CHAPTER 11
VISUALLY REACTIVE, MULTIMODAL, INTENTIONAL, AND PREDICTIVE MOVEMENTS: A SYNTHESIS

11.1. Avoiding Infinite Regress: Planned Movements Share Reactive Movement Parameters

In this Chapter, we tie together the design principles and mechanisms which we have introduced in the preceding chapters. This synthesis quantifies a developmental sequence in which obligatory saccadic reactions to flashing or moving lights on the retina are supplemented by attentionally mediated movements towards motivationally interesting or intermodal sensory cues. These movements are supplemented once again by predictive saccades which form part of planned sequences of complex movement synergies that can totally ignore the sensory substrate on which they are built.

The analysis of such a complex system would be hopeless in the absence of sufficient design constraints. The previous chapters illustrate our contention that an analysis of development and learning is essential to discover sufficiently many constraints. This chapter continues such an analysis until it discloses a small number of closely related global network designs that are capable of simultaneously satisfying all of the constraints.

The primacy of learning constraints is perhaps most urgently felt when one squarely faces the problem of infinite regress. On what firm computational foundation can adaptive calibrations be based? If parameters in several subsystems can all be changing due to different types of learning signals, then what prevents them from confounding each other's learning by causing a global inconsistency within the system? An analysis of how the problem of infinite regress may be solved shows how to hierarchically organize motor learning circuits so that each circuit can benefit from and build upon the learning of a more primitive learning circuit. This hierarchical organization of learning problems can be profitably compared to the hierarchical organization whereby simple motor synergies in the spinal cord can be temporally organized by descending cortical commands (Grillner, 1975; Grossberg, 1969d, 1974).

The need to solve the problem of infinite regress is evident in our method of calibrating a target position command in muscle coordinates at the head-muscle interface, or HMI (Chapter 4). We have suggested that an active target position map (TPM) population will sample whatever eye position corollary discharges are read-into the HMI after the eye comes to rest. These eye position signals may or may not correspond to the target position that is active within the TPM. The TPM will eventually sample the correct eye position only if an independent learning mechanism enables the eye to accurately foveate retinal lights.

We have suggested that this independent learning mechanism forms part of the RCN, or retinotopic command network (Chapter 3), whereas the HMI forms part of the VCN, or vector command network (Chapter 4).

One goal of this chapter is to complete our analysis of the interactions that occur between these two types of networks.

The RCN learning mechanism utilizes a multisynaptic pathway that probably includes the superior colliculus (SC) and the visual cortex (Huerta and Harting, 1984). This pathway is hypothesized to be activated by flashing and moving lights on the retina (Frost and Nakayama, 1983), and to be functional at an early stage of development. The VCN learning mechanisms, notably the TPM and HMI, are hypothesized to process attentionally modulated neocortical commands. These attentionally modulated VCN commands may be able to inhibit the relatively simple motion-sensitive signals of the RCN at a later developmental stage, but before they can do so, they first need to be calibrated using the motion-sensitive SC pathway as a reliable computational substrate. The first two problems that shall be considered in this chapter are how this calibration problem is solved and how attention can modulate cortical movement commands.

11.2. Learning and Competition from a Vector-Based Map to a Light-Based Map

The motion-sensitive system uses second light error signals to improve its foveating capability (Chapter 3). The HMI uses the learned accuracy of the motion-sensitive saccades in order to compute accurate difference vectors in muscle coordinates, which can then be transformed into a retinotopic map, or RM (Chapter 6). It remains to show how such an RM can generate accurate saccadic movements. Two main possibilities exist, and we suggest that both may be used *in vivo*. In one possibility, the RM learns to read-out commands to the saccade generator (SG) using the output pathways of the motion-sensitive RCN system. We suggest that this pathway is used by attentionally modulated commands from the parietal cortex, at least in monkeys and humans. In the other possibility, the RM controls its own output pathways through which it learns its own correct movement parameters. We suggest that this pathway is used by predictive commands from the frontal eye fields (FEF) in monkeys and humans. Each possibility implies a distinct constellation of functional and structural properties which can be empirically tested.

First we study how an RM in the VCN can learn to activate the RCN output pathways, and to thereby generate accurate saccades using the adaptive gains which had previously been learned within the RCN. This analysis resolves the following paradox: How does the VCN both learn from and inhibit the RCN? In order for the RM to become associated with the correct RCN output pathways, these pathways must be *excited* during the learning phase. In order for VCN commands to supplant RCN commands, the RM must be able to *inhibit* RCN output pathways during later performance trials. How can excitation during learning give rise to inhibition during performance?

The sampled RCN output pathways also form an RM. Thus our problem is to study learning of an $RM \rightarrow RM$ associative transform. In order

to avoid notational confusion, we denote the RM within the VCN by RM_V and the RM within the RCN by RM_R.

Before learning occurs, each RM_V population gives rise to conditionable pathways which are distributed nonspecifically across the RM_R, because it is not known *a priori* which RM_R population will be associated with which RM_V population. The LTM traces in these pathways are initially small. While the RCN is still learning how to make accurate saccades, the LTM traces within the $RM_V \rightarrow RM_R$ conditionable pathways remain broadly distributed and small, because each RM_V population can become associated with many different RM_R populations. Several reasons for this exist. Until accurate saccades persistently occur, the TPM cannot form with precision (Chapter 10) because each fixed combination of retinal position and initial eye position can sample a different final position whenever a foveation error occurs. Each TPM population can, in turn, sample many corollary discharge patterns at the HMI due to foveation errors. Thus, across learning trials, each *fixed* combination of TPM and corollary discharge inputs to the HMI can activate many different vectors there and, by extension, many different RM_V populations. Thus each RM_V population can sample many different RM_R populations until learning within the RCN stabilizes. Consequently the $RM_V \rightarrow RM_R$ associative transform will remain broadly distributed and weak until the RCN becomes capable of reliably generating correct foveations.

Competitive interactions are assumed to occur among the RM_R populations in order to choose a winning population for storage in short term memory (STM). Such a population's activity must be stored in STM until after the saccade terminates, so that it can sample any second light error signals that may then be registered (Chapter 2). During the learning trials when foveation errors still regularly occur within the RCN, the $RM_V \rightarrow RM_R$ signals, being weak, lose the competition for STM storage to the first light-activated movement signals. The RM_V consequently does not interfere with the initial stages of learning within the RCN.

By contrast, after the RCN becomes capable of generating correct foveations, the LTM traces of each RM_V population begin to sample the same RM_R population on all future learning trials. The LTM traces of these pathways become strong due to the cumulative effect of these consistent learning trials, while the LTM traces of other, less favored, pathways become weak. The $RM_V \rightarrow RM_R$ associative map thereby becomes more topographic as some of its LTM traces become large due to the cumulative effects of consistent learning on successive saccadic trials. After the LTM traces become sufficiently large, a conditioned signal from RM_V to RM_R can successfully compete with a first light-activated movement signal within the RCN. When this happens, attentionally amplified VCN commands can take precedence over motion-sensitive RCN commands.

In order for this scheme to work, the RM_R must occur at a stage that is prior to the RCN's error signal pathway to the adaptive gain (AG) stage, or cerebellum (Chapter 3). Then the $RM_V \rightarrow RM_R \rightarrow AG \rightarrow SG$ pathway can use the cerebellar gain that was previously learned in response to

second light error signals of RCN vintage (Figure 11.1).

This example shows in several ways how a suitably designed network hierarchy can avoid infinite regress by exploiting previously learned movement parameters. The RCN helps the VCN to learn its TPM, to calibrate its HMI, and to sample accurately calibrated movement commands using its AG stage adaptive gains.

11.3. Associative Pattern Learning and Competitive Choice: Non-Hebbian Learning Rule

This section mathematically demonstrates the properties which were intuitively explained within the previous section. Suppose that a network $F^{(1)}$ sends broadly distributed conditionable pathways to a disjoint network $F^{(2)}$. We have in mind, of course, the special case in which $F^{(1)} = RM_V$ and $F^{(2)} = RM_R$, but our discussion is more general. Denote the cells of $F^{(1)}$ by $v_j^{(1)}$ and the cells of $F^{(2)}$ by $v_i^{(2)}$. The cells $F^{(1)} = \{v_j^{(1)}\}$ are generically called *sampling cells* and the cells $F^{(2)} = \{v_i^{(2)}\}$ are generically called *sampled cells*.

For simplicity, let each cell population $v_j^{(1)}$ emit the sampling signal S_j within the axons e_{ji} from $v_j^{(1)}$ to all cells $v_i^{(2)}$ in $F^{(2)}$. Let a long term memory (LTM) trace z_{ji} be computed at the synaptic knobs of each pathway e_{ji}. Suppose that z_{ji} computes a time-average of the signal S_j times the postsynaptic potential $x_i^{(2)}$ of $v_i^{(2)}$. It has been proved, under general mathematical conditions, that the pattern of LTM traces $(z_{j1}, z_{j2}, \ldots, z_{jn})$ learns a weighted average of the STM patterns $(x_1^{(2)}, x_2^{(2)}, \ldots, x_n^{(n)})$ that occur across $F^{(2)}$ while S_i is active (Grossberg, 1968b, 1969c, 1972b, 1982a). This is the same theorem about associative pattern learning that we used to clarify the learning of an invariant TPM in Chapter 10. In the present example, a different version of the same learning laws shows how a topographic associative map can be learned.

To see how this can happen, suppose for definiteness that $F^{(2)}$ chooses at most one activity $x_i^{(2)}$ for storage in STM at any given time, as in Section 2.6. Suppose that the population $v_{i_j}{}^{(2)}$ with $i = i_j$ in $F^{(2)}$ is stored in STM at times when population $v_j^{(1)}$ is active in $F^{(1)}$. Since S_j is active during time intervals when the particular STM trace $x_{i_j}{}^{(2)}$ is chosen for STM storage, the relative LTM trace $z_{ji_j}(\sum_{k=1}^{n} z_{jk})^{-1}$ monotonically approaches 1 as learning proceeds. In other words, only the LTM trace z_{ji_j} becomes large, as all the other traces $z_{j1}, z_{j2}, \ldots, z_{jn}$ become small. This is a non-Hebbian form of learning, because both conditioned increases *and* conditioned decreases in LTM traces z_{ji} can occur due to pairing of

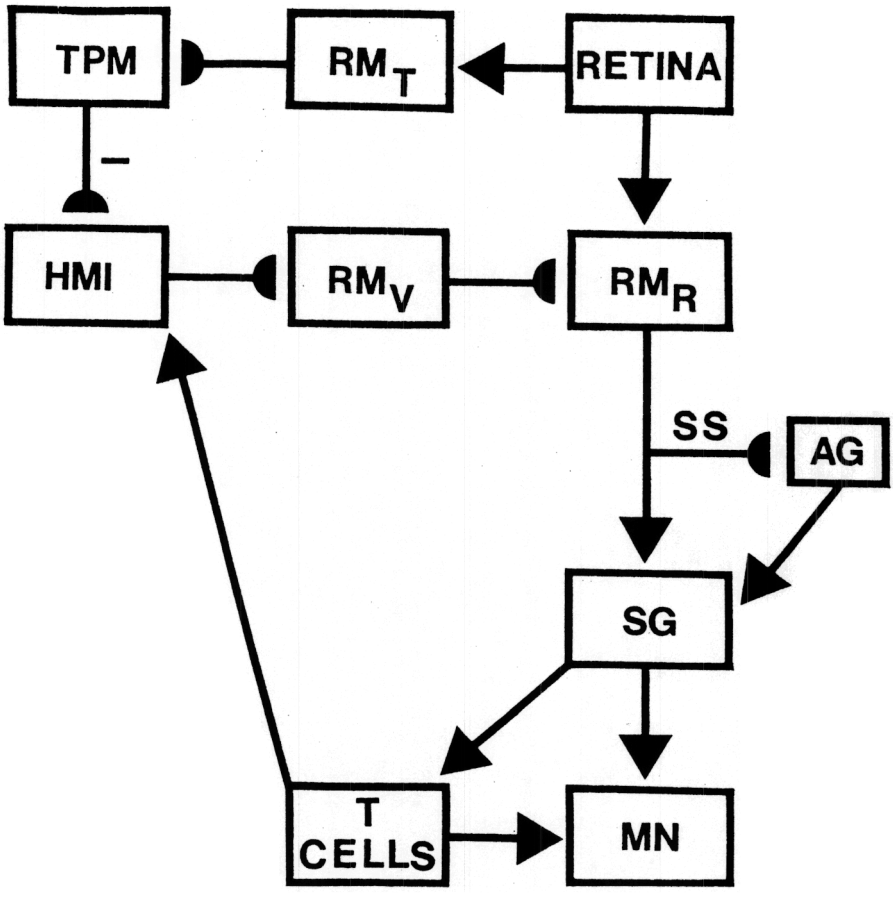

Figure 11.1. Interactions between the retinotopic command network (RCN) and the vector command network (VCN): The coupling between these subsystems enable the motorically encoded vectors in the head-muscle interface (HMI) of the VCN to be retinotopically recoded within the retinotopic map (RM_R) of the RCN. HMI vectors can then use the unconditioned movement pathway $RM_R \rightarrow$SG and the conditioned movement pathway $RM_R \rightarrow$AG$\rightarrow$SG of the RCN to generate accurate eye movements. Abbreviations: SS = sampling signals, AG = adaptive gain stage, SG = saccade generator, T cells = tonic cells, MN = motoneurons, TPM = target position map, RM_T, RM_V, and RM_R = retinotopic maps.

