

Sparse cerebellar innervation can morph the dynamics of a model oculomotor neural integrator

Thomas J. Anastasio · Yash P. Gad

Received: 19 June 2006 / Revised: 2 October 2006 / Accepted: 6 October 2006 / Published online: 4 November 2006
© Springer Science + Business Media, LLC 2007

Abstract The oculomotor integrator is a brainstem neural network that converts velocity signals into the position commands necessary for eye-movement control. The cerebellum can independently adjust the amplitude of eye-movement commands and the temporal characteristics of neural integration, but the percentage of integrator neurons that receive cerebellar input is very small. Adaptive dynamic systems models, configured using the genetic algorithm, show how sparse cerebellar inputs could morph the dynamics of the oculomotor integrator and independently adjust its overall response amplitude and time course. Dynamic morphing involves an interplay of opposites, in which some model Purkinje cells exert positive feedback on the network, while others exert negative feedback. Positive feedback can be increased to prolong the integrator time course at virtually any level of negative feedback. The more these two influences oppose each other, the larger become the response amplitudes of the individual units and of the overall integrator network.

Keywords Motor control · Oculomotor system · Vestibular system · Eye movements · Cerebellum · Dynamic systems · Adaptive systems · Genetic algorithm

Action Editor: Jonathan D. Victor

T. J. Anastasio (✉)
Department of Molecular and Integrative Physiology, Beckman
Institute for Advanced Science and Technology, University of
Illinois at Urbana-Champaign,
Urbana, IL 61801, USA
e-mail: tja@uiuc.edu

Y. P. Gad
Center for Biophysics and Computational Biology, Beckman
Institute for Advanced Science and Technology, University of
Illinois at Urbana-Champaign,
Urbana, IL 61801 USA

Introduction

The cerebellum powerfully regulates oculomotor subsystems, despite its sparse innervation of the brainstem neural networks that control eye movements. The mechanism of this sparse yet powerful regulation remains an enigma. The oculomotor integrator is a brainstem neural network that converts eye-velocity signals to eye-position commands (Robinson, 1989a). It is essential for eye-movement control by oculomotor subsystems such as the vestibulo-ocular reflex (VOR) (Robinson, 1989b). We used a combination of dynamic systems modeling and machine learning techniques to show how the cerebellum could regulate the neural integrator despite sparse cerebellar innervation of brainstem neurons.

The VOR provides a well-studied test bed for evaluating models of cerebellar regulation of the oculomotor neural integrator. The function of the VOR is to maintain retinal image stability by producing eye rotations that counterbalance head rotations (Wilson and Melville Jones, 1979). The brainstem VOR pathway is driven by head velocity signals from vestibular receptors, which are converted to position commands by the oculomotor integrator (Robinson, 1989b). The neural integrator is leaky (imperfect), and decays with a time constant of about 20 s in primates (Robinson, 1989a).

In mammals, the cerebellum is necessary for regulating both VOR gain and the integrator time constant. Lesions of the flocculus and ventral paraflocculus (referred to as the floccular complex) abolish VOR plasticity (Robinson, 1976; Nagao, 1983; Godaux and Vanderkelen, 1984; Rambold et al., 2002; Nagao and Kitazawa, 2003) and reduce the integrator time constant to a baseline value (Robinson, 1974; Zee et al., 1981; Chelazzi et al., 1990). Mice with genetic abnormalities of the cerebellum also show integrator deficits (Stahl and James, 2005) and reduced VOR plasticity

(Stahl, 2004; Katoh et al., 2005). With a normal, intact cerebellum, VOR gain and integrator time constant are independently adjustable (Tiliket et al., 1994). Because the position component dominates velocity in the VOR (Cheron et al., 1986; Cannon and Robinson, 1987), a significant proportion of VOR gain adjustment could be due to adjustment of the integrated component alone. Thus, integrator regulation is critical.

In mammals, both the VOR and the integrator are mediated by neurons in the medial vestibular (MVN) and prepositus hypoglossi (NPH) nuclei in the brainstem. Many neurons in the MVN/NPH carry eye position signals (McFarland and Fuchs, 1992; Stahl and Simpson, 1995; Cheron et al., 1996), and lesions of the MVN/NPH cause profound integrator deficits (Cheron et al., 1986; Cannon and Robinson, 1987). Estimates of the percentage of these neurons that receive inputs from floccular Purkinje cells range from 20% down to 1% (Babalian and Vidal, 2000; Sekirnjak et al., 2003). Thus, innervation of brainstem integrator neurons by the cerebellar flocculus is sparse.

Previous models of the mammalian neural integrator focused only on the brainstem (Cannon et al., 1983; Arnold and Robinson, 1991; Anastasio, 1998) and did not consider the effects of the cerebellar contribution. No previous modeling studies have addressed the issue of the sparseness of cerebellar innervation of brainstem neurons. Our model demonstrates how the cerebellum could independently regulate the gain and time constant of the mammalian oculomotor integrator, despite sparse innervation of MVN/NPH neurons by floccular Purkinje cells.

Methods

Our model is composed of Vunits, which represent neurons in the vestibular and prepositus nuclei, and Pcells, which represent Purkinje cells from the floccular complex. We constrained the general structure of the model according to neuroanatomical findings (Fig. 1), and ensured that the innervation of Vunits by Pcells was sparse. We preset the connectivity between Vunits in various patterns, and adjusted the connection weights so that integration by the Vunits alone (without Pcell connections) was very leaky. We then used the genetic algorithm to find specific connectivity patterns between the Vunits and Pcells that allow the Pcells, through adaptive changes in their connection weights, to regulate integrator time course and response amplitude independently. We used an efficient grid search to adjust connection weights during simulated evolution. We then used reinforcement learning in evolved networks to demonstrate that independent regulation by Pcells of integrator time course and response amplitude can occur in a neurobiologically plausible manner. Our model is linear, and we used linear systems techniques to analyze its dynamic behavior.

Network structure

Previous models of the mammalian neural integrator focused only on the connectivity among brainstem MVN/NPH neurons (Cannon et al., 1983; Robinson, 1989a; Arnold and Robinson, 1991; Anastasio, 1998). We extend these models by including connections with floccular Purkinje cells (Langer et al., 1985a,b; Büttner-Ennever, 1988; Epema et al., 1990; Tan and Gerrits, 1992). The brainstem portion of our model is similar to previous models of the mammalian integrator in that Vunits are arranged on either side of a bilateral network, and connected over the midline by reciprocal inhibition (Cannon et al., 1983; Robinson, 1989a; Arnold and Robinson, 1991; Anastasio, 1998). This reciprocal interaction is mediated by the commissural fibers that interconnect the MVN/NPH bilaterally (McCrea et al., 1987; Büttner-Ennever, 1988). Cutting these fibers produces severe integrator deficits (Anastasio and Robinson, 1991; Arnold and Robinson, 1997). The brainstem portion of our model varies both in the number of Vunits and in their connectivity patterns.

Our bilateral networks have equal numbers of Vunits on each side. Networks have either 20 Vunits (10 per side) or 400 Vunits (200 per side). Each Vunit on one side inhibits the Vunit directly opposite, and may also inhibit a fixed number of the neighbors of that Vunit. Networks of 20 Vunits have either uniform or local patterns of reciprocal inhibitory connectivity between Vunits. In uniform networks, each Vunit on one side inhibits all Vunits on the opposite side. Local 20-Vunit networks can have closed or open boundaries. Networks with closed boundaries are arranged as rings, so that Vunit-1 and Vunit-10 on each side are neighbors. Networks with open boundaries are arranged as a bar (Fig. 1(c)), so that Vunit-1 and Vunit-10 are not neighbors, and the Vunits near the ends do not have the full complement of connections. Locally connected 20-Vunit networks have neighborhood sizes from zero to four. The neighborhood extends along the Vunit array in both directions. For example, the local 20-Vunit network of Fig. 1(c) has a neighborhood size of two, so each Vunit on one side inhibits the Vunit directly opposite, as well as the two neighbors above and below that Vunit in the array. Four is the largest neighborhood size beyond which a closed 20-Vunit network becomes a uniform network. The computational complexity of the 400-Vunit network prohibited exploration of different Vunit connectivity patterns. Vunit connectivity in the 400-Vunit network is set arbitrarily to have a neighborhood size of ten and open boundaries.

The floccular complex receives heavy input bilaterally from brainstem vestibular neurons, but floccular Purkinje cells inhibit vestibular neurons on the same (ipsilateral) side only (Langer et al., 1985a,b; Büttner-Ennever, 1988; Epema et al., 1990; Tan and Gerrits, 1992). Our basic network structure respects this constraint (Fig. 1(c)). Pcells inhibit Vunits




Fig. 1 Patterns of innervation of neural integrator networks by cerebellum. The bilaterally symmetric networks are composed of brainstem vestibular neurons (Vunits, circles) and floccular Purkinje cells (P or Pcells, triangles). Vunit to Pcell (Vunit-Pcell) connections are excitatory (forks); all other connections are inhibitory (dots). For simplicity, we omit cell types such as inhibitory brainstem interneurons and cerebellar granule cells. Flocculus target neurons (FTNs) are Vunits that receive input from a Pcell. Pcells can receive from Vunits (including FTNs) bilaterally, but project to FTNs only ipsilaterally. (a) In the simplest case, there are only two Vunits (both FTNs) and two Pcells. The feedback

loop through the Pcells potentially augments the reciprocal feedback between the Vunits. Adjusting the strength of the Pcell feedback loop (which we model simply by tuning the weights from Pcells to FTNs) can adjust the integrator time constant, but independent control of time constant and gain is not possible. (b) Two Pcells (one per side) can also adjust the time constant in larger networks, but independent control of time constant and gain is again not possible. (c) With four Pcells (two per side) independent control of time constant and gain is possible. Only some of the potential Vunit-Pcell connections are shown in (b) and (c)

on the same side only, while Vunits can project to Pcells on either side. Vunits that receive input from Pcells are called flocculus target neurons (FTNs) (Lisberger and Pavelko, 1988). All networks have four Pcells (two per side), which is the minimum number sufficient for dynamic morphing (see Results).

