

Robust Persistent Neural Activity in a Model Integrator with Multiple Hysteretic Dendrites per Neuron

Mark S. Goldman^{1,4}, Joseph H. Levine², Guy Major³,
David W. Tank³ and H.S. Seung¹

¹Howard Hughes Medical Institute, Brain and Cognitive Sciences Department, MIT, Cambridge, MA 02139, USA, ²Electrical Engineering Department, MIT, Cambridge, MA 02139, USA and ³Departments of Molecular Biology and Physics, Princeton University, Princeton, NJ 08544, USA
⁴Present address: Department of Physics and Program in Neuroscience, Wellesley College, Wellesley, MA 02481, USA

Short-term memory is often correlated with persistent changes in neuronal firing rates in response to transient inputs. We model the persistent maintenance of an analog eye position signal by an oculomotor neural integrator receiving transient eye movement commands. Previous models of this network rely on precisely tuned positive feedback with <1% tolerance to mistuning, or use neurons that exhibit large discontinuities in firing rate with small changes in eye position. We show analytically how using neurons with multiple bistable dendritic compartments can enhance the robustness of eye fixations to mistuning while reproducing the approximately linear and continuous relationship between neuronal firing rates and eye position, and the dependence of neuron pair firing rate relationships on the direction of the previous saccade. The response of the model to continuously varying inputs makes testable predictions for the performance of the vestibuloocular reflex. Our results suggest that dendritic bistability could stabilize the persistent neural activity observed in working memory systems.

Introduction

The oculomotor neural integrator is a brain stem region that converts velocity-encoded eye movement commands into signals that encode eye position. The velocity commands include bursting inputs that cause saccades and head-velocity inputs that drive the vestibuloocular reflex (VOR). In the absence of command inputs, integrator neurons maintain persistent firing at a rate that is proportional to eye position. Thus, integrator neurons are said to maintain a ‘memory’ of eye position.

The first model of a neural integrator (Rosen, 1972) consists of a network of neurons with extremely strong self-excitation and weaker global recurrent excitation. The strong self-excitation makes the individual neurons bistable as a function of their external inputs, with both a quiescent (‘off’) state and a self-sustaining (‘on’) state of firing at saturation. The neurons have staggered thresholds for turning ‘on’ or ‘off’. The network response is characterized by the number of ‘on’ units. In the absence of the global excitatory coupling, an external input flips ‘on’ all units that have a threshold lower than this input. Repeated presentations of this input do not cause additional units to flip ‘on’. With recurrent excitatory coupling, repeated presentations of the input flip ‘on’ additional units that are brought closer to their firing threshold by the previous input’s effect on the network activity. This behavior was shown to result in temporal integration.

Rosen’s model was dismissed (Cannon *et al.*, 1983) as a model of the oculomotor neural integrator because real oculomotor neurons do not exhibit binary ‘on/off’ behavior. Subsequent models (Kamath and Keller, 1976; Cannon and Robinson, 1985; Seung, 1996; Seung *et al.*, 2000) substituted neurons with continuous firing rate versus injected current relationships for

Rosen’s bistable neurons. They also generalized Rosen’s recurrent excitatory network architecture to other architectures that lead to effective positive feedback between neurons, like reciprocal inhibition (Cannon *et al.*, 1983; Cannon and Robinson, 1985). These models successfully capture the threshold linear firing rate versus eye position relationships (Robinson, 1981; Aksay *et al.*, 2000) observed for oculomotor neurons and provide an explanation for why these relationships exhibit different slopes and thresholds (Seung *et al.*, 2000). However, they rely on fine-tuning of the network positive feedback to within a tolerance of <1% to achieve stable eye fixations. Too much feedback causes the network firing rates to grow to saturation, while too little causes firing rates to decay to zero.

Recently, Koulakov *et al.* (2002) showed how a network of coupled bistable units similar to Rosen’s model, but more realistic because the neurons exhibit a graded component in their firing response, could add robustness to eye fixations without requiring fine-tuning of parameters. In addition, these authors used conductance-based neurons and demonstrated how the bistability could be achieved by the voltage-dependent properties of the NMDA channel. This followed upon previous work that showed how bistability could make working memory network models more robust by providing the neurons with an intrinsic source of stability (Camperi and Wang, 1998; Lisman *et al.*, 1998).

Neuronal firing rates in the model of Koulakov *et al.* (2002), as in Rosen’s model, can exhibit large discontinuous jumps in steady-state firing rate with small smooth changes in eye position. These discontinuities reflect the neuronal bistability, with greater discontinuity tending to give greater robustness. Experimentally, large jumps in firing rate have not been observed to persist when current pulses have been injected at a variety of eye positions (Aksay *et al.*, 2001). Additionally, firing rates have not been observed to exhibit systematic jumps in firing rate with continuously varying command inputs as during the VOR (Pastor *et al.*, 1994).

Here, we propose a neural integrator model that attains robustness to mistuning from bistability but without large discontinuities in firing rate. This is achieved by placing the bistability in local dendritic compartments that are well isolated from the influence of the soma. Dendritic bistability has been observed experimentally in other systems under various conditions (Hounsgaard and Kiehn, 1993; Reuveni *et al.*, 1993; Yuste *et al.*, 1994; Lee and Heckman, 1998a,b; Schiller *et al.*, 2000). It has been attributed to voltage-dependent elements that can exhibit self-sustained activation, such as the NMDA channel and voltage-sensitive Ca^{2+} channels. Integrator neurons display extensive dendritic arborization (Aksay *et al.*, 2000), and theoretical work in other systems (Koch *et al.*, 1982; Poirazi *et al.*,

2003a,b) suggests that neurons can contain many approximately independent dendritic subunits.

The model's firing rate exhibits *quasi-continuous* behavior, defined in the limit of a large number of dendritic compartments as having discontinuities in firing rate versus eye position relationships that approach zero as the number of these compartments approaches infinity. Qualitatively, this means that any jumps in firing rate are sufficiently small as to be undetectable experimentally. This results from the firing rate of the neuron jumping only incrementally when a single dendritic compartment flips 'on'. Thus, this model achieves robust eye fixations without requiring fine-tuning of parameters or large discontinuities in firing rates.

Model Description

We construct a network of N recurrently connected neurons that controls the eye position in the positive half of its range. Each neuron has N bistable dendritic compartments, each of which receives recurrent input from one of the N neurons in the network (Fig. 1A). Each soma receives recurrent network input from the N dendrites, as well as direct (not mediated by the dendrites) constant-rate tonic input, saccadic burst input, and head velocity command input responsible for the VOR. We use a firing rate model of the soma and assume that the firing rate r_i of the i^{th} neuron is given by a threshold linear sum of its inputs (Aksay *et al.*, 2001):

$$r_i = \left[\sum_{j=1}^N W_{ij} D_{ij}(r_j) + r_{\text{ton},i} + r_{\text{com},i} \right]_+ \quad (1)$$

where $[\]_+$ denotes that all negative values are set to zero. Here, $r_{\text{ton},i}$ gives the contribution to the firing rate due to tonic back-

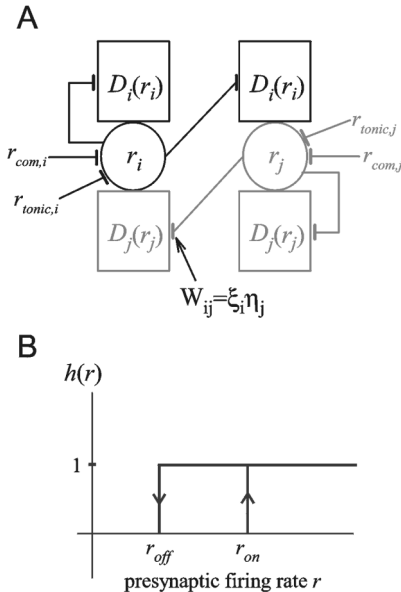


Figure 1. (A) Structure of the connections and inputs to the neurons in the analytic model network. Two of the N neurons in the network are shown. Neuron j projects to an identical dendritic compartment on every other neuron, characterized by the dendritic activation function $D_j(r_j)$. The contribution of this dendritic compartment to postsynaptic neuron i 's firing rate is $W_{ij} D_j(r_j)$. Neuron i 's firing rate also reflects contributions from tonic input, $r_{\text{ton},i}$, and command input, $r_{\text{com},i}$. (B) The hysteretic steady-state activation relation of a bistable dendrite receiving input from a presynaptic neuron that fires at rate r .

ground input, $r_{\text{com},i}$ gives the contribution due to eye movement commands (saccadic bursts and vestibular inputs), W_{ij} gives the maximum possible contribution from presynaptic neuron j , and $D_{ij}(r_j)$ gives the time-dependent activation of the dendrite that receives input from neuron j .

