibers fires a phasic burst of activity when the target is presented (this assumption is not necessary, but is included to illustrate another potential function for bistability in Purkinje cells). In cases such as this, when the target input is phasic, some short-term storage mechanism is needed to convert the phasic burst into a tonic effect on the Purkinje cell. The plateau potentials discussed earlier may be capable of doing this, since it has been suggested that local plateau potentials are produced by parallel fiber inputs to distal dendritic regions (Andersson et al, 1984). We further postulate that dendritic spines, which receive input from single parallel fibers, are individually bistable, and could thus be capable of a short-term storage of local responses.

The mechanism proposed in the previous paragraph converts a phasic visual input into a tonic local response. The latter then adds to the combined position input in Figure 5. As a consequence, the endpoint of the pattern generator will be adjusted to a shorter position, and the motor program will command a smaller movement. Now consider the set of visual parallel fibers, each of which fires to a different target. Each target will produce a different local event in a Purkinje cell's dendritic tree, and the strength of each event will depend on the synaptic weight of the particular parallel fiber. If we assume that the synaptic weights have already been adjusted to establish a correct correspondence between target position and final position of the movement, then different targets would automatically adjust the endpoint of the pattern generator so as to produce an appropriate amplitude of motion.

The previous example has shown how visual target information can adjust a pattern generator so that it controls movements to different target positions. Although this simple mechanism could function to control eye movements, it would not be effective for controlling limb movements. This is because target information is supplied in a retinal coordinate system that moves, depending upon the orientation of the eyes and head in space, whereas a pattern generator will control the end position of a limb movement in a coordinate system centered about the shoulder. However, if information about eye and head position were included in the input vector to the Purkinje domain, and if these signals were suitably calibrated by a training procedure, then these inputs would automatically adjust the endpoint of a pattern generator in order to accommodate different positions of a target in space. This gedanken experiment demonstrates the generality of the proposed mechanism for adjusting the final position of a movement. A large number of factors can easily be included in the computation, limited only by the 200,000 count on the number of parallel fibers.

5.2 Starting from different initial positions

A convenient feature of the present model is that it automatically adjusts motor programs to compensate for different initial positions of the limb. In essence, this is because any program, once triggered, will run until the combined position input reaches a particular endpoint. If the limb is initially farther from that endpoint, the program will simply run for a longer duration, and at a higher intensity, until the endpoint is reached.

The mechanisms that mediate these adjustments can be appreciated by extrapolating from Figure 5. To make this example concrete, let us assume that the pattern generator is wired to control elbow flexor muscles and that its proprioceptive input is from stretch receptors in the elbow extensor triceps. The targeted position is elbow flexion to, let us say, 45 degrees, and the movement illustrated in Figure 5 starts from 90 degrees. Instead, if the initial position were 135 degrees, the proprioceptive input would be proportionately smaller in the preselection interval. This would have two effects on the pattern generator.

The first effect is an increase in the duration of the motor program. When the initial position is 135 degrees, the proprioceptive component of the combined position input would be smaller than shown in Figure 5. Assuming for the moment that the movement is executed at the same velocity, it would take longer for the combined input to reach the on-threshold of Purkinje cells. Thus, the motor program would automatically have a longer duration in proportion to the increased distance between the initial limb position and its final position, as determined by the position of the target.

A second mechanism that compensates for different initial limb positions operates in the preselection period prior to the execution of a motor program. Recall that the velocity commanded by a motor program is regulated by the fraction of Purkinje cells in their off-state (Figure 4B). If more Purkinje cells were to be put in an off-state just before a trigger arrives, the trigger would produce a transition to more intense loop activity. The discharge rate of the rubrospinal fiber would be higher thus commanding a higher velocity of movement. This adjustment of commanded velocity is postulated to depend on an interplay between inhibitory and excitatory mechanisms that takes place in the preselection period. Inhibition from basket cells and other interneurons tends to set Purkinje cells to an off-state while excitation from parallel fibers tends to return them to an on-state. We will follow the convention of designating off-state Purkinje cells as being the selected ones, since it is specifically these cells that determine commanded velocity of the motor program.

Inhibitory input to Purkinje cells from basket cells is ideally suited for coordinating the preselection process. The distinctive branching pattern of basket axon terminals (Figure 2) tends to confine inhibition to individual Purkinje domains. The cells in a given domain are thus driven to the vicinity of their off-thresholds where they become sensitive to the balance between excitation and inhibition. The amount of excitation in this period depends on both initial and final position. Using the elbow flexion example introduced earlier, a target position of 45 degrees would produce less excitation than a target at 90 degrees, and this would result in more Purkinje cells being selected. Similarly, an initial limb position of 135 degrees would produce less excitation than would 90 degrees, which again would result in more Purkinje cells being selected. In this manner, the number of cells selected from a Purkinje domain will depend on the distance required to be traveled between the initial position of the limb and the desired final position designated by the target.

Ordinarily these two mechanisms for adapting motor programs to different initial positions would operate in combination. The preselection of more Purkinje cells would increase the commanded velocity. This would help compensate by making the movement faster. Assuming that the former does not fully compensate for a more distant initial position, the difference would automatically be made up by a longer duration command. The same mechanisms would work together to adjust the commanded velocity and duration to smaller values if the initial limb position were closer to the desired final position.

5.3 Changing velocity

A given movement can be made at a variety of velocities depending on the urgency of the situation. Since the commanded velocity depends on the fraction of Purkinje cells in the off-state, a relatively simple mechanism for varying this parameter would be to let a generalized urgency signal modulate the responsiveness of basket cells. This would potentiate the intensity of basket inhibition sent to all the cells in a Purkinje domain, thus increasing the number of selected Purkinje cells. Changing the number of selected cells in this manner would control the commanded velocity without affecting the endpoint of the movement. We propose that this modulating function is provided by one of the diffuse aminergic inputs to the cerebellum. Since these inputs branch to widespread areas in the cerebellar cortex, the urgency of movements could be set globally and thus influence all of the pattern generators that might take part in a particular action.

5.4 Moving in different directions

We assume that a given pattern generator targets a particular set of muscles as a consequence of the branching pattern of the descending rubrospinal fiber. Since neighboring rubrospinal fibers have different branching patterns, neighboring pattern generators will target different sets of muscles and control motion in a variety of directions. As a consequence, the way to make movements in different directions is to select different pattern generators, or more likely, different combinations of pattern generators that together produce a vector in the desired direction. This strategy fits well with the hypothesis that the cerebellum functions as an array of adjustable pattern generators (Houk, 1988b). The output of the array can be thought of as a composite motor program, with each pattern generator producing one element of the composite program. We have not yet dealt extensively with the problem of preselecting a whole set of pattern generators in order to control movements in space. However, it would seem that basket cells might be particularly important in this aspect of preselection.

6. TRAINING ADJUSTABLE PATTERN GENERATORS

A variety of evidence suggests that the cerebellum is important for sensorimotor learning. Patients with cerebellar lesions produce dysmetric movements, i.e., their movements are of the wrong size or in the wrong direction (Holmes, 1939). Normal subjects can be induced to make dysmetric movements by, for example, distorting visual input with a prism, but then they readjust their movements to make them normometric. Cerebellar patients apparently lack the ability to adaptively adjust motor programs so as to command accurate movements. Similarly, experimental lesions of the cerebellum in animals severely impairs adaptive mechanisms for coordinating several types of eye movement (Robinson & Optican, 1981). Cerebellar lesions also interfere with the formation of conditioned reflexes, an important model for sensorimotor learning in animals (Thompson, 1986; Yeo, 1987). It seems clear that the cerebellum is important for learning well-adapted sensorimotor functions.

In section 5 we developed a method for representing motor programs in adjustable pattern generators. We further indicated how the size of a movement controlled by one of these pattern generators will automatically accommodate different initial limb positions and different target positions, provided that the network is properly tuned. In the present section we discuss adaptive network strategies that might be used to tune an adjustable pattern generator to produce accurate movements. The adaptive strategies proposed in this section are, like the operational features proposed earlier, inspired by the anatomy and physiology of the cerebellum.