To enforce the sparseness constraint, each Pcell contacts only one Vunit (its FTN) on the same side. Because networks have either 20 Vunits (10 per side) or 400 Vunits (200 per side), Pcell innervation of Vunits is 20% and 1% in 20- and 400-Vunit networks, respectively. These percentage innervations span the range as estimated experimentally (Babalian and Vidal, 2000; Sekirnjak et al., 2003). Except that they receive input from Pcells, FTNs are treated in the same way as other Vunits. All Vunits, including FTNs, contribute to the output, and any Vunit, including FTNs, can potentially connect to any Pcell. For simplicity, all nonzero Vunit-Pcell connection weights take the same value, which is 1 in 20-Vunit networks, and 0.005 in 400-Vunit networks.

Even with these neuroanatomical constraints, there are a huge number of different network configurations. We in-

troduced several simplifications to reduce this huge number. Neither Vunits nor Pcells project to themselves. Pcells do not interconnect, whether ipsilaterally or contralaterally. Vunits receive input from outside the network, but Pcells do not. The main simplification is bilateral symmetry. The inhibitory commissural connections between Vunits on either side are symmetrical. The connections among Vunits and Pcells are the same on both sides. Vunit to Pcell connections are specified by four sets of connections, which are the same for Pcells on both sides: ipsilateral Vunit to Pcell-1, ipsilateral Vunit to Pcell-2, contralateral Vunit to Pcell-1, and contralateral Vunit to Pcell-2. FTN locations, and the single connection from Pcell-1 or Pcell-2 to its FTN, are the same on both sides.

There is an important distinction between network symmetry and bilateral symmetry. The integrator model with Pcell input is bilaterally symmetric because the connections between Vunits, the connections from Vunits to Pcells, and the connections from Pcells to Vunits are the same on both sides. The network itself is asymmetric because Pcells project to Vunits on the same side only, but Vunits can project

to Pcells on both sides. Network asymmetry is essential for dynamic morphing (see Results and Discussion). The bilateral symmetry constraint, which we impose to make genetic search more tractable, is not necessary for dynamic morphing. We also demonstrate dynamic morphing in networks that have randomized Vunit interconnections and bilaterally asymmetric FTN locations (see Results).

Network dynamics

Brainstem MVN/NPH neurons (Vunits) and floccular Purkinje cells (Pcells) are modeled as linear, first-order units. A network of such units can be described in state space form and analyzed in terms of its eigenvectors and eigenvalues. Eigenvectors determine the patterns of response amplitude among the units in the network, while the eigenvalues determine the time course over which their associated patterns will persist. This time course is expressed most naturally as a time constant, which is the opposite reciprocal of the real part of an eigenvalue. The dominant eigenvector is associated with the eigenvalue that corresponds to the longest time constant. Eigenvalues and eigenvectors can be real or complex. Complex eigenvalues (and eigenvectors) occur in complex conjugate pairs. These complex modes produce oscillatory responses, where the frequency of oscillation in radians/second is equal to the absolute value of the imaginary part of the complex eigenvalues (Luenberger, 1979).

Use of linear Vunits is justified because the firing rates of MVN/NPH neurons are linearly related to eye position and velocity over the oculomotor range (McFarland and Fuchs, 1992; Stahl and Simpson, 1995; Cheron et al., 1996). While the relationship between firing rate and eye-movement parameters is piecewise-linear (different slopes for ipsilateral and contralateral eye movements) for most floccular Purkinje cells (Miles et al., 1980a; Lisberger et al., 1994b), we approximate Pcell behavior as linear in this initial model. Vunit and Pcell dynamics are described by

$$\tau \dot{y}(t) + y(t) = vx(t) \quad (1)$$

where $y(t)$ represents the firing rate of the unit, $x(t)$ is its input, v its sensitivity, and τ is the single-neuron time constant. A single neuron acts as a leaky integrator, but its time constant is too short. The single-neuron time constant can be lengthened by positive feedback, as occurs in neural network models of the integrator (Cannon et al., 1983; Robinson, 1989a; Arnold and Robinson, 1991; Anastasio, 1998). Recent mouse brain slice experiments reveal that MVN/NPH neurons have single-neuron time constants near 10 ms (Sekirnjak and du Lac, 2006). In order to demonstrate the robustness of integrator models, the single-neuron time constant τ is traditionally set to the low value of 5 ms (Robinson, 1989a). We set τ to 5 ms to demonstrate the robustness of our model.

The network of linear, first-order Vunits and Pcells with single-neuron time constant τ can be described in matrix notation as:

$$\tau \dot{\mathbf{y}}(t) + \mathbf{y}(t) = \mathbf{W}\mathbf{y}(t) + \mathbf{v}x(t) \quad (2)$$

where \mathbf{W} is a matrix of weights interconnecting the units, and \mathbf{v} is a vector of input weights. The structures of \mathbf{W} and \mathbf{v} conform to the network structure described in the previous subsection. For simplicity, all reciprocal connections between Vunits have the same weight. In the absence of the cerebellum in mammals, the integrator time constant is reduced to 1 s or less (Robinson, 1974; Zee et al., 1981; Chelazzi et al., 1990). To demonstrate the robustness of the model, we set the time constant of the integrator network without Pcells to the low value of 0.2 s. We used a grid search (see below) to set the common weight value of the reciprocal connections between Vunits to produce the 0.2 s time constant.

The network can be represented in state space form (Luenberger, 1979) as:

$$\dot{\mathbf{y}}(t) = \mathbf{A}\mathbf{y}(t) + \mathbf{b}x(t), \quad \mathbf{A} = 1/\tau (\mathbf{W} - \mathbf{I}), \quad \mathbf{b} = 1/\tau (\mathbf{v}) \quad (3)$$

where \mathbf{I} is the identity matrix and τ is the single-neuron time constant. For simplicity, we do not scale input vector \mathbf{b} , so \mathbf{b} simply equals \mathbf{v} (Eq. (2)). The input $x(t)$ is a scalar function of time that is transmitted to all Vunits by input vector \mathbf{b} . When $x(t)$ is an impulse, the response of Eq. (3) is:

$$\mathbf{y}(t) = \sum_{i=1}^n \alpha_i \mathbf{e}_i e^{\lambda_i t}, \quad \alpha_i = \mathbf{f}_i^T \mathbf{b} / \mathbf{f}_i^T \mathbf{e}_i \quad (4)$$

where n is the number of units in the network. The λ_i are the eigenvalues, and the \mathbf{e}_i and \mathbf{f}_i are the associated eigenvectors and adjoint eigenvectors of the network, which satisfy:

$$\mathbf{A}\mathbf{e}_i = \lambda_i \mathbf{e}_i \quad \text{and} \quad \mathbf{f}_i^T \mathbf{A} = \lambda_i \mathbf{f}_i^T \quad (5)$$

The impulse response completely characterizes the dynamic behavior of the units. The amplitude of its impulse response is equivalent to the response gain of each unit.

Reciprocal inhibition in the model produces net positive feedback, which lengthens the time constants of the unit responses, but also rejects input spontaneous activity, and so integrates only differential (push-pull) input signals. The input vector \mathbf{b} (Eq. (3)) is set to excite left-side Vunits with weight 1, and to inhibit right-side Vunits with weight -1 (Fig. 3(a)). Vector \mathbf{b} thus provides a push-pull input to the Vunits, and the Vunit responses are therefore push-pull also. We compute the eye position command from the network

(Vcommand) as the sum of the Vunit responses on one side, minus the sum of the Vunit responses on the other side.

The long-duration behavior of the integrator network is determined by the eigenvalue (λ_1) and eigenvector (\mathbf{e}_1) associated with the longest time constant (τ_1). The dominant time constant τ_1 equals $-1/\lambda_1$ when λ_1 is real. It is necessary to scale dominant eigenvector \mathbf{e}_1 by α_1 (Eq. (4)). The gain (g) of the integrator (Vcommand gain) will mostly reflect the contribution due to the scaled dominant eigenvector ($\alpha_1\mathbf{e}_1$). We closely approximate g by finding the difference between the sums of the values of the elements of $\alpha_1\mathbf{e}_1$ that correspond to the Vunits on either side of the network. If there are n_v Vunits, with half on either side, then integrator (Vcommand) gain g can be computed as:

$$g = \alpha_1 \left(\sum_{j=1}^{n_v/2} \mathbf{e}_{1,j} - \sum_{j=(n_v/2)+1}^{n_v} \mathbf{e}_{1,j} \right) \quad (6)$$

Grid search

We used a grid search (Bevington, 1969) in bilaterally symmetric networks to tune the common value of the Vunit reciprocally inhibitory interconnection weights, and the Pcell to FTN weights (which we refer to simply as Pcell weights). To set the time constant of the brainstem portion of the model (without Pcells) using a grid search, we increment the absolute value of the Vunit interconnection weights in fixed steps until the desired time constant of 0.2 s is exceeded. We then remove the last increment from the weight value, reduce the increment by a factor of ten, and repeat the process until the desired time constant is attained within a tolerance of 0.0001. To set the time constant for a fixed value of the Pcell-2 weight, we tune the absolute value of the Pcell-1 weight, as described above, until the desired time constant (e.g. 2 s, 20 s, or 200 s) is attained within a tolerance of 0.001. To set the time constant and gain using a nested grid search, we increment the absolute value of the Pcell-2 weight using the grid search until the desired gain is attained within tolerance, but for each increment of the Pcell-2 weight, we tune the absolute value of the Pcell-1 weight using the grid search to attain the desired time constant within tolerance. Tolerance for both gain and time constant is 0.001. The nested grid search method is efficient, and reaches tolerance after two thousand iterations on average.