Recurrent input in the network is mediated by bistable dendritic compartments. The dendritic activation $D_{ij}(r_j)$ is governed by exponential approach of time constant τ_{dend} to a hysteretic steady-state activation relation $h_{ij}(r_j)$ (Fig. 1B). $h_{ij}(r_j)$ is like a function except that it is binary valued and history-dependent for presynaptic rates $r_{\text{off},ij} < r_j < r_{\text{on},ij}$: When r_j exceeds an activation threshold $r_{\text{on},ij}$, $h_{ij}(r_j)$ jumps up to a value of 1 (turns 'on'). $h_{ij}(r_j)$ then remains at a value of 1 until r_j drops below a deactivation threshold $r_{\text{off},ij}$ and jumps down to a value of 0 (turns 'off'). The $D_{ij}(r_j)$ dynamics may be summarized by the first order differential equation:

$$\tau_{\text{dend}} \frac{dD_{ij}}{dt} = -D_{ij} + h_{ij}(r_j) \quad (2)$$

Here, τ_{dend} sets the intrinsic biophysical timescale in the model. The simplified model of dendritic activation assumed by the above equation allows the main results of this paper to be derived analytically. Later in the paper, we consider more complicated but less analytically tractable variants of this model that include non-instantaneous synaptic input to the dendrite and background noise.

Results

Firing Rates during Fixations Are Threshold Linearly Related to Each Other and to Eye Position

We next show that the model outlined above can exhibit quasi-continuous, linear firing rate versus eye position relationships during fixations. This occurs when the recurrent input to each neuron has the form described below. Later, we show how modifying the form of the model described below can lead to experimentally observed effects such as history-dependence of the firing rate versus firing rate relationships of neuron pairs.

We suppose that the dendritic-activation thresholds, $r_{\text{on},ij}$ and $r_{\text{off},ij}$, depend only on the presynaptic neuron from which the dendrite receives input, $r_{\text{on},ij} = r_{\text{on},j}$ and $r_{\text{off},ij} = r_{\text{off},j}$. From this, it follows that $D_{ij}(r_j) = D_j(r_j)$. We suppose further that the neuron-dendrite coupling strengths W_{ij} can be written in the outer product form $W_{ij} = \zeta_i \eta_j$ (Seung *et al.*, 2000). These assumptions are made for computational convenience and reduce the network dynamics to a single equation (equation 5, below) that can be treated analytically. This equation can be thought of as modeling only the 'integrating mode' of a more general multi-dimensional network with many transiently decaying modes (Seung, 1996). η_j can be interpreted as the weight of activation of all dendrites to which neuron j projects, and ζ_i as the slope of neuron i 's firing rate versus total dendritic activation relationship. The outer product form of the recurrent feedback contributions implies that each of the N dendrites to which a given neuron projects has identical dynamics. Thus, the N^2 dendrites in the model function as N identically behaving groups (Fig. 1A).

Under these assumptions, equation 1 can be re-written as

$$r_i = [\zeta_i \hat{E} + r_{\text{ton},i} + r_{\text{com},i}] \quad (3)$$

where

$$\hat{E} = \sum_{j=1}^N \eta_j D_j(r_j) \quad (4)$$

Thus, during fixations ($r_{\text{com},i} = 0$), all firing rates are threshold linearly related to \hat{E} with slopes ζ_i and fixed thresholds $-r_{\text{ton},i}/\zeta_i$ (Fig. 2A). We refer to \hat{E} as the network representation of eye position. We do not model the readout of \hat{E} .

The linearity of the r_i versus \hat{E} relationships derives from the successive recruitment of 'on' dendrites with increasing eye position. Figure 2B illustrates the recruitment order of dendrites that receive input from three different neurons and have the same activation weights η_i . In the limit of very many independent dendritic compartments, the r_i versus \hat{E} relationships become not only linear, but also quasi-continuous.

Graphical Analysis of Fixations in a Well-tuned, Leaky or Unstable Integrator

Equations 2–4 can be combined to give a single equation governing the dynamics of \hat{E} . Multiplying equation 2 by η_i , summing over i , and substituting for r_i and E , gives

$$\tau_{\text{dend}} \frac{d\hat{E}}{dt} = -\hat{E} + \sum_{i=1}^N \eta_i h_i(\zeta_i \hat{E} + r_{\text{ton},i} + r_{\text{com},i}) \quad (5)$$

We have dropped the thresholding in the argument of $h_i(r_i)$ because we assume that $h_i(r_i) = 0$ when its argument is less than zero.

Stable eye fixations in the absence of eye movement commands ($r_{\text{com},i} = 0$) may occur at any eye position for which

$$\hat{E} = \sum_{i=1}^N \eta_i h_i(\zeta_i \hat{E} + r_{\text{ton},i} + r_{\text{com},i})$$




Figure 2. (A) Firing rate as a function of \hat{E} for a neuron with tonic background firing rate $r_{\text{ton},j}$ and with sensitivity to dendritic activation ζ_i . (B) Contribution to \hat{E} of the dendrites to which three neurons project. The weighted steady-state activations $\eta_i h_i(r_i)$ for each of these dendrites (solid, dashed and dotted lines) are plotted as a function of \hat{E} . The 'on' and 'off' thresholds as a function of \hat{E} are labeled for one dendrite (solid line).

This may be illustrated graphically [compare (Seung *et al.*, 2000; Koulikov *et al.*, 2002)] by plotting the left and right sides of this equation as a function of \hat{E} on the same set of axes (Fig. 3A). The \hat{E} term draws out a 45° line. The summed term traces out a ragged 'band' representing the summed contribution of the dendritic inputs as a function of \hat{E} . The intersections of the 45° line and band give the possible stable eye positions. For a particular value of \hat{E} , the dendritic hysteretic rectangles that appear below the horizontal line $y = \hat{E}$ represent dendritic compartments that are 'on', and those appearing above this line represent compartments that are 'off'. In the limit of many dendritic compartments, the ragged band of stacked rectangles approaches a smooth band, and the set of eye positions over which stable fixations can be achieved by a well-tuned integrator becomes quasi-continuous (Fig. 3B). This limit is achieved by scaling the dendritic weight factors as \hat{E}_{max}/N to maintain the total feedback to the neuron as the number of dendrites increases.

A 'well-tuned' integrator in this model does not require 'fine-tuning'. Stable fixations can be achieved over a quasi-continuous range of eye positions so long as the band intersects the 45° line throughout this range. Perturbations of the network parameters deform the shape of the band by changing the shapes and/or locations of the hysteretic dendritic rectangles (see Fig. 2B). As a result of such perturbations, the band may no longer enclose the 45° line throughout the full range of eye positions (Fig. 3C,D). For the eye positions at which the band does not enclose the line, fixations cannot be maintained. Instead, the eyes drift with a velocity proportional to the distance between the band and the line at the corresponding value of \hat{E} (as can be seen from equation 5). When there is insufficient feedback, the inte-

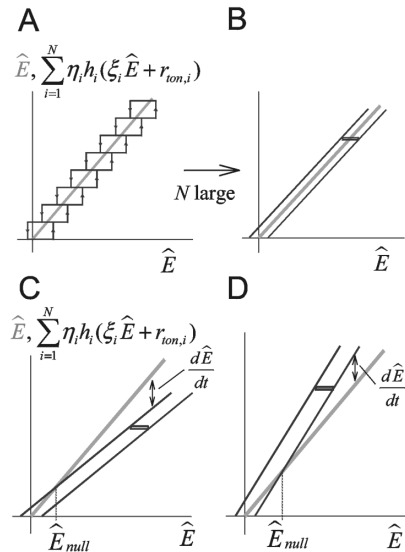


Figure 3. Graphical analysis of eye fixations. (A) Decay (gray line) and positive feedback (black hysteretic loops) contributions to eye velocity (equation 5) during fixations ($r_{\text{com},j} = 0$). Stable fixations can occur at any value for which the decay and positive feedback relations intersect. (B) In the limit of large N , the hysteretic loops form a nearly continuous positive feedback 'band'. The contribution of one dendrite is highlighted (black rectangle). (C,D) Positive feedback bands for leaky (C) and unstable (D) integrators. Stable fixations can occur at eye positions up to \hat{E}_{null} for which the band encloses the line. When the band does not enclose the line, \hat{E} drifts with a velocity proportional to the difference between the line and the near side of the band.

grator becomes ‘leaky’ and stable eye fixations are held only when \hat{E} is less than a ‘null’ eye position \hat{E}_{null} (Fig. 3C). For $\hat{E} > \hat{E}_{\text{null}}$, the eye position decays to the null position E_{null} . When the feedback is too large, the integrator becomes ‘unstable’, holding stable fixations only for $E > \hat{E}_{\text{null}}$ and growing until saturation when \hat{E} exceeds E_{null} .