6.1 Synaptic plasticity

From the standpoint of pattern recognition schemes, a particularly favorable site for adaptive modification is the synapse between a parallel fiber and a Purkinje cell (Marr, 1969; Albus, 1971). Each Purkinje cell is exposed to a large input vector consisting of 200,000 parallel fiber signals that pass through its fan-shaped dendritic tree. A Purkinje cell could be trained to process specific input patterns if there were an appropriate mechanism for adjusting the weights of parallel fiber synapses. In this section we summarize current hypotheses regarding cellular mechanisms that mediate plasticity at parallel fiber synapses. We then propose a training rule that can be applied to the problem of teaching a pattern generator to move to different targets.

There is now considerable evidence that conjunctive activation of climbing and parallel fiber inputs to Purkinje cells modifies the weights of parallel fiber synapses (Ito, 1989). These data provide important support for the proposed training function of climbing fiber signals. Figure 6 summarizes several of the steps that appear to be involved in this form of synaptic plasticity (Ekerot, 1984; Crepel & Krupa, 1988). It is helpful to trace these steps in some detail, since they provide considerable insight as to the likely training rules that apply to the cerebellar network.

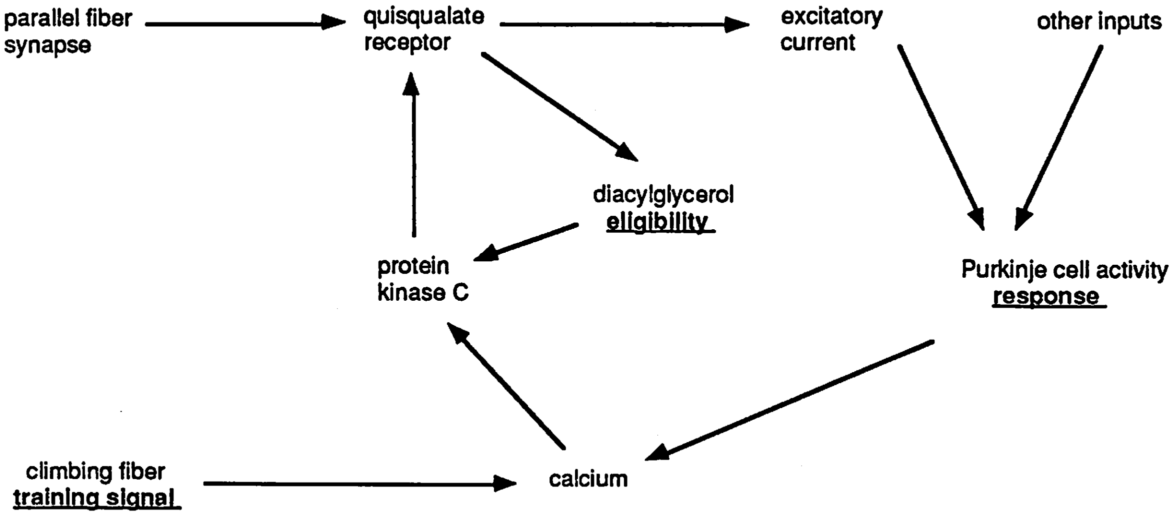

Figure 6: Flow diagram showing the electrical and chemical signals believed to mediate synaptic transmission and to regulate the synaptic weights for parallel fiber synapses onto Purkinje cells. The diagram also identifies (underlined labels) the correspondences between these cellular signals and the three factors in eqn(1), the weight adjustment equation proposed here. One factor is a training signal transmitted by a climbing fiber; it is postulated to signal when the movement falls short of a target. The second factor is the response of the Purkinje cell, as measured by its output activity. The third factor is the eligibility of each individual synapse for weight change; it is postulated to correspond with a local intracellular messenger, diacylglycerol.

Parallel fiber synapses occur on spines that extend from the dendrites of the Purkinje cell. Activity in a parallel fiber causes release of the neurotransmitter glutamate which binds to quisqualate receptors located in the adjacent dendritic spine, as indicted in Figure 6. This reaction between transmitter and receptor has two effects: (1) it opens ion channels leading to depolarizing current that excites the Purkinje cell and (2) it activates the intracellular messenger diacylglycerol within the membrane of the dendritic spine. While the activation of diacylglycerol is likely to be a local event, excitatory current has global effects on Purkinje cell activity. It combines with currents produced by many other synapses to regulate the level of activity of the Purkinje cell, designated the response in Figure 6.

We postulate that diacylglycerol marks synapses as being eligible for modification. Eligibility is a concept that has been useful in other models of learning that were inspired by biology (Klopf, 1982; Sutton & Barto, 1981). In these former models, eligibility was considered to have a short-term memory that decayed slowly after synaptic activity had ended. This trace feature might also apply to parallel fiber synapses, though we are not assuming this at present.

The same Purkinje cell that receives many parallel fibers is innervated by a single climbing fiber, designated as the training signal in Figure 6. The climbing fiber terminals spread over most of the cell body and stem dendrites. Although the climbing fiber does not fire very often, when it does, it produces a powerful global effect on the cell (Ekerot, 1984). This begins with an excitatory synaptic current which then activates calcium channels, the same channels discussed in section 4.1. Calcium channel activation produces a spike followed by a plateau potential with the latter spreading into the distal dendrites. These events cause substantial amounts of calcium ions to enter the cell, which is important since calcium also functions as an intracellular messenger. Since the duration of the plateau potential is very sensitive to the level of Purkinje cell activity, the amount of calcium that enters the cell becomes a function of the product (training signal) x (response) as illustrated in Figure 6.

The modification of parallel fiber synapses appears to be mediated by a substance called protein kinase C (Crepel & Krupa, 1988). The activation of this enzyme requires two factors, calcium and diacylglycerol (Nishizuka, 1986). This requirement for coactivation links synapse eligibility to the global components of the training rule described in the previous paragraph (Figure 6). The activated form of protein kinase C then modifies quisqualate receptors, causing them to be less sensitive to the neurotransmitter glutamate. These events thus produce a decrease in the synaptic weights of all eligible parallel fiber synapses. The experiments suggest that synaptic weight decrements can last for one to several hours.

Efficient training requires a mechanism for incrementing as well as decrementing synaptic weights. In the cerebellum, relatively little is known about incrementing mechanisms. However, there is some evidence for a potentiation of synaptic weight under conditions in which parallel fibers are activated without conjunctive climbing fiber activity (Sakurai, 1987). The parallel fiber activity would create eligibility in the form of diacylglycerol, and the absence of a climbing fiber discharge would result in modest or low levels of intracellular calcium. Presumably the lowest levels of calcium would occur if Purkinje cells were inactive in combination with an absence of climbing fiber input. In the training rule chosen here, we postulate that these lowest levels of intracellular calcium are required for increments in synaptic weight of eligible synapses. We further postulate than no appreciable changes in synaptic weight occur with modest levels of calcium.

The rules proposed for incrementing and decrementing can be combined into the following overall equation for adjusting synaptic weights:

$$\delta w_{ij} = \alpha (1.0 - c_t)(1.0 - y_i) \varepsilon_{ij} - \beta c_t y_i \varepsilon_{ij} \qquad (1)$$

where w_{ij} refers to the synaptic weight between the j^{th} parallel fiber and the i^{th} Purkinje cell and δw_{ij} represents the change in weight during one time step. Increments and decrements have learning rate coefficients α and β respectively. A climbing fiber response (designated by $c_t = 1.0$) is required for a decrement whereas its absence ($c_t = 0$) is required for an increment. Activity of Purkinje cells ($y_i = 1.0$) is assumed to inhibit increments and enable decrements. If the climbing fiber does not fire and the Purkinje cell is active, or if the climbing fiber does fire and and Purkinje cell is inactive, the weight is unchanged. Finally, the magnitudes of weight changes are assumed to be proportional to the eligibility of the particular synapse, where eligibility ε_{ij} can vary between 0 and 1.0 depending on the firing rate of the parallel fiber.

