CHAPTER 9
SACCADIC RHYTHM AND PREDICTIVE MOVEMENT SEQUENCES

9.1. Rhythmic Choices among Multiple Movement Sources

Chapter 8 provided a simple example of a functional rhythm in a sensory-motor system. During such a rhythm, movement and postural subsystems turn on and off intermittently through time, thereby supplementing and complementing each other's computational competences. The demands upon the postural subsystem are always the same: Hold the eyes in place no matter how they got there. The choice of movement subsystems is more varied. A visually reactive saccade may occur. Or a saccade that is triggered by an intermodality signal, such as a sound, may occur. Or a saccade may be volitionally produced, as in the rear view mirror problem (Section 1.13). Or a preplanned sequence of predictive movements may occur, as in a dance. As one ascends the evolutionary ladder from visually reactive to predictive movements, sensory cues play an increasingly limited role.

The remarkable capability of humans to override the influence of sensory cues has been experimentally analysed (Murphy, Haddad, and Steinman, 1974; Steinman, 1965). Such studies led Steinman (1976, p.121) to write: "Perhaps the most striking aspect of human oculomotor performance is its independence from stimulus variables. By this I mean that a normal human adult can look about in his visual world and attend whatever region catches his fancy undisturbed by this distinction of light on his retina, or, in perceptual terms, the way the visual world looks at a particular moment."

Special purpose networks are required to achieve a nonobligatory reaction to sensory cues. For example, patients with discrete frontal lobe removals for intractable epilepsy produce reflex-like saccades to visual stimuli (Buchtel and Guitton, 1980; Goldberg and Bushnell, 1981). In normal humans making intentional saccades, a slow negative shift of scalp potentials begins approximately 650 msec. before an eye movement and is largest in amplitude over the frontal region (Kurtzberg and Vaughan, 1982). This potential does not occur before a visually reactive saccade. The ability of normal humans and higher mammals to override momentary visual stimuli shows that the relationship between saccadic decisions and learned calibrations is not a simple one. As in our discussion of the rear view mirror problem, we will suggest how an intentional saccadic command can take advantage of learned calibrations due to visually reactive saccades even though the intentional command suppresses visually reactive commands. We will suggest that an intentional command does not both inhibit and sample the same visually reactive locus. Intentional commands need a separate subsystem to calibrate some parameters of their saccades, yet also rely on the visually reactive subsystem to calibrate other saccade-related parameters. We will relate this formal property with data

demonstrating the existence of separate saccade-generating subsystems in the superior colliculus and the frontal eye fields, such that each subsystem is capable of its own vector compensations (Schiller and Sandell, 1983; Schiller, True, and Conway, 1979).

The dichotomy between intentional saccades and visually reactive saccades does not imply that neurons within the intentional subsystem are insensitive to light signals. Quite the contrary must be true in order for visual signals to initiate the calibration of movement commands within this subsystem. We compare this property with the existence of visually responsive cells in the frontal eye fields, whose activities are enhanced only when the animal makes an eye movement to a stimulus in its receptive field (Goldberg and Bushnell, 1981; Wurtz and Mohler, 1976). Despite the ability of light to calibrate the intentional subsystem, this subsystem must be able to inhibit the visually reactive subsystem, say via inhibitory signals from the prefrontal eye fields to the SC (Goldberg and Bushnell, 1981; Holmes, 1938). The intentional subsystem must also be susceptible to activation by signals other than those elicited by visual cues, notably by internally generated commands capable of encoding sequential cognitive plans. The fact that lights may help to calibrate the intentional subsystem in no way implies that only lights can activate this subsystem.

In this chapter, we consider a network that is capable of encoding a predictive sequence of movements in long-term memory, and of reading-out this movement sequence at will. This model will put us in a better position to integrate the networks which process visually reactive, intermodality, intentional, and predictive movements into a single global processing scheme in Chapter 11.

9.2. Distinguishing Correct Predictive Saccades from Incorrect Individual Saccades

The basic nature of a predictive competence may be better appreciated through the following example. Consider the two light-and-movement sequences:

I. Incorrect Light-Activated Saccade

light 1—incorrect movement 1—light 2—movement 2

II. Correct Light-Activated Predictive Saccade

light 1—light 2—correct movement 1—correct movement 2

In both case (I) and case (II), two successive, nonfoveated lights strike the retina. In case (I), both lights are due to a single source that remains on continuously, whereas in case (II) each light is due to a different source that is flashed on only briefly. In case (I), we want the second nonfoveated light to be the source of an error signal which corrects the saccade caused by the first light. We also want this second light to be able to trigger a second saccade. These properties can be realized by a retinotopic command

network, or RCN (Chapter 3). In case (II), by contrast, we do not want the second nonfoveated light to be a source of an error signal, because the first predictive saccade is correct. However, we still want the second nonfoveated light to trigger a second predictive saccade.

How does the saccadic system know whether the second nonfoveated light in a light sequence should or should not trigger an error signal? Cases (I) and (II) differ in terms of how lights and movements are interspersed during visually reactive and predictive saccades, respectively. An analysis of these differences leads to the conclusion that each light activates two parallel pathways that are capable of triggering saccadic motions. The first pathway is the more direct one that leads through the network analog of the superior colliculus. The second pathway is more indirect, and leads through the network analog of the frontal eye fields. The second pathway subserves the predictive capability of the saccadic model.

This bifurcation of pathways is not the bifurcation between sampling pathways and error signal pathways that was introduced in Chapter 2. Both of the pathways in the present bifurcation are sampling pathways. We therefore need to explain how both of these sampling pathways work together to achieve the desired properties of both cases (I) and (II).

9.3. The Temporal Control of Predictive Saccades

In Section 1.16, we suggested that predictive saccades are encoded as target positions in a target position map (TPM) computed in head coordinates. Several target positions can simultaneously be stored in short term memory (STM) by such an TPM. The positions so stored are assumed to have already been influenced by attentional factors. We will review data which suggest that such TPMs may lie within, or closely interact with, the posterior parietal cortex or the frontal eye fields. These data can be interpreted using our explanation of how a spatial pattern of stored lights can be read-out as a temporal sequence of predictive saccades. Thus we are considering how a *parallel* internal representation is transformed into *serial* behavior. We will show how such an TPM interacts with a head-muscle interface, or HMI (Chapter 4), to generate the correct temporal sequence of vector commands. In the light of these results, we will consider in Section 9.10 the difficulties that need to be overcome by a system in which stored retinotopic positions are directly recoded as difference vectors, rather than first being recoded in a TPM before being compared with present position to generate a difference vector.

The minimal network design capable of controlling the STM storage and sequential read-out of predictive saccades uses two successive TPMs, denoted by TPM_1 and TPM_2. The spatial map TPM_1 initially stores all the relevant target positions in STM. The map TPM_2 selects and stores only the target position which controls the current saccade, and interacts directly with the HMI to generate output vectors. We call this entire control structure a *Predictive Command Network* (PCN).

The main functional insight that is embodied by PCN design is that matches and mismatches between target positions and present positions

at the HMI regulate the sequential performance and learning of predictive saccades (Grossberg, 1978a, 1980). In particular, a match at the HMI can cause the next target position command to be read into the HMI, thereby causing a mismatch between the new target position and the present position. This mismatch generates the next saccade. In other words, as the network attains its present target, it resets itself to attain its next intended target. Once the PCN circuit is nonspecifically activated by a volitional signal, this match-mediated reset scheme can automatically run off its intended saccades until the entire predictive sequence has been performed.