Dendritic Hysteresis, Not Binary Dendritic Responses, Leads to Robustness of a Quasi-continuous Set of Stable Eye Fixations

The tolerance of stable fixations to mistuning of parameters comes about primarily because of the hysteresis in the functions $h_i(r_i)$, rather than their binary (‘on’ or ‘off’) nature (Pouget and Latham, 2002). We illustrate this for a specific example below.

Consider a network of N neurons with constant parameter values $r_{\text{on},i} = r_{\text{on}}$, $r_{\text{off},i} = r_{\text{off}}$, and $W_{ij} = W = \eta\zeta$, and equally spaced tonic inputs

$$r_{\text{ton},i} = \frac{N-i+0.5}{N} \bar{r}$$

where $\bar{r} = (r_{\text{on}} + r_{\text{off}})/2$ (Fig. 3). These values were chosen to form a band with parallel edges whose bottom is centered at the origin. For this model, stable fixations across the entire range of eye positions can be achieved for values of W within a range ΔW defined by (see Appendix for derivation):

$$\frac{\Delta W}{W^*} = \frac{r_{\text{on}} - r_{\text{off}}}{\bar{r}} + O\left(\frac{1}{N}\right) \quad (6)$$

where W^* gives the midpoint of the range ΔW . $\Delta W/W^*$ defines the fractional tolerance of the network to mistuning of the W .

The first term of equation (6) gives the effect of the hysteresis: the fractional tolerance of the network due to hysteresis equals the width of each dendrite’s hysteretic loop divided by its midpoint. Equivalently, this equals the width of the hysteretic band in the limit $N_j \rightarrow \infty$ (Fig. 3B) divided by the maximum eye position. The fact that $\Delta W/W^*$ remains finite in the continuum limit means that the network remains robust. This is seen in Figure 4A,B: when the weights are mistuned by 10% (Fig. 4A), a network of 100 neurons with hysteretic dendrites still maintains stable firing across the entire range of \hat{E} (Fig. 4B).

The effect of having discrete fixed points is given by the second term of equation (6) and, for large N , decays as $1/N$. In the absence of hysteresis, the positive feedback ‘band’ becomes a jagged ‘line’ that occupies some width due to its nonlinear shape. For an excitatory network, the feedback line is a monotonically non-decreasing function. This creates a tradeoff between the number of discrete eye positions that can be maintained and the stability of these eye positions (Fig. 4C): Wider deviations of the feedback line from the 45° decay line allow greater tolerance to mistuning of parameters that distort the shape of the feedback line. However, wide deviations from the 45° line imply further separated, and hence more discrete, stable eye positions (Fig. 4C, left). Given that stable eye positions are observed without noticeable digitization, it is unlikely that discreteness alone is a major source of stability in the oculomotor neural integrator.

The tradeoff discussed above implies that, to achieve a quasi-continuum of potentially stable eye positions without hysteresis, the feedback band must be nearly perfectly linear and never deviate far from the 45° line. To make this idea concrete, let us compare the results obtained above for a hysteretic and binary

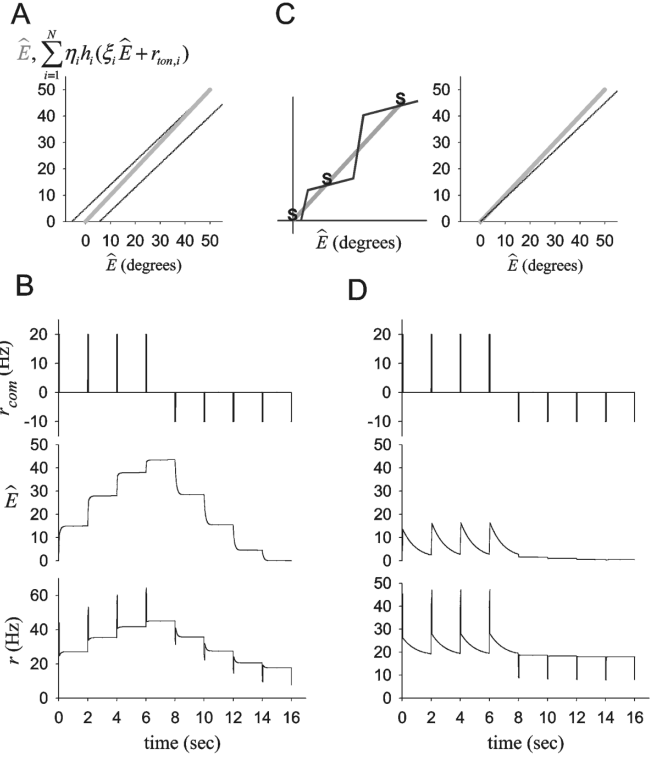


Figure 4. Hysteresis is the primary contributor to the robustness of stable fixations. (A) Decay line (gray) and positive feedback band (black) for a network of 100 neurons with hysteresis and coupling strengths $W = \eta\zeta$ decreased by 10% from W^* (equation 15) for a network that tolerates $\pm 10\%$ mistuning. (B) The hysteretic network tolerates the 10% mistuning, producing stable firing rates over a quasi-continuous range of \hat{E} . Top, middle, and bottom panels, respectively, show the eye movement command received by each neuron, \hat{E} , and the firing rate of one neuron in the network as a function of time. (C) A recurrent excitatory network without hysteresis has a tradeoff between having a few discrete fixed points that are robust to mistuning (left, stable fixed points denoted by s) or a quasi-continuum of fixed points that are highly sensitive to mistuning (right, drawn for a 10% decrease of W from optimal). (D) The non-hysteretic network with band pictured in C (right) shows severe leak. Panels are as described in B. In all simulation figures, $N = 100$, $\eta_i = \hat{E}_{\text{max}}/N$ where the maximum eye position $\hat{E}_{\text{max}} = 50^\circ$, and the midpoint of the hysteretic rectangles $\bar{r} = 35$ Hz. r_{on} and r_{off} are chosen to set the tolerance to mistuning of the network according to equation 6 (for parallel-edge shaped bands) or equation 7 (for cone-shaped bands). Because the η_i are constant, uniform changes across the ζ_i determine the mistunings of W from W^* . $r_{\text{ton},i} = ((N-i+0.5)/N)\bar{r}$ for the parallel-edge shaped bands. $\tau_{\text{dend}} = 100$ ms in all figures except Figure 6.

$h(r)$ (equation 6; Fig. 4A,B) with the case where $h(r)$ is binary but not hysteretic. In the latter case, the first term of equation (6) is zero. Therefore, taking the continuum limit $N_j \rightarrow \infty$ abolishes the robustness of the model. This is demonstrated in Figure 4C (right) and D for a network with identical parameters to that of Figure 4A,B except for the removal of the hysteresis. When the feedback is fine-tuned, the network can maintain stable firing at N values of \hat{E} (not shown). However, when the weights are mistuned by 10%, the non-hysteretic network exhibits rapid decay in \hat{E} for nearly all values of \hat{E} (Fig. 4D). Only at extremely low values of \hat{E} does the network maintain stable fixations. This small amount of remaining stability is due to the jaggedness of the feedback line. Similar stability due to discreteness has been observed in previous models (Seung *et al.*, 2000) and is a general

consequence of the nonlinear shape of $h(r)$ rather than the particular binary form of $h(r)$ used here.

Although we have only discussed uniform changes in W here, changes in particular weights will change only the shape of the corresponding hysteretic rectangles (Fig. 2B). The tuning criterion described in equation (6) is strictest for the N th dendrite because tilting the band by uniformly changing W values will most easily push the top of the band outside the line. The methods used in the Appendix to determine the tuning can be applied to changes in particular sets of weights, or random sets of weights, with the general result that random weight changes will have less effect than equal-magnitude uniform weight changes for which tolerances were defined here.

The Shape of the Band Relates to the r_i versus \hat{E} Relationships and Determines the Network Response to Strong Perturbations

In the previous example, we considered the special case of uniform dendritic parameters and equally spaced tonic inputs. We next consider the effect of having nonuniform dendritic parameter values. For simplicity, we here restrict the analysis to the large N limit and to band shapes that have linear slopes. A more detailed discussion is provided in the Appendix.

The shape of the positive feedback band is determined by the sizes and positions of each of the hysteretic rectangles (Fig. 2B). These sizes and positions as a function of \hat{E} are given by the individual terms in the sum of equation (5): the hysteretic response as a function of presynaptic firing rates, $h_i(r_i)$, is scaled vertically by η_i , shifted horizontally by $r_{\text{ton},i}$, and scaled horizontally by ζ_i^{-1} . In general, there are many ways to tune a band through partly redundant variations in the parameters η_i , ζ_i , $r_{\text{on},i}$, $r_{\text{off},i}$, and $r_{\text{ton},i}$. Throughout this work, we fix η_i at the constant value $\eta_i = E_{\text{max}}/N$. Except for the simulations of history-dependent effects (Fig. 6), we also hold constant the values of $r_{\text{on},i} = r_{\text{on}}$ and $r_{\text{off},i} = r_{\text{off}}$.