S_j and $x_i^{(2)}$. A simple learning equation with this property is

$$\frac{d}{dt}z_{ji} = \epsilon S_j(-z_{ji} + x_i^{(2)}). \tag{11.1}$$

Performance of the learned pattern of LTM traces occurs as follows. Suppose that the signal that reaches each population $v_i^{(2)}$ from $v_j^{(1)}$ equals $S_j z_{ji}$. In other words, S_j is multiplicatively gated by the LTM trace z_{ji} before the gated signal $S_j z_{ji}$ perturbs $v_i^{(2)}$. Since only z_{ji_j} among all the LTM traces $z_{j1}, z_{j2}, \ldots, z_{jn}$ is large, the only large gated signal is $S_j z_{ji_j}$. Consequently, only $v_{i_j}{}^{(2)}$ is significantly activated by $v_j^{(1)}$.

To illustrate how all the inputs to $F^{(2)}$ influence its dynamics through time, let us approximate the choice-making behavior of $F^{(2)}$ by an algebraic rule. This is permissible in the present case because the LTM traces, which are the variables of primary interest, fluctuate through time much more slowly than the STM traces $x_i^{(2)}$ or their inputs I_i. Thus, let

$$J_i = I_i + \sum_{j=1}^{m} S_j z_{ji} \tag{11.2}$$

equal the total input to $v_i^{(2)}$ at any time, where I_i is a light-induced motion-sensitive input and $\sum_{j=1}^{m} S_j z_{ji}$ is the total conditioned input from $F^{(1)}$. Also let

$$x_i^{(2)} = \begin{cases} 1 & \text{if } J_i > \max\{\delta, J_k : k \neq i\} \\ 0 & \text{otherwise} \end{cases} \tag{11.3}$$

be the rule whereby $F^{(2)}$ activities are chosen and stored in STM. In (11.3), parameter $\delta > 0$ estimates the quenching threshold (QT) of $F^{(2)}$ (Section 2.6). By (11.3), $v_i^{(2)}$ is chosen and stored with the normalized activity 1 if its input exceeds the QT and is the largest total input received by $F^{(2)}$. Equation (11.3) approximates a network with broad lateral inhibitory interactions and signal functions selected to make a global choice (Section 2.6).

Before learning occurs, all the LTM traces are small and uniformly distributed across $F^{(2)}$. Consequently the maximal J_i is determined by the maximal I_i. The choice behavior of $F^{(2)}$ is then controlled by the motion-sensitive inputs I_i, as in Chapter 3. Suppose that at most one signal S_j is large at any time (because only one RM_V population is activated by the HMI at any time). Also suppose, as above, that activity in $v_{i_j}{}^{(2)}$ is paired with activity in $v_j^{(1)}$, so that only z_{ji_j} becomes large due to associative

learning. In particular, equation (11.1) implies that z_{ji_j} approaches 1 as a result of learning and all z_{ji} approach zero, $i \neq i_j$. These conclusions follow from the conjoint action of several properties. For one, $S_j = 0$ implies that $\frac{d}{dt} z_{ji} = 0$; no learning occurs unless the sampling signal S_j is active. If $S_j > 0$ when $x_{i_j} = 1$, then (11.1) implies that z_{ji_j} approaches 1. If $S_j > 0$ when $x_i = 0$, $i \neq i_j$, then (11.1) implies that all z_{ji} approach zero, $i \neq i_j$.

Suppose, moreover, that each S_j becomes larger than the maximal possible input size I_i whenever S_j is positive. The size of S_j may, in particular, be amplified by attentional gain control or by incentive motivational signals (Section 2.6). After learning occurs, consider a performance trial during which the sampling signal S_j is on while a motion-sensitive input I_i is on in a different $F^{(2)}$ channel ($i \neq i_j$). Because $z_{ji_j} \cong 1$ and all $z_{ji} \cong 0$, $i \neq i_j$, it follows by (11.2) that

$$J_{ij} \cong S_j z_{ji_j} \cong S_j \tag{11.4}$$

and that

$$J_i \cong I_i, i \neq i_j. \tag{11.5}$$

Consequently,

$$J_{ij} > J_i, i \neq i_j. \tag{11.6}$$

By (11.3), it follows that

$$x_{ij} = 1 \tag{11.7}$$

whereas all

$$x_i = 0, i \neq i_j. \tag{11.8}$$

In other words, after learning occurs, the VCN-activated population $v_{i_j}{}^{(2)}$ wins the STM competition over the RCN-activated population $v_i^{(2)}$ despite the fact that $v_i^{(2)}$ wins the STM competition over $v_{i_j}^{(2)}$ before learning occurs.

By winning the STM competition, population $v_{i_j}{}^{(2)}$ also stabilizes the memory within its LTM trace z_{ji_j}. This property also follows from the learning equation (11.1). When $S_j = 0$, it follows that $\frac{d}{dt} z_{ji_j} = 0$, so that no learning *or forgetting* occurs unless the sampling signal S_j is active. Suppose that when $S_j = 0$, z_{ji_j} starts out equal to 1 due to prior learning. If z_{ji_j} can cause $x_{i_j}{}^{(2)}$ to quickly win the STM competition, then also $x_{i_j}{}^{(2)} = 1$ at these times. By (11.1), even if $S_j > 0$,

$$\frac{d}{dt} z_{ji_j} = S_j(-1 + 1) = 0. \tag{11.9}$$

Consequently, once LTM trace z_{ji_j} learns the value 1, it reinforces its memory of this value during subsequent performance. Note, however, that the LTM trace has not lost its plasticity. It is in a dynamical equilibrium with the STM activation that it causes.

In summary, after learning of the associative transform $F^{(1)} \rightarrow F^{(2)}$ occurs, $F^{(1)}$ can control STM decisions within $F^{(2)}$, even though $F^{(1)}$ had no influence over these STM decisions before learning occurred. Moreover, $F^{(1)}$ can maintain its control in a stable fashion due to the feedback exchange between STM and LTM that occurs in equations (11.1) and (11.3).

This model describes a critical period that ends when its LTM traces achieve a dynamical equilibrium with its STM traces. In the absence of further learning-contingent structural changes within the network, such a model predicts that new learning can occur in the adult if the correlations between sampling populations $v_j^{(1)}$ and sampled populations $v_i^{(2)}$ are systematically changed using artificially large inputs I_i, despite the absence of any apparent plasticity before such a change takes place. Artificially large direct inputs to $F^{(2)}$ are needed to overcome the competitive advantage of $F^{(1)}$ signals after learning takes place. Additional memory buffering interactions, modulated by the catecholaminergic transmitter norepinephrine, have been predicted to occur in more complex associative mapping models (Grossberg, 1976b, 1980, 1984) and have received some experimental support (Kasamatsu and Pettigrew, 1976; Pettigrew and Kasamatsu, 1978). In such models, a learned feedback map $F^{(2)} \rightarrow F^{(1)}$ is also predicted to exist. No data are yet available on which to base the judgment of whether such a dynamic LTM buffer also exists between the RM_V and the RM_R networks.

11.4. Light Intensity, Motion, Attentional, and Multimodal Interactions within the Parietal Cortex

The previous two sections have shown how the VCN can feed into the RCN. We begin this section by noting that the RCN can also feed into the VCN, as in Figure 11.1. Indeed, when the target position map (TPM) of the VCN is trying to gain control over the RM_R, it does so by correlating the *same* retinotopic position within the RM_V and the RM_R. In other words, from an early stage of development, the TPM, no less than the RM_R, is also sensitive to light-induced intensity and motion changes on the retina. The RM that feeds into the TPM, which is denoted by RM_T in Figure 11.1, registers these light-sensitive activations and passes them on to the TPM. The high correlation between activations of the *same* retinotopic positions within the RM_T and the RM_R leads to the growth of large LTM traces within the $RM_V \rightarrow RM_R$ associative transform.

The TPM never loses its sensitivity to light-induced changes on the retina. As development progresses, however, other modulatory influences

can also begin to play upon the RM_T and the TPM. These include multimodal interactions, notably auditory inputs (Section 1.15), and attentionally modulated signals. In Section 11.2, for example, we discussed how a TPM-activated population v_{ij} can win the STM competition within the RM_R over a population v_i, $i \neq i_j$, that is directly activated by a light-sensitive input I_i. This can only occur if multimodal or attentional inputs can bias the TPM to activate a target position that does not correspond to the retinotopic position v_i. Such mismatches must not be allowed to frequently occur until after the $RM_V \rightarrow RM_R$ map is learned, because they would undermine the transform learning process. Fortunately, such mismatches *cannot* occur until after the $RM_V \rightarrow RM_R$ transform is learned, because this transform must be learned before intentional movements can generate a mismatch at the RM_R via the $RM_V \rightarrow RM_R$ transform. Infinite regress is thus once again averted.

In this section, we describe some of the multimodal and attentional mechanisms that can influence which TPM population will be chosen to activate the $TPM \rightarrow HMI \rightarrow RM_V \rightarrow RM_R$ pathway. Concerning the important problem of multimodal interactions, we will make only some brief observations.

A light-activated invariant TPM is computed in head coordinates. So too is a TPM based on auditory signals. Once both types of TPMs have been formed, strong correlations between sounds and lights can be the basis for learning an associative transform from sound-activated target positions to light-activated target positions, and possibly an inverse transform from vision to audition. This transform can be learned using the same associative laws that enabled the $RM_V \rightarrow RM_R$ transform to be learned.

After this transform is learned, the target position of a sound can compete for STM activation with the discordant target position of a light at all TPMs where both types of signals are registered (Figure 11.2). The target position of a sound can also amplify the activation of the target position of a light into which it is associatively mapped. Meredith and Stein (1983) have reported both competitive and cooperative interactions due to light and sound stimuli within the deeper laminae of the superior colliculi of cats and hamsters. Within our theory, such an interaction manifests itself via the circuit

$$(\text{auditory } TPM) \rightarrow (\text{visual } TPM) \rightarrow HMI \rightarrow RM_V \rightarrow RM_R \quad (11.10)$$

in Figures 11.1 and 11.2.

As in the associative transform described in Section 11.3, an intermodal associative transform tends to stabilize its prior learning via a dynamically maintained critical period (Knudsen, 1983, 1984). Although this discussion does not address the complex issue of how binaural signals are used to localize an object in space, any more than it describes how binocular disparity cues are used for visual localization, the following conclusion is robust: Whatever be the fine structure of these TPMs, if they are both computed in head coordinates, then multimodal associative

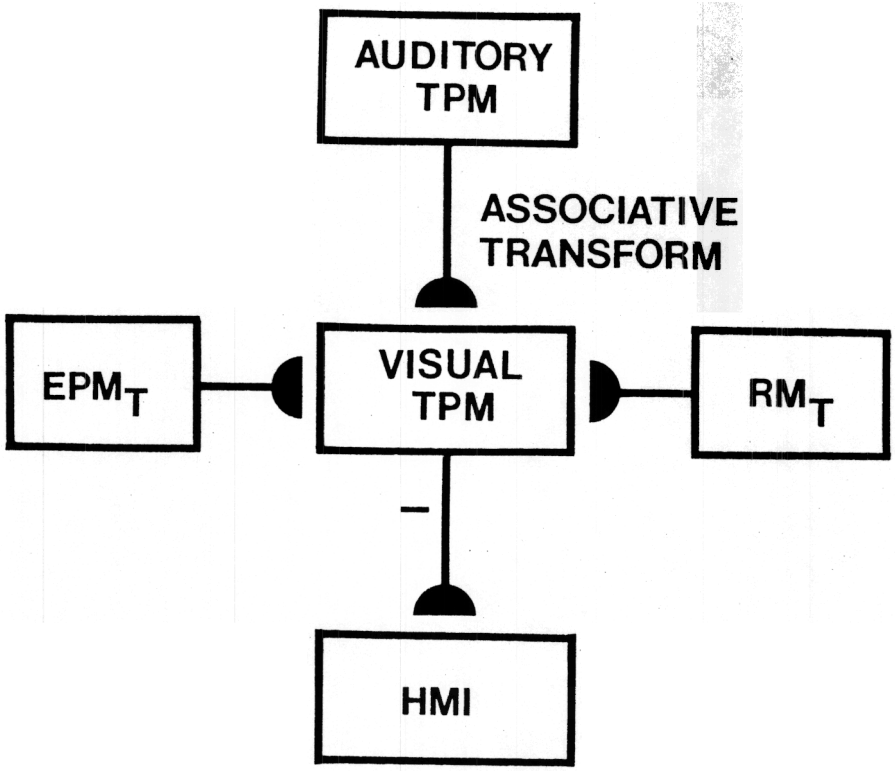

Figure 11.2. An auditory target position map (TPM) can be associatively mapped upon a visual TPM because the two maps are dimensionally consistent.

transformations can form with a precision that covaries with the degree to which their sampling and sampled map positions are correlated.