6.2 Simulation model

In section 5.1 we indicated how target information transmitted by parallel fiber input can be used to adjust the endpoint of a pattern generator, thus controlling movements to different final positions. As a specific example, we outlined how a set of visual targets, each designated by activity in a different parallel fiber, could be made to specify endpoints that would produce movements to appropriate target positions, provided the network were properly tuned. The problem of tuning the network reduces to the selection of an appropriate set of synaptic weights for the target lines. On what basis might these synaptic weights be adjusted to match particular endpoints with particular targets in space, and how might a pattern generator correct dysmetria? In the present section we set up a simulation model of an adjustable pattern generator, and in the following section we describe the results of simple simulation experiments which suggest that these modules can be trained to command accurate movements.

The simulation model of an adjustable pattern generator consists of one cerebellorubral loop (Figures 1 & 4B) innervated by a domain of 12 Purkinje cells (Figures 2 & 4A). The output of the pattern generator is a motor program that commands movement velocity. This command is zero prior to the arrival of a trigger signal, and then jumps to a discharge rate R given by

$$R = V_{max} (1 - \Sigma y_i / N) \qquad (2)$$

where N represents the total number of Purkinje cells in the domain, 12 in this example, and y_i is the output of the ith Purkinje cell. This output can be either 0 or 1.0, and the synaptic effect on the loop is assumed to be the same for all of the Purkinje cells. R thus assumes values between 0 and a maximum commanded velocity V_{max}, depending on how many Purkinje cells are preselected.

At each time step t, Purkinje cells receive an input vector $X_t = [x_0(t), x_1(t), x_2(t), x_3(t), x_4(t)]$. Lines 0, 2, 3 and 4 represent parallel fibers whereas line 1 represents a basket cell axon. This vector was assumed to be identical for all cells in the Purkinje domain. Figure 2 shows how an input vector with two parallel fibers and one basket axon could be distributed to an entire domain of Purkinje cells.

Although Purkinje cells receive several sources of position information, for the purposes of the present simulation we lumped these into a single, real valued, continuously updated and undelayed position signal, and presented it on line x_o with a fixed weight of 0.01. The combined use of accurate but delayed, together with less accurate but undelayed, position inputs will be important to explore in the future.

The inhibitory input from basket cells was presented on line x_1. Basket inhibition is postulated to be a transient input that occurs in the preselection period. It drives Purkinje cells near to their off-threshold where they become sensitive to the other parallel fiber inputs. Depending on these other inputs, a cell will be switched to its off-state (i.e., selected for participation in the movement) or left in its on-state. We further assumed that a stochastic process added noise to the off-threshold of Purkinje cells. The importance of this noise term is that it randomizes the preselection of Purkinje cells in successive trials, helping to insure that all cells in the domain are exposed to the full training regimen. We found empirically that using a basket input weight of -0.6 counterbalanced the excitatory target and position lines such that approximately half the Purkinje cells would be preselected for average length movements.

The remaining lines of the input vector were used to indicate the presence or absence of 3 different targets. When a given target was present, the parallel fiber line was set to a value of 1.0, whereas it was 0 when the target was absent. The pattern generator was taught to command movements that acquired these targets by a training regimen that adjusted the synaptic weights of the target inputs, as will be described later.

We chose a 100 unit long, one-dimensional external world in which to run our simulations. Movement position along this domain was determined by numerically integrating the velocity command. In the future it will be interesting to explore more complex transformations between the velocity command and the movement. We arbitrarily selected values of 28, 52 and 95 to represent the desired endpoints for the 3 targets. One of the target lines was randomly chosen and set to 1.0 at the start of a preselection period. Similarly an initial position somewhere in the range between 0 and the target position was randomly chosen. Then the basket cell fired and a certain number of Purkinje cells was selected, based on whether or not the net input fell below the cell's off-threshold. At this point the pattern generator was fully adjusted, i.e., all of its parameters were set by the preselection process. Upon the arrival of a trigger signal, it then executed its program.

After exposing this circuit to the training regimen described in the next section, the pattern generator began to command quite accurate movements to any of the 3 targets, starting from any initial position smaller than the target position. Initial positions larger than the target position resulted in zero output. Thus, like the muscles they control, adjustable pattern generators are capable of commanding movements only in one direction. Figure 7 shows examples of movements performed from two initial positions to a target located at position 95. Note how velocity and duration both increase when the larger

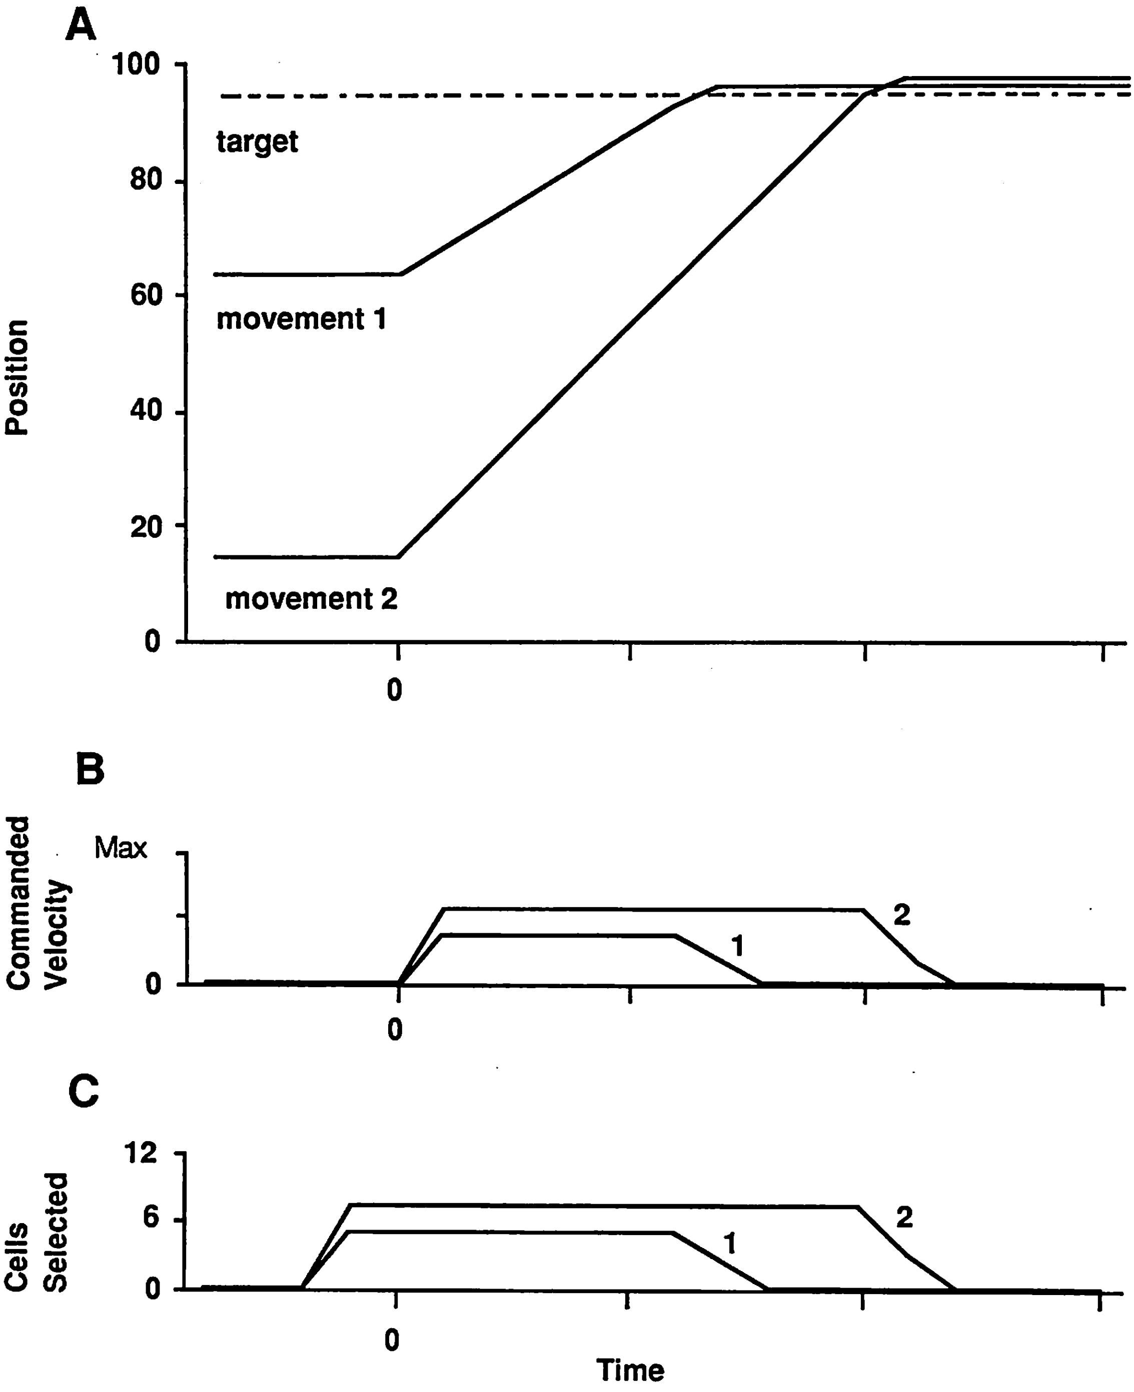
A
100
80
60
40
20
0
Position
target
movement 1
movement 2
0
B
Commanded Velocity
Max
0
2
1
0
C
Cells Selected
12
6
0
2
1
0
Time