Although many parameters affect the shape of the band, because ζ_i and $-r_{\text{ton},i}/\zeta_i$ give the slopes and firing rate onset thresholds of the r_i versus \hat{E} relationships, the shape of the band is intimately related to these relationships. The band with parallel edges (Fig. 3A,B; Fig. 5B), or ‘parallel-edge shaped’ band, was constructed from r_i versus \hat{E} relationships having constant slopes but staggered thresholds that reflect different levels of tonic background input (Fig. 5A). Experimentally, neurons differ in both their slopes and thresholds, with a tendency for higher threshold neurons to have higher slopes (Aksay *et al.*, 2000). This arrangement can be captured in a simple model in which the dendritic parameters $r_{\text{on},i} = r_{\text{on}}$, $r_{\text{off},i} = r_{\text{off}}$, and $\eta_i = \eta$ and tonic inputs $r_{\text{ton},i} = r_{\text{ton}}$ are assumed to be equal, but the ζ_i are staggered according to (Fig. 5D)

$$\zeta_i = \frac{\bar{r} - r_{\text{ton}}}{i\eta}$$

This gives rise to a cone-shaped band (Fig. 5E).

Stable fixations for the cone-shaped band network tolerate perturbations that do not cause the \hat{E} line to fall outside the band. This can be shown to occur over a range ΔW of W 's given by (Appendix)

$$\frac{\Delta W}{W^*} = \frac{r_{\text{on}} - r_{\text{off}}}{\bar{r} - r_{\text{ton}}} + O\left(\frac{1}{N}\right) \quad (7)$$

where W^* gives the midpoint of the range ΔW . The form of the tolerance of this network to mistuning is similar to that for the




Figure 5. Relationship between firing rate versus \hat{E} relationships, shape of the band, and network behavior. (A) Firing rate versus \hat{E} relationships of five neurons for a network with constant couplings $W_{ij} = \eta\zeta$ and uniformly spaced $r_{\text{ton},i}$. (B) Shape of the band for a network with firing rate relationships as depicted in A, and mistuning of weights by 12% for a band width that tolerates $\pm 10\%$ mistuning. (C) Performance for a leaky integrator with band shape as in B. \hat{E} decays rapidly towards a null position \hat{E}_{null} for $\hat{E} > \hat{E}_{\text{null}}$. (D) Firing rate versus \hat{E} relationships of five neurons for a network with constant $r_{\text{ton},i} = r_{\text{ton}}$ and staggered slopes $\zeta_i = (\bar{r} - r_{\text{ton}})/(i\eta)$ where $\bar{r} = (r_{\text{on}} + r_{\text{off}})/2$. (E) Shape of the band for a network with firing rate relationships as depicted in D, and mistuning of weights by 12% for a band width that tolerates $\pm 10\%$ mistuning. (F) Performance for a leaky integrator with band shape as in E. \hat{E} decays with a longer time constant than in C towards a null position $\hat{E}_{\text{null}} = 0$. For other parameter values, see caption of Figure 4.

network with parallel-edge shape band (equation 6), except that the rates are shifted by r_{ton} .

For more severe perturbations, the networks behave much differently. The network with the parallel-edge shaped band responds to large decreases in coupling strengths with a rapid decay in \hat{E} towards a fixed nonzero null eye position \hat{E}_{null} (Fig. 5C). In contrast, the network with the cone-shaped band decays towards a zero null position, but with a much slower decay time (Fig. 5F). Aggregated experimental measurements of r_i versus eye position relationships across many animals appear to fall somewhere between the two extreme cases shown here (Aksay

et al., 2000). This suggests that the real network behavior would be intermediate between these two cases.

The null point to which \hat{E} decays and the time constant of this decay can be generalized to a network with band edges that are arbitrary linear functions of \hat{E} . From equation (5), it can be shown that the dynamics of such a network describe exponential growth towards or leak (decay) away from a null eye position

$$\hat{E}_{\text{null}} = \text{intersection with } \hat{E} \text{ line} \quad (8)$$

with a time constant

$$\tau_{\text{network}} = \frac{\tau_{\text{dend}}}{|1 - \text{slope of band}|} \quad (9)$$

where the 1 in the denominator gives the slope of the \hat{E} line. In the equations above, the upper side of the band should be used to calculate \hat{E}_{null} and τ_{network} for a leaky integrator, and the bottom side of the band for an unstable integrator.

The results above have all been computed in the absence of noise. The presence of noise can cause additional drift and, in a network in which the band is not perfectly centered around the decay line, this drift may be biased. Thus, for example, if noise were included in the simulation of Figure 5C, there would be a less discontinuous increase in eye drift with increasing eye position. Examples of fixations in the presence of noise are given in the section ‘Network dynamics during the VOR and effects of noise’.

History Dependence of Firing Rate versus Firing Rate Relationships Between Neurons

To a first approximation, the firing rates of oculomotor integrator neurons are threshold linearly related to one another. However, recent work reported in this issue (Aksay *et al.*, 2003) shows that the firing rate relationships of neuron pairs are systematically different following saccades that increase the recorded neurons’ firing rates (‘ON-direction’ saccades, corresponding to eye movements in a direction ipsilateral to the side of the recorded neurons) versus those that decrease the neurons’ firing rates (‘OFF-direction’ saccades). In particular, if the firing rates of neuron pairs are plotted with the higher firing-rate threshold neuron as the ordinate and the lower firing-rate threshold neuron as the abscissa, then the points following ON-direction saccades will preferentially lie above those following OFF-direction saccades. That is, the higher firing-rate threshold neurons have relatively higher firing rates following ON-direction saccades than following OFF-direction saccades. We refer to this previous saccade direction-dependent relationship in the neurons’ firing rate versus firing rates relationships as ‘ON-OFF hysteresis’.

The integrator model described in the previous sections (equations 2–4) does not exhibit this behavior because the outer product form of the model implies that, at all times, the neurons’ firing rates have a fixed relationship relative to one another. However, if the model is modified so that the outer product form is broken in an appropriate manner, the model can exhibit ON-OFF hysteresis.

We choose to break the outer product form of the model by relaxing the assumption $D_{ij} = D_j$ that the dendritic-activation thresholds for a particular dendrite depend only on the presynaptic neuron projecting to a given dendrite. We also, for added realism, include a synapse whose activation decays with a time

constant τ_s . The inclusion of the synapse does not by itself produce rate-rate hysteresis, although the value of τ_s does affect the relative positions of the ON-direction and OFF-direction sets of points in the rate-rate plots (see below). The equations governing the modified model are:

$$\tau_s \frac{ds_i}{dt} = -s_i + \alpha r_i \quad (10)$$

$$\tau_{\text{dend}} \frac{dD_{ij}}{dt} = -D_{ij} + h_{ij}(s_j) \quad (11)$$

$$r_i = \left[\sum_{j=1}^N W_{ij} D_{ij}(r_j) + r_{\text{ton},i} + r_{\text{com},i} \right] \quad (12)$$

where the steady-state dendritic activation $h_{ij}(s_j)$ is defined as previously, except that the dendritic-activation thresholds for turning ‘on’ and ‘off’ are now defined in terms of synaptic activations, $s_{\text{on},ij}$ and $s_{\text{off},ij}$, rather than in terms of the presynaptic neurons’ firing rates, $r_{\text{on},ij}$ and $r_{\text{off},ij}$. Because this network can no longer be written in an outer-product form, it cannot be reduced to a one-dimensional equation as was the case previously (equation 5).

Figure 6 shows how a network with different dendritic-activation functions for low firing rate threshold neurons (defined, in this example, as neurons with large tonic background activities $r_{\text{ton},i}$) and high firing rate threshold neurons (defined, in this example, as neurons with small tonic background activities $r_{\text{ton},i}$) can lead to ON-OFF hysteresis. For a given firing rate of the low firing rate threshold neuron (Fig. 6A, top), the high firing-rate threshold neuron (Fig. 6A, bottom) attains a different firing rate depending on whether the fixation follows an ON-direction or OFF-direction saccade (Fig. 6A). When the steady-state firing rate versus firing rate relationships of the two neurons are plotted, they form two approximately linear sets of values, with the points following ON-direction saccades lying above those following OFF-direction saccades (Fig. 6B). This behavior is a result of arranging the dendritic-activation thresholds $s_{\text{on},ij} = s_{\text{on},i}$ and $s_{\text{off},ij} = s_{\text{off},i}$ of the different neurons such that the dendrites on the higher firing-rate threshold neurons both turn ‘on’ and turn ‘off’ before those on the lower firing-rate threshold neurons (Fig. 6C, network parameters are given in the Fig. 6 caption). As a result, when the eyes change direction, the dendrites on the high firing-rate threshold neuron change their activation states before the corresponding ones on the low firing-rate threshold neuron. This implies that the high firing-rate threshold neuron will change its rate by relatively more than the low firing-rate threshold neuron immediately following a change in eye movement direction (Fig. 6A, gray arrows shows one example).