Reinforcement learning

We used reinforcement learning (Sutton and Barto, 1998) to tune Pcell weights in both bilaterally symmetric and bilaterally asymmetric networks. In bilaterally symmetric networks, the Pcell-1 weight is the same on both sides, and similarly for the Pcell-2 weight. In bilaterally asymmetric

networks, all four Pcell weights are independently modifiable. On each cycle of reinforcement learning in bilaterally symmetric networks, we randomly perturb the Pcell-1 weights by the same amount on both sides, and similarly for the Pcell-2 weights. In bilaterally asymmetric networks we independently perturb all four Pcell weights. In either case, all perturbations are retained if the combined error between the actual and desired time constant and gain decreases. Perturbing all modifiable weights simultaneously is neurobiologically more plausible than perturbing one at a time, but still allows learning to occur. Pcell weight perturbations are drawn from a zero-mean Gaussian. To improve performance in high-gain cases, we decrease the variance of the perturbations as learning proceeds, and apply unequal scaling to the perturbations of the weights of Pcell-1 and Pcell-2. We also apply unequal scaling to the components of the error due to time constant and gain. The reinforcement strategy is inefficient, and requires more than ten thousand iterations to produce a good fit.

Genetic algorithm

We used the genetic algorithm (Holland, 1975; Goldberg, 1989) to search the space of possible, bilaterally symmetric network configurations. Binary chromosomes encode the pattern of Vunit-Pcell connectivity (20- and 400-Vunit networks) and FTN locations (20-Vunit networks only). Each bit in the population (composed of 20 chromosomes) initially takes value 1 with probability 0.5. Populations evolve for 100 or 5000 generations for 20- and 400-Vunit networks, respectively. On every generation, each bit in the population mutates with probability 0.002. Also on every generation, the algorithm chooses four chromosomes for mating to produce four offspring, and chooses four other chromosomes for replacement. Mating pairs copy then recombine their chromosomes at random cross-over points. The algorithm chooses mating and replacement pairs with probabilities proportional to their fitness. The fitness of a chromosome depends on the ability of the network it encodes to adjust gain by changing Pcell weights, while keeping the time constant at 20 s. We determine fitness by holding the Pcell-2 weight at a series of preset values (-0.1 to -2), and tuning the Pcell-1 weight, using the grid search, until a time constant of 20 s is attained within a tolerance of 0.001. We compute integrator (Vcommand) gain (Eq. (6)) from the response using the preset value for the Pcell-2 weight, and the grid search result for the Pcell-1 weight, provided that these weights also produce a real dominant eigenvalue corresponding to a time constant of 20 s. We then compute fitness as the difference between the maximal and minimal gains over the series. To ensure sharp response onsets, we discount fitness by a factor related to the ratio of the absolute real parts of the dominant and first subdominant eigenvalues.

All simulations and analyses were carried out using MATLAB (The MathWorks).

Results

We used our model to demonstrate how the cerebellum could independently regulate the gain and time constant of the mammalian oculomotor integrator, despite sparse innervation of MVN/NPH neurons by floccular Purkinje cells. Model units representing MVN/NPH neurons and floccular Purkinje cells are called Vunits and Pcells, respectively. V-units that receive input from Pcells are called flocculus target neurons (FTNs). Vunits and Pcells are modeled as linear, first-order units, and the networks composed of Vunits and Pcells are described in state space form and analyzed in terms of their eigenvectors and eigenvalues (see Methods). The dominant eigenvalue determines the time constant of the integrator network. The networks operate in push-pull, and the eye position command from a network (Vcommand) is taken as the sum of the Vunit responses on one side, minus the sum of the Vunit responses on the other side (see Methods). The gain of the integrator (Vcommand gain) will mostly reflect the contribution due to the scaled dominant eigenvector (Eq. (6) in Methods).

Inhibitory Pcells could regulate integrator networks by augmenting the net positive feedback among Vunits. The simplest bilateral integrator network with Pcells has two V-units (both FTNs) and two Pcells (Fig. 1(a)). If the Pcells receive a preponderance of input from the FTN on the opposite (contralateral) side then, by tuning their weights to the FTNs, they could adjust the value of the dominant time constant (for discussion of sites of learning see Discussion). In the two-Vunit case, independent adjustment of time constant and integrator gain is not possible. Each Pcell weight setting will be associated with a unique dominant time constant, and a unique, scaled dominant eigenvector and integrator gain. This situation is not changed by increasing the number of Vunits (Fig. 1(b)). Regardless of the number of Vunits, or the connectivity among the Vunits and Pcells, each Pcell weight setting will correspond to a unique dominant time constant, scaled dominant eigenvector, and integrator gain. Independent adjustment of time constant and gain is again not possible. It is possible to achieve independent adjustment of integrator time constant and gain by adding just one more pair of Pcells (Fig. 1(c)). The problem is determining how to configure the network.

Exploring configurations with 20% cerebellar innervation

Our goal was to find network configurations, characterized by sparse Pcell innervation of Vunits, in which Pcells can independently adjust integrator time constant and gain. To ensure sparseness, each network has only four Pcells (two

Pcells per side), and each Pcell contacts only one FTN. The networks are bilaterally symmetric. Thus, networks with 20 Vunits (10 Vunits per side) have a 20% innervation of Vunits by Pcells. The weights of the connections between the Vunits take the same value, and the Vunit to Pcell (Vunit-Pcell) weights also take the same value (see Methods). While these constraints on network structure greatly simplify the problem, there remain a huge number of possible network configurations. In the 20-Vunit network there are 2^{40} ways in which the Vunits can project to the Pcells, and 100 ways the Pcells can project to the Vunits, for a total of over 10^{14} possible network configurations. We used the genetic algorithm (see Methods) to search this vast configuration space.

We encoded the phenotype (configuration) of each 20-Vunit network using a chromosome (binary string) specifying the FTN locations and Vunit-Pcell connectivity pattern. The genetic algorithm operated on a population of these chromosomes. We subjected ten sets of 200 populations to simulated evolution. Each set has its own pattern for the reciprocal connections between Vunits (uniform, or closed or open, with neighborhood size from zero to four; see Methods). The fitness of a chromosome depends on the range over which the network it encodes can change integrator gain, through changes in Pcell to FTN weights (Pcell weights), while holding the time constant at 20 s. In the absence of network connections, all units would have a gain of 1, and the integrator gain (Vcommand gain) for a 20-Vunit network would be 20. Since mammals can adaptively halve or double their VOR gains (Robinson, 1976; Nagao, 1983; Godaux and Vanderkelen, 1984; Rambold et al., 2002; Nagao and Kitazawa, 2003), a minimally fit network should have a fitness of 80, so it can match the minimal gain of 20 and double it twice. For robustness, we set the criterion fitness at 100. Almost 30% of populations in each set produced a fittest chromosome that attained fitness of 100 or greater. Thus, a mechanism for independent adjustment of time constant and gain evolves in our model integrator networks with Pcell innervation. This mechanism, which we term dynamic morphing, is described in detail in a later subsection, following analysis of the configurations of fit evolved networks.

The genetic algorithm found fit networks in all sets, regardless of the pattern of reciprocal connections between Vunits. We analyzed the configurations of fit networks and took statistics on FTN locations and Vunit-Pcell connections. The evolved configurations of fit networks varied. On average, FTN-1 and FTN-2 had a separation between them of two Vunits, and were located at a distance of three Vunits from either end in networks with open configurations. Fewer than 1% of the fit networks actually had FTN-1 and FTN-2 at identical locations. Thus, while at least two Pcells per side are needed for independent adjustment of time constant and gain, both can work through the same FTN. Pcells could receive inputs from FTNs and other Vunits

alike, both ipsilaterally and contralaterally, though Pcell-1 and Pcell-2 received a preponderance of contralateral and ipsilateral input, respectively. Half of the potential Vunit-Pcell connections were utilized in fit networks.

The performance of a fit example network, with a Vunit interconnection neighborhood of two, is shown in Fig. 2(a). This open network has FTN-1 and FTN-2 at locations four and seven, and therefore represents the average configuration found by genetic search. To achieve specific combinations of time constant and gain, we nested the grid search to tune both the Pcell-1 and the Pcell-2 weights (see Methods). Figure 2(a) depicts a range of gain and time constant adjustment that encompasses the range observed experimentally (Robinson, 1976; Nagao, 1983; Godaux and Vanderkelen, 1984; Tiliket et al., 1994; Rambold et al., 2002; Nagao and Kitazawa, 2003).

The grid search used to tune Pcell weights is efficient (see Methods) but implausible neurobiologically. To explore a more plausible alternative, we trained the 20-Vunit example network using a reinforcement strategy (see Methods), and found Pcell weights to accurately reproduce all of the combinations of time constant and gain that were found using the grid search (Fig. 2(a)). The details of cerebellar learning are outside the focus of this study (see Discussion), but the results show that Pcell weights can be learned via adaptive mechanisms in the model, and it is reasonable to assume that they can be learned in the real brain also.

Altering evolved configurations

Genetic search found many ways to solve the problem of independent adjustment of time constant and gain, but solutions were specific. We examined network specificities by altering their configurations. Alterations involve moving an FTN location or changing (making or breaking) a Vunit-Pcell connection bilaterally. These alterations sometimes result in instability, and almost always change the dominant time constant. To ease comparison between them, we use the grid search on all altered networks to tune the Pcell-1 weight to yield a dominant time constant of 20 s, while holding the Pcell-2 weight at the mid-gain value. Changes in network configuration usually degrade performance, either by decreasing gain, or by causing network responses (of all Vunits, FTNs, Pcells, and the Vcommand) to oscillate.