Figure 7: Responses of the simulation model obtained after a period of training. Both of the simulated movements shown in Panel A were to the same target, located at position 95, but they started from two different initial positions that had been randomly selected. Panel B shows the corresponding commanded velocities. The model used a higher velocity and a longer duration in an appropriate combination to make the larger movement. Panel C plots the number of Purkinje cells selected (ie, the number switched to their off-state) for the two movements. The preselection process that occurred before the trigger arrived (at time zero) automatically selected more Purkinje cells when the initial position was farther from the target position, and this made the commanded velocity higher. The commanded duration also automatically increased. This was because it took longer for the position input to reach the on-threshold of the Purkinje cells.

movement is required. The overshoot is due to slight inaccuracy in training this target input. In some simulation runs, training inaccuracy resulted in small undershoots instead.

6.3 Learning target positions

The training signal in the postulated learning rule is carried by climbing fibers. As discussed earlier, these fibers appear to detect unexpected somatic events. Here we speculate on how the observed properties of climbing fibers might detect situations in which movements are dysmetric. Then we demonstrate that a signal sensitive to hypometric movements (ones that are too short) is effective for training the simulation model described in the previous section.

One type of climbing fiber is predominantly responsive to cutaneous input whereas another type is predominantly responsive to proprioceptive input (section 3.1). The responsiveness of both types is modulated by efference copy signals, but in different ways. The cutaneous fibers are inhibited toward the end of a movement, whereas the proprioceptive ones are inhibited during movement (Gellman, Gibson & Houk, 1985). As a consequence, the cutaneous fibers fire when a limb bumps into an object in the course of a movement. We speculate that these climbing fibers would also fire when a movement command is hypermetric, since in these circumstances the limb would tend to bump into the target before the movement's intended completion. In contrast, proprioceptive climbing fibers are depressed during movement, and responsiveness resumes as a movement is completed. Fibers (presumably of this type) were found to detect instances in which the object encountered at the end of a movement yielded (Andersson & Armstrong, 1987). We speculate that a failure to find a solid object at the end of the commanded movement is what triggers this type of climbing fiber response, and that the significance of the response is that it detects when movement commands are hypometric.

By this reasoning, we postulate that some climbing fibers would tend to detect hypermetric movement commands whereas others would tend to detect hypometric commands. In the present example, we used hypometric climbing fibers to train the pattern generator. When the movement fell short of the target, the training signal c_t in eqn(1) was set equal to 1.0 for one time interval after the motor program ended. All of the synapses made by the active target line received decrements, since these synapses were the eligible ones and since all Purkinje cells were active by this time. At all other times c_t was kept at 0. The incrementing part of the learning rule operated throughout each movement.

kept at 0. The incrementing part of the learning rule operated throughout each movement. Eligible synapses that innervated inactive Purkinje cells (preselected ones) received increments at each time interval. Good results were obtained with α=0.0005 and β=0.01.

A typical learning curve for the model is shown in Figure 8A. Absolute error drops quickly but then remains at a low level indefinitely. It takes approximately 600 trials for the asymptote to be reached. The indefinite presence of an absolute error is due to fluctuations of the signed error about zero. These fluctuations are caused by variations in synaptic weight, apparent in the weight-convergence graph shown in Figure 8B. This graph shows how the weights of three synapses formed by one of the target lines progressively converge toward the correct weight of 0.48 (value calculated analytically) and then fluctuate about this value.

The three curves in Figure 8B are representative of the 36 adjustable synapses in the model. Each of these synapses was initially assigned a random weight between 0 and 1.0. For cases in which the initial weight was high, the excitatory input was large, and there was a high probability that the innervated Purkinje cell would either not be selected, or, if selected, would return to its on-state early in the movement. Both of these factors reduced the likelihood of receiving weight increments without affecting the likelihood of receiving weight decrements. Thus, the weight moved downward toward the correct value. In contrast, when the initial weight was low, the excitatory input from the target was small, and it was more likely that the Purkinje cell would be selected for participation in the movement and more likely that the synapse would receive an increment. In addition, the lower the weight, the longer the cell was likely to remain inactive, which further increased the likelihood of receiving increments. Thus, the weight moved upward toward the correct value.

6.4 Interpretation of the learning procedure

How does this learning procedure fit into the categories usually associated with network learning methods? Is it supervised learning, where the climbing fiber specifies a desired response or a signed error signal; is it reinforcement learning, where the climbing fiber plays the role of a penalty signal (Barto, Sutton & Anderson, 1983; Barto, this volume); or is it some other form of learning? The method we have described differs sufficiently from the familiar learning methods that it is not immediately clear how to characterize it. There are two ways one can view this learning process: a local view focusing on individual Purkinje cells, and a global view emphasizing its logic at the level of the pattern generator module.

Viewed at the level of an individual Purkinje cell, the learning process is not easily characterized. According to eqn(1), if a cell fires, climbing fiber activity instructs it to decrease its tendency to fire under similar input conditions. In this case, the climbing fiber signal might be thought of as a teaching signal specifying a desired response of "off". On the other hand, when a cell is not firing, its weights always change so as to increase its tendency to fire under similar input conditions (unless climbing fiber activity occurs, in which case nothing happens: a case never exercised in the simulations reported here). Consequently, for a cell that is not firing, the absence of climbing fiber activity might be thought of as a teaching signal specifying a desired response of "on". However, given that the instructive significance of climbing fiber activity depends on the output of the cell, the learning rule does not conform to the usual logic of supervised-learning. On the other hand,