In order to gradually develop this insight into a precise circuit design, a functional description of network operations will first be given. Then a specific network realization, along with possible cell-type identifications, will be suggested. Several closely related network realizations can be envisaged, and more than one version may be used across different species.

9.4. Storage of Temporal Order Information

A. *Storage of Temporal Order, Target Match, and Memory Reset*

Stage TPM_1 is assumed to store in STM the temporal order information of attended target positions. The temporal order information is encoded by a spatial pattern of activation across the target positions, with more intensely activated positions tending to be performed first. Grossberg (1978a, 1978c, 1986) has derived STM codes whereby temporal order information can be laid down in a way that permits its stable long term memory (LTM) encoding as a unitized motor plan. It turns out that a suitably designed shunting competitive network can do the job (Section 2.6). A discussion of these mechanisms lies outside the scope of this chapter. A temporal order code that is based upon relative STM activity can take into account both the order of item occurrences and the motivational salience of individual items. Motivationally more salient events can be looked at at earlier times, other things being equal. Grossberg (1982b, 1982c, 1984) has described mechanisms whereby motivational salience can modulate the attention that is given to a subset of sensory cues. A discussion of these mechanisms also lies outside the scope of this chapter. For present purposes, it suffices to note that earlier occurring items will often, but not always, be stored in STM with larger activities.

B. *Read-Out and STM Storage of a Target Choice*

When a nonspecific output gate opens between TPM_1 and TPM_2, the most active target position stored within TPM_1 is read into TPM_2. This target position is chosen by a competitive interaction among the active output pathways from TPM_1 to TPM_2. Several things now happen (Figure 9.1a).

C. *HMI Mismatch, Output Gate Closure, and Target Self-Inhibition*

Consider the state of the HMI prior to the moment of target choice at TPM_2. A match may exist there due to correct execution of a prior saccade, or the HMI may be inactive. If target read-out from TPM_2 then occurs, it nonspecifically activates the HMI (Chapter 4) as it instates a new

target position there in agonist-antagonist muscle coordinates. This new target mismatches the eye position corollary discharge. Thus a primary effect of target read-out is to cause a mismatch within the HMI (Figure 9.1b).

A second effect of target read-out is to prevent the read-out of other active target locations from TPM_1 until the next saccade is over. The chosen target position in TPM_2 is stored in STM until this saccade is over. During this time, it blocks further read-out there from TPM_1 and maintains the target position within the HMI.

A third effect of target read-out from TPM_1 is to prepare TPM_1 to read-out its next saccadic command after the present saccade is over. A topographically-specific inhibitory signal from the chosen target position to its source in TPM_1 accomplishes this reset event (Figure 9.1a). Although its STM source in the temporal order code of TPM_1 is inhibited, the chosen target position can remain active in STM within TPM_2 due to the internal positive feedback loops within this network.

D. *Read-Out, Reset, and STM Storage of Retinotopic Commands*

The mismatch at the HMI generates a vector difference which generates an output pattern to an RM (Chapter 4). When this happens, the eye is still at rest. This output pattern is stored in STM as a spatial location within the RM after inhibiting any command that may have previously been active within the RM.

The RM command, in turn, activates three output pathways: an unconditioned excitatory pathway to the saccade generator (SG), a prewired inhibitory pathway to the superior colliculus (SC), and a conditioned pathway to the adaptive gain (AG) stage. These signals together can generate a saccade.

While the saccadic motion is taking place, output from the HMI to the RM is prevented by offset of an output gate. The stored movement command in the RM is thus undisturbed during the saccade.

E. *LTM Printing*

After the saccade terminates, a gate opens that enables the LTM traces in the TPM_2 $\rightarrow$HMI pathway to print the final eye position in muscle coordinates (Chapter 4).

F. *Match-Induced Reset of the TPM*

After a correct saccade takes place, a match between target position and present eye position is generated within the HMI (Figure 9.2a). Whenever a match occurs at the HMI, the stored target position in the TPM_2 is inhibited. Its target position signals to the HMI are thereby also inhibited (Figure 9.2b), a mismatch is caused at the HMI (Figure 9.2c), and the inhibitory signal from the TPM_2 to the TPM_1 output pathways is eliminated. In all, a match at the HMI induces a mismatch at the HMI by inhibiting the TPM_2 $\rightarrow$HMI target position command.

After the TPM_2 target position is inhibited, the target positions that are active in STM at the TPM_1 stage can compete to choose the next target light and store it in TPM_2. Then the cycle can automatically

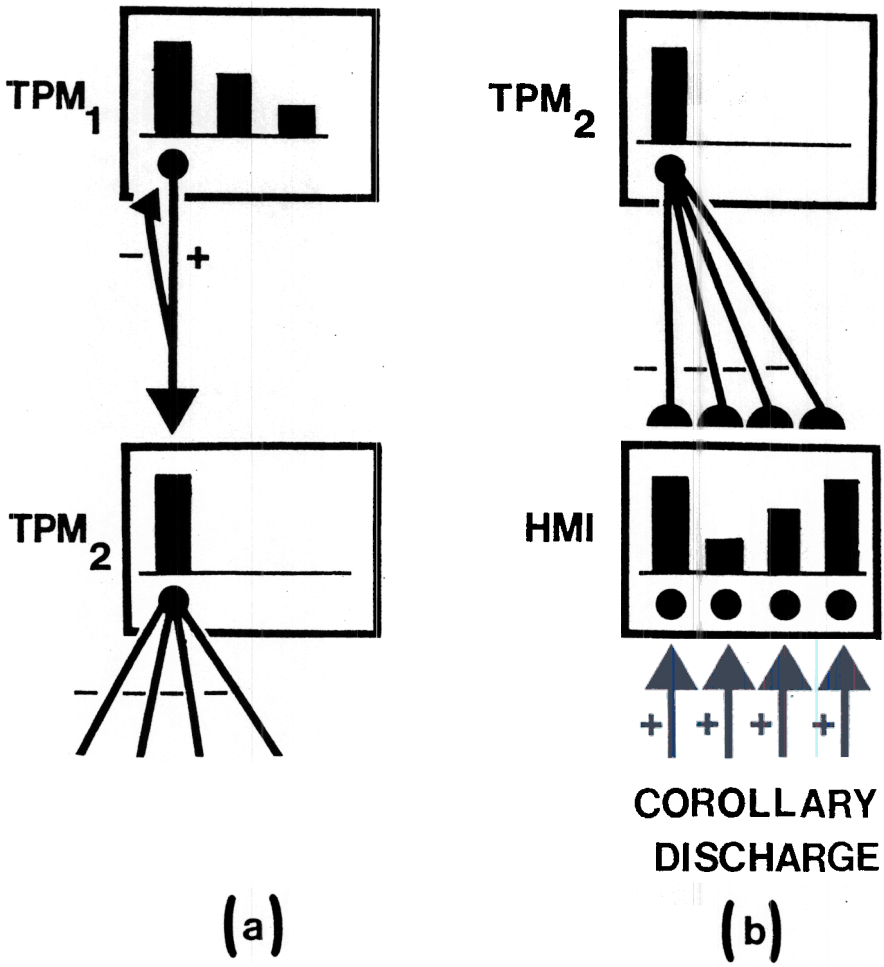

Figure 9.1. Read-out of the most active population in the first stage of the target position map (TPM_1): (a) When the most active target position is read-out of TPM_1, it self-inhibits its activity as it is stored in the second stage of the target position map (TPM_2); (b) Then TPM_2 activates its inhibitory conditioned pathways to the head-muscle interface (HMI), at which excitatory corollary discharges are also received.