11.5. Nonspecific and Specific Attentional Mechanisms

Two types of attentional mechanisms are suggested to occur in the neocortical regions that control saccadic movements. Both types of mechanisms were derived from postulates concerning classical conditioning and instrumental conditioning (Grossberg, 1973, 1975; Grossberg and Levine, 1975) before relevant cortical data began to be reported (Motter and Mountcastle, 1981; Mountcastle, Anderson, and Motter, 1981; Wurtz, Goldberg, and Robinson, 1982). This is thus another area where theory and data seem to be rapidly converging.

The most elementary attentional mechanism is a nonspecific alteration of the gain, or sensitivity, across a neural network. This operation leads to either amplification or attenuation of all responses within the network. Such a nonspecific modulation of sensitivity is controlled by a parameter called the *quenching threshold* (QT) of the network (Section 2.6). Any of several network parameters can, in principle, cause state-dependent changes of the QT. A large QT imposes a high criterion for cell activation. A very large QT can totally desensitize the network to all inputs. A low QT enables even small inputs to activate their cells. A typical mechanism for controlling the QT is a nonspecific signal that shunts the interneuronal pathways within the network.

A nonspecific type of attentional reaction has been reported in the posterior parietal cortex (Mountcastle, Anderson, and Motter, 1981). If, for example, the QT of the auditory TPM became large and/or the QT of the visual RM_T became small during attention to visual cues, then auditory inputs would have great difficulty competing to activate the HMI. By contrast, if the QT of the visually-activated RM_T became large and the QT of the auditory TPM became small during attention to auditory cues, then auditory inputs could more easily activate the HMI.

A more specific type of attentional modulation is also needed to effectively control classical conditioning and instrumental conditioning. These specific attentional signals are learned incentive motivational signals from midbrain reward and punishment centers (Section 2.6). The conditioning model developed in Grossberg (1975, 1982b, 1982c, 1984) shows how the sensory representations of motivationally charged cues can be differentially enhanced by incentive motivational signals before the locations of the enhanced representations are encoded in a TPM. Analogous cue-specific enhancements have been reported in the posterior parietal cortices of monkeys who were trained to perform eye movement tasks to receive juice as a reward (Wurtz, Goldberg, and Robinson, 1982).

If a target position receives a large incentive motivational signal, then it can more effectively compete with factors like retinal light intensity or retinal movement for STM storage within a cortical TPM. If a target position receives a multimodal combination of auditory and visual signals, then it can more effectively compete with a source of only visual signals.

In addition, nonspecific gain change may alter the salience of different modalities through time. At any given moment, a complex interplay of retinal light intensity signals, retinal motion signals, specific intermodal signals, specific motivational signals, and nonspecific modality-selective attentional signals will help to determine which activated target positions within the TPM will win the competition to activate the HMI. In addition to these input factors, the learned factors which group familiar sensory patterns into object representations will also influence the intensity with which these representations activate their corresponding target positions (Carpenter and Grossberg, 1985b; Grossberg, 1980, 1984). As in life, the more interesting and complex are the objects being viewed, the more complex will be the activity patterns across the TPM that enable us to visually explore them.

11.6. Multiple Retinotopic Maps

We can now begin to gather together constraints derived in earlier chapters to refine the macrocircuit diagram of Figure 11.1. We will do this in stages using figures that omit processes which are not immediately being discussed. Where important circuit variations are consistent with functional requirements, we will note them in order to stimulate experiments capable of deciding between them.

Figure 11.3 depicts a refinement of Figure 11.1. Figure 11.3 contains two macrostages that were not included in Figure 11.1: a pair of retinotopic maps, RM_{R1} and RM_{R2}, instead of a single retinotopic map RM; and a pair of target position maps, TPM_1 and TPM_2, instead of a single target position map TPM. The retinotopic map RM_{R2} plays the role of RM_R and the target position map TPM_1 plays the role of TPM. First, we describe the role of the new retinotopic map RM_{R1}.

Figure 2.4 summarizes the functional role of the RM_{R1}. In Figure 2.4, two successive network stages are used to process light-activated signals from the retina. The first stage, the RM_{R1}, chooses the retinotopic position that receives the most favorable combination of retinal intensity and motion factors at any time. As in Figure 2.4, the RM_{R1} gives rise to second light error signals (ES) which act at the adaptive gain (AG) stage, or cerebellum. In both Figures 2.4 and 11.3, the RM_{R1} activates the next retinotopic map, the RM_{R2}, at which the light that is chosen by the RM_{R1} is stored in STM. The RM_{R2} gives rise to the conditioned pathways along which sampling signals (SS) learn adaptive gains within the AG stage in response to second light error signals. The chosen light must thus be stored in STM until after the saccade is over, so that its conditioned pathways can sample the second light error signals.

As in Figure 11.1, the RM_V in Figure 11.3 samples the RM_{R2}, so that it can learn to transmit an attentionally modulated choice of retinotopic position to the RM_{R2}. Due to $RM_V \rightarrow RM_{R2}$ signals, the choice which is stored within the RM_{R2} might differ from the choice which was made by the RM_{R1}. The RM_{R2} chooses the largest signal which it receives before storing it in STM, as in equation (11.3). Attentional and intentional

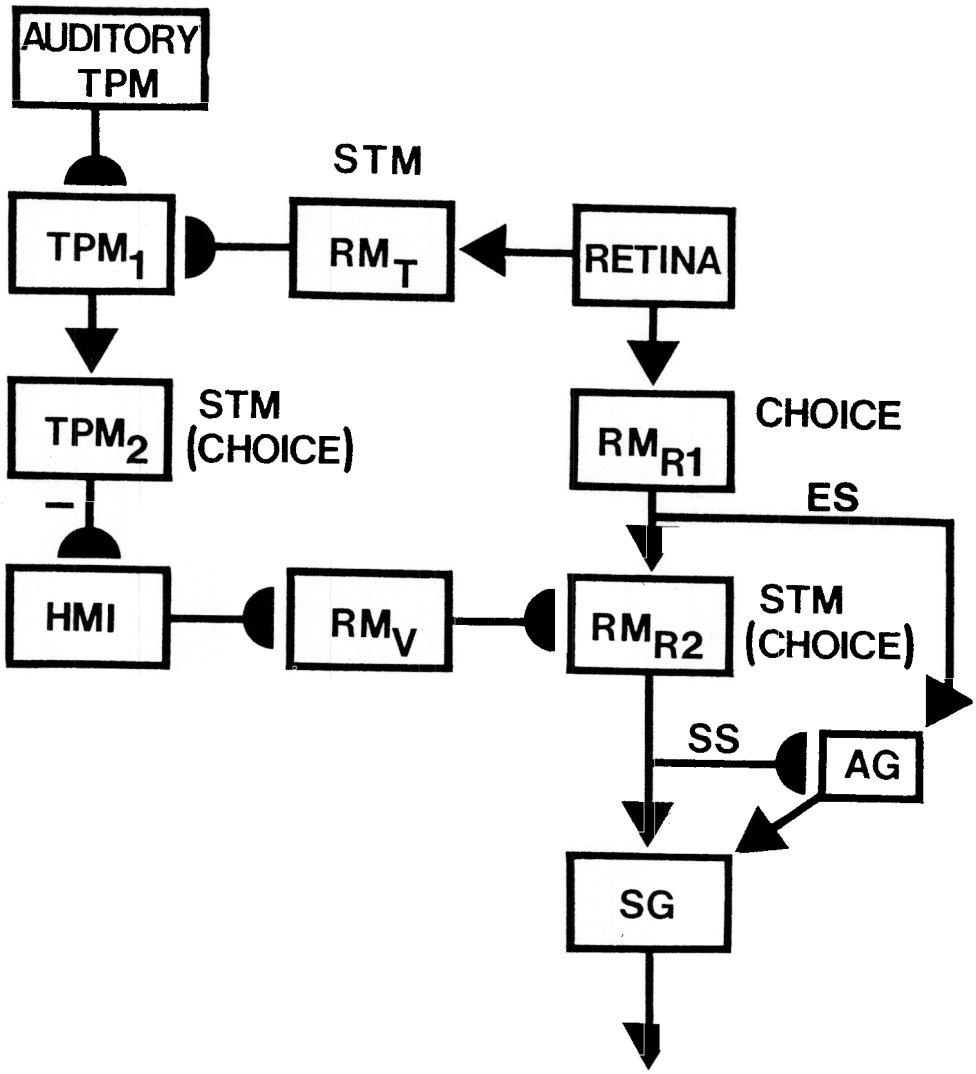

Figure 11.3. Refinement of the circuit in Figure 11.1: A pair of retinotopic maps RM_{R1} and RM_{R2} replace RM_R and a pair of target position maps TPM_1 and TPM_2 replace TPM. See text for details. Abbreviations: ES = error signal.

factors may hereby override visually reactive factors within the RM_{R2}, as in Section 11.3. However, sufficiently vigorous reaction to a moving stimulus may override an insufficiently amplified attentional focus. Despite this fact, visual factors remain the exclusive source of second light error signals from the RM_{R1} to the AG stage.

Figure 11.4 describes a variation on Figure 11.3 that is consistent with functional requirements. In Figure 11.4, a multisynaptic pathway $RM_{R1} \rightarrow RM_T \rightarrow RM_{R2}$ replaces the direct $RM_{R1} \rightarrow RM_{R2}$ pathway of Figure 11.3. This multisynaptic pathway plays the following functional role. In both Figures 11.3 and 11.4, the RM_T is a retinotopic map at which attentional modulation of visual signals can occur. In Figure 11.4, the $RM_{R1} \rightarrow RM_T$ pathway can selectively enhance the retinotopic position that was chosen within the RM_{R1}. As a result, the favored RM_{R1} population is also "attentionally" amplified within the RM_T. The TPM_1 thus receives from the RM_T a spatial pattern of activated retinotopic positions such that the favored RM_{R1} position can effectively compete with positions that are amplified by other attentional factors.

In Figure 11.4, the RM_{R2} receives the entire spatial pattern of positions from the RM_T, not just the position chosen by the RM_{R1}, as in Figure 11.3. Thus if a retinotopic position other than the one chosen by the RM_{R1} is favored by the RM_T, then, in the absence of intermodal influences, both the RM_V and the RM_{R2} will tend to choose the same retinotopic position even before the $RM_V \rightarrow RM_{R2}$ transform is learned. By contrast, in Figure 11.3, the RM_{R2} and the RM_V can more easily choose different positions whenever attentional factors within RM_T favor a position other than the one chosen by the RM_{R1}.

In summary, the interactions in Figure 11.4 strengthen the tendency for both the RM_V and the RM_{R2} to choose the same positions whenever only light-activated cues enter the movement decision, whether or not the attentionally favored positions are the retinally favored ones. This property of Figure 11.4 tends to further stabilize learning of the $RM_V \rightarrow RM_{R2}$ associative transform.

11.7. Interactions between Superior Colliculus, Visual Cortex, and Parietal Cortex

We interpret the stages RM_{R1}, RM_{R2}, and RM_T in Figures 11.3 and 11.4 in terms of interactions between superior colliculus (SC), visual cortex, and parietal cortex. These retinotopic maps interact with other network stages that possess unambiguous anatomical interpretations. Such linkages constrain the anatomical interpretation of the retinotopic maps themselves. For example, the TPM_1 is assumed to occur in the parietal cortex. Both the RM_{R1} and the RM_{R2} project to the AG stage, which is assumed to occur in the cerebellum. The RM_{R1} gives rise to visually activated error signals, which are assumed to reach the cerebellum via climbing fibers. The RM_{R2} gives rise to visually activated sampling signals, which are assumed to reach the cerebellum via mossy fibers. The