Altering (making or breaking) a Vunit-Pcell connection bilaterally in the example network either decreases integrator gain (22 of 40 Vunit-Pcell connections), or causes oscillation (18 of 40). The integrator model with Pcell input, like the real network, is susceptible to oscillation because its connectivity is asymmetric (see Discussion). The presence of oscillation is determined from the numeric type of the dominant eigenvalue. If it is real then the dominant mode is purely exponential, but if it is complex then the dominant mode

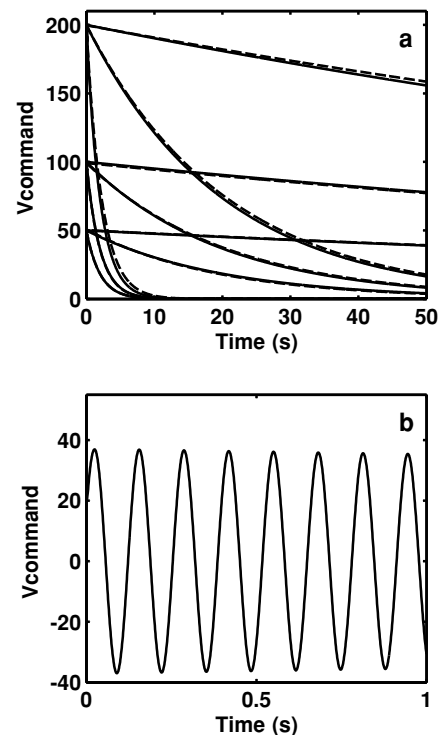
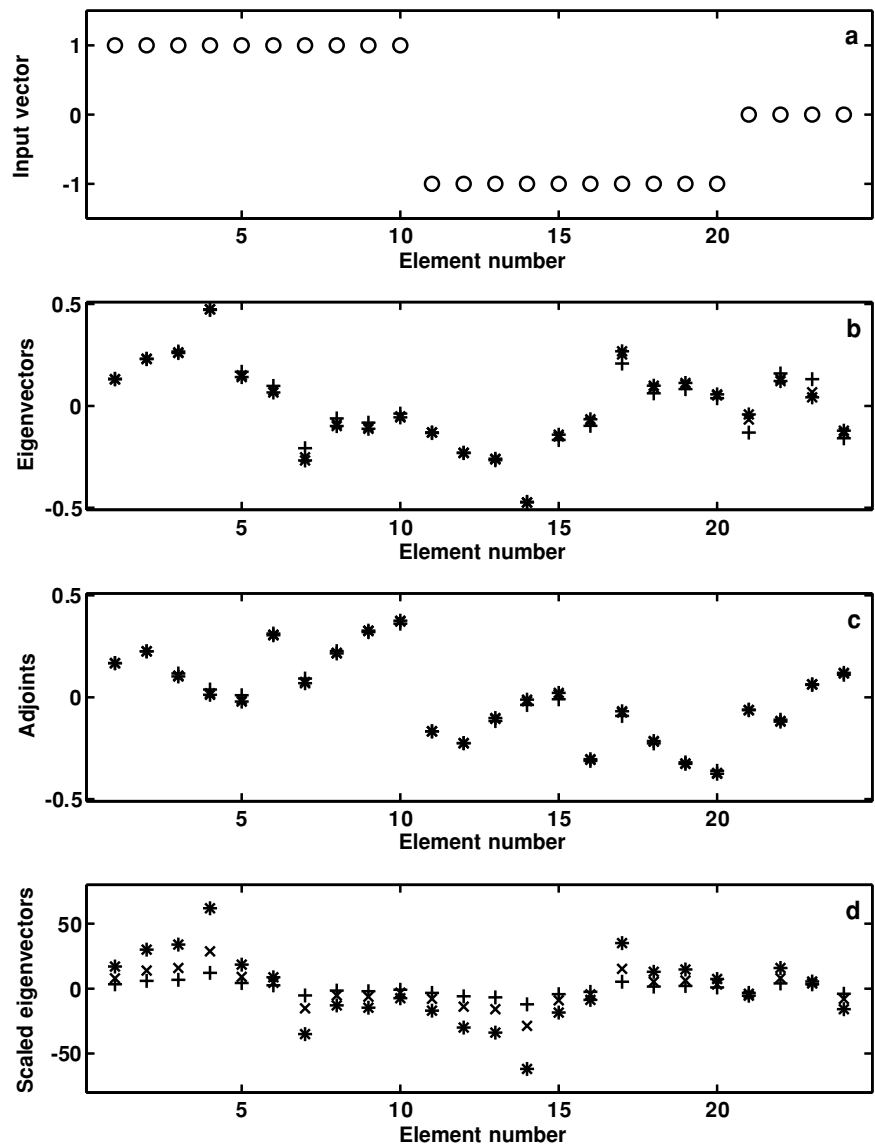


Fig. 2 Network response (Vcommand), under normal and pathological conditions, in the example, 20-Vunit network with four Pcells (as in Fig. 1(c)). (a) Through changes in the Pcell to FTN weights (Pcell weights), the normal network can achieve any time constant, and independently adjust integrator gain (Vcommand gain) over a broad range (time constants: 2 s, 20 s, and 200 s; gains: 50, 100, and 200). We trained Pcell weights using either a grid search (solid curves) or a reinforcement strategy (dashed curves). (b) Altering the normal network, by introducing an ipsilateral connection from Vunit-1 to Pcell-2 bilaterally, causes the network response to oscillate

is oscillatory. Making a normally absent connection from Vunit-1 to Pcell-2 causes the example network to oscillate at 8 Hz (Fig. 2(b)). Oscillations due to altering a Vunit-Pcell connection bilaterally in the example network range in frequency from 6 Hz to 38 Hz. These alterations correspond to point mutations of the chromosomes specifying network configurations (see Discussion).

In the integrator model with Pcell input, the numeric type of the dominant eigenvalue is determined both by the pattern of Vunit-Pcell connectivity and by the value of the Pcell to FTN weights (Pcell weights). As explained below, integrator gain is regulated by the Pcell-2 weight. In most fit networks, integrator gain is limited by the change in the dominant eigenvalue from real to complex as the Pcell-2 weight is increased past a certain value, which itself depends on network configuration. Conversely, oscillation due to alteration in Vunit-Pcell connectivity can be removed through adjustment of the Pcell weights in some networks, again depending on their specific connectivity. Presumably, the ability to plastically modify the Vunit-Pcell weights should prevent integrator oscillation in normal mammals (see Discussion).

Fig. 3 Analysis of dynamic morphing in the 20-Vunit network. The response amplitudes of the units in the network, and so the integrator (Vcommand) gain, are determined by the dominant eigenvector after it has been scaled by coefficient α_1 (d). Coefficient α_1 depends (Eq. (4)) on the input vector (a), the dominant eigenvector (b), and the dominant adjoint eigenvector (c). Dynamic morphing produces small rotations in both the dominant eigenvector and its adjoint vector that result in large changes in α_1 . The input vector is the same at all Vcommand gains. The other vectors are shown at three different levels of Vcommand gain (50, +; 100, \times ; 200, *) at a time constant of 20 s (the analysis at other time constants is similar). The first ten elements of each vector correspond to left-side Vunits, the next ten to right-side Vunits, and the last four to Pcells. Input weights are 1 and -1 to left and right Vunits, respectively, and 0 to Pcells (a)



The mechanism of dynamic morphing

Dynamic morphing is a mechanism by which the integrator gain and the integrator time constant can be independently adjusted. Because the integrator networks considered here are linear, the dynamic morphing phenomenon can be described using the results of linear analysis. The behavior of a linear network is determined mostly by the dominant eigenvalue and the scaled dominant eigenvector (see Methods). The dominant eigenvalue in fit integrator networks is real, and the integrator time constant is equal to its opposite reciprocal. The gain of the integrator (Vcommand gain), which is proportional to the sum of the Vunit responses on the left side minus the sum on the right (Eq. (6)), depends on the amplitudes of the responses of the units. The unit response amplitudes are mostly determined by the dominant eigenvector after it has been scaled by its coefficient α_1 (Eq. (4)).

Coefficient α_1 is a function of the input vector, the dominant eigenvector, and its associated adjoint eigenvector (Eq. (4)).

Figure 3 shows how relatively small rotations of the dominant eigenvector and its adjoint can lead to relatively large changes in α_1 and in the scaled dominant eigenvector (Fig. 3(d)). The vectors are those of the 20-Vunit example network, at a time constant of 20 s and at three different levels of integrator gain (50, +; 100, \times ; and 200, *). The elements of the input vector (Fig. 3(a)) are the same at all gain levels. Dynamic morphing causes both the dominant eigenvector (Fig. 3(b)) and its adjoint vector (Fig. 3(c)) to rotate slightly as it changes integrator (Vcommand) gain.

The amount of rotation can be quantified as the length of the vector that is the difference between the eigenvectors, or their adjoints, before and after rotation. Because the dominant eigenvector and its adjoint are unit vectors, and because a vector rotated by the maximal amount will point in the

direction opposite to its original direction, the maximal length of the difference vector would be two. The actual difference vectors for dominant eigenvector or adjoint rotations range in length from 1% to 6.5% of maximal. Thus, the rotations are small, but they produce large changes in α_1 .

Coefficient α_1 is computed as the scalar product of the input vector and the adjoint vector, divided by the scalar product of the dominant eigenvector and the adjoint vector (Eq. (4)). The reason that small rotations in the dominant eigenvector and its adjoint cause large changes in α_1 is that the dominant eigenvector and its adjoint are nearly orthogonal. Therefore, their scalar product is close to zero. Because the scalar product of the dominant eigenvector and its adjoint occurs in the denominator of the relation determining α_1 (Eq. (4)), small rotations in those vectors cause large changes in α_1 , which in turn cause large changes in the overall integrator (Vcommand) gain (Eq. (6)). Thus, the dynamic morphing mechanism increases Vcommand gain by making the dominant eigenvector and its adjoint more orthogonal. The near orthogonality of the dominant eigenvector and its adjoint, or indeed any angular separation between eigenvectors and their adjoints, is possible only in asymmetric systems (Luenberger, 1979). Thus, the dynamic morphing mechanism depends on the asymmetry of the integrator network with Pcell input (see Discussion).

The gain changes analyzed in Fig. 3 were made while holding the time constant at 20 s, but similar gain changes can be made while holding the time constant at any value over a broad range (e.g. Fig. 2(a)). Alternatively, the gain can be held constant while the time constant is varied over a broad range (e.g. Fig. 2(a)). In holding gain constant, the dominant eigenvector and its adjoint are not rotated (not shown), even as the dominant eigenvalue, and so the time constant, are varied over several orders of magnitude. Thus, dynamic morphing involves rotating the dominant eigenvector, and its adjoint, while holding the dominant eigenvalue constant, or keeping the directions of the dominant eigenvector and its adjoint constant while changing the dominant eigenvalue.