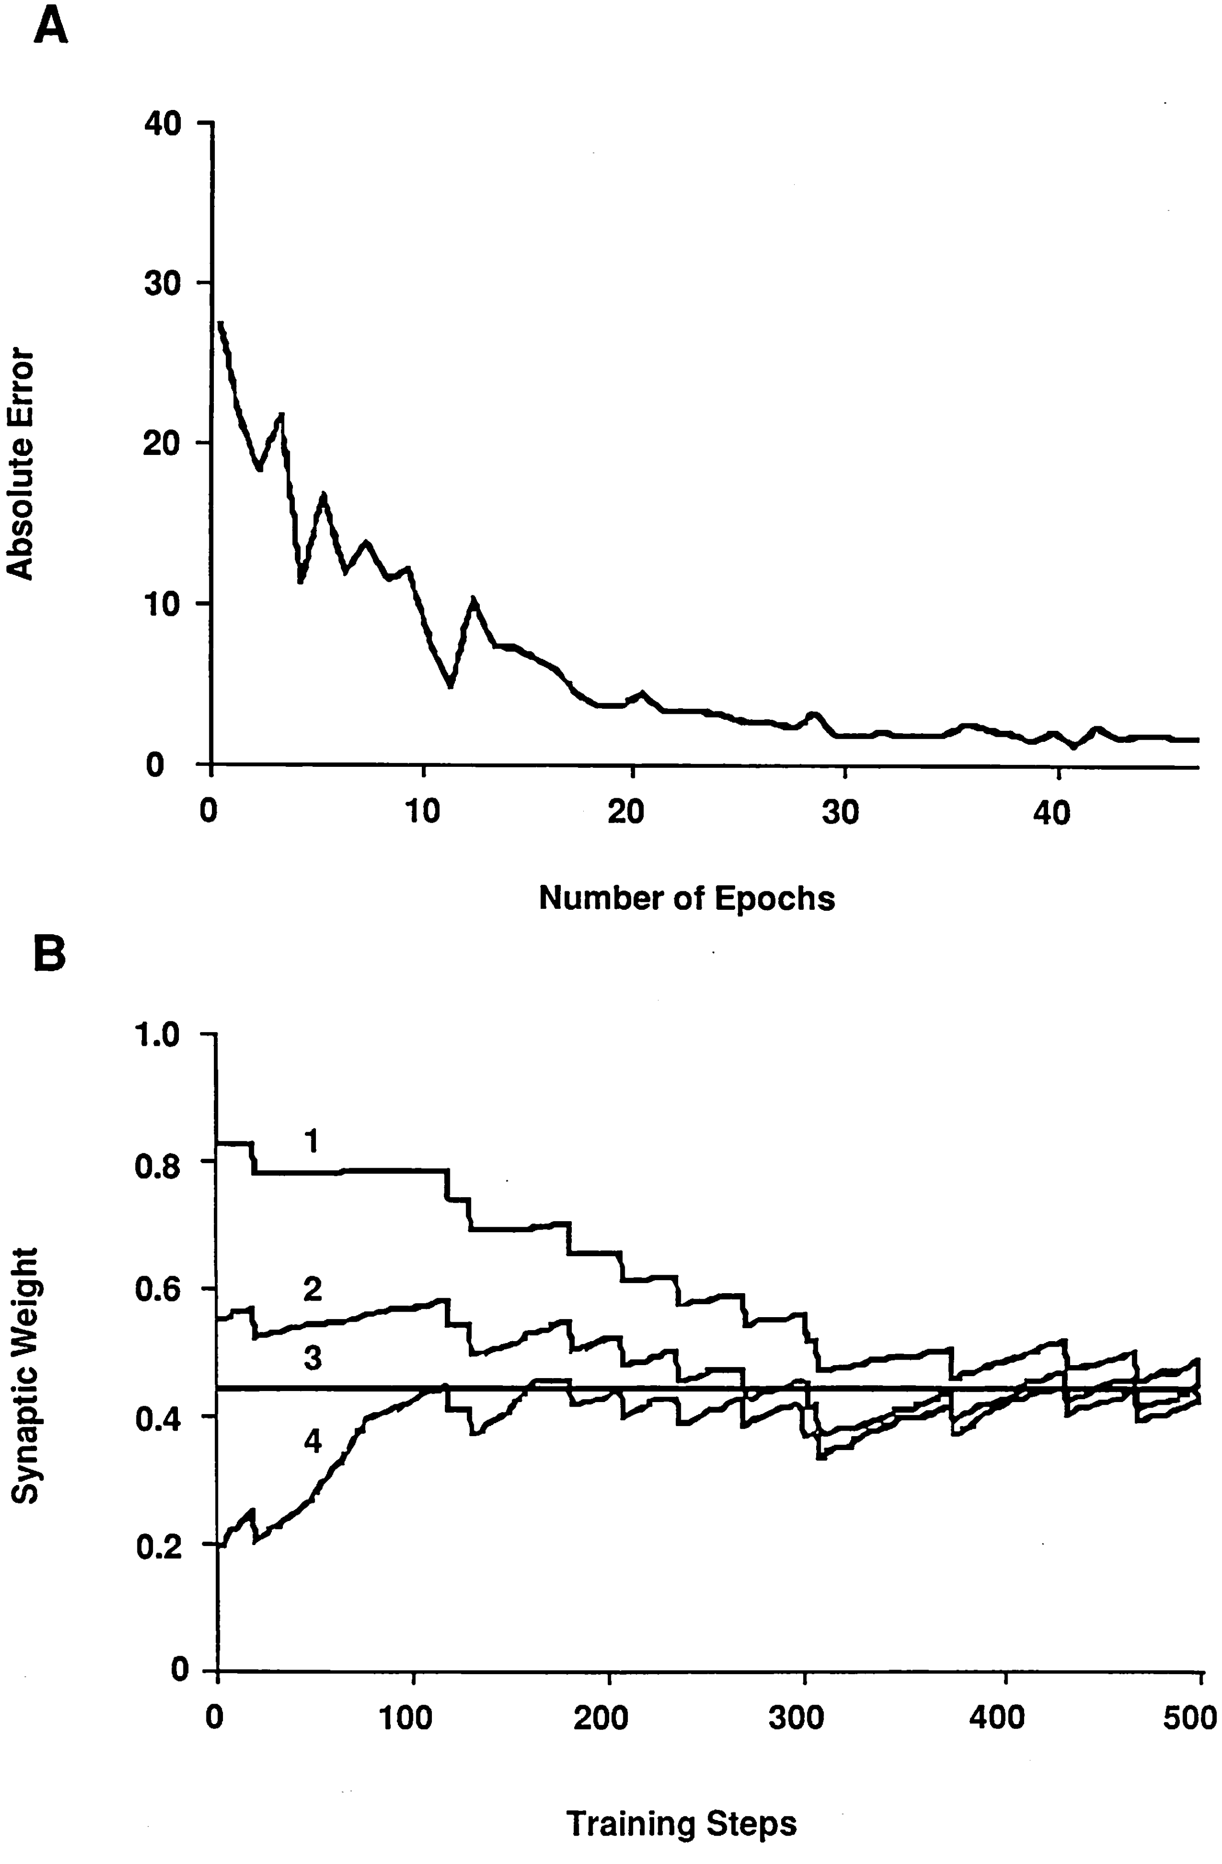
A
Absolute Error
40
30
20
10
0
0
10
20
30
40
Number of Epochs
B
Synaptic Weight
1.0
0.8
0.6
0.4
0.2
0
1
2
3
4
0
100
200
300
400
500
Training Steps

Figure 8: Learning performance of the simulation model. Panel A shows a learning curve, expressed as a plot of the average absolute error in successive 20 trial epoches. Performance reached an asymptote after about 30 epoches. Panel B shows three examples of the convergence of synaptic weight toward a precisely correct value that was calculated analytically (the horizontal line labeled 3). A synapse that started with a large weight (curve 1) initially experienced decrements, whenever the climbing fiber signalled that the movement was too short, in combination with very few increments, so the synaptic weight progressively became smaller. In contrast, a synapse that started with a small weight (curve 4) experienced many increments in addition to decrements, so the synaptic weight progressively became larger. After having experienced several hundred training steps, all of the synapses had weights that fluctuated about the correct value.

it also does not conform to the usual logic of reinforcement-learning in which climbing fiber activity would be interpreted as a "punishment signal" (Widrow, Gupta & Maitra, 1973; Barto et al., 1983; Barto, 1985, 1989, and this volume). Climbing fiber activity would play the role of a punishment signal if it told the cell to do the opposite of what it was currently doing (Widrow et al., 1973, called this "negative bootstrapping"). In the case of eqn(1), however, climbing fiber activity does not tell an "off" cell to be "on". It may be most appropriate to regard a climbing fiber in the present model as playing the role of a kind of "one-sided" punishment signal. According to the model, a Purkinje cell autonomously increases its tendency to fire; climbing fiber activity counters this tendency by 1) suspending weight increases of non-firing cells, and 2) instructing firing cells to stay off if they happen to fire in inappropriate circumstances.