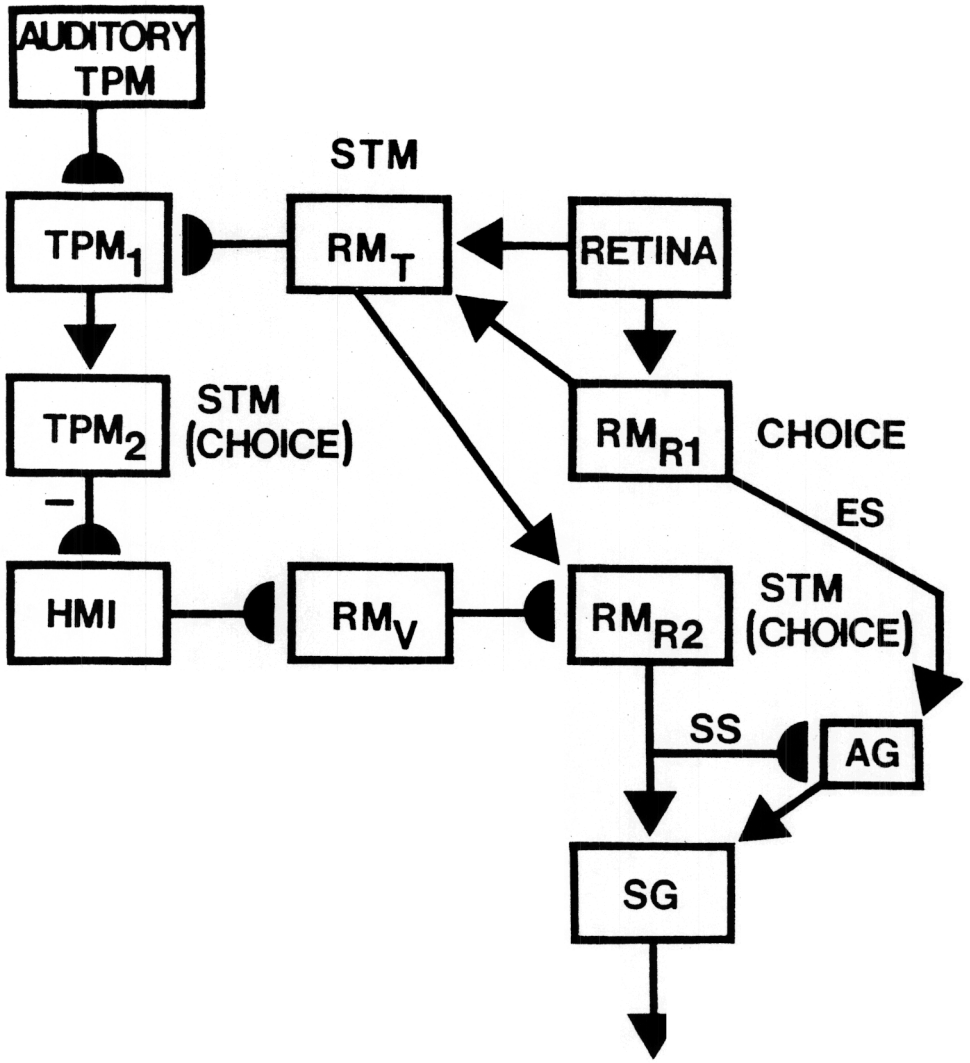

Figure 11.4. A variation on the circuit in Figure 11.3: The direct $RM_{R1} \rightarrow RM_{R2}$ projection of Figure 11.3 is replaced by an indirect projection in which the RM_{R1} projects to the RM_T which, in turn, projects to the RM_{R2}. The text describes the functional benefits of using such an indirect pathway.

RM_{R1} activates the RM_{R2}, whether monosynaptically or multisynaptically, but not conversely. The RM_{R2}, but not the RM_{R1}, stores a choice in STM from a time before a saccade begins until a time after a saccade ends. The RM_{R2}, but not the RM_{R1} or RM_T, projects to the saccade generator (SG). The RM_{R2} also receives inputs that are sensitive to corollary discharges. The RM_T, but not the RM_{R1}, processes retinotopic position signals that are attentionally modulated. Last but not least, each of the regions RM_{R1}, RM_{R2}, and RM_T is a *retinotopic* map; hence, it possesses enough internal topography to differentially process many distinct map positions.

These properties suggest that the RM_T is either part of the parietal cortex, or is in a part of the prestriate cortex that directly projects to parietal cortex. Because all of the maps are retinotopic maps, the deep layers of the SC, which are coded in motor coordinates, seem to be ruled out as a possible interpretation. Thus any SC role in computing RM_{R1}, RM_T, or RM_{R2} would seem to be restricted to the superficial layers of the SC (Huerta and Harting, 1984; Wurtz and Albano, 1980). The deeper SC layers, including the quasi-visual cells (Section 4.2), are interpreted to occur somewhere between the RM_{R2} and the SG in Figures 11.3 and 11.4. A likely location of the RM_{R1} is within the superficial SC layers. The following facts are consistent with this interpretation. The superficial SC layers receive a topographically organized retinal input. The superficial SC layers also project to several areas of visual cortex, as well as to the parietal cortex (compare $RM_{R1} \rightarrow RM_T$ in Figure 11.4). If this interpretation of the RM_{R1} is correct, then a positional choice based upon retinal light intensity and motion factors is made somewhere along the pathways between the superficial SC, the visual cortex, and the parietal cortex.

Given that the RM_T is interpreted to be in parietal cortex or in an adjacent prestriate cortical region, the reader may wonder why the retina is shown projecting directly to the RM_T in Figure 11.4. The projection from the retina to the RM_T is recognized to be multisynaptic *in vivo*, but we are not herein analysing the many processing stages used for visual form, depth, and color perception (Cohen and Grossberg, 1984a, 1984b; Grossberg and Mingolla, 1985a, 1985b) or for visual object recognition (Carpenter and Grossberg, 1985b; Grossberg and Stone, 1985a). The circuit in Figure 11.4 is a lumped representation of the visual factors that now play a role in our eye movement theory. The minimal network can and will be expanded as recent models of visual cortex are used to generate locations of movement commands in the theory.

It remains to anatomically interpret the RM_{R2} stage. Although this stage is the most functionally distinctive one—receiving as it does retinotopically recoded signals due to corollary discharges, projecting via mossy fibers to the cerebellum, and storing chosen lights in STM throughout a saccade—the available neural data do not seem to force an unambiguous anatomical interpretation. We consider the anatomical localization of the RM_{R2} to be a problem of major importance for neurobiologists interested in the saccadic system. Just as several anatomical stages intervene between the retina and the RM_T *in vivo*, the single macrostage RM_{R2} may

consist of several closely interacting nuclei of cells *in vivo*.

11.8. Multiple Target Position Maps within Parietal Cortex and Frontal Eye Fields

The single TPM in Figure 11.1 is expanded into two successive TPMs, namely TPM_1 and TPM_2, in Figure 11.4. These successive TPMs play a role that is similar in important respects to the successive TPMs that were used to regulate sequences of predictive saccades in Figure 10.3. Map TPM_1 in both Figures 10.3 and 11.4 is an invariant TPM which can simultaneously encode several active target positions. In Figure 11.4, the activity level at each target position is the resultant of a combination of light-activated, multimodal, incentive motivational, and attentional gain control signals. In Figure 10.3, such factors interact with predictive mechanisms to read-out sequences of learned movement synergies. In both Figure 10.3 and Figure 11.4, the map TPM_2 chooses the target position from TPM_1 which achieves the largest activity based upon all of these influences. The chosen target position is stored in STM until after the saccade is over, so that it can then sample the corollary discharge at its HMI. In Figure 11.4, we interpret the TPM_1 on-center cells to be formal analogs of the light-sensitive attentionally modulated cells of the posterior parietal cortex (Yin and Mountcastle, 1977), and the on-center cells in TPM_2 to be formal analogs of the saccade cells in the posterior parietal cortex (Lynch, 1980; Motter and Mountcastle, 1981).

The TPMs of the parietal cortex (Figure 11.4) are assumed, however, to differ from the TPMs that control predictive saccades (Figure 10.3) in a crucial way: The parietal TPMs are incapable of storing temporal order information in STM. As we saw in Chapter 10, the ability to effectively store temporal order information in STM requires specialized auxiliary mechanisms that enable a match-mediated nonspecific rehearsal wave to sequentially read-out and reset the STM activity pattern within the TPM_1. We assume that this type of reset mechanism is found in the frontal eye fields (FEF) and is a version of a frontal cortex architecture that is specialized to organize "the temporal ordering of recent events" (Milner and Petrides, 1984, p.403). In addition to the machinery that is needed to regulate storage and sequential read-out of STM in a predictive system, we suggest that the predictive STM patterns may be encoded in long term memory (LTM) as motor plans, or chunks. These motor plans may be read-into STM from LTM in the absence of sensory cues. Both the regulatory STM machinery and the LTM chunking and read-out machinery are assumed to be part of the predictive design. Applications of this type of temporal ordering machinery to the processing of language and other types of motor control sequences are described in Grossberg (1978a, 1978b, 1985c) and Grossberg and Stone (1985a, 1985b).

11.9. Learning Multiply-Activated Target Position Maps

The previous sections have suggested that one TPM system is needed at which intermodal, motivational, and light-related factors can compe-

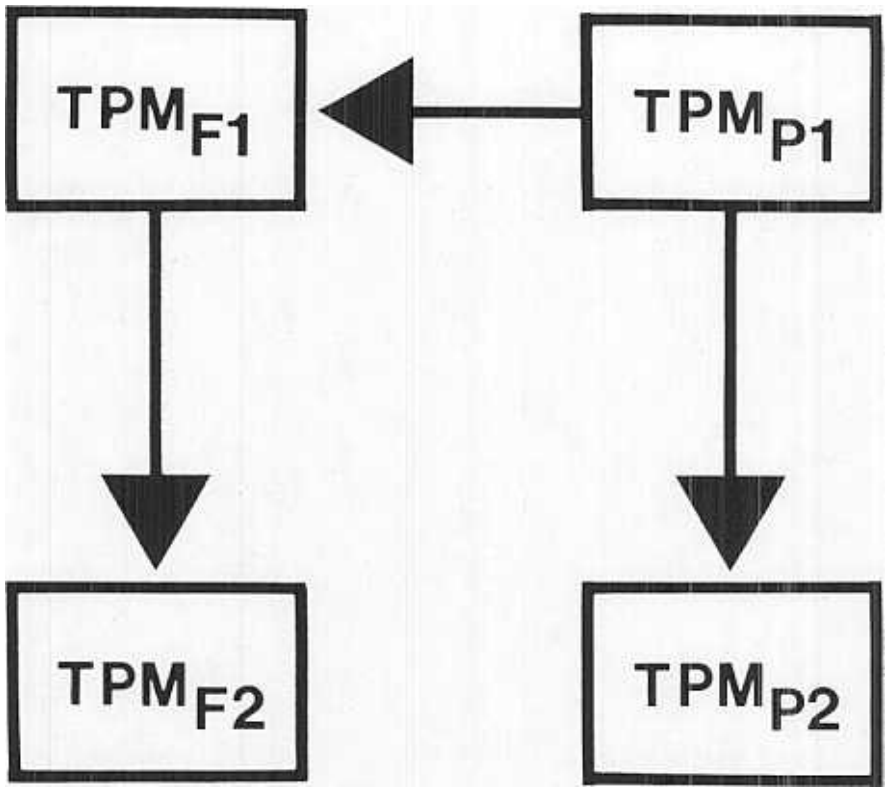

Figure 11.5. A possible circuit for connecting an attentionally modulated target position map (TPM_{P1}) with a predictive target position map (TPM_{F1}): Using this arrangement, the TPM_{F1} can automatically inherit both the target position coordinates and the attentional modulations of the TPM_{P1}. Both TPMs could, in principle, be implicitly built up from their own separate sources of retinotopic and eye position signals.

titively struggle to define an attentional focus capable of directing an eye movement. A second TPM system is needed at which predictive sequences of eye movements can be performed and can be learned, perhaps as part of more general motor plans. Neural data suggest, moreover, that the former TPM system is centered in the parietal cortex, whereas the latter TPM system is centered in the frontal eye fields (FEF). Consideration of how these two types of TPM systems interact leads to further constraints on system design. In order to unambiguously discuss these two TPM systems, we denote the parietal TPM stages by TPM_{P1} and TPM_{P2} and the frontal TPM stages by TPM_{F1} and TPM_{F2}.

A major issue concerns how these TPMs form. In Chapter 10, we considered several ways in which an invariant TPM can arise. In the subsequent discussion, we will fix ideas by assuming that the TPM_{P1} forms through a process of self-organization, as in Section 10.6. *Whatever* mechanism gives rise to the TPM_{P1}, once the TPM_{P1} develops, it can *automatically* generate the TPM_{F1}. If a topographic map exists from the TPM_{P1} to the TPM_{F1}, then the cells in the TPM_{F1} inherit the invariant target positions that are encoded by the corresponding cells in TPM_{P1} (Figure 11.5). Moreover, all of the attentional factors that weight the importance of target positions in TPM_{P1} can also weight the target positions in TPM_{F1}. Thus the specialized intermodal and motivational mechanisms that feed directly into TPM_{P1} can automatically influence the TPM_{F1} without necessitating the replication of these mechanisms at the TPM_{F1}. Although we consider this an attractive possibility, the TPM_{F1} could also implicitly form from combinations of RM_T and EPM signals, as in Chapter 10. Such a TPM would not, however, benefit from parietal forms of attentional modulation.

It is also conceivable that retinotopic inputs to the FEF are directly recoded into difference vectors. Such a direct recoding would, however, have to surmount the functional problems described in Section 9.10. In the present exposition, we explicitly consider the case in which the FEF possesses its own TPM_{F1}. With this exposition in hand, the above variations can also be understood.

From the TPM_{P1}, the target position with the largest weight, or activity, can generate a movement command via the TPM_{P2}. From the TPM_{F1}, the whole pattern of active target positions can be sequentially read-out in order of decreasing activity, and can be gradually encoded on successive performance trials into long term memory as part of a unitized motor plan for future predictive performances. Once it is recognized that these two types of TPM systems exist, several design issues come into focus concerning how it is decided which of these TPM systems will learn or perform eye movements at any given time.