Figure 6 demonstrates just one possible way of generating ON-OFF hysteresis. We note that as long as the model does not have an outer product form, there are many possible parameter choices that could lead to ON-OFF hysteresis. For example, we have found that the locations of the ON and OFF points in Figure 6B depend on the dynamics of the synaptic and dendritic activations and the form of the saccadic burst as well as the parameters determining the shape of the hysteretic dendritic rectangles.

Network Dynamics during the VOR and Effects of Noise

In addition to integrating transient saccadic burst commands, the goldfish oculomotor neural integrator also receives contin-

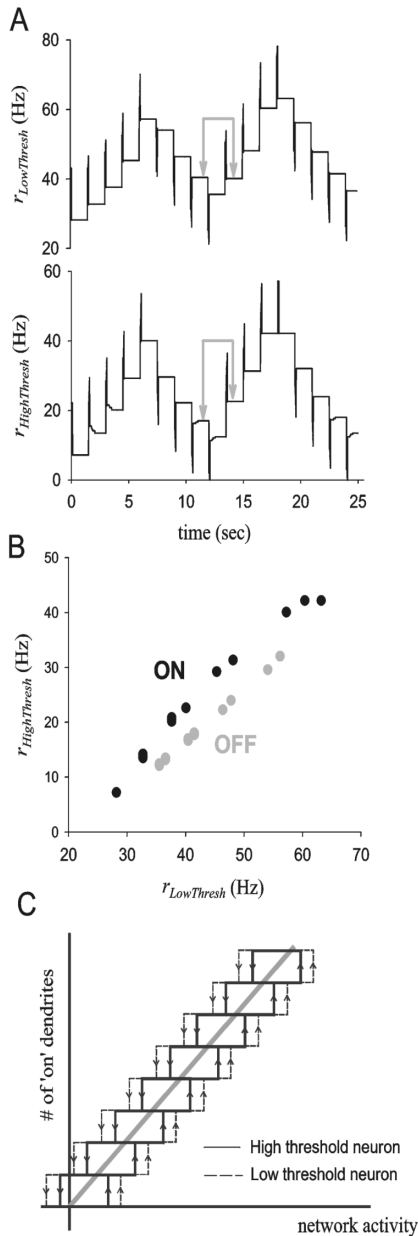


Figure 6. Dendritic hysteresis can lead to history dependence in rate–rate relationships. (A) Firing rates of a low firing-rate threshold (top) and high firing-rate threshold (bottom) neuron in a well-tuned network with parallel-edge shaped band and with outer product form broken by having a linear dependence of the ‘on’ and ‘off’ dendritic-activation thresholds on the neuron’s firing-rate threshold. Parameters are as described in Figure 4 except that $\alpha=1$, the $s_{on,j}/\alpha$ are equally spaced from 42 Hz for the lowest-threshold neuron down to 36 Hz for the highest firing-rate threshold neuron, and the $s_{off,j}/\alpha$ are equally spaced from 28 Hz for the lowest firing-rate threshold neuron up to 34 Hz for the highest firing-rate threshold neuron. For the same value of firing rate of the low firing-rate threshold neuron (top, gray arrows), the high firing-rate threshold neuron has a relatively larger value following ON-direction saccades (bottom, gray arrows). Plots correspond to the 20th and 80th neurons in the network when neurons are ordered by firing-rate threshold. (B) Firing rates of the high firing-rate threshold neuron of panel A as a function of the low firing-rate threshold neuron’s firing rate following ON-direction (black) or OFF-direction (gray) saccades. Note the hysteretic relationship. (C) Qualitative explanation for ON–OFF hysteresis. Dendrites on the high firing-rate threshold neurons turn ‘on’ before those on the low firing-rate threshold neurons as the network activity increases (corresponding to firing rates and eye position increasing) and turn ‘off’ before those on the low firing-rate threshold neurons as the network activity decreases. This leads to relatively higher firing rates of the high firing-rate threshold neurons relative to the low firing-rate threshold neurons following ON-direction saccades and relatively lower rates following the OFF-direction saccades.

uous eye velocity command inputs involved in the vestibulo-ocular and optokinetic reflexes (VOR and OKR, respectively). Here, we consider the integration of eye movement commands that drive the VOR. The firing rates of vestibular neurons, which provide the input to the oculomotor neural integrator, primarily encode head velocity. These velocity signals are then converted by the neural integrator to eye position signals that are responsible for stabilizing gaze during head movements.

Dendritic bistability lends stability to eye fixations because it causes the system to be relatively insensitive to small perturbations. This suggests that the network might not be responsive to small velocity commands. Using sinusoidal head velocity input, we show that a network in the absence of noise is indeed insensitive to small command inputs but that this insensitivity may be largely overcome by the presence of noise in the neuron’s firing rates. When the network is mistuned, the noise causes significant drift in the VOR and a lesser drift in fixations. However, neither of these drifts is as large as would occur in the absence of hysteresis. Thus, hysteresis in the presence of noise appears to add robustness to fixations (although not as much robustness as in a network without noise) and a lesser amount of robustness to the VOR.

The effect of a nonzero command input $r_{com,i}$ is to shift the i^{th} hysteretic rectangle in the previous figures to the left by an amount $r_{com,i}/\zeta_i$. This can be seen from equation (5), in which $r_{com,i}$ appears in the argument of the bistable function h . For sufficiently large $r_{com,i}$, the band may no longer intersect the 45° line at the current value of \hat{E} . \hat{E} is then driven with a velocity determined from equation (5) by the difference between the 45° line and the near side of the band.

The response of the hysteretic network to continuously varying velocity commands is illustrated in Figure 7. We consider the network defined by equations (10–12) with a parallel-edge shaped band (Fig. 7A,B) that tolerates $\pm 10\%$ mistuning in the absence of noise. Head velocity commands to each neuron are modeled by identical slowly varying sinusoidal inputs $r_{com,i} = A \sin(2\pi ft)$ of amplitude A and frequency f (Fig. 7B,E,H top). In all simulations, $f = 0.1$ Hz, $\tau_{dend} = 100$ ms, and $\tau_{syn} = 5$ ms. Because $\tau_{dend} \gg \tau_{syn}$, the model behaves qualitatively like the analytic model (equations 2–4), and we define \hat{E} as in that model (equation 4).

In the absence of noise, the sinusoidal input periodically cycles through very small velocities that are insufficient to push the band far enough to be outside the line, i.e. to reach the ‘on’ or ‘off’ activation thresholds of any dendrites. Therefore, no dendrites are flipped ‘on’ or ‘off’ at these times, although dendrites that are flipped ‘on’ previous to these times continue to relax towards a steady state for a time τ_{dend} . Thus, if the input causes the line to remain in the band for times much greater than τ_{dend} , \hat{E} will remain essentially unchanged for these times (Fig. 7B, bottom, flat sections of gray lines). For sufficiently low amplitude inputs (Fig. 7B, top, dashed line), the input is never large enough to push the band outside the line and \hat{E} does not change at all (Fig. 7B, bottom, dashed gray line). It should be noted that, although \hat{E} does not change, the firing rates of the neurons (not shown) still do show some modulation that reflects the direct influence of the velocity input on the firing rate (equation 12).