The significance of this weight-adjustment process becomes more clear at the global level of a pattern generator module. A pattern generator module essentially conducts a search for the endpoint of a motor program that is calibrated to produce a movement to the specified target position. It does this by continuing to try, on each trial with a given position command, shorter programs (the increment part of the learning rule given by eqn(1)) until the attempted movement is too short (indicated by climbing fiber activity), at which point it slightly lengthens the endpoint of the program (the decrement part of eqn(1)). Thereafter, for the simple example we have simulated, the endpoint of the program oscillates, on a trial-by-trial basis, around the correct value with an amplitude depending on the constants α and β and the magnitudes of the eligibility terms ε_{ij} in eqn(1). By dynamically adjusting the magnitudes of α and β, the magnitude of this oscillation can be reduced, although at the cost of adaptability to future changes in task requirements. Alternatively, averaging the output of many pattern generator modules will reduce the oscillation.

After learning, a target command arriving via parallel fibers sets the pattern generator to produce, upon triggering, a program with an endpoint that moves the limb very near to the position of the target. In other words, after learning, the search aspect of the process is no longer required. Results of previous searches have been retained so that a program of appropriate endpoint is immediately selected. This is an example of the process called "associative search" by Barto, Sutton, and Brouwer (1981) and illustrated by Barto, Sutton, and Anderson (1981, 1983). It is an instance of reinforcement learning because desired program endpoints are learned even though the training information provided by the climbing fiber is not based on direct knowledge of the desired endpoint.

Instead of directly specifying this endpoint, as would be necessary in a supervised-learning module, the climbing fiber effectively terminates a search for the desired endpoint. The computational significance of this process is that the system can learn to produce correctly calibrated motor programs without requiring a teacher which already knows the correct endpoints for a set of training examples. Although there are other ways to learn under these conditions, the procedure outlined above is a natural consequence of the architecture of the pattern generator modules we have described. Future research will address the question of how efficient this learning process can be when movements are the result of the activity of many pattern generator modules executing their motor programs simultaneously.

7. DISCUSSION

We have reviewed recent findings regarding the anatomy and physiology of the cerebellum and have attempted to apply this knowledge to the design of an adaptive neural network for controlling limb motion. The proposed network is comprised of an array of modules that are modeled after the neural circuitry found in the cerebellum. Each module functions as an adjustable pattern generator that uses semi-realistic neurons to process information in a manner consistent with single unit studies in animals. Individual pattern generators learn to combine target information with limb position signals to produce velocity commands. The modules are thus capable of storing, retrieving and executing simple motor programs. The simple programs discussed here should be thought of as elements of more complete programs that would be required to control robot or actual limbs. We have not yet dealt with the problem of coordinating a set of pattern generators in order to produce such composite motor programs.

The sensorimotor network presented here is related to several previous models that were also inspired by the anatomy and physiology of the cerebellum. The most obvious similarity concerns our treatment of climbing fiber input as a training signal that adjusts the synaptic strengths of parallel fiber inputs to Purkinje cells, as originally proposed by Marr (1969) and Albus (1971). In subsequent publications this basic idea has been applied to the control of robot manipulators (Albus, 1975; Miller, 1987), the adaptive adjustment of filters (Fujita, 1982), the identification of dynamics and inverse dynamics (Kawato, Furukawa & Suzuki, 1987), adaptive gain control (Grossberg & Kuperstein, 1987) and novel associative memories (Kanerva, 1988). The proposed network also shares some features with other previous models of the cerebellum (Kornhüber, 1971; Boylls, 1975; Pellionisz & Llinás, 1979; Braitenberg, 1987; Moore, Desmond & Berthier, in press).

While these prior models have yielded interesting and useful ideas, they incorporate little of the new biological knowledge reviewed in this chapter. For example, the nature of information processing in the models is relatively unrelated to that deduced to occur in the cerebellum on the basis of single unit recordings in behaving animals (section 3). The networks were built from static elements that compute weighted sums of inputs rather than exploring the unusual dynamic properties resulting from special ion channels in neural membranes (section 4). Also, the training rules used in these former studies bear little relation to the biology of synaptic plasticity (section 6). While the model proposed here remains unrealistic in many of its details, we have attempted to devise simplifying abstractions that incorporate most of the salient features of this new information. Undoubtedly changes will have to be made, so it is probably more appropriate to view the proposed network as an inspiration based on biology rather than as a realistic model of the cerebellum. Given this caveat, let us review some of the potential merits of this tentative model.

7.1 Novel functional features related to persistent internal states

The design for an adjustable pattern generator developed here owes many of its unique features to a capacity for undergoing persistent transitions in internal state. Pattern generator modules have several stable internal states because they are composed of elements with persistent on- and off-states. Inputs to a pattern generator can produce transitions in internal state that then alter the basic manner in which the pattern generator subsequently operates. While this feature could lead to chaotic behavior, instead, state transitions are used in a controlled manner to regulate the preselection and execution of motor programs. It is worth reviewing the two types of element that have persistent internal states, pointing out how they give rise to special functional features of pattern generator modules as a whole.

One form of persistent state is a direct consequence of a recurrent pathway from red nucleus output neurons back to cerebellar nuclear cells (Figure 1). In our model, this positive feedback loop has two stable states -- an off-state of quiescence, in which output is zero independent of the level of Purkinje cell activity, and a on-state in which discharge rate is a variable regulated by the amount of Purkinje cell inhibition (Figure 4B).

The existence of a stable off-state of the loop is very important for the process of preselection. While the loop is in this state, Purkinje cells can be freely turned on and off so as to preset the parameters of the pattern generator, without having any immediate effect on pattern generator output. The preselection process can be thought of as the retrieval of a motor program from memory, in preparation for its execution.

The ability of extrinsic inputs to the loop to produce transitions to an on-state of loop activity is the basis for the triggering feature of the model. This transition is analogous to switching on the power to a function generator that has already been adjusted to produce a desired patterned output. Triggering is important since it allows a control process to separate the problem of deciding when to take an action from the problem of determining what kind of action to take. The kind of action (the particular motor program) is specified by the cerebellar module whereas the timing of the action can be decided by a separate process. There may be many circumstances when it is desirable to use the same motor program. Different subsystems can be designed to evaluate these circumstances and transmit their decision by sending a simple trigger signal to any of the neurons in the loop. This will cause a transition from off- to on-state, thus initiating the execution of a motor program. Motor cortical mechanisms that might contribute to this triggering function are discussed in another article (Houk, 1989).

Another form of persistent state derives from the dynamic properties of special ion channels present in the membranes of Purkinje cells. The bistable input-output curve assumed for these neurons (Figure 4A) is an abstraction inspired by non-inactivating sodium and calcium channels (section 4.1). However, Purkinje cells definitely show more than two stable activity states, so this model is clearly provisional. The actual properties of Purkinje cells are likely to be multistable, as a consequence of bistability in several dendritic processes, but this hypothesis needs to be tested experimentally. The present use of a bistable model is meant to be a conservative choice, since this simplification is expected only to diminish the computational capabilities of adjustable pattern generators.

Bistability of Purkinje cells is responsible for the unique ability of pattern generator modules to operate in a semi-autonomous manner during the execution of a motor program, and yet respond to a rich interplay of inputs during the preselection period. The absence of responsiveness to inputs during most of the execution period is due to the flatness of the input-output function, except at two values of input, the on- and off-threshold points for state transitions. This nonresponsiveness is important since it allows the time course of pattern generator output to be specified independently of the time course of its inputs. In fact, one can have a time varying output under conditions in which all of the inputs remain constant. The independent specification of output time course permits a great deal of flexibility in controller design, such as the possibility that the same program can be used in many different circumstances, even though the different circumstances supply appreciably different input patterns to Purkinje cells.