On a more qualitative level, morphing involves changing the shape of the scaled dominant eigenvector. In the integrator networks with Pcell innervation, morphing produces changes in the response amplitudes of the Vunits. The Pcells change the Vunit response amplitudes by pulling their responses in opposite directions. Basically, Pcell-1 pulls up the responses of FTN-1 and its Vunit neighbors, while Pcell-2 pulls down the responses of FTN-2 and its Vunit neighbors. This pulling increases the response amplitudes of FTN-1 (and its Vunit neighbors) more than it decreases the response amplitudes of FTN-2 (and its Vunit neighbors). Thus the Vcommand is positive, and the more strongly the Pcells oppose each other (i.e. the harder they try to pull apart the responses of the Vunits), the larger the Vcommand gain (integrator gain) becomes.

Regardless of their FTN positions or Vunit-Pcell connectivities, all fit networks show the same pattern of Pcell responses. Neural responses in bilateral VOR pathways can be described as type I or type II, according to whether they increase (type I) or decrease (type II) their firing rate when their ipsilateral input increases (Wilson and Melvill Jones, 1979). In all fit networks, Pcell-1 has a type II response, while Pcell-2 has a type I response (Fig. 4(a)). Pcells are inhibitory. Thus, FTN-1 is type I because it receives input from Pcell-1, which is type II. FTN-2 is likewise type II (Fig. 4(c)).

The effects of Pcell feedback are reflected by the scaled dominant eigenvector, which determines the response amplitudes of the units. Figure 4(b) shows the elements of the scaled eigenvector. Each element of the scaled eigenvector largely determines the gain, and so the amplitude of the impulse response, of its corresponding Vunit. For example, the fourth element of the scaled eigenvector (within the gray bar) corresponds to FTN-1 (Fig. 4(b)). Figure 4(b) shows the elements of the scaled eigenvector at each of three levels of integrator gain. The crosses correspond to a gain of 100. The responses of the Vunits (including the FTNs) at an integrator gain of 100 are shown in Fig. 4(c). Note in Fig. 4(b) that the fourth element of the scaled eigenvector (corresponding to FTN-1) at an integrator gain of 100 (the fourth cross from the left) has a value near 30. This largely determines the gain, and so the amplitude of the impulse response, of FTN-1, which is almost 30 at an integrator gain of 100 (Fig. 4(c)). The other elements of the scaled eigenvector at an integrator gain of 100 (crosses in Fig. 4(b)) likewise determine the response amplitudes of the other Vunits (including FTN-2) at an integrator gain of 100 (Fig. 4(c)). Figure 4(b) shows the elements of the scaled dominant eigenvector at three different levels of integrator gain (50, 100, and 200) but, for clarity, Fig. 4(c) only shows the Vunit (including FTN) impulse responses at an integrator gain of 100.

Although response amplitudes change less for Pcell-1 than for Pcell-2, increases in integrator gain are associated with increases in the response amplitudes of both Pcell-1 and Pcell-2 (Fig. 4(a)). Increases in the response amplitudes of Pcell-1 and Pcell-2 in turn cause increases in the response amplitudes of both FTN-1 and FTN-2 (Fig. 4(b)). Because of Vunit interconnectivity, FTN responses are transmitted throughout the network. Some Vunits have their responses pulled up by FTN-1, while others have their responses pulled down by FTN-2 (Fig. 4(b)). The amplitudes of both type I and type II FTNs and Vunits increase together, but the type I increase more than the type II. The result is a net increase in integrator gain (Vcommand gain).

Dynamic morphing results from an interplay of opposites. Because Pcells are inhibitory, Pcell-1 (type II) and Pcell-2 (type I) exert positive and negative feedback, respectively. Positive feedback by Pcell-1 must overcome negative feedback by Pcell-2 to achieve time constant lengthening. This

tug-of-war is what changes Vunit (and FTN) response amplitudes and produces adjustment of integrator gain. Basically, positive feedback can be increased to achieve virtually any time constant at any level of negative feedback, but the more these two influences oppose each other, the more the unit response amplitudes are increased, and the larger the overall integrator gain becomes.

The Pcells evolve to be of specific types because of the way their fitness is evaluated. To evaluate fitness, the Pcell-2 weight is held fixed at various levels, while the Pcell-1 weight is tuned to adjust the time constant. Thus, Pcell-1 evolves to be type II because it is responsible for producing the positive feedback that is needed to lengthen the integrator time constant. Pcell-2 evolves to be type I so that it can exert the negative feedback necessary to adjust the unit response amplitudes and produce gain changes. Pcell type depends on the input received from Vunits. Evolved networks tend to have more type I than type II Vunits. Thus, Pcell-1 (type II) and Pcell-2 (type I) receive the preponderance of their input from the contralateral and ipsilateral sides, respectively.

Reinforcement learning in bilaterally asymmetric networks

We used a regular and bilaterally symmetric structure to facilitate evolution and comparison of network configurations and analysis of network dynamics, but these constraints are by no means necessary for network performance. It is also possible to achieve independent adjustment of time constant and gain in bilaterally asymmetric networks. The robustness of the dynamic morphing phenomenon is illustrated using the 20-Vunit example network (Fig. 5). First, the reciprocal connections between the Vunits are randomly perturbed, independently on either side (Fig. 5(a) and (b)), and scaled to produce a dominant time constant of 0.2 s without Pcell feed-

back. Then the FTNs are positioned asymmetrically (FTN-1 is located at position three and four, and FTN-2 at seven and eight, on the left and right sides, respectively). Finally, the four Pcell weights are trained using reinforcement learning (see Methods) to achieve three different levels of Vcommand gain at a time constant of 20 s (Fig. 5(c)). The behavior of the Pcells, FTNs, and other Vunits is similar in the randomly perturbed, bilaterally asymmetric network (Fig. 5(d)–(f)) and in the regular, bilaterally symmetric network (Fig. 4).

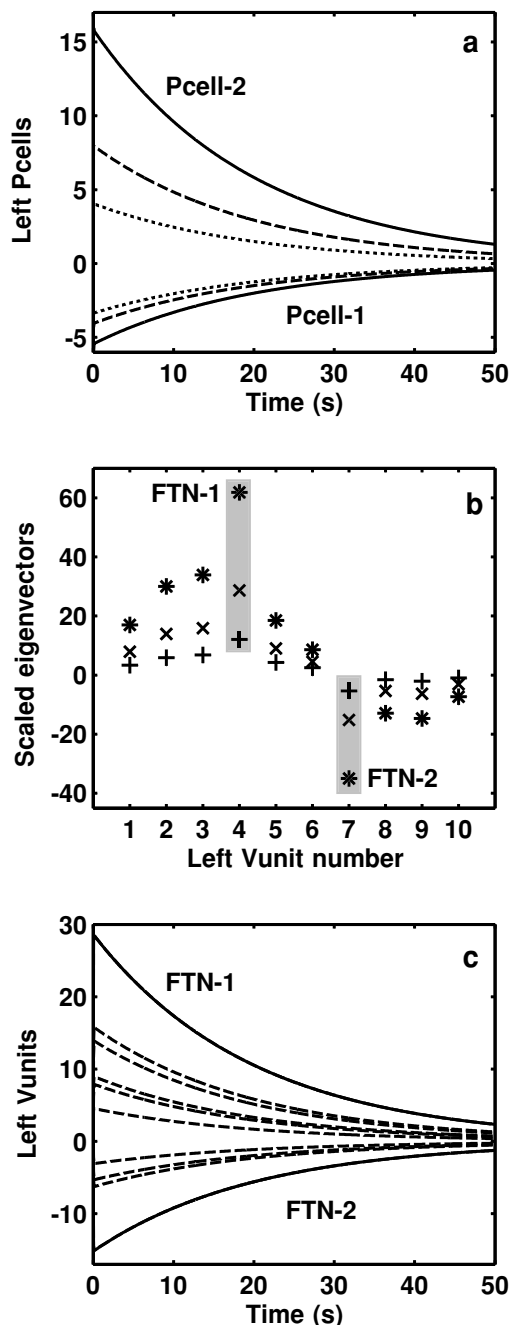


Fig. 4 Response analysis in the example, 20-Vunit integrator network at a time constant of 20 s. (a) The response amplitudes of Pcell-1 (type II, lower negative curves) and Pcell-2 (type I, upper positive curves) increase as integrator (Vcommand) gain increases. The responses of left-side Pcells at three different Vcommand gains are shown (gains: 50, dotted; 100, dashed; 200, solid). Right-side Pcell responses (not shown) are the additive inverses of those on the left side. (b) The scaled dominant eigenvector, which determines the unit response amplitudes, changes as Vcommand gain increases. Elements 4 and 7 (within gray bars) correspond to FTN-1 and FTN-2, respectively, while the other elements correspond to the Vunits, on the left side (gains: 50, +; 100, x; 200, *). (c) The impulse responses of the FTNs (solid) and other Vunits (dashed) are almost completely determined by the dominant eigenvalue (corresponding to a 20 s time constant) and the scaled dominant eigenvector (shown in (b)). The example responses are for FTNs and Vunits on the left side at the middle Vcommand gain of 100. Right-side FTN and Vunit eigenvector elements and unit responses (not shown) are the additive inverses of those on the left side. The convention of showing only responses on the left side will be followed in subsequent figures