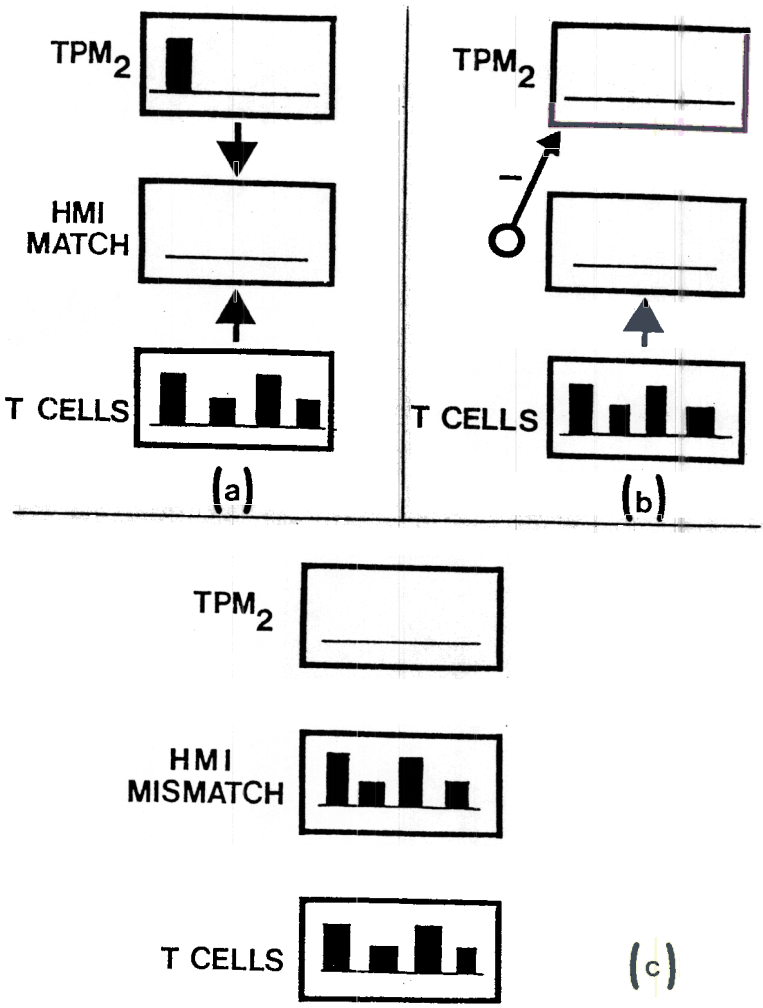

Figure 9.2. Interactions between target position map (TPM_2), head-muscle interface (HMI), and tonic cell (T) source of corollary discharges: (a) When target position matches present position in motor coordinates at the HMI, activity across the HMI is inhibited; (b) Inhibition of HMI activity causes the stored target position in the TPM_2 to be inhibited; (c) Inhibition of the TPM_2 also inhibits its inhibitory efference copy to the HMI. The corollary discharges can then register a mismatch at the HMI.

repeat itself until no temporal order information is stored any longer in STM within the TPM_1.

9.5. Design of a Predictive Command Network

Figures 9.3 and 9.4 together describe a formally competent microcircuit that is built up from neurally plausible components. Some cellular stages of these figures possess properties of known neuron types. The circuit thus suggests predictions about the functional roles of such cells and their interactions with other cells. Figure 9.4 provides details of HMI gating processes that could not be drawn into Figure 9.3.

The networks of TPM_1 and TPM_2 possess similar cellular architectures, even though they carry out different functional tasks. In particular, both TPM_1 and TPM_2 contain shunting on-center off-surround feedback networks (Chapter 2.6). For example, pathways 1 and 4 in TPM_1 and TPM_2, respectively, are both part of positive feedback loops in an on-center (Figure 9.3). Pathway 5 in TPM_2 illustrates a negative feedback loop in an off-surround. Despite these similarities, network TPM_1 is designed to simultaneously store a spatial pattern of temporal order information in STM. Network TPM_2, by contrast, is designed to store only one target light at a time in STM.

The subsequent discussion will show that the cells in the TPM_1 have properties similar to those of light-sensitive cells in the parietal cortex, whereas the cells in TPM_2 have properties similar to those of the light-insensitive saccade neurons in the parietal cortex (Lynch, 1980; Motter and Mountcastle, 1981). The visuomovement cells and movement cells in the frontal eye fields (Bruce and Goldberg, 1984) also have similar properties (Chapter 11). The formal properties in question were independently derived as part of a theory of of how motor commands are read-out from STM in a prescribed temporal order (Grossberg, 1978a, Section 52; 1978b). This theory predicts that similar types of cells will be found wherever this functional task is performed. The discovery of consistent parietal and frontal eye field data suggests that a convergence of theory and experiment is reflected in these network examples.

To start our discussion of the PCN microcircuit, consider a time at which a pattern of temporal order information is being stored in STM and the eye is at rest. In other words, the network has attended some lights and stored them in STM, but has not yet translated them into overt movements. At such a time, network TPM_1 is active, but network TPM_2 is inactive and no target position commands are being read by TPM_2 into the HMI. Consequently, all output gates from the HMI are closed and there exists a mismatch between corollary discharge and target position inputs to the HMI.

The activity within the TPM_1 cannot activate the TPM_2 because of the nonspecific inhibitory gating action which pathway 2 exerts on outputs from the TPM_1 (Figure 9.3). Such a gating action may *in vivo* take place at an interneuronal stage between TPM_1 and TPM_2, but we have avoided all such extra complexities to draw a simpler circuit. The cells which