While it is useful for Purkinje cells to be minimally responsive to inputs during the execution phase, they must respond in a sensitive fashion to inputs present during the preselection period. This is how an appropriate pattern generator is chosen and how the parameters of the pattern generator are adjusted to accommodate different target distances, initial positions of the limb and degrees of urgency. In our model, basket cell inhibition is used to drive Purkinje cells into the vicinity of their off-thresholds where, due to bistability, they become extremely sensitive to small differences in input. Purkinje cells are thus made to respond to an interplay between basket inhibition and the constellation of excitatory input from parallel fibers, with some cells being driven into their off-state and others remaining in their on-state.

The on-threshold region of bistability is also used to advantage in this model. It serves as a sensitive switching point for stopping a motor program. In our model, position signals are continuously fed back to Purkinje cells, but Purkinje cells do not respond to these inputs until their on-threshold is reached. At this point the cells quickly turn on whereupon they again cease responding to position feedback. Because the position loop is operative only at the switching point, the usual fear that feedback will cause instability is circumvented. This gives rise to the quasi-feedforward mode of operation discussed in the next section.

Bistability was also seen as a potential mechanism for prolonging the persistence of target position information, a simple form of short-term memory. We postulated that local bistability, due to calcium channels in the distal dendrites and spines, is capable of converting phasic representations of target position into tonic responses that then influence the final position commanded by a motor program. It is important to note that pattern generator modules do not require this mechanism, since tonic inputs representing target positions may be provided by other sources of mossy fiber input.

In summary, the use of elements capable of transitions in internal state gives rise to many advantageous features in the overall properties of pattern generator modules, and it seems likely that additional useful features await future exploration. While our present implementation deals with individual movements, other investigators have used networks capable of transitions in internal state to generate action sequences (Jordan, 1986; Kleinfeld, 1986). This is likely to be a fruitful topic to explore when we begin to deal with multiple pattern generators in a larger sensorimotor network.

7.2 The quasi-feedforward mode of operation

While feedforward has many advantages over feedback as a general control strategy, on-line compensation for errors in performance and prompt adjustments to changing conditions are unique capabilities of feedback (cf. Houk, 1988a). Conventional feedback loops are very sensitive to time delays, and their tendency to oscillate when gain is increased so as to improve regulatory performance is a big problem (Robinson, 1987). Therefore, the use of bistability to circumvent this problem, as discussed in sections 4.2 and 7.1, seems to be a very nice feature of the proposed pattern generator module. Recall that stability problems are avoided because position feedback is discontinuous -- it is operative only momentarily, at the switching point of Purkinje cell bistability. We refer to this feature as quasi-feedforward processing.

The present model is related to two previous models of motor command generation. These previous models used efference copy signals driving internal feedback loops as a strategy for avoiding the delay associated with feedback from the movement. This strategy first appeared in a model of the saccadic eye movement control system (Robinson, 1975) and was later put to detailed experimental test and refined (Gisbergen et al, 1981). The same basic strategy was applied recently in vector format to limb control problems (Bullock & Grossberg, 1988).

In Robinson's model for saccadic burst generation, an efference copy position signal is subtracted from a target position signal to compute an error, and this error signal is then processed by an amplifier with a high gain that rapidly saturates. The saturating amplifier is needed to reproduce the rapid rise and fall of discharge rate that is recorded from burster neurons in the brainstem. Bullock and Grossberg instead postulate that a similar error signal is fed into a time-varying gain function specifically tailored to produce symmetrical velocity profiles. The time-varying gain starts out at zero and grows in an uneven manner that was chosen to produce slowly rising and slowly falling velocity profiles. The purpose of doing this is not entirely clear, since recordings from brain cells that transmit motor commands generally show abrupt onsets and offsets. The slowly rising and slowly falling features of actual velocity profiles probably arise as the feedback mechanisms in the spinal cord translate the movement commands into actual motion.

Adjustable pattern generators differ from these prior models in several respects. First they use position information from proprioceptive feedback in addition to efference copy sources. The presence of proprioceptive feedback allows a pattern generator to adjust itself automatically to accommodate different initial limb positions (section 5.2), and it also appears to provide some measure of compensation for disturbances that prevent a movement from being completed (section 4.2). Proprioceptive feedback would not be feasible in the other two schemes, since they use continuous feedback, and the delay in the feedback pathway is sufficiently long to cause stability problems. This delay has some undesirable effects on the performance of pattern generator modules as well, but these effects are relatively minor compared with a threat of instability. Delay will present no problem for the pattern generators when movements are slow. However, at high velocities the distance traveled during the delay period will become appreciable. Under the latter conditions, it would be desirable for the pattern generator to depend more heavily on efference copy, and this possibility needs to be investigated.

A more radical difference between adjustable pattern generators and previous models concerns the process that combines target and limb position information to produce a motor program. As pointed out earlier, the pattern generators use persistence of internal state in module elements and plasticity in parallel fiber synapses to provide several features not available in the alternative models. These features include the storage of motor programs and a preselection mechanism for retrieving these programs from memory. For the designation of target position, pattern generator modules are thought to be capable of directly combining a variety of signals in several coordinate systems, whereas the other models typically postulate a sequential set of coordinate transformations that recode target information into the coordinates of the movement that is controlled by the command generator. Further advantages may result from the fact that individual Purkinje cells receive such a large input vector. This suggests that it may be possible to incorporate situation effects and other contingencies without a necessity for invoking additional stages of processing. In future research, we hope to show that the learning mechanisms described here can be extended so that this variety of information expressed in different coordinate systems can be combined and motor programs can be calibrated to refine the level of skill exhibited by the final motor output.

Clearly more work is needed to assess the genuine advantages of building sensorimotor networks from arrays of adjustable pattern generators. Adaptive neural networks of this type might also be studied for their potential utility in the control of other types of process.

References

Albus, J. (1981) Brains, Behavior and Robotics. Petersborough, NH, Byte Books

Albus, J.S. (1971) A theory of cerebellar function. Math Biosci **10**: 25-61

Andersson, G., Armstrong, D.M. (1987) Complex spikes in Purkinje cells in the lateral vermis (b zone) of the cat cerebellum during locomotion. J. Physiol. London **385**: 107-134

Andersson, G., Campbell, N.C., Ekerot, C.-F., Hesslow, G., Oscarsson, O. (1984) Integration of mossy fiber and climbing fiber inputs to Purkinje cells. Exp. Brain Res **Suppl. 9**: 145-150

Andersson, G., Ekerot, C.-R., Oscarsson, O., Schouenborg, J. (1987) Convergence of afferent paths to olivo-cerebellar complexes. In: Glickstein, M., Yeo, C., Stein, J. (eds) Cerebellum and Neuronal Plasticity. London, Plenum Press, pp 165-173

Barto, A.G., Sutton, R.S., Anderson, C.W. (1983) Neuronlike elements that can solve difficult learning control problems. IEEE Transactions on Systems, Man, and Cybernetics **13**: 835-846

Barto, A.G., Sutton, R.S., Brouwer, P.S. (1981) Associative search network: reinforcement learning associative memory. IEEE Transactions on Systems, Man, and Cybernetics **40**: 201-211

Barto, A.G., Sutton, R.S. (1981) Landmark learning: illustration of associative search. Bio. Cybernetics **42**: 1-8

Barto, A.G. (1985) Learning by statistical cooperation of self-interested neuron-like computing elements. Human Neurobiology **4**: 229-256

Barto, A.G. (1989) From chemotaxis to cooperativity: exercises in neuronal learning strategies. In: R. Durbin, R. Maill, G. Mitchison (eds) The Computing Neurone. Addison-Wesley, pp

Boylls, C.C. (1975) A Theory of Cerebellar Function with Applications to Locomotion, Vol 76-1. Amherst, MA, COINS Tech. Rept.