When noise is included in the model, the insensitivity to small inputs can be overcome. To model the effects of noise, we add low-pass filtered Gaussian noise into each dendrite. This is

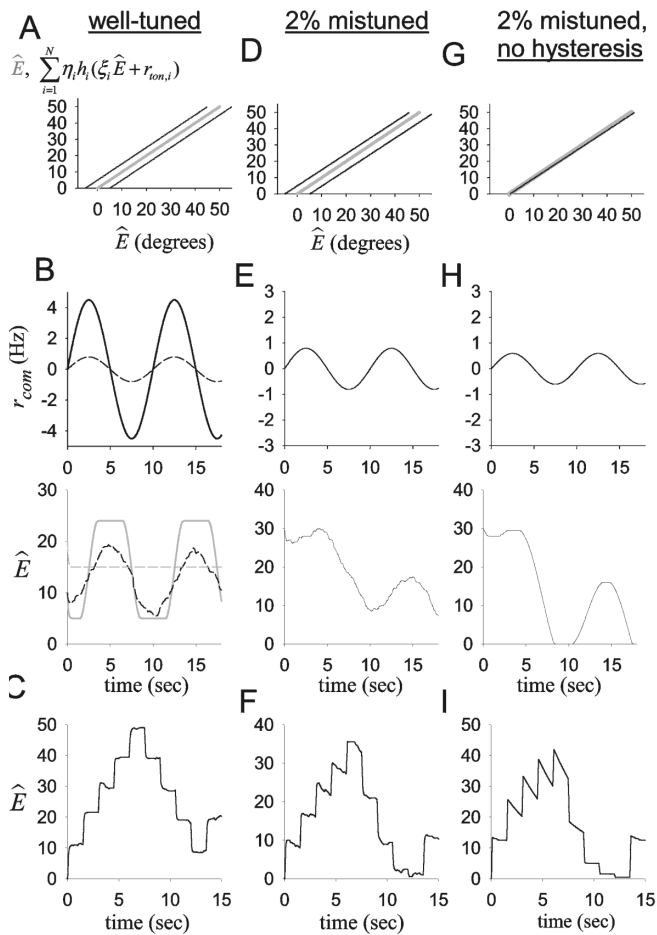


Figure 7. Performance of a network with dendritic hysteresis during the sinusoidal VOR. (A) Band shape for a well-tuned network with band centered about the line. (B) The network response to 0.1 Hz sinusoidal input without noise shows a truncated sinusoidal behavior for larger input ($A = 4.5$ Hz, solid lines) and is unresponsive for smaller input ($A = 0.8$ Hz, top, dashed line; bottom, gray dashed line). With noise and the smaller ($A = 0.8$ Hz) input, the network is responsive throughout the sinusoidal cycle (bottom, black dashed line). $\sigma = 4$ Hz in all panels that include noise. (C) Fixations for the same amount of noise as in B. (D,G) Band shape for a network mistuned towards decay by 2% for a band width that produces stable fixations for $\pm 10\%$ mistuning of weights with noise (D) or for a non-hysteretic band and no noise (G). (E,H) The network response is more sensitive to negative velocity commands than to positive commands, causing \hat{E} to decay. The decay is less for the band with hysteresis (E; $A = 0.8$ Hz) than for one without (H; $A = 0.6$ Hz). (F,I) Fixations with hysteresis and noise show some decay (F) but less than is shown by the VOR (E) and much less than would occur in the absence of hysteresis (I). Saccadic amplitudes were set to give approximately equal ranges of eye position. Flat portions of traces in H and I reflect the effects of discreteness.

accomplished by adding to each neuron's firing rate an independently generated Gaussian noise value every 1 ms. This firing rate noise is then filtered by the synaptic time constant $\tau_{\text{syn}} = 5$ ms so that each of the dendrites to which a neuron projects receives low-pass filtered Gaussian noise. Modeled in this manner, the noise acts as an external input that randomly jiggles each of the hysteretic rectangles in the positive feedback band left and right. If a noise fluctuation is larger than a minimum size equal to the width of the band, then it can cause \hat{E} to drift. The width σ of the Gaussian noise distribution was chosen so that

the characteristic size of the fluctuations in the noise input to each dendrite is slightly smaller than the width of the band, in particular just large enough to occasionally turn dendrites 'on' or 'off'. This causes a realistic amount of random drift in \hat{E} fixations (Fig. 7C). Considerably larger noise values cause much larger drift in fixations and considerably smaller values cause fixations to be perfectly stable as in the simulations without noise.

Adding even small amounts of command input on top of this noise causes the random drift to obtain a bias. In this way, the network can be driven (Fig. 7B, bottom, black dashed line) by small VOR inputs (Fig. 7B, top, dashed line) that, in the absence of noise, would not cause any change in \hat{E} (Fig. 7B, bottom, gray dashed line).

Both fixations and the VOR show some drift or inaccuracy, respectively, when noise is added to the model and the coupling strengths or tonic inputs are mistuned by amounts less than the width of the band. This is because, when the band is not centered about the 45° line, the response to positive and negative command inputs or noise fluctuations are unequal. For example, if the band is slightly mis-centered towards reduced positive feedback (Fig. 7D, 2% mistuning towards decay from perfect centering of the band), noise fluctuations are more likely to cause negative than positive eye drifts. In this way, unlike in the simulations with parallel-edge shaped band but no noise (Fig. 5C), small amounts of mistuning can lead to small amounts of drift. Nevertheless, the amount of drift is much smaller than would occur in the absence of hysteresis (compare to Fig. 7I, which shows \hat{E} for the same amount of mistuning in a model without hysteresis or noise; results are similar to Fig. 7I if noise is added in the absence of hysteresis).

Relative to fixations, the VOR is less robust because its response reflects both an asymmetry in the response to noise and an asymmetry in the response to the VOR command inputs. Again consider an integrator with feedback that is insufficient to center the band about the line (Fig. 7D). For sinusoidal command inputs (Fig. 7E) with the same noise level and same command input as was considered in the perfectly centered band of Figure 7B, the response of the network is now greater to negative inputs than to positive inputs. This asymmetry in response reflects that additional command inputs and noise are more likely to push the band outside the decay line in the negative-velocity direction than in the positive-velocity direction. Although these deviations are larger than those for the fixations, they are nevertheless somewhat smaller than would be the case for the same amount of mistuning in the absence of hysteresis (Fig. 7H, amplitude of input is scaled to make amplitude of response comparable to that in Fig. 7E).

Discussion

We have shown how a recurrent network model with multiple bistable dendrites per neuron can maintain stable persistent neural activity without precise tuning of model parameters and without large discontinuities in firing rate with small changes in \hat{E} . The robustness of stable fixations to perturbations in parameters arises because the bistable dendritic compartments are unresponsive to small changes in their inputs. The main purpose of including multiple dendritic subunits in this model is to reduce the large jumps in firing rate that would typically be associated with a somatic mechanism of bistability. Linear quasi-

continuous firing rate versus \hat{E} relationships can result from the successive recruitment of many dendritic compartments with increasing \hat{E} . In this manner, our model maintains the robustness that was attributable to bistability in other models (Camperi and Wang, 1998; Koullakov *et al.*, 2002) while distributing the jumps in firing rate across multiple compartments. Further, neuron pair firing rate versus firing rate relationships can be hysteretic in the manner seen in experiments (Aksay *et al.*, 2003) if the parameters of the bistable dendritic compartments are assumed to differ systematically with the neuron's firing rate threshold (Fig. 6).

We have not assumed any particular mechanism of dendritic bistability. Previous work has implicated, among others, the voltage-dependent properties of NMDA and voltage-sensitive Ca^{2+} channels as sources of dendritic bistability (Hounsgaard and Kiehn, 1993; Reuveni *et al.*, 1993; Yuste *et al.*, 1994; Lee and Heckman, 1998a,b; Schiller *et al.*, 2000). The graphical analysis of our model (Fig. 3) suggests that any mechanism that provides a hysteretic response to inputs can add robustness to fixations. The amount of robustness relates to the width of the hysteretic response. Robust quasi-continuous performance is not obtained if the dendrites in our model are made binary but not hysteretic. This implies that hysteresis, rather than the binary nature of the dendritic compartments, is the crucial feature of our model for robustness.

Specific mechanisms of achieving bistability might have different advantages when considered in detail. For example, we have found that the ability of the network to maintain stable fixations is not particularly sensitive to whether the dendrite or synapse provides the slow time constant underlying recurrent feedback. However, differences in synaptic dynamics can have notable differences on the exact form of ON-OFF hysteresis in the rate-rate relationships and in the amplitude of the response to command inputs or noise. Future models should also include both hysteretic and non-hysteretic components for the dendritic responses.

The weighted dendritic activations $W_{ij}D_{ij} = \zeta_i \eta_j D_j(r_j)$ in the analytic model were taken to be a perfect outer product. This approximation should be regarded as a computational convenience that simplifies the network behavior to a one-dimensional set of possible eye positions governed by equation 5 and makes the analysis of the model analytically tractable. More realistically, integrator networks display multi-dimensional behavior – the outer-product form can be considered a trick for reducing this more complicated behavior to a single 'integrating mode'. However, the outer-product form also implies that there is no hysteresis in the firing rate versus firing rate relationships during fixations. Relaxing this assumption in an appropriate manner can lead to more interesting effects as a result of corresponding dendrites on different neurons being recruited at different times. Such effects can include ON-OFF hysteresis in rate-rate relationships (Fig. 6) or other history-dependent trends in rate-rate relationships [such as the ON1-ON2 hysteresis observed in Aksay *et al.* (2003)]. However, it is not yet clear experimentally whether the ON-OFF hysteresis or other history-dependent effects are due to an actual bistable hysteretic element as was used in this model or are due to more complicated dynamics that may be present in multi-dimensional models [for more on this issue, see the discussion in Aksay *et al.* (2003)].