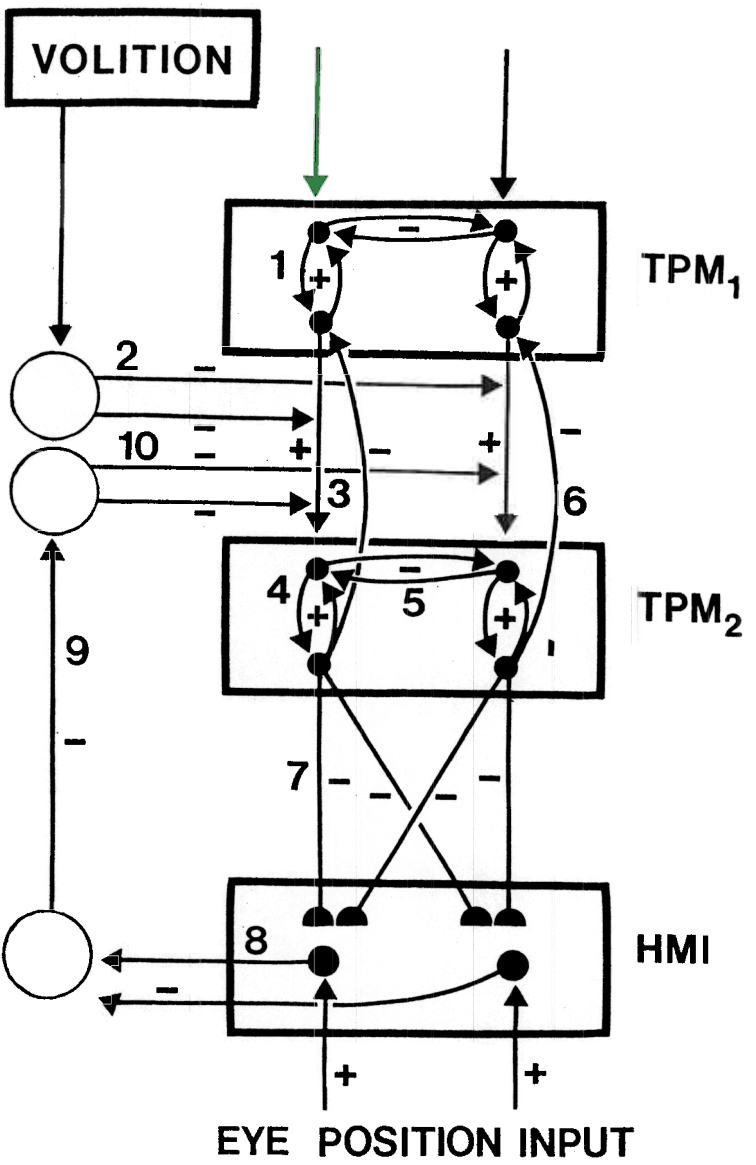

Figure 9.3. Part of the predictive command network (PCN) circuitry for storing, resetting, and sequentially reading out a series of predictive movement commands: The text describes how temporal order information that simultaneously stores several target lights in the TPM_1 can be read-out as a series of vectors from the HMI. The TPM_2 stores just one target position at a time in short term memory (STM). A match between target position and eye position at the HMI triggers read-out of the most active target position command in the TPM_1, which is then stored in the TPM_2 until its saccade is over.

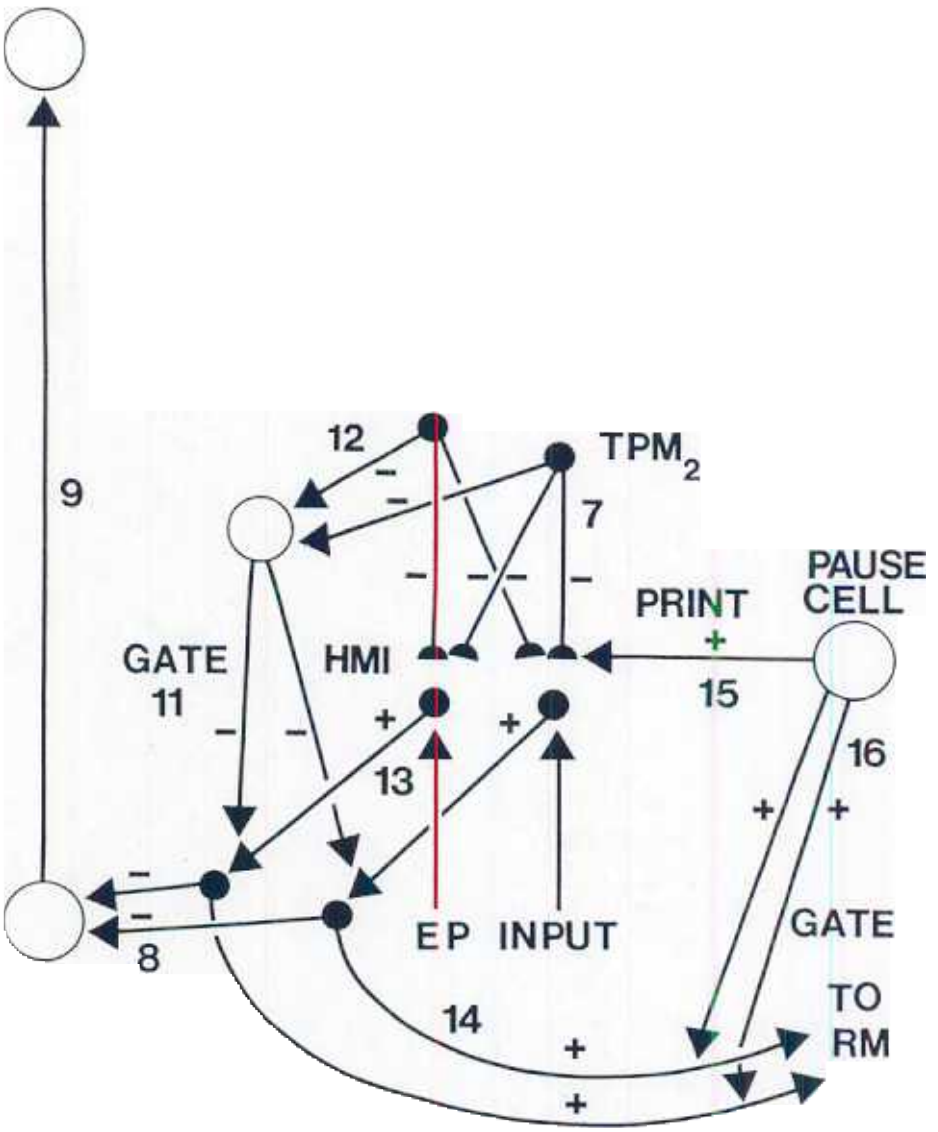

Figure 9.4. Microcircuitry of the predictive TPM_2 and its HMI: The text describes how light-mediated gating actions control read-out of reset signals to the TPM_1, read-in of target positions in motor coordinates at the HMI, and read-out of HMI vectors to the retinotopic map (RM). The RM recodes HMI vectors as focal activations within a spatial map. Abbreviation: EP = eye position.

activate the inhibitory pathway 2 are tonically active. (Tonic cells are depicted as open circles, phasic cells as closed circles.)

Pathway 10 can also gate shut the output pathways from TPM_1 to TPM_2. Pathway 10 is, however, inactive at this time. It is inhibited by pathway 9. Pathway 9 is active because pathway 8 from the HMI is inactive. Pathway 8 from the HMI is inactive because the inhibitory pathway 11 in Figure 9.4 is active. Pathway 11 is active because the TPM_2 is inactive, and thus cannot activate the inhibitory pathway 12. In all, the multisynaptic pathway 12 → 11 → 8 → 9 → 10 is shut off whenever the TPM_2 is inactive. Thus the TPM_2 does not block its own activation by the TPM_1 at times when it is inactive. The only impediment to activation of TPM_2 by TPM_1 at such a time is the inhibitory pathway 2.

Suppose that volition acts to inhibit the nonspecific pathway 2. Volition does not need to have specific consequences of any kind in order for the PCN to respond by rapidly reading out its sequence of stored commands. We now explain how this can happen.

When pathway 2 shuts off, excitatory signals from the active TPM_1 populations can reach the TPM_2 along topographically organized pathways such as pathway 3. The largest signal causes the most rapid and vigorous activation of its receptive population. The on-center (pathway 4) off-surround (pathway 5) feedback interactions within the TPM_2 contrast enhance this input pattern until only the positive feedback loop corresponding to the largest input is active. The activity within this loop is also stored in STM by the positive feedback interaction.

This active TPM_2 population gives rise to three types of output signals. First, a topographic inhibitory feedback signal (pathway 6) from the TPM_2 to the TPM_1 inhibits its source of feedforward TPM_1 →TPM_2 signals (pathway 3). The stored activity in TPM_2 hereby resets the temporal order information across TPM_1 by deleting the target position that was just stored in TPM_2. This reset event prepares TPM_1 to read-out the next-most-active population during the next saccadic cycle.

Second, activation of the TPM_2 causes the nonspecific inhibitory gating pathway 10 to turn on, thereby inhibiting further outputs from TPM_1. Thus, as one target position is being stored in TPM_2, it prevents other target positions from being read-into TPM_2. This gating action is accomplished by the multisynaptic pathway 12 → 11 → 8 → 9 → 10. The multiple links in this pathway are needed because this pathway also controls another important property. The multisynaptic pathway 12 → 11 also prevents tonic read-out of eye position signals from the HMI when no target positions are registered there (Section 4.6D). Pathway 12 → 11 does this by gating shut all HMI outputs—including gating signals of TPM_1 outputs and inputs to the RM—whenever the TPM_2 is inactive. Pathway 8 links the HMI output gate 12 → 11 to the TPM_1 output gate 9 → 10. Pathway 12 → 11 → 8 → 9 → 10 is thus an assemblage of two functionally distinct but synergetic subsystems.