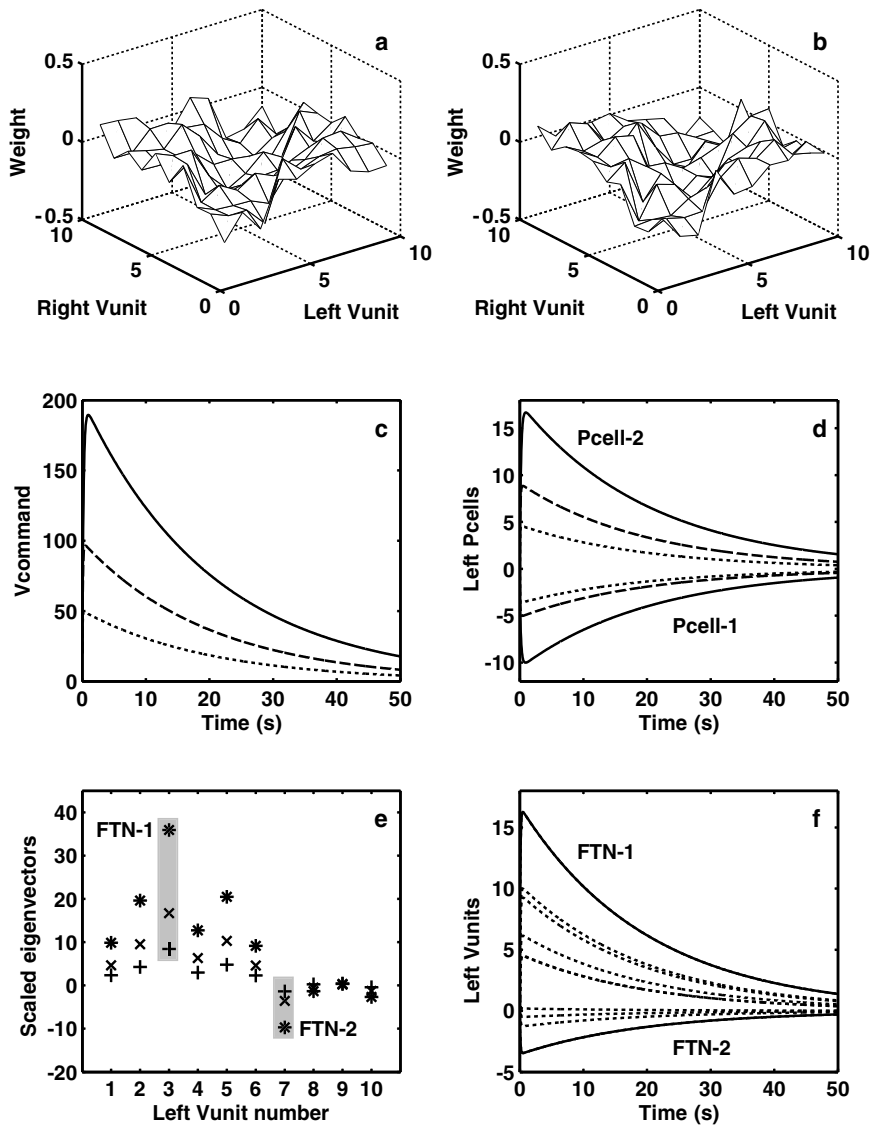


Fig. 5 Response analysis in the example 20-Vunit integrator network with randomly perturbed Vunit interconnections and asymmetric FTN locations. FTN-1 is located at position 3 and 4 on left and right, and FTN-2 is located at position 7 and 8 on left and right, respectively. (a) Perturbed weights of connections from Vunits on the right to Vunits on the left. (b) Perturbed weights of connections from Vunits on the left to Vunits on the right. Both Vunit weight arrays, if unperturbed, would have a value of -0.21 along a diagonal band down the center and zero elsewhere. (c) The network is able to independently adjust time constant and gain despite randomly perturbed Vunit interconnections and asymmetric FTN locations. Shown are three different Vcommand gains (50, dotted; 100, dashed; 200, solid) at a time constant of 20 s. The same gains are attainable over a range of different time constants

(not shown). (d) As in the unperturbed network, Pcell-2 is type I while Pcell-1 is type II. The responses of left-side Pcells at the different Vcommand gains are shown (50, dotted; 100, dashed; 200, solid). (e) Opposing Pcell feedback alters the scaled dominant eigenvector. The elements of the scaled eigenvector corresponding to Vunits (including FTNs) on the left side are shown at the three Vcommand gains (50, +; 100, \times ; 200, *). Gray bars mark the FTNs. (f) The scaled eigenvector determines the Vunit (and FTN) response amplitudes. The responses of FTNs and other Vunits on the left side are shown at the middle Vcommand gain of 100 (FTNs, solid; other Vunits, dashed). Thus, the behavior of the perturbed, bilaterally asymmetric network is essentially the same as that of the bilaterally symmetric network

As previously mentioned, alterations in evolved network structure, which involve moving an FTN location or changing (making or breaking) a Vunit-Pcell connection bilaterally, often degrade network performance. These alterations almost always change the dominant time constant, which can be reset by tuning the Pcell-1 weight

alone. In many altered networks, it is also possible to independently adjust time constant and gain by tuning both the Pcell-1 and Pcell-2 weights, as in the normal network. The ability to adaptively tune the Pcell contribution makes dynamic morphing by sparse Pcell innervation a robust phenomenon.

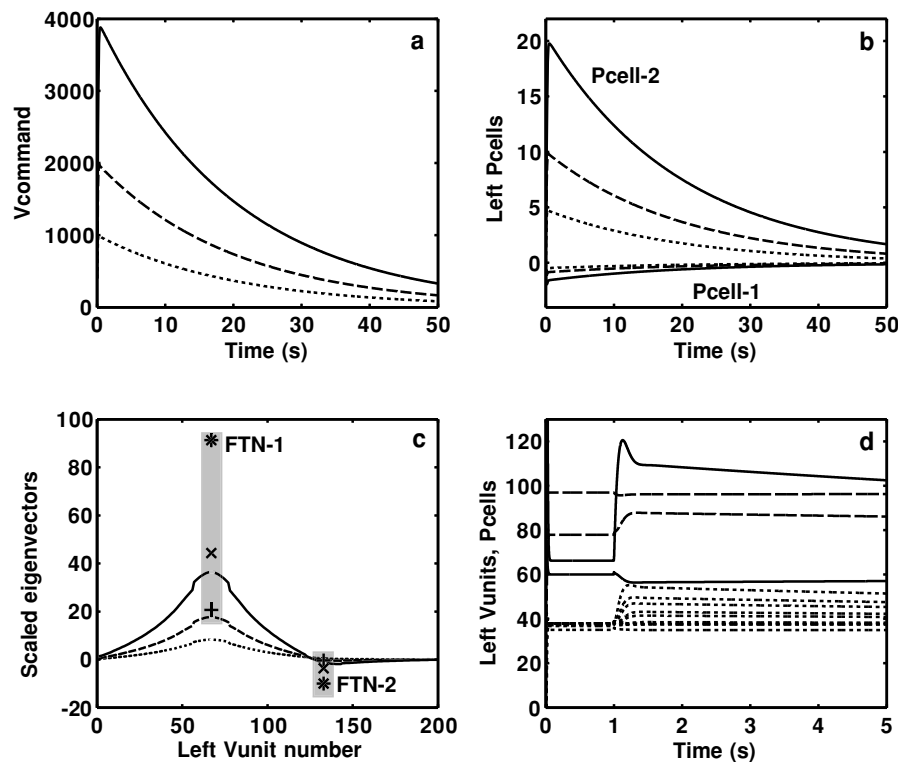


Fig. 6 Response analysis of a 400-Vunit integrator network at a time constant of 20 s. (a) Independent adjustment of time constant and gain is possible in the 400-Vunit network, despite very sparse (1%) Pcell innervation. Shown are three different Vcommand gains (1000, dotted; 2000, dashed; 4000, solid) at a time constant of 20 s. The same gains are attainable over a range of different time constants (not shown). (b) As in 20-Vunit networks, Pcell-2 is type I while Pcell-1 is type II. The responses of left-side Pcells are shown at the three Vcommand gains (1000, dotted; 2000, dashed; 4000, solid). (c) Adjustment of Pcell weights alters the scaled dominant eigenvector. Element 67 corresponds

to FTN-1, and element 133 to FTN-2 (within gray bars) at the three Vcommand gains (1000, +; 2000, ×; 4000, *). The other elements correspond to left-side Vunits at the three Vcommand gains (1000, dotted; 2000, dashed; 4000, solid). (d) We simulated spontaneous rates by adding bias values to the Pcells, FTNs, and Vunits. Left-side pulse responses correspond to the middle gain of 2000 (FTNs, solid; Pcells, dashed; selected Vunits, dotted). Right-side pulse responses (not shown) are the same as those on the left, but rotated about axes defined for each unit by its bias

Dynamic morphing with 1% cerebellar innervation

The 20-Vunit networks, with four Pcells, demonstrate how Purkinje cells could regulate oculomotor integrator dynamics despite only 20% innervation of integrator neurons by the flocculus. To investigate 1% innervation, we expanded the network to 400 Vunits, keeping the number of Pcells at four. To make simulated evolution of the larger networks more manageable, we fixed the locations of the FTNs at positions 67 and 133 on each side. This corresponds to the spacing that evolved in the 20-Vunit networks. Fixing the FTN locations still leaves almost 10^{240} possible configurations for Vunit-Pcell connections. The minimal Vcommand gain in 400-Vunit networks is 400, so a minimally fit network would have a fitness of 1600, but we set the criterion fitness at 2000 for robustness. Genetic search found fit 400-Vunit networks despite the complexity of the expanded problem.

We used the nested grid search (see Methods) to tune the Pcell weights in a fit 400-Vunit network to achieve

various combinations of integrator time constant and gain. Figure 6(a) depicts a realistic range of gain adjustment at a time constant of 20 s. The same adjustments in gain are possible over a broad range of time constants (not shown). As in the 20-Vunit networks, Pcell-1 is type II and Pcell-2 is type I in the 400-Vunit network (Fig. 6(b)). In both 20- and 400-Vunit networks, the response amplitudes of the Pcells, FTNs, and other Vunits, whether type I or type II, all increase as overall integrator (Vcommand) gain increases (Fig. 6(b) and (c)). These results are consistent with the real oculomotor system (see Discussion).