Further consideration needs to be given to how the TPM_{P1} is formed. In Chapter 10, we considered a self-organization model in which a single retinal position (RM) and a single initial eye position (EPM_1) can sample a single target eye position (EPM_2) on every learning trial. Eventually, pairs of signals from the RM and the EPM_1 to the EPM_2 learn to generate

individual target position activations at the EPM_2, thereby converting it into a TPM. In the adult TPM_{P1}, by contrast, we assume several target positions can simultaneously be active at any time. In order to ensure the applicability of the model in Chapter 10, one could alternatively suppose that just one target position can win the attentional competition at any time. This assumption is equivalent to lumping TPM_{P1} and TPM_{P2} into a single stage. This may, in fact, be true in sufficiently simple organisms.

This assumption is inadequate for our purposes, however. Then the TPM_{P1} could not map to the TPM_{F1}, because the TPM_{F1}, as a predictive mechanism, must be able to simultaneously store several target positions in STM. We would therefore have to ask how the TPM_{F1} could self-organize in response to multiple retinotopic sampling inputs. We would thus have to solve the same functional problem for the TPM_{F1} if we try to escape this problem at the TPM_{P1}. Moreover, if the TPM_{P1} can store only one target position at a time, then we would have lost the desirable property that attentional modulation within the TPM_{P1} alters the temporal ordering of predictive movements from the TPM_{F1}. Consequently, we now face the issue of how the TPM_{P1} can self-organize in the presence of multiple retinotopic sampling signals. The manner in which this issue is forced upon us illustrates how available design constraints can severely limit the number of hypotheses that one can reasonably entertain.

Two types of solutions are known. They are, moreover, experimentally testable. In the first solution, we suggest that the problem causes no new difficulties for the following reason. On any given saccadic trial, a single initial eye position (EPM_1) can combine with several retinotopic positions (RM) to sample a single target eye position (EPM_2). Only one combination of retinotopic position with initial eye position is the correct one, and this combination will consistently sample its correct target eye position on all learning trials, thereby generating a cumulative effect on the corresponding LTM traces. On the average, the other retinotopic positions will be randomly chosen with respect to the correct retinotopic position across sampling trials. Their LTM traces will tend to encode as many erroneous LTM increments as decrements, thereby tending to cancel their net LTM change due to erroneous correlations across learning trials. If only this property of random error distribution prevents map relearning errors from accumulating, then an electrode input that persistently activates the same false target eye position at EPM_2 after each saccade should cause a massive learned distortion in the adult TPM_{P1}.

In the second solution, TPM_{P1} map learning occurs only during a developmentally plastic critical period. This critical period terminates before such diverse attentional influences as multimodal and motivational factors can begin to override retinal light intensity and motion factors. During such a critical period, the RM_{R1} in Figure 11.4 can strongly bias the RM_T to differentially amplify its preferred light position. The input from the RM_T to the TPM_{P1} then approximates the choice of a single retinotopic position. After the TPM_{P1} is formed in this way, the critical period ends, and future multiple activations of the RM_T cannot cause map errors to be learned. If this model is valid, then approximate choices should

occur within the RM_T at an early developmental stage, and electrode activations of an adult TPM_{P1} should not recode its target positions.

11.10. Multiple Parietal and Frontal Eye Field Vector Systems

Given that two TPM systems exist, it is necessary to consider whether they compute neural vectors using their own head-muscle interfaces (HMI) or whether they share a single HMI. The results in Chapter 4 showed how a TPM can convert its target positions into vectors at an HMI via a learning process. These results imply that only a single target position at a time can sample an HMI. If two or more target positions simultaneously sampled an HMI, then each target position would learn only a fractional part of the target eye position, by equations (4.1) and (4.2). Thus either each TPM system has its own HMI (Figure 11.6a), or each TPM system activates the same HMI at different times (Figure 11.6b), or only the TPM_{F2} projects to an HMI, which is reached by the TPM_{P1} via the TPM_{F1} (Figure 11.6c). In the last case, the TPM_{P1} would have to equal the TPM_{P2} in order for the TPM_{P2} to be able to activate a single (nonpredictive) vector. Then we would be reduced once again to considering a network in which the TPM_{F1} would have to self-organize in response to multiple sampling inputs from a retinotopic map (RM) and a source of initial eye movements (EPM_1). The parietal TPM_{P1} system could not then attentionally weight the order of predictive FEF saccades. We therefore reject this option and consider the two cases in which each TPM system can directly access an HMI.

Experimental data which support this conclusion were reported by Schiller and Sandell (1983), who showed that vector compensation using a Mays and Sparks (1980) light-electrode arrangement (Section 4.2) can occur after extirpation of the FEF, or of the SC, while stimulating the other. In our theory, the SC-controlled vector compensation is interpreted to include a TPM system, namely TPM_{P1} in the parietal cortex.

It remains to decide whether the TPM_P and TPM_F systems share a single HMI or whether each possesses its own HMI. We favor the latter alternative on general grounds because it provides the predictive movement system with a greater independence from interference by irrelevant sensory cues. We also favor this latter alternative for specific reasons. For example, if only one HMI existed, then the TPM_P and TPM_F systems would need to be able to competitively inhibit each other's output signals before these signals reached the shared HMI, since only one TPM can be allowed to sample the HMI at any time. Given that the TPM_{P2} and the TPM_{F2} are most likely to include large neocortical or subcortical regions, this would necessitate the ability of two large and topographically organized neural regions to nonspecifically shut each other off. Whereas this is not formally impossible, it is much easier for the two systems to compete nonspecifically at a smaller neural region, such as the SC or the SG. It is not inconceivable that, as a predictive movement capability evolved through phylogeny, a single HMI-based system evolved into a multiple HMI-based system. By considering the more difficult multiple HMI case

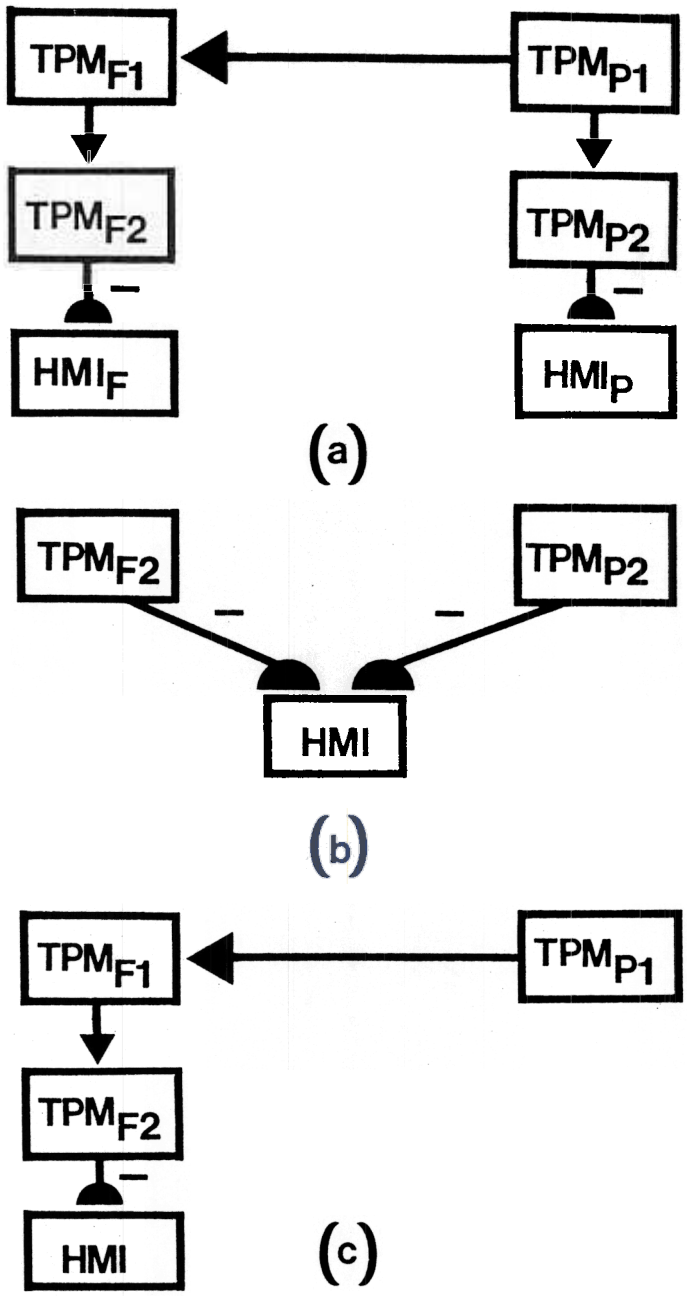

Figure 11.6. Some possible circuits for attaching a head-muscle interface (HMI) to a target position map (TPM): The text discusses some advantages in using separate HMIs in the attentional (HMI_P) and predictive (HMI_F) systems.

herein, many of the properties which arise using a simpler single HMI case are also implicitly analysed.

A multiple HMI-based system is also supported by the data of Mays and Sparks (1980) and of Schiller and Sandell (1983). The data of Mays and Sparks (1980) show that HMI vectors get read out through the quasi-visual (QV) cells of the SC. Suppose that the RM_R stage in Figure 11.4 feeds into the QV cells, or includes the QV cells. Then extirpation of the SC would prevent the HMI in Figure 11.4 from reading-out its movement commands. If the TPM_F system fed its signals into this same HMI, then SC lesions would prevent the read-out of FEF saccades. The data of Schiller and Sandell (1983) suggest that this is false, at least in monkeys. Bruce and Goldberg (1984) have provided additional evidence that the FEF can generate vector commands. They showed that electrical stimulation of certain FEF cells could generate saccades of a fixed size and direction, independent of the initial eye position.

11.11. Learning Neural Vectors and Adaptive Gains in a Predictive Movement System

Henceforth we assume that two HMI systems exist, corresponding to the two TPM systems TPM_P and TPM_F. We denote these HMIs by HMI_P and HMI_F, respectively, as in Figure 11.6a. We now consider what new learning issues are raised by the existence of a separate HMI_F system. The main issues concern how target positions are converted into muscle coordinates by the learned $TPM_{F2} \rightarrow HMI_F$ transform, and how vector commands from the HMI_F gain control of adaptive gains at the AG stage that can generate accurate foveations.

The problem in its most severe form arises if the TPM_F system, unlike the TPM_P system of Figure 11.4, does not sample the RM_{R2}. Then the TPM_F system cannot use the adaptive gains which the RM_{R2} controls via the AG stage. It is theoretically possible for the TPM_F system to sample the RM_{R2}, but it would have to do so in a way that does not also inhibit sampling by the TPM_P system. In particular, the TPM_F system would have to sample the RM_{R2} at a stage subsequent to the stage at which the TPM_P system samples the RM_{R2}, so that TPM_F signals would not competitively preempt sampling of the RM_{R2} by the TPM_P system. Such a convergence of visually reactive, TPM_P, and TPM_F commands at the RM_{R2} could, however, cause a high probability of erroneous command choice during sequences of predictive saccades. Erroneous command choices during predictive saccades can be minimized if the predictive movement system can nonspecifically inhibit all visually reactive and TPM_P signals. This cannot happen if the TPM_F system can sample the RM_{R2}, because the sampled pathways must be activated by visually reactive signals. In order for the predictive movement system to nonspecifically inhibit other sources of movement commands, it must feed into the SG at a stage subsequent to the RM_{R2}, so that it can nonspecifically inhibit $RM_{R2} \rightarrow SG$ movement signals without inhibiting its own movement signals.

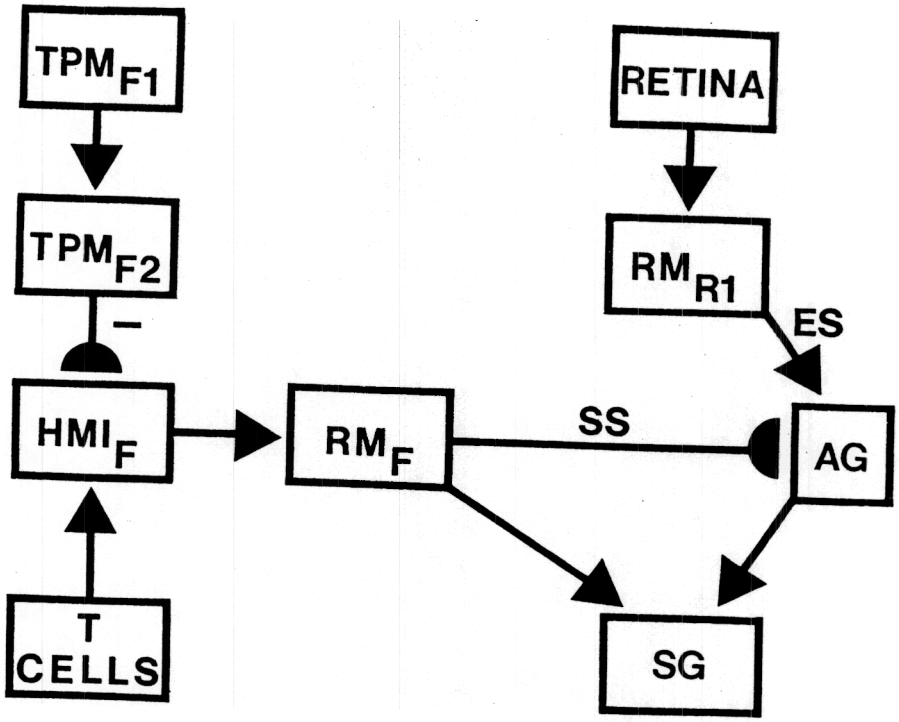

Figure 11.7. A circuit in which the predictive system controls its own unconditioned and conditioned movement pathways. Abbreviations: ES = error signal, SS = sampling signal.