Due to the importance of pathway 12 → 11 → 8 → 9 → 10, we now consider in greater detail how it inhibits TPM_1 outputs. When TPM_2 gets

activated, the inhibitory pathway 12 shuts off the inhibitory gating signal in pathway 11, thereby disinhibiting pathway 8. When the disinhibitory pathway 8 fires, it shuts off pathway 9, which disinhibits pathway 10, thereby shutting off outputs from the TPM_1. Pathway 8 can fire if it receives an excitatory signal from the HMI. Such a signal arises from the HMI if there is a mismatch between target position and eye position there. This property leads to a consideration of the third type of output signal that is controlled by the TPM_2.

The active population in the TPM_2 reads a target position, computed in muscle coordinates, into the HMI. This target position was learned using the mechanisms described in Chapter 4. Should this target position be different from eye position, then a mismatch will be registered in the HMI and the output pathways from the TPM_1 to the TPM_2 will be gated shut. By contrast, should this target position equal the eye position that is read into the HMI, then a match will be registered in the HMI, and the target position will cancel the eye position. Even if the inhibitory gating pathway 11 is off, no output can activate pathway 8 when a match occurs within the HMI. At such times, pathway 9 will inhibit pathway 10, thereby enabling the next target position to be read-into the TPM_2 from the TPM_1.

Thus, output along pathway 8 is contingent upon two events: activation of the TPM_2 and mismatch within the HMI. This contingency enables new inputs from the TPM_1 to reset an active TPM_2 should its target position have already been realized by prior eye movements. Were it not for this property, the sequential reset mechanism could sometimes get "stuck" before it could read-out its entire command sequence.

If the target position does not equal the eye position within the HMI, then no more signals can pass from the TPM_1 to the TPM_2, and a target position will be stored within the TPM_2, until a match can be generated within the HMI. Such a match cannot be generated by a change in target position, since the same target position will continue to be read into the HMI by the TPM_2 until a match occurs. A match can be created only if the eye moves until the eye position pattern matches the stored target position pattern.

9.6. Saccade Generation by Predictive Commands

In order for the eye to move in this way, the vector difference computed within the HMI must be able to generate a saccade, in the manner described by Chapter 4. Such a saccade can be elicited due to the second role played by the multisynaptic pathway $12 \rightarrow 11$. We have already summarized the first role of this pathway: When the TPM_2 is inactive, the tonically active gating pathway 11 inhibits the target cells at which the HMI vector is registered via pathway 13. At such times, the HMI vector encodes only eye position and would generate ceaseless saccade staircases were it not for pathway 11 (Chapter 7). By contrast, when the TPM_2 is activated, pathway 12 inhibits pathway 11. The target cells can then

TPM_2	HMI	12 → 11 → 8 → 9 → 10 → 14					
OFF	mismatch	OFF	ON	OFF	ON	OFF	OFF
ON	mismatch	ON	OFF	ON	OFF	ON	ON
ON	match	ON	OFF	OFF	ON	OFF	OFF

Table 9.1. Path activities during different target storage and matching phases of the predictive circuit in Figures 9.3 and 9.4.

register the HMI vector, which now compares target position with eye position. If a mismatch occurs, pathway 8 maintains this difference vector within the HMI by preventing STM reset of the TPM_2.

The second role of pathway 12 → 11 is to enable pathway 14 to carry muscle-coded vector signals to a retinotopic map (RM), where they are recoded as an active map position and stored in STM (Chapter 6). Output signals from this RM to the saccade generator (SG) update the eye position update network (EPUN) and generate a new saccade (Chapter 7). If this saccade is correct, it causes eye position to match target position in the HMI. If the saccade is incorrect, another mismatch is caused by a vector difference in the HMI, and another corrective saccade is generated. This description thus shows how temporally discrete target position information and temporally continuous eye position information are continuously matched at the HMI.

When eye position finally comes close to matching target position within the HMI, pathway 8 becomes inactive, despite the fact that the gating pathway 11 is still inhibited. Consequently pathway 10 is inhibited via the multisynaptic pathway 8 → 9 → 10, and output signals from the TPM_1 can once again reach the TPM_2. Otherwise expressed, a match at the HMI triggers a "rehearsal wave" that enables read-out of a new target position to occur from the TPM_1 (Grossberg, 1978a, 1980). The most active target position within the TPM_1 can then begin to compete with the stored target position within the TPM_2. The new target position wins this competition because the stored target position has only its positive feedback pathway within TPM_2 with which to maintain its activity, whereas the new target position has both the input from TPM_1 plus its positive feedback pathway with which to become instated. As the new target position inhibits the old one within the TPM_2, it can generate a new target position input to the HMI. This new target position causes a mismatch within the HMI, which in turn causes activation of pathway 10. Pathway 10 inhibits the TPM_1 output pathways and stabilizes the STM storage of the new target position within the TPM_1.

The next saccade in the predictive sequence can then be generated by the HMI. This cycle continues until all of the stored target positions within the TPM_1 have been sequentially inhibited. After all of the saccades have been performed, the TPM_1 is inactive, the TPM_2 is left storing the last target position command in STM, and the HMI is inactive because its target position and eye position match.

Several variations on this microcircuit design can be contemplated which compute the same basic functional properties but differ in terms of testable details. One of these variations is suggested by consideration of Table 9.1. Table 9.1 summarizes how different branches of the multisynaptic pathways 12 → 11 → 8 → 9 → 10 and 12 → 11 → 8 → 14 are activated during different phases of the saccadic cycle. This summary suggests that these pathways are designed rather parsimoniously. Inspection of Table 9.1 also shows that the pathways 8 and 10 are always on and off in-phase with each other. Consequently, pathway 8 could directly

inhibit the TPM_1 output pathways, instead of acting via the multisynaptic pathway 8 → 9 → 10, without altering any functional properties. We have included the tonic pathways 9 and 10 to call attention to a possible homology that may exist *in vivo* between pathway 10 and pathway 2. In Figure 9.3, both volitional signals and HMI signals play upon a general gating system whose function is to inhibit TPM_1 outputs. Pathways 2 and 10 simply mark the target cells where volitional and HMI-activated signals happen to feed into this general gating system. Thus Figure 9.3 would be "more parsimonious" than a figure in which pathway 8 directly inhibits TPM_1 outputs if such a general gating system were to exist. On the other hand, if pathway 2 inhibits not only TPM_1 →TPM_2 outputs but also the TPM_2 cells, then pathways 2 and 10 could not possibly be part of the same functional system, because STM storage of activity within the TPM_2 must be able to occur while pathway 10 is active. Enabling pathway 2 to inhibit TPM_2 cells achieves a small advantage, since shutting off the volitional signal after a saccade sequence terminates would then inhibit the stored final target postion within TPM_2, and thereby completely reset the TPM_2 in preparation for the next predictive saccadic sequence. Given such an arrangement, the rationale for the multisynaptic pathway 8 → 9 → 10 would collapse, and a direct inhibitory pathway 8 onto the TPM_1 output pathways might be expected to exist.