The outer product form of the analytic model also implies that all of the dendrites to which a neuron projects turn 'on' and 'off' simultaneously. This leads to N groups of identically behaving

dendritic compartments, and N stable eye positions. If, however, these dendrites have a staggered recruitment order then, in principal, the number of states of the network could be much larger than N . If there is additionally some graded component to the neuron's input, rather than the binary steady-state dendritic activation assumed for simplicity here, then each neuron could additionally have many more than the M states corresponding to the number M of independent dendritic subunits. Although it is not presently known how many independent dendritic subunits may be contained in a single integrator neuron, modeling work (Koch *et al.*, 1982; Poirazi *et al.*, 2003a,b) suggests that neurons in other systems may contain as many as ~ 30 approximately independent subunits.

We have assumed that the dendritic compartments act independently of each other and of the somatic voltage. This is consistent with the studies of Poirazi *et al.* (Poirazi *et al.*, 2003a,b), which find that a neuron's firing rate can be well-approximated by a sum of independent dendritic contributions followed by a static nonlinearity. We have found that isolation of the dendritic voltage from the somatic voltage can be achieved if the dendritic tree is large compared with the soma and the locus of plateau generation is distal, so that somatic currents are effectively shunted (Goldman *et al.*, 2002b). If the dendritic bistability is mediated by the NMDA receptor, the dependence of NMDA activation on presynaptic activity could provide additional isolation from the soma (Lisman *et al.*, 1998; Schiller and Schiller, 2001; Koullakov *et al.*, 2002).

The strongest predictions of this model relate to the VOR. Dendritic bistability provides robustness for fixations by being unresponsive to small inputs. This implies that there is a threshold value below which vestibular input cannot activate the VOR and below which noise will not affect the activation state of the dendrites. We have shown that when the amplitude of noise is close to this threshold value, the network becomes responsive to any additional inputs such as head velocity commands associated with the VOR. An implication of this modeling is that there is a tradeoff: with little noise, the network is insensitive to small vestibular inputs; with more noise, the fixations show random drift and become more sensitive to mistuning (Fig. 7F). This result suggests that the amplitude of the noise should assume an intermediate value roughly equal (in appropriate units) to the size of the hysteresis band. If this is the case, one might be able to indirectly estimate the width of the hysteresis band in a given system by measuring the noise in the firing rates and synaptic transmission properties.

Even with noise, the sensitivity to mistuning for fixations is still far less than would occur in the absence of hysteresis (see Fig. 7D). For the VOR, the hysteresis band also contributes some robustness (Figs 7E,H) although the robustness of the VOR is less than that of the fixations. With less noise (not shown), the difference between the VOR and fixation performances are even larger – this can be understood easily by considering the limit of no noise, for which the fixations are stable so long as the band encloses the 45° decay line but the response to the VOR is nevertheless asymmetric when the command inputs move the band far enough that it no longer encloses the decay line. Previous work (Goldman *et al.*, 2002a) comparing the fixation and VOR time constants following a cumulative series of permanent lesions found that the fixation time constant did drop less than the VOR time constant following the early lesions, consistent with this prediction, but the data is quite noisy.

For working memory systems that do not continuously integrate inputs, constant-rate persistent neural activity might be the only relevant neuronal response. Dendritic bistability provides a biophysically plausible mechanism for stably maintaining constant-rate activity in the absence of external input. Thus, this work more generally proposes a hypothesis for how memory ‘fixations’ may be held robustly by neurons with discernibly continuous firing rate responses.

Notes

This work was supported by HHMI, NIH EY06558 and RR00166, and NSF 9986 022. We thank Ila Fiete for helpful discussions and Emre Aksay for helpful discussions and critical reading of the manuscript.

Address correspondence to Mark Goldman, Department of Physics, Wellesley College, 106 Central St., Wellesley, MA 02481, USA. Email: mgoldma1@wellesley.edu.

Appendix: Model Analysis and Solution

Conditions for Uniqueness of the Positive Feedback Band Shape

In the main text, we graphically represent the positive feedback term

$$\sum_{i=1}^N \eta_i h_i(\xi_i \hat{E} + r_{\text{ton},i} + r_{\text{com},i})$$

of equation (5) as a static ‘band-like’ function of \hat{E} (Figs 3–6). Given this band, the eye velocity is determined as the difference between this band and the 45° line representing \hat{E} . In general, the shape of this band is not static and depends on the relative strengths and histories of the inputs to the integrator. Here, we explain how this history-dependence arises and derive conditions under which the positive feedback can be represented by a static, history-independent band.

The essential assumption that allows the consideration of static bands in the main text is that the order of dendrites turning ‘on’ is history-independent and exactly the reverse of the order of dendrites turning ‘off’. This is certainly the case under two conditions: (1) the ordering of ‘on’ thresholds is the same as the ordering of the ‘off’ thresholds. That is, if $r_{\text{on},i} < r_{\text{on},j}$ then $r_{\text{off},i} < r_{\text{off},j}$. (2) The inputs are arranged such that dendrite i always turns on before dendrite $i + 1$ and turns ‘off’ after dendrite $i + 1$. This condition would be violated if, for example, a large input to dendrite i caused it to turn on before dendrite $i - 1$. Condition 2 can be guaranteed so long as positive inputs to dendrites with lower ‘on’ thresholds are at least as large as those to dendrites with higher ‘on’ thresholds, and negative inputs to dendrites with higher ‘off’ thresholds are at least as large in magnitude as negative inputs to dendrites with lower ‘off’ thresholds.

We illustrate these conditions by examining the possible states of each of the dendrites that make up the band as a function of \hat{E} . In the absence of external inputs ($r_{\text{com},i} = 0$ for all i), $h_i(r_i)$ is determined by the value of its input r_i relative to its ‘on’ and ‘off’ thresholds $r_{\text{on},i}$ and $r_{\text{off},i}$. From equation 3, $h_i(r_i)$ equals:

$$\begin{aligned} 0 \text{ (off)} & \quad \text{if } \xi_i \hat{E} + r_{\text{ton},i} \leq r_{\text{off},i} \\ 1 \text{ (on)} & \quad \text{if } \xi_i \hat{E} + r_{\text{ton},i} \geq r_{\text{on},i} \\ 0 \text{ or } 1 \text{ (history-dependent)} & \quad \text{if } r_{\text{off},i} < \xi_i \hat{E} + r_{\text{ton},i} < r_{\text{on},i} \end{aligned} \quad (13)$$

where by ‘history-dependent’ it is meant that the dendrite could be either ‘on’ or ‘off’ depending on the previous history of the input.

Possible band structures that violate assumptions 1 and 2 above are depicted in Figure 8. Assumption 1 can be violated when two dendrites do not obey the same recruitment order for increasing and decreasing eye positions (Fig. 8A, top two dendritic rectangles). One dendrite (the ‘wider’ dendrite) has both a greater ‘on’ threshold as a function of \hat{E} and a lower ‘off’ threshold. For increasing eye positions, the ‘wider’ dendrite turns ‘on’ after the ‘narrower’ dendrite (Fig. 8A, top). For decreasing eye positions, the ‘wider’ dendrite remains ‘on’ after the ‘narrower’ dendrite turns ‘off’ (Fig. 8B, bottom). Thus, the band shape differs depending on whether the eye position is increasing or decreasing.




Figure 8. Sources of non-uniqueness in the band structure. (A) When some dendritic hysteretic rectangles are wider than others (top two rectangles), the order of turning ‘on’ dendrites can be different from the order of turning ‘off’ dendrites. This can result in a different band structure for a given value of \hat{E} (dash on y-axis) depending on whether the input arrived at this value from a lower value of \hat{E} (top panel) or from a higher value (bottom panel). (B) When the input to a higher activation-threshold dendrite is greater than that to a lower activation-threshold dendrite (middle neurons of band), the higher activation-threshold dendrite may turn ‘on’ before the lower activation-threshold dendrite. Analogous behavior can occur for dendrites turning ‘off’.

Even when the ‘on’ and ‘off’ thresholds have the same order (assumption 1), the neurons in the ‘history-dependent’ state (equation 13) can be recruited in a different order from these thresholds (violating assumption 2). Figure 8B shows a possible band structure resulting from a dendrite with a higher ‘on’ threshold receiving a larger input than a dendrite with a lower ‘on’ threshold. The band structure reflects that any dendrite with bistable region straddling \hat{E} could be either ‘on’ or ‘off’ depending on the history of the inputs. Of course, if the hysteretic rectangles are of different heights then there may be only one arrangement of inputs that strictly leads to a particular value of \hat{E} , but for many thousands of dendrites and realistic amounts of noise, such precision would not be detectable.