We therefore consider a model in which the TPM_F system controls its own unconditioned $HMI_F \rightarrow RM_F \rightarrow SG$ and conditioned $HMI_F \rightarrow RM_F \rightarrow AG \rightarrow SG$ movement pathways to the SG, as in Figure 11.7. This model faces and meets the full force of the infinite regress problem. If the adaptive gains controlled by the RM_{R2} are not used to achieve accurate predictive saccades, then how does the predictive command network (PCN) learn its adaptive gains? No learning difficulties arise if the vectors used by the PCN via the HMI_F are accurately calibrated. Any adequate solution of how sequences of predictive saccades are controlled must include a solution of this problem. If HMI_F vectors are accurately calibrated, then predictive commands from the PCN can cause unconditioned movements via the pathway $HMI_F \rightarrow RM_F \rightarrow SG$ in Figure 11.7. Second light error signals, registered via the Retina $\rightarrow RM_{R1} \rightarrow AG$ pathway, can then alter the adaptive gains within the $RM_F \rightarrow AG \rightarrow SG$ pathway, using the mechanisms of Chapter 3. Thus the retinotopic command network (RCN) helps to calibrate the PCN using its second light error signal subsystem. We now indicate how the RCN also helps to calibrate the vectors computed within the HMI_F of the PCN. Thus the RCN acts in two ways to overcome the infinite regress problem that might otherwise be faced by the PCN.

The main requirements are (1) the $TPM_{F2} \rightarrow HMI_F$ transform is correctly learned before the HMI_F can begin to generate predictive commands, and (2) the $TPM_F \rightarrow HMI_F$ transform is not unlearned due to performance errors caused by HMI_F-generated saccades before the $HMI_F \rightarrow RM_F \rightarrow AG \rightarrow SG$ pathway can learn the correct adaptive gains. Both requirements are achieved by making an assumption concerning the type of saccade-driven gate (Section 9.7) that turns on learning within the HMI_F.

The best performance is achieved if the SG is broken up into two subsystems: One subsystem SG_R receives unconditioned and conditioned movement signals from the RM_{R2}, and thus also from the HMI_P (Figure 11.8). The other subsystem SG_F receives unconditioned and conditioned movement signals from the HMI_F via the RM_F. Both subsystems are assumed to possess their own complement of burster cells and pauser cells that send convergent signals to tonic cells and motoneurons (Section 7.8). We assume that the learning gate within the HMI_F is turned on either by onset of pauser activity or by offset of burster activity within the SG_R. We also assume that the same class of SG_R cells turns on the learning gate within the HMI_P (Figure 11.8). Thus the learned transformation of target position into motor coordinates is assumed to be regulated by RM_{R2}-controlled gating signals in *both* HMIs. Since the RM_{R2} can learn to generate correct saccades using second light error signals (Chapter 3), it can generate correct target eye positions for learning at the HMIs.

In order for the TPM_{F2} to sample the RM_R-generated target positions at the HMI_F, the TPM_{F1} must be capable of activating the TPM_{F2} before RM_R-generated saccades occur. Thus the TPMs of the FEF must be sensitive to lights at a developmental stage prior to the emergence of predictive saccades. This property implies that factors other than volition

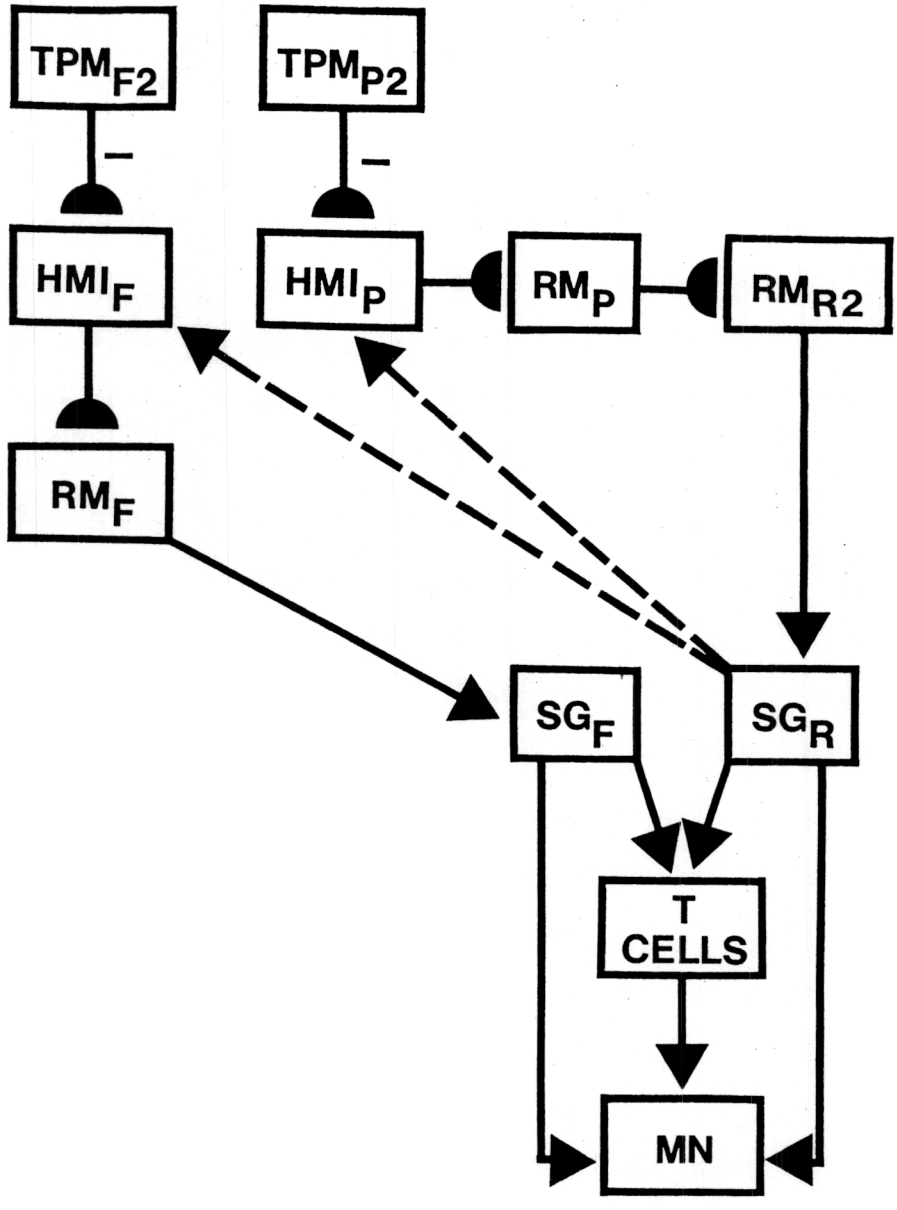

Figure 11.8. Regulation of target position learning within both the parietal head-muscle interface (HMI_P) and the frontal head-muscle interface (HMI_F) by the same source of gating signals. This source is the saccade generator (SG_R) that is activated by the retinotopic command network (RCN). This scheme enables vectors in the HMI_F to be correctly calibrated before the predictive command network (PCN) begins to compete with the RCN. Abbreviations: SG_F = frontally activated saccade generator. The gating pathways are marked with dashed lines.

can enable the TPM_{F1} to activate the TPM_{F2} (Section 9.4). In particular, we assume that specific output signals from the TPM_{P1} to the TPM_{F1} give rise to inhibitory interneurons capable of shutting off the tonic gating cells which inhibit TPM_{F1} output signals (Figure 11.9). A volitional signal can also shut off these cells, as in Section 9.4. We assume furthermore that such a volitional signal opens an output gate which enables the HMI_F to activate the SG_F (Figure 11.9).

11.12. Frontal Eye Field Control of Voluntary Saccadic Eye Movements and Posture: Cell Types

These assumptions imply that the FEF system can learn correct vectors in response to light-sensitive inputs, but can only generate saccades in response to an act of volition. These properties are consistent with the conclusion of Bruce and Goldberg (1984, p.439) that "frontal eye field cells discharge before purposive saccadic eye movements." Bruce and Goldberg (1984) described *visuomovement cells* which begin to discharge in response to a visual stimulus and continued to respond until after a saccade began, and *movement cells* which responded just before and during a saccade. These cells were studied during an experiment in which the monkey was trained not to respond to a peripheral target until the fixation light was turned off over a second later. The visuomovement cells continued to discharge throughout the presaccadic interval even if the target presentation was brief (100 msec.). Thus these visuomovement cells were capable of storing the target position in STM until after the saccade began.

These experimental properties are consistent with the assumptions that visuomovement cells occur in the TPM_{F1}, movement cells occur in the TPM_{F2}, and a volitional signal enables the stored command in the TPM_{F1} to be read into the TPM_{F2}, thereby generating a voluntary saccade.

The FEF also contains cells that discharge after saccadic eye movements (Bizzi, 1968; Bizzi and Schiller, 1970; Goldberg and Bushnell, 1981). These post-saccadic cells respond only after saccades of fixed dimensions. Bruce and Goldberg (1984, p.440) noted that "Because they respond after both spontaneous and visually elicited saccades, it seems likely that they are driven by an efference copy from the brainstem saccade generator." Our theory suggests that three types of post-saccadic cell may be found in the FEF, or in closely associated neural regions. One type may be part of the eye position map (EPM) within the tension equalization network (TEN) of Chapter 8. These EPM cells are activated by tonic cells of the SG after a saccade terminates. They are driven by an efference copy from the brainstem saccade generator, and respond only after saccades of fixed dimensions are over. Their functional role is to read out the learned postural gains that prevent postsaccadic drift.

A second type of post-saccadic cell may exist in the FEF if the HMI_F is housed within the FEF. Such cells would relay gating signals from the SG_R to the HMI_F (Figure 11.8). These cells would be active after both spontaneous and visually elicited saccades. They could, however, respond after saccades of variable size, since their function is to enable the TPM_{F2}

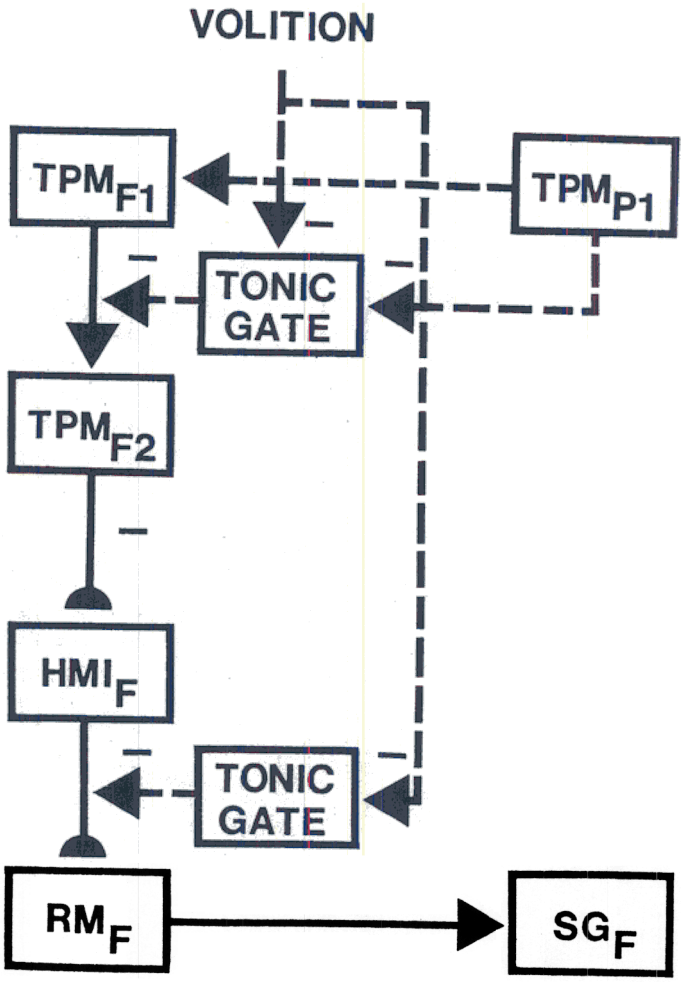

Figure 11.9. Gating signals which regulate the predictive, or frontal, command system: Volitional signals can enable read-out of the TPM_{F1} and of the HMI_F signals to occur. Attentionally modulated sensory signals, say from the TPM_{P1}, can also enable read-out of the TPM_{F1} to occur. In this way HMI_F vectors can be calibrated even if movement commands from the HMI_F are not released. The gating pathways are marked with dashed lines.

to learn whatever target eye position was attained by the previous saccade. It may be that the SG_R directly gates the HMI_F without the intervention of HMI_F interneurons. Then the activity of axons, rather than cell bodies, would be predicted to possess these functional properties.