Purkinje cells and MVN/NPH neurons have spontaneous firing rates about which their responses are modulated. To simulate this, we added bias values to the unit responses in the 400-Vunit network and computed the pulse response, following a period of constant input (Fig. 6(d)). The biases were set to make the unit spontaneous rates commensurate with those observed for real eye-movement-related Purkinje cells and brainstem neurons (McFarland, 1992; Stahl and Simpson, 1995; Cheron et al., 1996; Lisberger and Pavelko, 1988; Lis-

berger et al., 1994a; Keller and Precht, 1979; Lisberger and Miles, 1980; Miles et al., 1980a,b; Lisberger et al., 1994b; Blazquez et al., 2003). The transient onsets of the FTN pulse responses are similar to those observed experimentally (Stahl and Simpson, 1995; Lisberger and Pavelko, 1988; Lisberger et al., 1994a). Bias does not alter network dynamics.

Discussion

Our main finding is that adaptive cerebellar feedback can independently adjust the time constant and gain of a model oculomotor integrator, and that this adjustment is possible using very sparse Purkinje cell innervation of integrator units. Previous modeling work suggests that the Purkinje cells that regulate VOR might become specialized for specific temporal patterns of head rotation (Anastasio, 2001). Since few Purkinje cells are required to regulate the integrator, the cerebellum could use separate and differently specialized Purkinje cells, perhaps working through the same FTNs, to morph the dynamic behavior of the oculomotor neural integrator to produce combinations of time constant and gain that are tailored for specific circumstances. This conjecture is supported by recent experimental findings that both integrator function and Purkinje cell responses depend on context (Belton and McCrea, 2004; Chan and Galiana, 2005).

While the details of cerebellar learning are outside the focus of this study, we show that Pcell weights can be learned via adaptive mechanisms in the model, and it is reasonable to assume that they can be learned in the real brain also. We restrict synaptic plasticity in the model to Pcell to FTN (Pcell-FTN) synapses for simplicity. There is evidence for synaptic plasticity both in cerebellum and brainstem (Ito, 1982; du Lac et al., 1995). Changing the weight of a Pcell-FTN synapse is equivalent to changing the sensitivity of the Pcell, or to scaling by the same amount the weights of all Vunit connections to the Pcell. Restricting synaptic plasticity to Pcell-FTN synapses facilitates comparison between networks of their Vunit to Pcell connectivity patterns, and demonstrates network robustness.

We studied evolved networks by altering their structures. These alterations correspond to point mutations of the chromosomes specifying network configurations. Alterations generally degrade network performance and often produce oscillations. Oscillatory dynamics are produced by complex modes that can only occur in asymmetric systems, though asymmetric systems need not have complex modes (Luenberger, 1979). The connectivity (system) matrices describing integrator networks with Pcell feedback are essentially asymmetric, because of the neuroanatomical fact that Pcells receive from Vunits bilaterally but project to Vunits only ipsilaterally (Langer et al., 1985a,b; Büttner-Ennever, 1988; Epema et al., 1990; Tan and

Gerrits, 1992). Fit networks evolve to have real (non-oscillatory) dominant modes, but point mutations could make the dominant mode complex (oscillatory). Oscillation frequencies due to mutations of network configurations range between 1 Hz and 40 Hz, and encompass the range observed for oscillatory eye-movement disorders (Leigh and Zee, 1983). Certain congenital eye-movement disorders have been attributed to a leaky or unstable neural integrator (Leigh and Zee, 1983). Our model can be made unstable or leaky by increasing or decreasing the Pcell-1 weights. Our results further suggest that oscillatory eye-movement disorders may also involve integrator abnormalities.

The genome probably does not specifically encode the Vunit-Pcell connectivity in real brains. It is more likely that the genome sets up a basic structure that is tuned by synaptic plasticity (Ito, 1982; du Lac et al., 1995). Plasticity of Vunit-Pcell synapses could prevent oscillatory responses by correcting errors in Vunit-Pcell connectivity. We predict that oscillatory neural responses and eye movements could occur in mammals lacking plasticity in Vunit-Pcell synapses. Recent work in mutant mice deficient in the glutamate receptor $\delta 2$ subunit, which lack plasticity of synapses onto Purkinje cells, reveals oscillations of about 10 Hz in Purkinje cell responses and eye movements (Yoshida et al., 2004). Oscillatory eye movements cease in these mutant mice following flocculectomy (*ibid.*).

Network asymmetry, which is responsible for the abnormal, oscillatory dynamics of mutated networks, plays a critical role in normal networks. In the integrator network with Pcell input, gain increases as the scalar product of the dominant eigenvector and its adjoint decreases (see Results). That scalar product can be made arbitrarily small, because the dominant eigenvector and its adjoint are nearly orthogonal, so small rotations of those vectors can produce large gain changes. In symmetric systems, eigenvectors and their adjoints are the same; angular separation of the dominant eigenvector and its adjoint can occur only in asymmetric systems (Luenberger, 1979). Thus, asymmetry in the integrator network with Pcell input is essential for dynamic morphing.

The mechanism of dynamic morphing is similar in 20- and 400-Vunit networks, which have cerebellar innervation of 20% and 1%, respectively. In both, the response gains of the Pcells, FTNs, and other Vunits, whether type I or type II, all increase or decrease as integrator gain increases or decreases. We predict similar behavior for real Purkinje cells and MVN/NPH neurons that are involved in producing neural integration. This prediction has not been addressed specifically for integrator neurons, but findings on sensitivity to head rotation are in accord.

MVN neurons, whether type I or type II, change their gain in the same direction as the adapted VOR in cats and monkeys (Keller and Precht, 1979; Lisberger and Miles, 1980). Keller and Precht (1979), working in cat, further divided

their type I and type II populations into high-gain and low-gain groups, and found that the change in gain was larger for high-gain than for low-gain neurons, regardless of response type (I or II). This corresponds precisely with our model. FTNs, and the Vunits near the FTNs, have the highest gains (whether their response polarity is type I or type II), and their gains change the most with changes in overall integrator gain (see Figs. 4(b), 5(e), and 6(c)). Studies that have focused on FTNs find that, whether the response is type I or type II, the sensitivity of FTNs to head rotation is considerably higher in monkeys with up-adapted as compared with down-adapted VORs (Lisberger and Pavelko, 1988; Lisberger et al., 1994a).

Studies on floccular Purkinje cells provide an intriguing set of data. Purkinje cells are inhibitory and only project ipsilaterally onto vestibular nucleus neurons (see above). This lead researchers to assume that type II Purkinje cells should change gain in the same direction as the adapted VOR, but that type I Purkinje cells should change gain in the direction opposite the adapted VOR (e.g. Miles et al., 1980b). Although some type I floccular Purkinje cells in rabbits change their gain in the direction opposite to that of the adapted VOR (Watanabe, 1985), most others of both types change gain in the same direction as that of the adapted VOR (Watanabe, 1985; Dufossé et al., 1978; Nagao, 1989). Miles and coworkers (1980a,b) made similar findings in monkeys, leading them to conclude that the gains of type I Purkinje cells change “in the wrong direction” to account for VOR plasticity.

It is possible that some Purkinje cells, including the type I rabbit Purkinje cells that change gain opposite the adapted VOR, but also including some type II Purkinje cells that change gain with the adapted VOR, regulate the velocity component of the VOR rather than the integrated component. Our model of the integrator offers an explanation for the others: Purkinje cells that are involved in controlling the integrator should change their gain in the same direction as the adapted VOR, whether their response polarity is type I or type II.

Recent findings are consistent with this explanation. Gaze velocity Purkinje cells, found in monkeys, appear to carry an eye velocity command during pursuit and a head rotation response during cancellation of the VOR. When normalized by their eye velocity gains, the gains of the head rotation responses of monkey gaze velocity Purkinje cells change in the same direction as that of the adapted VOR (Lisberger et al., 1994b; Blazquez et al., 2003). These findings support the dynamic morphing hypothesis.

Our model of the oculomotor neural integrator is the first to include the crucial contribution of the cerebellar floccular complex. It incorporates recent findings on the sparseness of floccular innervation of brainstem integrator neurons. The structure of the model reflects the known asymmetry in the connections between flocculus and brainstem. That asym-

metry, which may underlie the ability of the flocculus to independently adjust time constant and gain in normal integrator networks, can also lead to oscillations in altered networks that may underlie certain eye-movement disorders. The model is consistent with findings that floccular Purkinje cells and brainstem integrator neurons, regardless of their response polarity (type I or type II), change their gain in the same direction of the adapted VOR. Our approach provides a new way of analyzing neural integrator models, in terms not only of the dominant eigenvalue (and dominant time constant) but also in terms of the scaled dominant eigenvector. It shows how floccular Purkinje cells, in contacting only a few integrator neurons, could shape the scaled dominant eigenvector and change the pattern of response amplitude among the neurons, thereby changing overall integrator gain, while holding the dominant time constant at virtually any value. Our model provides a new perspective on how regulatory structures could morph the dynamics of motor control networks.

Acknowledgments We thank J. Bronski for mathematical consultation and J. Malpeli for suggestions. We also thank the two anonymous reviewers for constructive comments.