9.7. Two Types of Output Gates: Target-Driven Gates and Saccade-Driven Gates

To complete our description of PCN design, we need to consider in finer mechanistic detail several functional issues that were raised in Chapter 4. What happens if a predictive saccadic command generates a saccadic error due to inaccurate translation of a vector command into a movement, as in the Shebilske paradigm? How can a vector command be stored within the RM throughout a saccade? What prevents this stored vector command from being reset by eye positions attained during a saccade? How is learning within the HMI prevented from occurring except after a saccade terminates? We suggest that all of these functions are controlled by a single gating system. This gating system is, however, different from the *target-driven* gating system (pathway 12 → 11) that is modulated by storage of target positions within the TPM_2. We call the other gating system a *saccade-driven* gating system because its activity is modulated by the saccade generator (SG).

Figure 9.4 describes one possible version of the saccade-driven gating system. This version uses a pause cell input source. Such a pause cell is tonically on except during a saccade. Before and after a saccade occurs, when the pause cell population is active, it can open two gating pathways. A Now Print gate (pathway 15) enables pathway 7 from the TPM_2 to learn the eye position that was attained by the last saccade. A read-out gate (pathway 16) enables the HMI to read-out a vector command along pathway 14 to the RM, where it can be stored in STM. As soon as a saccade begins, both gates shut. Learning within the HMI therefore

cannot occur during a saccade, and new vectors from the HMI cannot be read-out during a saccade. Consequently the RM command can be stored until after a saccade is over. It can thereby sample second light error signals using its conditionable pathways within the adaptive gain (AG) stage, or cerebellum (Chapter 3). After a saccade terminates, the pause cell turns on and both gates open again. It takes some time, however, for the target position within the TPM_2 to be reset, for a new vector to be computed by the HMI, for this new vector to activate the RM via pathway 14, and for the new RM input to competitively inhibit the stored RM position. During this time interval, the old RM position can sample second light error signals at the AG stage.

Reset of the RM position is accomplished in the same way that an updated target position from the TPM_1 resets the TPM_2. This type of competitive reset of the RM follows from the ability of the RM to contrast enhance and store retinotopic inputs using a shunting on-center off-surround feedback network (Chapter 6).

The conjoint action of target-driven gating and saccade-driven gating help to refine our explanation of the Shebilske (1977) data in Section 4.12. Shebilske found that incorrect saccades in response to a briefly flashed light could correct themselves in the dark. In Figure 9.4, the target-driven gate remains closed after such an incorrect saccade, because a mismatch of eye position with target position occurs within the HMI. After the incorrect saccade terminates, however, the saccade-driven gate opens, and enables the new HMI vector to instate an updated movement command in the RM. A corrective saccade ensues.

Several alternative saccade-driven gating schemes all possess the functional properties that we need, and may therefore exist *in vivo*. For example, pause cells could inhibit tonic cells that inhibit learning and vector read-out, thereby causing a disinhibitory gating reaction, rather than a direct excitatory gating reaction, in pathways 7 and 14. A saccadic burst cell population could directly inhibit learning in pathway 7 and read-out to the RM in pathway 14 (Chapter 7). To test whether one has discovered such gate-controlling cells *in vivo*, lesion experiments of several types can be carried out. For example, suppose that pathway 11 in the target-driven gating system were cut. Then outputs from the HMI could occur even in the absence of target position inputs. Saccade staircases would therefore ensue. Suppose that pathway 16 in the saccade-driven gating system were cut. Then HMI outputs could continually reset the RM during a saccade. Consequently, by the time a correct saccade was over, the RM would have received an input corresponding to the smallest vectors that it can encode. This, in itself, is not a bad property. However, if the RM were then required to relearn its conditioned AG gains, it could not do so, because it would have "forgotten" the vector which caused the saccade by the time the saccade was over.

9.8. Parietal Light-Sensitive and Saccade Neurons

The on-center cells within the TPM_1 and the TPM_2 have properties

that are analogous to those of light-sensitive neurons and saccade neurons in the parietal cortex, respectively. The TPM_1 cells are light-sensitive because lights are an important information source in building up a TPM in head coordinates. If these cells in the TPM_1 store activity in STM without triggering a saccade, then their activity can persist for a long time. However, as soon as they initiate a saccade by activating the TPM_2, their activity is rapidly inhibited by negative feedback from the TPM_2 (pathway 6 in Figure 9.3). Thus these TPM_1 cells have the curious property that they shut off just when the saccade begins that their stored light might have been expected to control. Yin and Mountcastle (1977, p.1383, Figure 2A) have reported the existence of a light-sensitive cell type with similar properties in their parietal lobe data.

Supraliminal activation of TPM_2 cells, by contrast, causes these cells to generate a saccade and to persistently fire due to their STM storage until the saccade is over. These TPM_2 cells are not as sensitive to light because they are activated by TPM_1 cells, and only when the control processes which regulate pathways 2 and 10 in Figure 9.3 allow. These cell properties resemble those of the saccade neurons in the parietal cortex (Lynch, 1980; Motter and Mountcastle, 1981). In Chapter 11, we will summarize the similar properties of visuomovement cells and movement cells in the frontal eye fields (Bruce and Goldberg, 1984).

Experimental tests can be made of whether parietal light-sensitive and saccade neurons or frontal visuomovement and movement neurons are examples of TPM_1 and TPM_2 cells, respectively. Electrical stimulation of a saccade neuron should inhibit its light-sensitive source via pathway 6 in Figure 9.3. Cutting or otherwise inactivating this inhibitory feedback pathway should enable the light-sensitive neuron to remain on during the ensuing saccade. The light-sensitive cell should also respond to certain eye positions, in particular the eye positions which together with appropriate retinal positions give rise to the target position of a TPM_1 cell. The saccade neuron should remain on when an incorrect saccade occurs in the dark until after a corrective saccade of Shebilske type is made.

In addition to cells with properties of light-sensitive neurons and saccade neurons, the PCN also contains other cells with interesting characteristics. We have already noted that the target-driven and saccade-driven gating systems can be controlled by pause cells or by burst cells. In fact, several types of pause cells are included in Figure 9.4. The pause cell which controls pathway 11 will shut off just before and during a correct PCN-generated saccade and will then turn back on until the next TPM_2 command is instated. However, after an incorrect PCN-generated saccade occurs, such a pause cell will remain off until a corrective saccade of Shebilske type can be made. By contrast, the pause cells which control pathways 15 and 16 will turn off during any saccade, whether or not it is generated by the PCN and whether or not it is correct. The cells which control the "rehearsal wave" pathway 9 are also pause cells. These cells turn off only when the TPM_2 is active and a mismatch is registered within the HMI. Thus they turn off only for saccades that are generated by the PCN, and remain off during corrective saccades (Section 5.5). In addition

to these differences, the onset and offset times of these different types of pause cell types are not synchronous.

Thus a neurophysiological study of a circuit such as Figure 9.4 would reveal several populations of pause cells whose local cell properties, as such, would provide an insufficient index of the cells' functional role in the network. This very difficulty also illustrates the more optimistic evolutionary theme that a small number of local cell properties can give rise to a complex repertoire of network functions.

9.9. Switching between Movement and Postural Eye Position Maps: Frontal Eye Field Control

Our analysis of the PCN has focused upon how a pattern of stored target positions can give rise to a sequence of predictive saccades. We have shown how such a stored activity pattern at the TPM_1 can sequentially activate an RM and can correct Shebilske-type errors due to mistranslation of a saccadic command between the RM and the eye muscles. This analysis calls attention to the following issues.

Before a predictive saccade takes place, the RM of the PCN and an EPM that registers initial eye position must both send their movement commands to the SG. The RM and the EPM are part of the eye position update network (EPUN) that was described in Chapter 7. *After* a predictive saccade is over, it takes some time for the next TPM_1 command to be instated within the TPM_2 and then give rise to the next RM output. In order to prevent postsaccadic drifts from occurring between these saccades, the EPM of the tension equalization network (TEN) must switch on to read-out its conditioned postural gains.