Braitenberg, V. (1987) The cerebellum and the physics of movement: some speculations. In: Glickstein, M., Yeo, C., Stein, J. (eds) Cerebellum and Neuronal Plasticity. London, Plenum Press, pp 193-207

Brodal, A. (1981) Neurological Anatomy in Relation to Clinical Medicine. New York, Oxford University Press

Bullock, D., Grossberg, S. (1988) Neural dynamics of planned arm movements: emergent invariants and speed-accuracy properties during trajectory formation. Psych. Rev **95**: 49-90

Crepel, F.C., Krupa, M. (1988) Activation of protein kinase C induces a long-term depression of glutamate sensitivity of cerebellar Purkinje cells. An in vitro study. Brain Res **458**: 397-401

Eccles, J.C., Ito, M., Szentagothai, J. (1967) The Cerebellum as a Neuronal Machine. Berlin, Springer-Verlag

Ekerot, C.-F. (1984) Climbing fibre actions of Purkinje cells -- plateau potentials and long-lasting depression of parallel fibre responses. In: Bloedel, J., et al. (eds) Cerebellar Functions. Berlin, Springer-Verlag, pp 268-274

Fujita, M. (1982) Adaptive filter model of the cerebellum. Biol. Cybern **45**: 195-206

Gellman, R., Gibson, A.R., Houk, J.C. (1985) Inferior olivary neurons in the awake cat: Detection of contact and passive body displacement. J. Neurophysiol **54**: 40-60

Gibson, A.R., Houk, J.C., Kohlerman, N.J. (1985) Relation between red nucleus discharge and movement parameters in trained macaque monkeys. J. Physiol. (Lond) **358**: 551-570

Gielen, C.C.A.M., Houk, J.C. (1987) A model of the motor servo: Incorporating nonlinear spindle receptor and muscle mechanical properties. Biol. Cybern **57**: 217-231

Gisbergen, J.A.M. van, Robinson, D.A., Gielen, S. (1981) A quantitative analysis of generation of saccadic eye movements by burst neurons. J. Neurophysiol **45**: 417-442

Grossberg, S., Kuperstein, M. (1986) Neural Dynamics of Adaptive Sensory-motor Control: Ballistic Eye Movements. Amsterdam, Elsevier, pp 72-79

Harvey, R.J., Porter, R., Rawson, J.A. (1979) Discharges of intracerebellar nuclear cells in monkeys. J. Physiol **297**: 559-580

Holmes, G. (1939) The cerebellum of man. (The Hughlings Jackson Memorial Lecture). Brain **62**: 1-30

Houk, J.C., Gibson, A.R. (1987) Sensorimotor processing through the cerebellum. In: King, J.S. (ed) New Concepts in Cerebellar Neurobiology. New York, NY, Alan R. Liss, Inc., pp 387-416

Houk, J.C. (1987) Model of the cerebellum as an array of adjustable pattern generators. In: Glickstein, M., Yeo, C., Stein, J. (eds) Cerebellum and Neuronal Plasticity. New York, NY, Plenum Press, pp 249-260

Houk, J.C. (1988a) Control strategies in physiological systems. FASEB J **2**: 97-107

Houk, J.C. (1988b) Schema for motor control utilizing a network model of the cerebellum. In: Anderson, D.Z. (ed) Neural Information Processing Systems. New York, NY, Amer. Inst. Physics., pp 367-376

Houk, J.C. (1989) Cooperative control of limb movement by the motor cortex, brainstem and cerebellum. In: Cotterill, R.M.J. (ed) Models of Brain Function. Cambridge, Cambridge Univ. Press, in press

Ito, M. (1984) The Cerebellum and Neural Control. New York, Raven Press

Ito, M. (1989) Long-term depression. Ann. Rev. Neurosci **12**: 85-102

Jordan, M.I. (1986) Serial order: a parallel, distributed processing approach. Tech. Rept. 8604, UCSD Inst. Cognitive Science

Kanerva, P. (1988) Sparse Distributed Memory. Cambridge, MA, MIT Press

Kawato, M., Furukawa, K., Suzuki, R. (1987) A hierarchical neural-network model for control and learning of voluntary movement. Biol. Cybern **56**: 1-17

Kleinfeld, D. (1986) Sequential state generation by model neural networks. Proc. Natl. Acad. Sci **83**: 9469-9473

Kornhüber, H.H. (1971) Motor functions of cerebellum and basal ganglia: the cerebellocortical saccadic (ballistic) clock, the cerebellonuclear hold regulator, and the basal ganglia ramp (voluntary speed smooth movement) generator. Kybernetik **8**: 157-162

Kuperstein, M. (1988) An adaptive neural model for mapping invariant target position. Beh. Neurosci **102**: 148-162

Kuypers, J.G.J.M. (1981) Anatomy of the descending pathways. In: Brooks, V.B. (ed) Motor Control, Vol II of Sect 1 of Handbook of Physiology, Vol Part 1. Bethesda,MD, American Physiological Society, pp 597-666

Llinás, R., Sugimori, M. (1980) Electrophysiological properties of in vitro Purkinje cell dendrites in mammalian cerebellar slices. J. Physiol. London **305**: 197-213

Marr, D. (1969) A theory of cerebellar cortex. J. Physiol **202**: 437-470

Miller, W.T. (1987) Sensor-based control of robotic manipulators using a general learning algorithm. IEEE J. Robotics & Automation **RA-3**

Moore, J.W., Desmond, J.E., Berthier, N.E. (in press) Adaptively timed conditioned responses and the cerebellum: a neural network approach. Bio. Cybernetics

Nishizuka, Y. (1986) Studies and perspectives of protein kinase C. Sci **233**: 305-312

Pellionisz, A., Llinás, R. (1979) Brain modeling by tensor network theory and computer simulation. The cerebellum: distributed processor for predictive coordination. Neuroscience **4**: 323-348

Robinson, D.A., Optican, L.M. (1981) Adaptive plasticity in the oculomotor system. In: Flohr, H., Precht, W. (eds) Lesion Induced Neuronal Plasticity in Sensorimotor Systems. Berlin, Springer-Verlag, pp 295-304

Robinson, D.A. (1975) Oculomotor control signals. In: Lennerstrand, G., Bach-y-rita, P. (eds) Basic Mechanisms of Ocular Motility and their Clinical Implications. Oxford, Pergamon Press, pp 337-374

Robinson, D.A. (1987) Why visuomotor systems don't like negative feedback and how they avoid it. In: Arbib, M.A., Hanson, A.R. (eds) Vision, Brain and Cooperative Computation. Cambridge, MIT Press, pp 89-107

Sakurai, M. (1987) Synaptic modification of parallel fibre-Purkinje cell transmission in in vitro guinea-pig cerebellar slices. J. Physiol. London **394**: 463-480

Stein, J.F. (1986) Role of the cerebellum in the visual guidance of movement. Nature **323**: 217-221

Thompson, R.F. (1986) The neurobiology of learning and memory. Science **233**: 941-947

Van Kan, P.L.E., Houk, J.C., Gibson, A.R. (1987) Response characteristics of mossy fiber afferents during forelimb movements. 2nd World Congr. Neurosci., Budapest

Voogd, J., Bigaré, F. (1980) Topographical distribution of olivary and corticonuclear fibers in the cerebellum. A review. In: Courville, J., Montigny, C. de, Lamarre, Y. (eds) The Inferior Olivary Nucleus. New York, Raven Press, pp 207-234

Widrow, B., Gupta, N.K., Maitra, S. (1973) Punish/reward: learning with a critic in adaptive threshold systems. IEEE Transactions on Systems, Man, and Cybernetics **5**: 455-465

Yeo, C.H. (1987) Cerebellum and classical conditioning. In: Glickstein, M., Yeo, C., Stein, J. (eds) Cerebellum and Neuronal Plasticity. London, Plenum Press, pp 321-338