The TPM_{F2} cells which we have identified with FEF movement cells (Bruce and Goldberg, 1984) may also occasionally respond like a type of post-saccadic cell. Suppose that the TPM_{F1} has been storing several target positions in STM while spontaneous or visually elicited saccades are taking place. The TPM_{F2} and the HMI_F can process such target positions because the TPM_{F2} must be able to learn its target eye positions at the HMI_F in response to visually activated movement commands (Section 11.11). Thus after a saccade takes place, a rehearsal wave may occur enabling the TPM_{F1} to read a new target position into the TPM_{F2}. Such a TPM_{F2} cell would occur after the saccade was over. Hence the cell would appear to be a "post-saccadic cell." This cell might not cause the next saccade unless the volitional gate is opened (Figure 11.9). Hence its role in controlling a future voluntary saccade could be mistaken for a reaction to a past visually reactive saccade. Such a cell would not encode an efference copy from the SG. Rather, it would encode the target position of a possible voluntary saccade. To distinguish this type of "post-saccadic cell" from an EPM cell which controls the read-out of postural gain, one could use the Bruce and Goldberg (1984) paradigm to discover whether the cell can also function as a movement cell in an experiment that requires voluntary saccades.

11.13. Coupled Vector and Adaptive Gain Learning

A possible variation of the circuit design described in Figures 11.8 and 11.9 is also worthy of experimental study. If the gating rules summarized in Figure 11.9 are imposed, then the PCN can achieve a significant competence even if the SG is not broken into two subsystems. The gating rules in Figure 11.8 enable the TPM_{F2} to calibrate accurate vectors at the HMI_F before the HMI_F can begin to generate predictive saccades. Thus, when HMI_F-generated saccades emerge, the vectors computed within the HMI_F are already accurately calibrated, whereas the adaptive gains of the conditioned pathway $HMI_F \rightarrow RM_F \rightarrow AG \rightarrow SG$ are not yet accurately calibrated. Consequently saccadic errors will occur. In Figure 11.8, we prevented these errors from causing the relearning of inaccurate vectors by using the SG_R to shut off vector learning in the HMI_F when predictive saccades are being made.

We now observe that new vector learning may be permitted to occur in the predictive mode. The learning of incorrect target positions within the HMI_F which is caused by foveation errors will tend to be compensated by learning within the AG stage *if* the HMI_F starts out with accurately calibrated vectors. For example, suppose that an undershoot error occurs. Then the vector which caused the saccade will learn a larger adaptive gain (Chapter 3), but the TPM_{F2} will learn to generate a smaller vector in response to the same retinal position on a later saccadic trial (Chapter

4). On successive undershoot trials, progressively smaller vectors and progressively larger gains will be learned hand-in-hand. Finally, a time is reached at which either the saccades cannot get shorter, or the larger adaptive gains by themselves will start to cause longer, and therefore more correct, saccades. At such a time, the adaptive gains are larger than they were at the outset, and are hence better able to prevent undershoot errors, but the vectors read-out by the TPM are too small. As large vectors begin to be learned, the adaptive gains will grow even larger to prevent further undershoot errors. Thereafter, relearning of correct vectors and of large enough adaptive gains to prevent undershoot errors go hand-in-hand until accurate saccades are generated.

This coupled vector-gain learning process is not yet fully understood. It may help to stabilize learning within the PCN even in organisms wherein HMI_F learning is gated by the SG, but not by a specialized SG_R subsystem.

11.14. Gating of Learning, Movement, and Posture

The network to which our analysis of sensory-motor learning has led us is now a large and complex one. It includes circuits whose physiological and anatomical properties help to explain and predict data from the retina, superior colliculus, peripontine reticular formation, oculomotor nuclei, cerebellum, visual cortex, parietal cortex, and frontal eye fields. Underlying this complexity are a few general circuit designs—such as head-muscle interfaces, eye position maps, target position maps, and adaptive gain stages—which are utilized in multiple network locations. The placement of these circuits into different network positions endows them with different functional capabilities. The reader need only recall the many uses of eye position maps, including their possible role in generating an invariant TPM (Chapter 10), to appreciate this fact. Thus a significant part of the intelligence embodied within these large circuits lies in the genetic and developmental rules which control the correct placement of a few standardized circuit designs.

Another part of the intelligence of such a system lies in the gating rules which synchronize the switching on and off of the correct subsystems through time. Such gating rules have arisen repeatedly during the preceding chapters. They must be simple enough to generate correct switching decisions based upon locally computable signals. They give rise to global intelligence due to the manner in which they are connected to network subsystems. In particular, they are organized to achieve synchronous switching between movement and postural subsystems, and to modulate learning within each of these subsystems (Chapter 9). The present section summarizes and further develops the gating rules that are needed to prevent contradictory movement or postural signals from being generated during the saccade cycle.

Figures 11.10–11.13 describe the main types of gates that are needed in our theory. In all of the figures, gating pathways are denoted by dashed lines. The gates that regulate sequential read-out of target positions

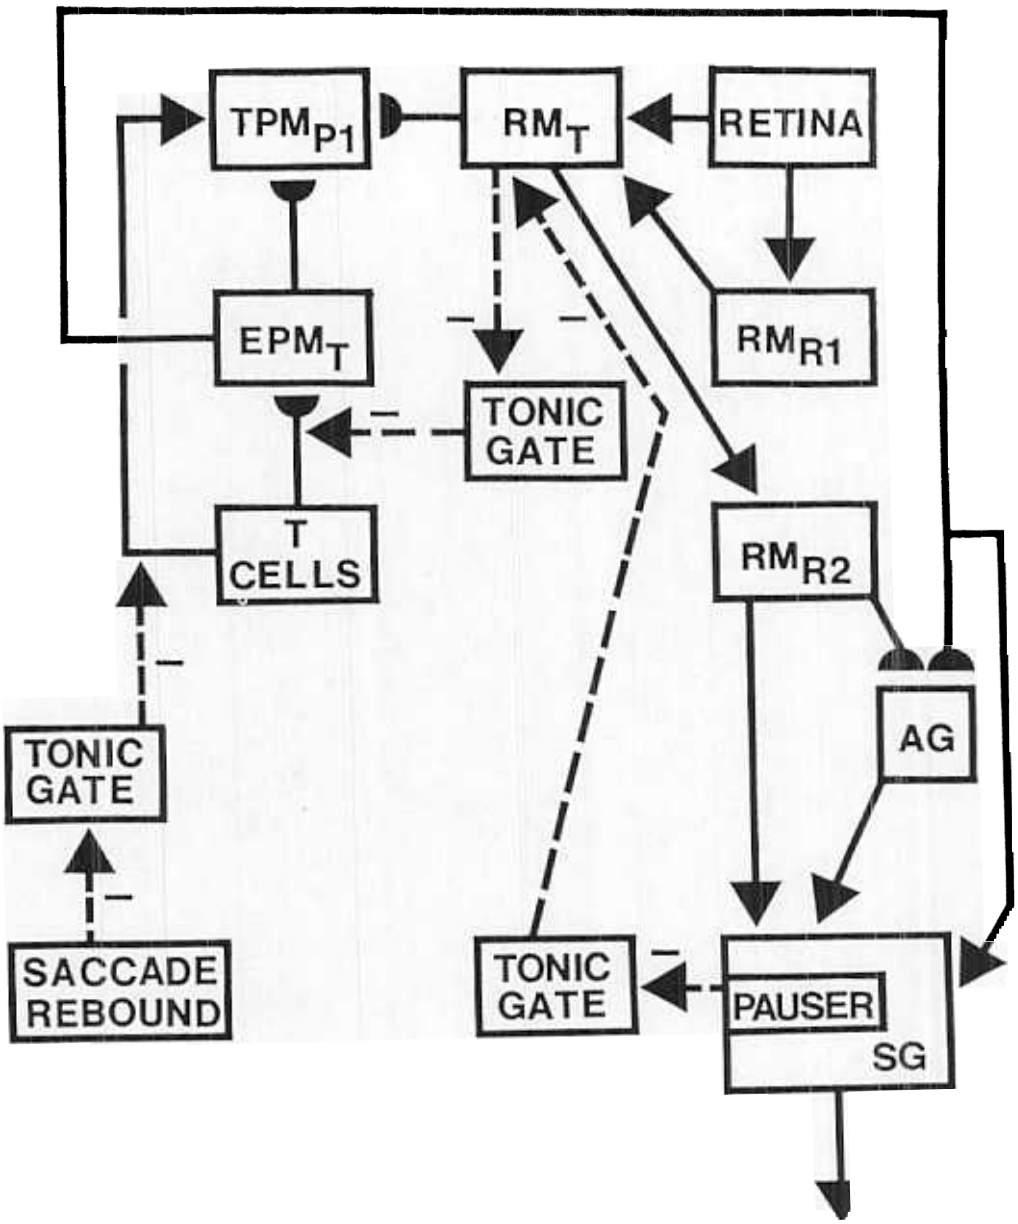

Figure 11.10. Gating of map reset events by signals from the saccade generator (SG). See text for details. Gating pathways are marked by dashed lines.

within the PCN will not be discussed, because they were already dealt with in Chapter 9. Figures 11.10–11.12 show that just a few types of gating actions suffice to regulate the entire network.

A. *Read-in, Reset, and Storage of Movement Commands*

Figure 11.10 indicates how movement commands activate the network, how they reset previous commands, and how they are stored in STM long enough after a saccade terminates to sample adaptive gains at the SG stage (Chapter 2).

Suppose that the network starts out at rest. The pauser population in the SG is then on. This is an omnipauser population that is on between all saccades and is turned off by all saccades. Because these pausers are on, they inhibit the inhibitory tonic gate. The RM_T is therefore free to respond to lights that impinge upon the retina. (In the subsequent discussion, we will suppose for definiteness that inhibitory tonic gates modulate network stages. Excitatory phasic gates could, in principle, do the same job, as Figure 11.13 will illustrate. Such thematic variations will be determined experimentally.)

Activation of the RM_T has three major effects. First, the RM_T can instate a new retinotopic position into the RM_{R2}. By so doing, it inhibits any retinotopic position that was previously stored by the RM_{R2}, as well as the output signals from the RM_{R2} to the SG and to the AG stage.

Second, the RM_T enables a new initial eye position to be read into the EPM_T, by inhibiting the tonic gate that kept the T cells from doing so before. This eye position is stored in STM within the EPM_T. Due to this sequence of events, the RM_{R2} is inhibited before the EPM_T can be reset. Consequently, the new eye position at the EPM_T cannot initiate a saccade staircase by inputting a new eye position signal to the SG before the previous retinotopic signal is shut off (Section 7.6).

Third, both the new RM_T position and the new EPM_T position can now input to the TPM_{P1}, which leads to choice and STM storage of a target position at the TPM_{P2}.

These events prepare the network to initiate a new saccade by instating new retinotopic (RM_{R2}), initial eye position (EPM_T), and target position (TPM_{P2}) commands. These reset events also occur with a sufficient delay after a previous saccade has ended to allow the RM_{R2}, EPM_T, and other systems which sample the AG stage to benefit from any second light error signals that the RM_{R1} might generate there. All of these events are regulated by two tonic gates, both of which are controlled, directly or indirectly, by populations in the SG.

Figure 11.10 also notes that a different type of gate is needed in order to learn an invariant TPM_{P1}. Such a gate is opened only after a saccade occurs, in response to the activation of saccade rebound cells. This gate enables the RM_T retinotopic position and the EPM_T initial eye position to sample the target eye position that is attained by the saccade (Section 10.5). Saccade rebound cells can arise as the off-cells of a gated dipole opponent process (Grossberg, 1972c, 1980, 1984), whose on-cell input source

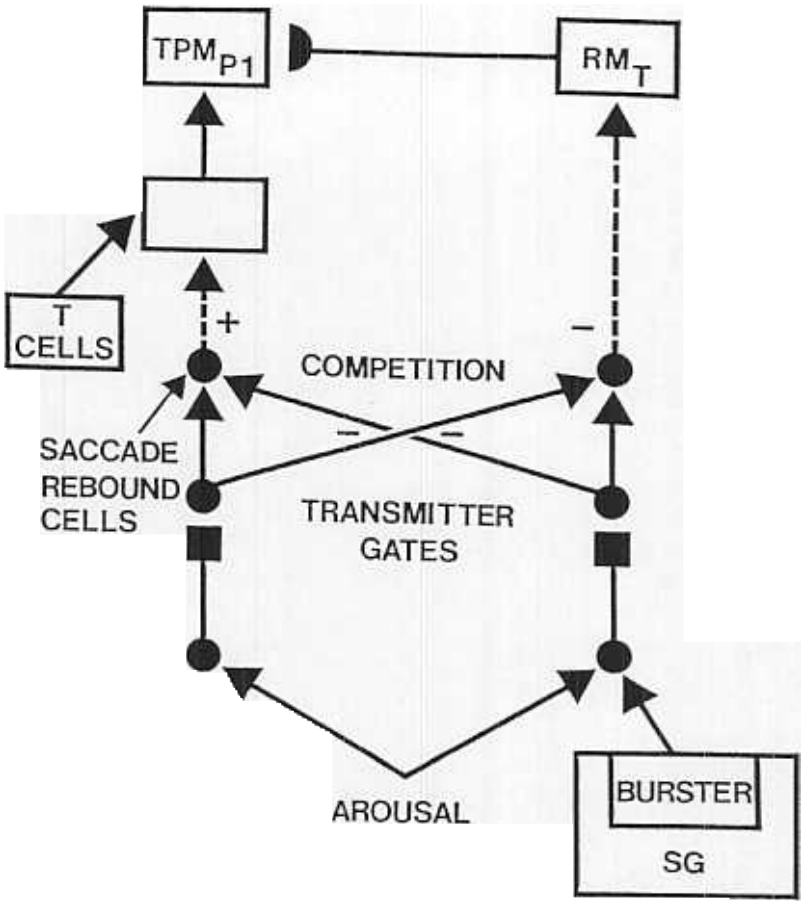

Figure 11.11. A possible mechanism giving rise to the saccade rebound cells of Figure 11.10. A tonically-aroused opponent process, called a gated dipole, can inhibit the retinotopic map RM_T via its on-channel and excite the gate in the T cell→TPM_{P1} pathway via its off-channel. A burster population in the saccade generator (SG) can excite the on-channel. Offset of the burster generates a transient output burst from the off-channel. Habituation of the slowly accumulating transmitter gates within the gated dipole regulates its on and off responses.

is an SG burster cell population (Figure 11.11). In fact, the on-cells of such a gated dipole can replace the disinhibitory pauser gate in Figure 11.10 by an inhibitory burster gate (Figure 11.11).