This discussion emphasizes that two different types of EPMs are turned on and off at different phases of the saccadic cycle in order to guarantee the accuracy of movements and the stability of postures. The EPM which cooperates with the RM to generate the SG input is called the *movement EPM*, or EPM_M. The EPM which controls postural gains is called the *postural EPM*, or EPM_{Pos}.

Figure 9.5 schematizes the timing of RM, EPM_M, and EPM_{Pos} onsets and offsets during a series of saccades. The durations of the EPM_M signals are not certain. They may be as long as the sillouettes drawn in Figure 9.5, or they may be more synchronous with RM durations, as illustrated by the cross-hatched area. We will develop some implications of the latter hypothesis.

Suppose that the RM and the EPM_M are synchronously switched on and off, while the EPM_{Pos} is being switched off and on. The simplest way to accomplish these properties is to suppose that the gating system which turns the RM on and off also turns the EPM_M on and off while it turns the EPM_{Pos} off and on. This gating system is a saccade-driven gating system.

It remains to suggest where the cells that carry out these EPM functions may be found. We hypothesize the the computation and storage

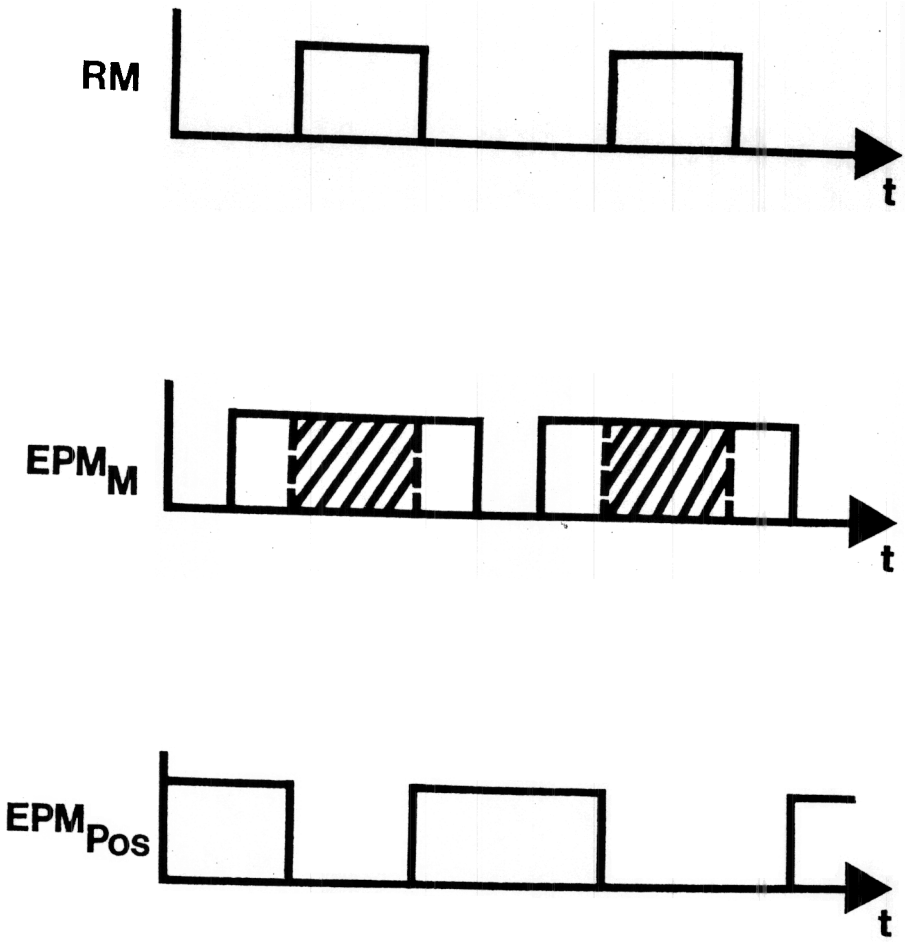

Figure 9.5. Timing of activity in the retinotopic map (RM), the movement eye position map (EPM_M), and the postural eye position map (EPM_{Pos}): The RM and EPM_M activities are out-of-phase with the EPM_{Pos} activities. The EPM_M activities may occur only during the hatched times or during the broader time intervals that alternate with EPM_{Pos} activity.

of temporal order information by the TPM_1 takes place in the frontal eye fields (FEF) as part of the frontal lobes' role in cognitive and motor planning. All of the subsequent PCN stages—TPM_2, HMI, and RM-could also, in principle, take place in the FEF, or in regions to which the FEF projects. Let us consider the hypothesis that the EPM_M and/or the EPM_{Pos} are also computed within the FEF. Is there any evidence for such an assumption? Cells within an EPM have properties analogous to those of fixation neurons, which exist in the FEF (Suzuki and Azuma, 1977). Lesions of the FEF can cause fixation deficits in monkeys (Latto and Cowey, 1971). Such a deficit could be caused, for example, by destruction of the EPM_{Pos}, since the conditioned gains needed to maintain a stable posture could not then be read into the MN cells. By contrast, both humans and monkeys can maintain visual fixation after posterior parietal lesions (Hyvärinen, 1982). Hyvärinen compares the fixity of gaze after posterior parietal damage in humans to the somatic dystonia that follows extrapyramidal cortical ablation. He concludes that "the role of the posterior parietal cortex is to interrupt fixation when an interesting visual stimulus appears in the visual periphery" (Hyvärinen, 1982, p.1112). We will discuss possible interactions between the posterior parietal cortex and the frontal eye fields in Chapter 11.

9.10. Direct Computation of Predictive Difference Vectors from Stored Retinotopic Positions?

We now consider some functional problems that would need to be solved if a neural system were to directly transform stored retinotopic positions $R_1, R_2, \ldots, R_m$ into difference vectors.

A. *Getting Started*

The first problem is to get started. Consider the first light R_1 in a stored predictive sequence of retinotopically coded positions. How is R_1 encoded into a difference vector? With what stored quantity R_0 does R_1 get compared in order to compute the correct difference vector $R_1 - R_0$? To emphasize the nature of this difficulty, suppose that the last saccade was made in response to a sound, not to a light. How does the system know how to subtract R_1 from the "retinotopic" position R_0 of the sound?

These considerations suggest that a brain capable of solving this problem would encode all sources of movement commands, including auditory sources, within a *universal* retinotopic map, or URM (Figure 9.6a). Such a URM would store only the retinotopic position of the last movement command to be executed. When R_1 was read-out to be converted into a difference vector, so too would the previously stored command in R_0 in the URM (Figure 9.6b). Both retinotopic commands would be read-into a region where the difference vector $R_1 - R_0$ could be computed.

As R_1 was read-out of its RM, it would self-inhibit its RM representation to prevent perseverative performance of its movement command. It would also inhibit R_0 at the URM in order to instate itself into STM storage there. Once stored in the URM, R_1 could later be read-out along with R_2 to compute the difference vector $R_2 - R_1$ (Figure 9.6b).