References

- Anastasio TJ (1998) Nonuniformity in the linear network model of the oculomotor integrator produces approximately fractional-order dynamics and more realistic neuron behavior. *Biol. Cybern.* 79: 377–391.
- Anastasio TJ (2001) A pattern correlation model of vestibulo-ocular reflex habituation. *Neural Netw.* 14: 1–22.
- Anastasio TJ, Robinson DA (1991) Failure of the oculomotor neural integrator from a discrete midline lesion between the abducens nuclei in the monkey. *Neurosci. Lett.* 127: 82–86.
- Arnold DB, Robinson DA (1991) A learning network model of the neural integrator of the oculomotor system. *Biol. Cybern.* 64: 447–454.
- Arnold DB, Robinson DA (1997) The oculomotor integrator: Testing of a neural network model. *Exp. Brain Res.* 113: 57–74.
- Babalian AL, Vidal PP (2000) Floccular modulation of vestibuloocular pathways and cerebellum-related plasticity: An in vitro whole brain study. *J. Neurophysiol.* 84: 2514–2528.
- Belton T, McCrea RA (2004) Context contingent signal processing in the cerebellar flocculus and ventral paraflocculus during gaze saccades. *J. Neurophysiol.* 92: 797–807.
- Bevington PR (1969) *Data Reduction and Error Analysis for the Physical Sciences*. McGraw-Hill, New York.
- Blazquez PM, Hirata Y, Heiney SA, Green AM, Highstein SM (2003) Cerebellar signatures of vestibulo-ocular reflex motor learning. *J. Neurosci.* 23: 9742–9751.
- Büttner-Ennever JA (ed.) (1988) *Neuroanatomy of the Oculomotor System*. Elsevier, Amsterdam.
- Cannon SC, Robinson DA (1987) Loss of the neural integrator of the oculomotor system from brain stem lesions in monkey. *J. Neurophysiol.* 57: 1383–1409.
- Cannon SC, Robinson DA, Shamma S (1983) A proposed neural network for the integrator of the oculomotor system. *Biol. Cybern.* 49: 127–136.

- Chan WWP, Galiana HL (2005) Integrator function in the oculomotor system is dependent on sensory context. *J. Neurophysiol.* 93: 3709–3717.
- Chelazzi L, Ghirardi M, Rossi F, Strata P, Tempia F (1990) Spontaneous saccades and gaze holding ability in the pigmented rat. II. Effects of localized cerebellar lesions. *Eur. J. Neurosci.* 2: 1085–1094.
- Cheron G, Escudero M, Godaux E (1996) Discharge properties of brain stem neurons projecting to the flocculus in the alert cat. I. Medial vestibular nucleus. *J. Neurophysiol.* 76: 1759–1774.
- Cheron G, Godaux E, Laune JM, Vanderkelen B (1986) Lesions in the cat prepositus complex: Effects on the vestibulo-ocular reflex and saccades. *J. Physiol.* 372: 75–94.
- Dufossé M, Ito M, Jastreboff PJ, Miyashita Y (1978) A neuronal correlate in rabbit's cerebellum to adaptive modification of the vestibulo-ocular reflex. *Brain Res.* 150: 611–616.
- Epema AH, Gerrits NM, Voogd J (1990) Secondary vestibulocerebellar projections to the flocculus and uvulo-nodular lobule of the rabbit: A study using HRP and double fluorescent tracer techniques. *Exp. Brain Res.* 80: 72–82.
- Godaux E, Vanderkelen B (1984) Vestibulo-ocular reflex, optokinetic response and their interactions in the cerebellectomized cat. *J. Physiol.* 346: 155–170.
- Goldberg E (1989) *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Boston.
- Holland H (1975) *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor.
- Ito M (1982) Cerebellar control of the vestibulo-ocular reflex – around the flocculus hypothesis. *Ann. Rev. Neurosci.* 5: 275–296.
- Katoh A, Yoshida T, Himeshima Y, Mishina M, Hirano T (2005) Defective control and adaptation of reflex eye movements in mutant mice deficient in either the glutamate receptor $\delta 2$ subunit or Purkinje cells. *Eur J Neurosci* 21: 1315–1326.
- Keller EL, Precht W (1979) Adaptive modification of central vestibular neurons in response to visual stimulation through reversing prisms. *J. Neurophysiol.* 42: 896–911.
- du Lac S, Raymond JL, Sejnowski TJ, Lisberger SG (1995) Learning and memory in the vestibulo-ocular reflex. *Ann. Rev. Neurosci.* 18: 409–441.
- Langer T, Fuchs AF, Scudder CA, Chubb MC (1985a) Afferents to the flocculus of the cerebellum in the rhesus macaque as revealed by retrograde transport of horseradish peroxidase. *J. Comp. Neurol.* 235: 1–25.
- Langer T, Fuchs AF, Chubb MC, Scudder CA, Lisberger SG (1985b) Floccular efferents in the rhesus macaque as revealed by autoradiography and horseradish peroxidase. *J. Comp. Neurol.* 235: 26–37.
- Leigh RJ, Zee DS (1983) *The Neurology of Eye Movements*. Davis, Philadelphia.
- Lisberger SG, Miles FA (1980) Role of primate medial vestibular nucleus in long-term adaptive plasticity of vestibuloocular reflex. *J. Neurophysiol.* 43: 1725–1745.
- Lisberger SG, Pavelko TA (1988) Brain stem neurons in modified pathways for motor learning in the primate vestibulo-ocular reflex. *Science* 242: 771–773.
- Lisberger SG, Pavelko TA, Broussard DM (1994a) Neural basis for motor learning in the vestibuloocular reflex of primates. I. Changes in the responses of brain stem neurons. *J. Neurophysiol.* 72: 928–953.
- Lisberger SG, Pavelko TA, Bronte-Stewart HM, Stone LS (1994b) Neural basis for motor learning in the vestibuloocular reflex of primates. II. Changes in the responses of horizontal gaze velocity Purkinje cells in the cerebellar flocculus and ventral paraflocculus. *J. Neurophysiol.* 72: 954–973.
- Luenberger G (1979) *Introduction to Dynamic Systems*. John Wiley, New York.
- McCrea RA, Strassman A, May E, Highstein SM (1987) Anatomical and physiological characteristics of vestibular neurons mediating the horizontal vestibulo-ocular reflex of the squirrel monkey. *J. Comp. Neurol.* 264: 547–570.
- McFarland JL, Fuchs AF (1992) Discharge patterns in nucleus prepositus hypoglossi and adjacent medial vestibular nucleus during horizontal eye movement in behaving macaques. *J. Neurophysiol.* 68: 319–332.
- Miles FA, Fuller JH, Braitman DJ (1980a) Long term adaptive changes in primate vestibuloocular reflex. III. Electrophysiological observations in flocculus of normal monkeys. *J. Neurophysiol.* 43: 1437–1476.
- Miles FA, Braitman DJ, Dow BM (1980b) Long term adaptive changes in primate vestibuloocular reflex. IV. Electrophysiological observations in flocculus of adapted monkeys. *J. Neurophysiol.* 43: 1477–1493.
- Nagao S (1983) Effect of vestibulo-cerebellar lesion upon dynamic characteristics and adaptation of the vestibuloocular and optokinetic responses in pigmented rabbits. *Exp. Brain Res.* 53: 36–46.
- Nagao S (1989) Behavior of floccular Purkinje cells correlated with adaptation of vestibulo-ocular reflex in pigmented rabbits. *Exp. Brain Res.* 77: 531–540.
- Nagao S, Kitazawa H (2003) Effects of reversible shutdown of the monkey flocculus on the retention of adaptation of the horizontal vestibulo-ocular reflex. *Neuroscience* 118: 563–570.
- Rambold H, Churchland A, Selig Y, Jasmin L, Lisberger SG (2002) Partial ablations of the flocculus and ventral paraflocculus in monkeys cause linked deficits in smooth pursuit eye movements and adaptive modification of the VOR. *J. Neurophysiol.* 87: 912–924.
- Robinson D (1974) The effect of cerebellectomy on the cat's vestibulo-ocular integrator. *Brain Res.* 71: 195–207.
- Robinson DA (1976) Adaptive gain control of vestibuloocular reflex by the cerebellum. *J. Neurophysiol.* 39: 954–969.
- Robinson DA (1989a) Integrating with neurons. *Annu. Rev. Neurosci.* 12: 33–45.
- Robinson DA (1989b) Control of eye movements. In: V.B. Brooks ed. *Handbook of Physiology, Sect. 1: The Nervous System, Vol. II part 2*. American Physiological Society, Bethesda, pp. 1275–1320.
- Sekirnjak C, du Lac S (2006) Physiological and anatomical properties of mouse medial vestibular nucleus neurons projecting to the oculomotor nucleus. *J. Neurophysiol.* 95: 3012–3023.
- Sekirnjak C, Vissel B, Bollinger J, Faulstich M, du Lac S (2003) Purkinje cell synapses target physiologically unique brainstem neurons. *J. Neurosci.* 23: 6392–6398.
- Stahl JS (2004) Eye movements of the murine P/Q calcium channel mutant Rucker, and the impact of aging. *J. Neurophysiol.* 91: 2066–2078.
- Stahl JS, James RA (2005) Neural integrator function in murine CACNA1A mutants. *Ann. NY Acad. Sci.* 1039: 580–582.
- Stahl JS, Simpson JI (1995) Dynamics of rabbit vestibular nucleus neurons and the influence of the flocculus. *J. Neurophysiol.* 73: 1396–1413.
- Sutton RS, Barto AG (1998) *Reinforcement Learning: An Introduction*. MIT Press, Cambridge.
- Tan H, Gerrits NM (1992) Laterality in the vestibulo-cerebellar mossy fiber projection to flocculus and caudal vermis in the rabbit: A retrograde fluorescent double-labeling study. *Neuroscience* 47: 909–919.
- Tiliket C, Shelhamer M, Roberts D, Zee DS (1994) Short term vestibulo-ocular reflex adaptation in humans. I. Effect on the ocular motor velocity-to-position neural integrator. *Exp. Brain Res.* 100: 316–327.

- Watanabe E (1985) Role of the primate flocculus in adaptation of the vestibulo-ocular reflex. *Neurosci. Res.* 3: 20–38.
- Wilson VJ, Melvill Jones JG (1979) *Mammalian Vestibular Physiology*. Plenum Press, New York.
- Yoshida T, Katoh A, Ohtsuki G, Mishina M, Hirano T (2004) Oscillating Purkinje neuron activity causing involuntary eye movement in a mutant mouse deficient in the glutamate receptor $\delta 2$ subunit. *J. Neurosci.* 24: 2440–2448.
- Zee S, Yamazaki A, Butler PH, Gücer G (1981) Effects of ablation of flocculus and paraflocculus on eye movements in primate. *J. Neurophysiol.* 46: 878–899.