B. *Read-Out and Competition of Movement Commands*

Figure 11.12 describes gating actions that regulate the flow of movement commands to the saccade generator (SG). The critical gate in this figure is an inhibitory tonic gate which acts at a stage somewhere between (or within) the RM_{R2} and the SG. We hypothesize that this gate exists within the substantia nigra (Hikosaka and Wurtz, 1983) and therefore interpret its site of action as the deep layers of the superior colliculus (SC). We also assume that several different movement sources can act upon this gate.

The RM_{R2} can inhibit the gate, and thereby enable its stored retinotopic position to activate the SG, other things being equal. Two properties govern our choice of the RM_{R2} as a possible source of gate inhibition. First, the RM_{R2} stores a choice in STM, and hence is active throughout a saccade. Thus the gate stays off throughout an RM_{R2}-activated saccade and thereby enables the SG to fully react to the RM_{R2} movement command. Second, the RM_{R2} is part of the retinotopic command network (RCN). The RM_{R2} can therefore be activated by attended lights before the developmental stage occurs during which the invariant TPM_{P1} forms. The RCN can thus generate and correct the visually reactive saccades on which invariant TPM formation, vector calibration within the HMI, and other later learning processes are based.

The TPM_{P2} is also assumed to inhibit the tonic gate. This enables attentionally modulated multimodal inputs, in addition to visually reactive visual inputs, to generate saccades via the $TPM_{P2} \rightarrow HMI_P \rightarrow RM_P \rightarrow RM_{R2} \rightarrow SG$ pathway.

Finally, we assume that the RM_F can powerfully excite the tonic gate. This excitation can counteract inhibition from the RM_{R2} and the TPM_{P2}. Thus, when the volitional gate in the FEF opens (Chapter 9) and a vector from HMI_F activates the RM_F, this voluntary command can inhibit other movement commands as it activates its own pathway to the SG_F. This hypothesis is consistent with the fact that the FEF projects to the SC via the substantia nigra (Hikosaka and Wurtz, 1983).

Figure 11.12 also describes another possible way in which the FEF movement system can compete with the SC movement system. Suppose that RM_F commands excite the bursters of SG_F and the pausers of SG_R, and possibly that RM_{R2} commands excite the bursters of SG_R and the pausers of SG_F. If the pausers of SG_R could be kept on by RM_F excitation even in the presence of inhibition from SG_R long lead bursters, then the FEF could dominate SC commands even within the peripontine reticular formation (Section 7.8). We do not know any data relevant to this possibility. The circuit can work without this mechanism.

One general conclusion from Figures 11.10–11.12 is that SG cells such as tonic cells, bursters, and pausers may interact with far-flung neocortical cells, no less than with nearby cells of the SG circuit.

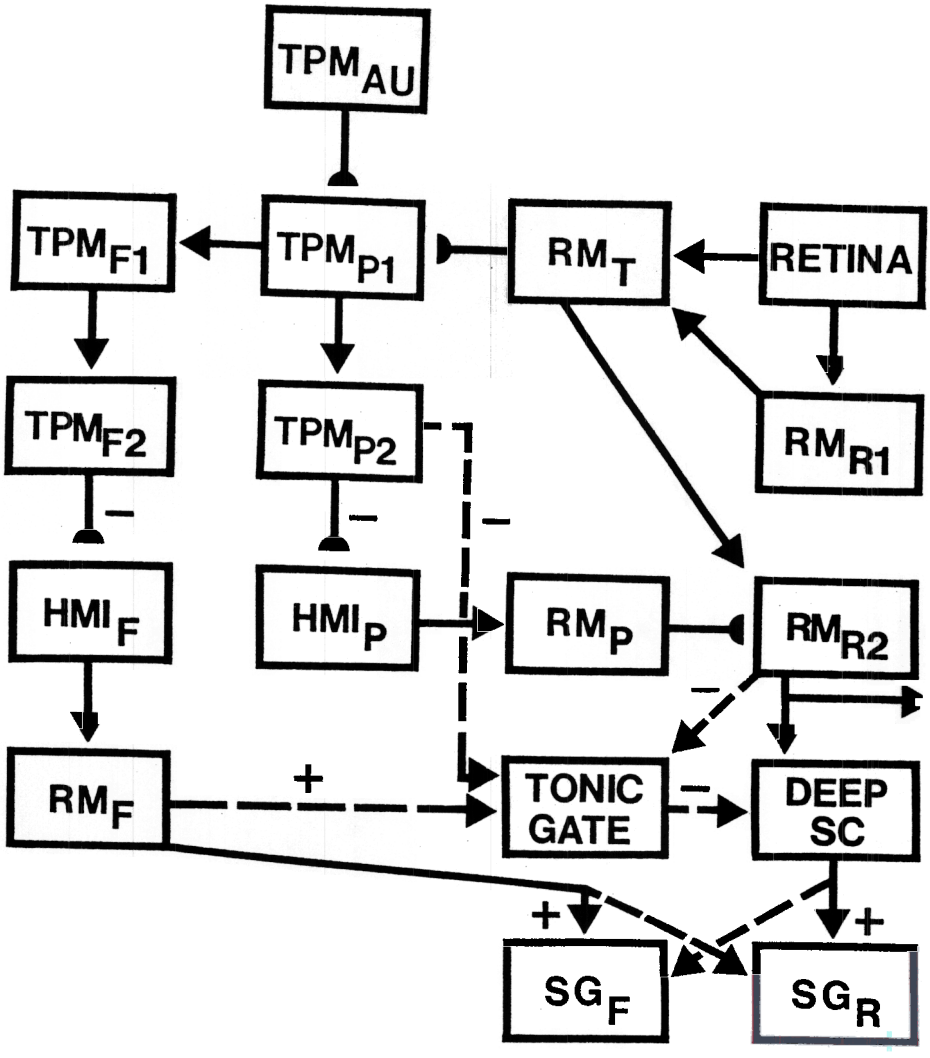

Figure 11.12. Gating processes that regulate the flow of movement commands to the saccade generator (SG). See text for details. Gating pathways are indicated by dashed lines.

C. *Gating of Posture and Learning*

In Figure 11.13, the postural eye position map (EPM_{Pos}) is allowed to store eye positions only between saccades. While a saccade is being performed, the tonic gate inhibits the EPM_{Pos}. Thus the EPM_{Pos} reads out the adaptive gains which prevent postsaccadic drift only during the postural state (Chapter 8).

In Figure 11.13, we have assumed for simplicity that the same tonic gate acts upon the EPM_{Pos} and upon vector outputs from the HMI_P and the HMI_F. Outputs leave each HMI for STM storage within its retinotopic map RM only during the postural state (Chapter 4). Thus before a saccade begins, each HMI can compute a vector difference of target position and initial eye position and store it retinotopically in its RM. As soon as a saccade begins, the $HMI \rightarrow RM$ pathways are gated shut. Consequently the stored retinotopic positions are not changed due to the saccadic motion. After the saccade terminates, the output gates open again and enable new vectors to be retinotopically stored.

Figure 11.13 also includes a tonic gate that prevents learning from occurring within the HMI except during the postural state (Chapter 4). In networks wherein a saccade generator subsystem like SG_R does not exist, all the gating actions in the figure could, in principle, be governed by a single SG-controlled gating system.

11.15. When Saccade Choice May Fail: Saccadic Averaging and Partial Vector Compensation

In order to emphasize the dynamic nature of the interactions which we have described, we end this chapter by considering some circumstances during which saccadic choice may partially fail.

In Section 11.3, we suggested that choice of one retinotopic position occurs due to the broad lateral inhibitory signals and signal functions at the stage that is caricatured by equation (11.3). Such a network may not, however, make a choice if two electrodes simultaneously activate a pair of its cells with equal intensity, or with intensities that are greater than the strength of its lateral inhibitory signals. Under these circumstances, both cells may retain partial activity and may thus generate a saccade that is a compromise, or weighted average, of the individual saccadic commands (Schiller and Sandell, 1983).

In Chapter 9 and Figures 11.3, 11.4, 11.8, and 11.9, we suggested that only one target position at a time can normally be stored within a TPM_2; then only one vector at a time can be computed at the corresponding HMI. By contrast, the process of STM reset and STM storage at a TPM_2 takes some time. It may be possible for one light to be partially stored within the TPM_2 when another, more salient, light occurs. Such a second light may be able to instate itself in the TPM_2 before the previous light can be fully stored. In this case, a partial vector compensation may occur that is rapidly followed by a second, more complete, vector compensation (Young, 1981).

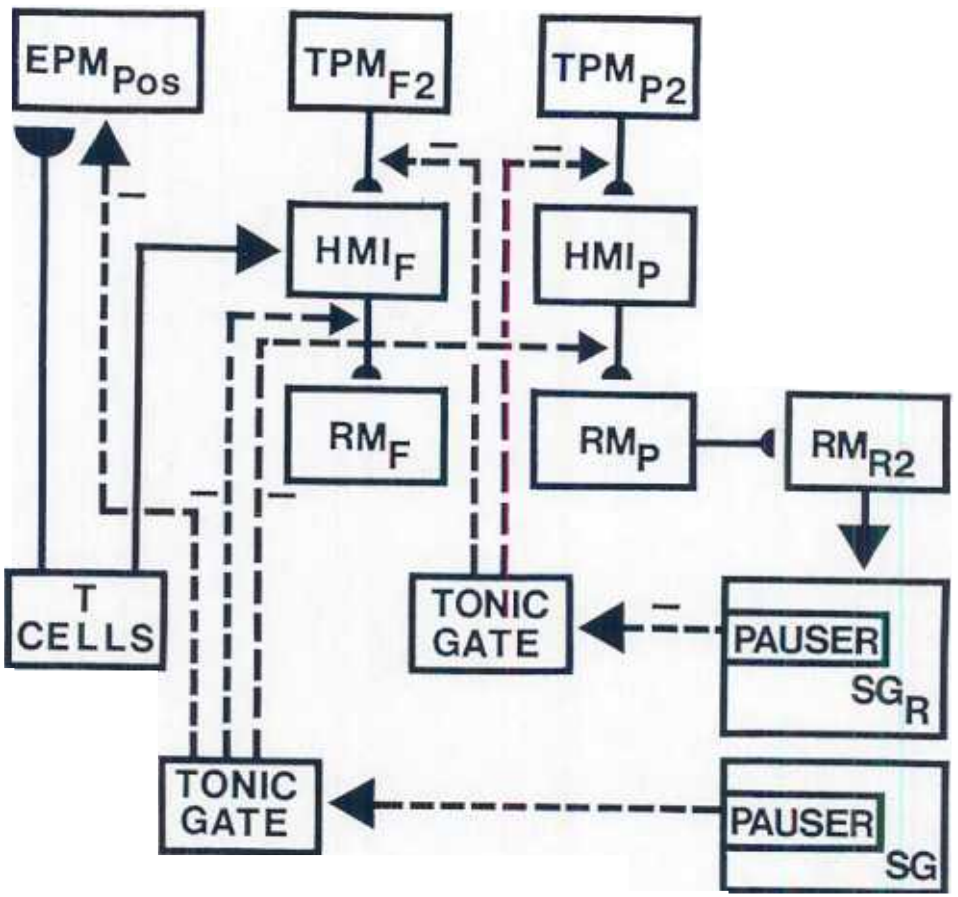

Figure 11.13. Gating processes that regulate reset events during posture. See text for details. Gating pathways are indicated by dashed lines.

Finally, only one saccade is usually generated by the retinotopic command network (RCN), or the predictive command network (PCN), but not both. On the other hand, simultaneously active electrodes placed in both subsystems may either bypass their inhibitory gates (Figure 11.12) or may override the inhibitory signals due to these gates. Then a saccade may occur that is a weighted average of both saccadic commands (Schiller and Sandell, 1983).