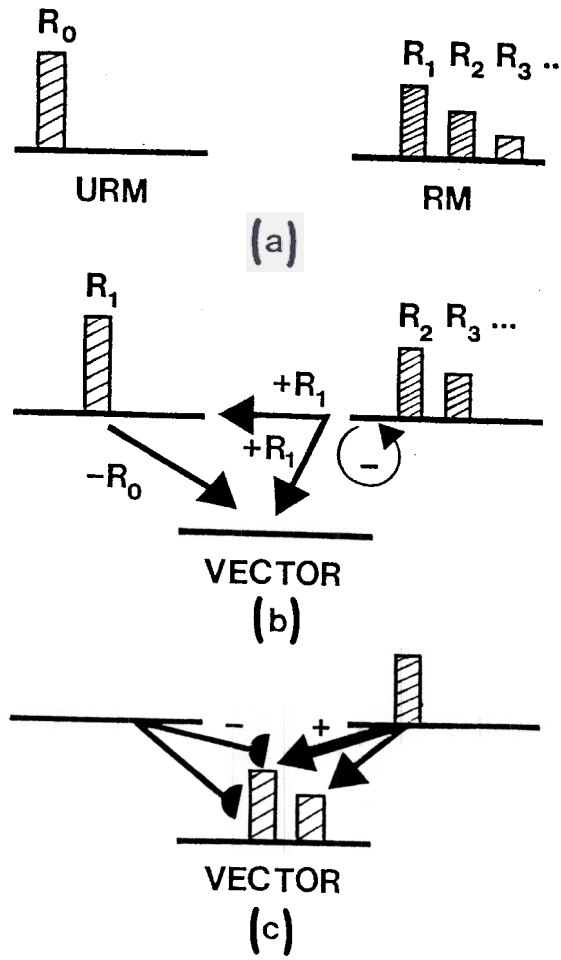

Figure 9.6. A possible circuit for direct recoding of retinotopic positions R_1, R_2, R_3, ..., into difference vectors: (a) Temporal order information is stored in short term memory at the retinotopic positions R_1, R_2, R_3, ..., of the retinotopic map (RM). A previously stored position R_0 is also active at the universal retinotopic map (URM); (b) Activation of these maps transforms R_1 and R_0 into a vector representation. R_0 is replaced in URM by R_1, as R_1 self-inhibits its activity at RM; (c) The transform from RM to the vector representation can be prewired, whereas the transform from the URM can be learned as in the head-muscle interface (HMI).

In summary, the existence of a URM and of coordinated STM reset and storage operations between the URM and all other RMs is implicit in the apparently simple idea that retinotopic commands are directly recoded into difference vectors. These constraints are relatively simple to realize compared with the following ones.

B. *Vector Sign Reversal*

Consider a pair R_1 and R_2 of retinotopic positions corresponding to horizontal eye motions. For simplicity, suppose that these positions are encoded within a 1-dimensional RM whose positions are labeled by real numbers. Consider the case in which both R_1 and R_2 lie to the right of the fovea. Suppose moreover that R_2 lies to the left of R_1. If we label the fovea with the number 0, then $R_1 > R_2 > 0$ and $R_2 - R_1 < 0$. The problem can now be stated: How can retinotopic positions with one sign (in this case positive) generate a difference vector with the correct length and opposite sign (in this case negative)? The need for such a property is strikingly illustrated by the data of Mays and Sparks concerning the quasi-visual (QV) cells of the superior colliculus (Section 4.2).

We have not succeeded in discovering any neurally plausible mechanism whereby these properties can be achieved through a mapping of an RM and a URM into another *spatial* map (Chapter 6). No plausible combination of narrow or broad excitatory or inhibitory spatial gradients seems capable of generating a spatial encoding which can simultaneously represent the length and the direction of all physically realizable difference vectors. In particular, a single difference vector can be realized by many pairs of retinotopic positions. We see no way to neurally accomplish this many-to-one transformation using signals between spatial maps.

C. *Motor Recoding and Dimensional Inconsistency*

From this perspective, one can better appreciate the type of solution that we have suggested for computing difference vectors. In Chapter 4, we introduced the head-muscle interface, or HMI, and suggested that vector differences are computed in muscle coordinates, not within a spatial map. We also showed that the same process which learns the transformation into muscle coordinates can automatically compute difference vectors. Motor coordinates are ideal for computing difference vectors because the balance of excitation and inhibition across their agonist-antagonist populations can easily encode a reversal of sign, as in $R_2 - R_1 < 0$.

Retinotopic coordinates are, however, dimensionally inconsistent with motor coordinates. By contrast, target position coordinates are dimensionally consistent with motor coordinates. That is why recoding of retinotopic coordinates into target position coordinates was necessary before the target position coordinates could be recoded into motor coordinates at the HMI.

D. *Opponent Recoding and Linearity*

Another possibility is suggested by the HMI example. Perhaps the opponency of agonist-antagonist populations can be used without their motor interpretation. Perhaps each RM is recoded into a vector field that

is encoded with opponent populations, albeit opponent populations that are not derived from motor outflow or inflow signals (Figure 9.6c). Such a recoding into opponent populations could be achieved by spatial gradients from the RM to the vector field. The URM could, in turn, learn these opponent patterns using the learning mechanism which is found within the HMI (Figure 9.6c).

This scheme also faces a formidable obstacle. Unless the spatial gradients from the RM to the vector field vary *linearly* with RM position, then retinotopic positions R_1, R_2, and R_3 which generate the same *formal* difference vector $R_3 - R_2 = R_2 - R_1$ could generate different patterns of opponent activation at the vector field. Moreover, the same pattern of opponent activation at the vector field could be generated by different vector differences of RM and URM positions. Such a vector field could not generate consistent movement commands.

One can now better appreciate how the HMI solves this problem. Corollary discharges derived from motor outflow signals are used to compute present position signals at the HMI. These present position signals are linearly related to present eye position signals due to the action of the muscle linearization network, or MLN (Chapter 5). Target position signals learn these corollary discharges at the HMI, and the HMI computes a vector whose suprathreshold activities vary linearly with the difference of target position and present position. Then these vectors can be consistently mapped back into retinotopic coordinates to generate movement commands which are corrected by second light error signals (Chapter 3).

We have not discovered any neurally plausible mechanism whereby direct recoding of retinotopic maps RM and URM into difference vectors can achieve these linearity properties. Due to the combined impact of these failures, we have reached the tentative conclusion that retinotopically coded positions are not directly recoded as difference vectors. If this conclusion is correct, then retinotopically coded cells in a predictive region, such as the frontal eye fields, may elsewhere be recoded into target positions before being recoded once again into difference vectors in motor coordinates at an HMI.

Such a recoding into target position coordinates may be implicit rather than explicit. As we will show in Chapter 10, combinations of retinotopic and eye position signals can be distributed across a network in such a way that certain cells act *as if* they are encoded in target position coordinates, even though they may also be excited by many different combinations of retinotopic and eye position signals.

It cannot yet be ruled out that other neural designs than the ones we have considered are capable of directly transforming retinotopic positions into difference vectors. If such a design does exist, then we might anticipate that it solves the functional problems which we have identified, such as coordinated reset of all RMs with a URM, and vector sign reversal using opponent processing. In addition, such a network would most likely inhibit all the stored positions in its predictive RM each time a new pattern of lights was instated there. Otherwise, sequential scanning of a fixed

set of light sources could encode each light with a new retinotopic position every time the eye moved. After n eye movements, each nonfoveated light would have n retinotopic encodings within the RM, thereby creating massive confusions in the choice of which light should be foveated next. It therefore seems unlikely that any RM which is not totally reset by a new light input is directly recoded into difference vectors. This STM reset property can be used to indirectly test the nature of the retinotopic recoding into difference vectors in brain regions where all the transformations leading to difference vectors are not under direct experimental control.

If an RM is recoded into an invariant target position map (TPM) before the TPM is used to compute difference vectors, then total reset of the RM by each new light is not necessary, but also is not forbidden. We now turn to an analysis of how an RM may be implicitly recoded into an invariant TPM.