Derivation of Tolerance to Mistuning for a Network with Parallel-stripe Shaped Band

We here derive equation (6) for the tolerance to mistuning of the coupling strengths W in a network with the parallel-edge shaped band defined by $r_{\text{on},i} = r_{\text{on}}$, $r_{\text{off},i} = r_{\text{off}}$, $W_{ij} = W = \zeta \eta$, and equally spaced tonic inputs

$$r_{\text{ton},i} = \frac{N - i + 0.5}{N} \bar{r}$$

where $\bar{r} = (r_{\text{on}} + r_{\text{off}})/2$.

For an eye position $\hat{E} = m\eta$ corresponding to m 'on' dendrites to be stably maintained in the absence of external inputs, the presynaptic rates of all neurons must satisfy:

$$\begin{aligned} Wm + r_{\text{ton},m+1} &< r_{\text{on}} \\ Wm + r_{\text{ton},m} &> r_{\text{off}} \end{aligned} \quad (14)$$

The first inequality gives the condition that the network feedback must be sufficiently small that it does not flip 'on' the $(m+1)$ th dendrite. The second inequality gives the condition that the network feedback must be large enough to not flip 'off' the m th dendrite. These conditions are most strict for $m = N - 1$ in the first inequality and $m = N$ in the second inequality. Therefore, the range of W s that will guarantee stable fixations at all values of \hat{E} are

$$\frac{r_{\text{off}} - r_{\text{ton},N}}{N} < W < \frac{r_{\text{on}} - r_{\text{ton},N}}{N-1} \quad (15)$$

where we have substituted in the values of $r_{\text{ton},i}$ in deriving this relationship. Finally, defining W^* to be the midpoint of the above range of W s and ΔW to be the width of the allowed range, we get after some manipulation that the fractional tolerance of the network to perturbations, $\Delta W/W^*$, is given by

$$\frac{\Delta W}{W^*} = \frac{r_{\text{on}} - r_{\text{off}}}{\bar{r}} + \frac{(2r_{\text{on}} - \bar{r})r_{\text{off}} + r_{\text{on}}\bar{r}\left(1 - \frac{1}{N}\right)}{\bar{r}\left[\left(2N - 1 + \frac{1}{2N}\right)\bar{r} - r_{\text{off}}\right]} \quad (16)$$

Derivation of Tolerance to Mistuning for a Network with Cone-shaped Band

We derive the tolerance to mistuning of coupling strengths for a network with the cone-shaped band pictured in Figure 5E (with constant parameters $r_{\text{on},i} = r_{\text{on}}$, $r_{\text{off},i} = r_{\text{off}}$, $\eta_i = \eta$, and $r_{\text{ton},i} = r_{\text{ton}}$, and with slopes ζ_i staggered according to the relation

$$\zeta_i = \frac{\bar{r} - r_{\text{ton}}}{i\eta}$$

where $\bar{r} = (r_{\text{on}} + r_{\text{off}})/2$.

For an eye position $\hat{E} = m\eta$ corresponding to m 'on' dendrites to be stably maintained in the absence of external inputs, the presynaptic rates of all neurons must satisfy:

$$\begin{aligned} W_{m+1}m + r_{\text{ton}} &< r_{\text{on}} \\ W_m m + r_{\text{ton}} &> r_{\text{off}} \end{aligned} \quad (17)$$

where $W_m = \zeta_m \eta$. In the limit that the number of neurons $N \rightarrow \infty$ so that the weights form a continuum, we may approximate $W_m \approx W_{m+1}$. Then equation (17) reduces to

$$\frac{r_{\text{off}} - r_{\text{ton}}}{m} < W_m < \frac{r_{\text{on}} - r_{\text{ton}}}{m} \quad (18)$$

Defining W_m^* to be the midpoint of this range and ΔW_m to be the width of this range, the fractional tolerance to mistuning $\Delta W_m/W_m^*$ is independent of m and equals the expression given in equation (7).

References

Aksay E, Baker R, Seung HS, Tank DW (2000) Anatomy and discharge properties of pre-motor neurons in the goldfish medulla that have eye-position signals during fixations. *J Neurophysiol* 84:1035-1049.
Aksay E, Gamkrelidze G, Seung HS, Baker R, Tank DW (2001) In vivo intracellular recording and perturbation of persistent activity in a neural integrator. *Nat Neurosci* 4:184-193.
Aksay E, Major G, Goldman MS, Baker R, Seung HS, Tank DW (2003) History dependence of rate covariation between neurons during persistent activity in an oculomotor integrator. *Cereb Cortex* 13:1173-1184.

Camperi M, Wang XJ (1998) A model of visuospatial working memory in prefrontal cortex: recurrent network and cellular bistability. *J Comput Neurosci* 5:383-405.
Cannon SC, Robinson DA (1985) An improved neural-network model for the neural integrator of the oculomotor system: more realistic neuron behavior. *Biol Cybern* 53:93-108.
Cannon SC, Robinson DA, Shamma S (1983) A proposed neural network for the integrator of the oculomotor system. *Biol Cybern* 49:127-136.
Goldman MS, Kaneko CR, Major G, Aksay E, Tank DW, Seung HS (2002a) Linear regression of eye velocity on eye position and head velocity suggests a common oculomotor neural integrator. *J Neurophysiol* 88:659-665.
Goldman MS, Levine JH, Major G, Aksay E, Tank DW, Seung HS (2002b) Dendritic bistability increases the robustness of persistent neural activity in a model oculomotor neural integrator. *Soc Neurosci Abstr* 28:266.14.
Hounsgaard J, Kiehn O (1993) Calcium spikes and calcium plateaus evoked by differential polarization in dendrites of turtle motoneurons *in vitro*. *J Physiol (Lond)* 468:245-259.
Kamath BY, Keller EL (1976) A neurological integrator for the oculomotor control system. *Math Biosci* 30:341-352.
Koch C, Poggio T, Torre V (1982) Retinal ganglion cells: a functional interpretation of dendritic morphology. *Philos Trans R Soc Lond B Biol Sci* 298:227-263.
Koulakov AA, Raghavachari S, Kepecs A, Lisman JE (2002) Model for a robust neural integrator. *Nat Neurosci* 5:775-782.
Lee RH, Heckman CJ (1998a) Bistability in spinal motoneurons in vivo: systematic variations in persistent inward currents. *J Neurophysiol* 80:583-593.
Lee RH, Heckman CJ (1998b) Bistability in spinal motoneurons in vivo: systematic variations in rhythmic firing patterns. *J Neurophysiol* 80:572-582.
Lisman JE, Fellous JM, Wang XJ (1998) A role for NMDA-receptor channels in working memory. *Nat Neurosci* 1:273-275.
Pastor AM, De la Cruz RR, Baker R (1994) Eye position and eye velocity integrators reside in separate brainstem nuclei. *Proc Natl Acad Sci USA* 91:807-811.
Poirazi P, Brannon T, Mel BW (2003a) Arithmetic of subthreshold synaptic summation in a model CA1 pyramidal cell. *Neuron* 37:977-987.
Poirazi P, Brannon T, Mel BW (2003b) Pyramidal neuron as two-layer neural network. *Neuron* 37:989-999.
Pouget A, Latham P (2002) Digitized neural networks: long-term stability from forgetful neurons. *Nat Neurosci* 5:709-710.
Reuveni I, Friedman A, Amitai Y, Gutnick MJ (1993) Stepwise repolarization from Ca^{2+} plateaus in neocortical pyramidal cells: evidence for nonhomogeneous distribution of HVA Ca^{2+} channels in dendrites. *J Neurosci* 13:4609-4621.
Robinson DA (1981) Control of eye movements. In: *Handbook of physiology. The nervous system. visual motor control* (Desmedt JE, ed.), pp. 1275-1320. Bethesda, MD: American Physiology Society.
Rosen MJ (1972) A theoretical neural integrator. *IEEE Trans Biomed Eng* 19:362-367.
Schiller J, Schiller Y (2001) NMDA receptor-mediated dendritic spikes and coincident signal amplification. *Curr Opin Neurobiol* 11:343-348.
Schiller J, Major G, Koester HJ, Schiller Y (2000) NMDA spikes in basal dendrites of cortical pyramidal neurons. *Nature* 404:285-289.
Seung HS (1996) How the brain keeps the eyes still. *Proc Natl Acad Sci USA* 93:13339-13344.
Seung HS, Lee DD, Reis BY, Tank DW (2000) Stability of the memory of eye position in a recurrent network of conductance-based model neurons. *Neuron* 26:259-271.
Yuste R, Gutnick MJ, Saar D, Delaney KR, Tank DW (1994) Ca^{2+} accumulations in dendrites of neocortical pyramidal neurons: an apical band and evidence for two functional compartments. *Neuron* 13:23-43.