Mechanisms of neural integration: recent results and relevance to nystagmus modeling

Andrea K. Barreiro
Department of Applied Mathematics, University of Washington

I. Introduction

In many systems in the brain, computation requires the persistence of information beyond the time of sensory input. Other computations require information to be accumulated over a period of time before a decision or action can be made. Both of these situations can be thought of as examples of integration in the sense of calculus; to be concrete, the transformation of a signal by

$$S(t) \to \int_0^t S(t')dt' \tag{1}$$

or more generally integration against a kernel

$$S(t) \to \int_0^t S(t')K(t-t')dt'. \tag{2}$$

Many working memory tasks can be interpreted as "integration" operations. For example, the neural correlate of a remembered sensation can be thought of as the integration of a brief sensory input which is thereafter stored, as has been found in primate pre-frontal cortex (Romo *et al.* 1999). In area LIP during motion discrimination tasks, firing rates increase at a rate proportional to motion strength (Shadlen and Newsome 2001). In the oculomotor system, a mechanism is needed to maintain eye position after a shift in gaze to an object of interest. This same system also helps to compensate for involuntary head movements, stabilizing our vision during head rotations,

breathing, and walking. In all of these cases, this system must translate sensations of eye *velocity* into estimates of eye *position*; this is precisely performed by the operation of mathematical integration.

This system is known as the *oculomotor integrator* and is present in all vertebrates. It has perhaps been studied most extensively in the goldfish, which performs a sequence of automatic saccades while swimming in order to scan its environment, providing a convenient experimental protocol. Two brainstem nuclei, area I and area II, have been identified as coding for eye position or velocity; area II modulates with eye velocity whereas area I codes for eye position. Single neuron recordings in area I show that the firing rates of neurons in this area have a threshold linear relationship to eye position, with the position threshold and slope varying among individual neurons (Aksay *et al.* 2000).

The situation in mammals is similar; in this case two brainstem nuclei, the nucleus prepositus hypoglossi and medial vestibular nucleus, have been identified as involved in integrator function. A subpopulation of these neurons have a threshold-linear relationship with eye position (McFarland and Fuchs 1992), as in goldfish area I. In addition, connections with the mammalian cerebellum are necessary; lesions in the flocculus region of the cerebellum are known to compromise integrator function (Nagao and Kitazawa 2003; Zee *et al.* 1981).

Although the nuclei that are involved in oculomotor integration have been identified in lesion studies, the mechanisms by which it is performed are still unresolved. Authors have proposed a variety of mechanisms that range from single cells to large

networks. The most common idea is that integration is performed by excitatory recurrent connections in a network, which effectively lengthen the time constant (equivalently, decrease the rate of decay) so that a signal is held for a longer time. Alternatively, a network connected by bilateral inhibition can perform the same function, with the added bonus that it explains the bilateral push-pull behavior observed in integrator neurons, while not integrating a background input (Cannon *et al.* 1983). Whether the integration occurs via excitatory or inhibitory connections, the principle is the same: tune the dominant eigenvalue of the linear network so that it is very close to zero and the firing rate persists for a time that is long relative to the intrinsic membrane time constant of the individual neurons. These networks are prone to instability, because a small change in network parameters could push the eigenvalue over the boundary of stability and produce an exponentially growing response. This has been proposed as a mechanism for the slow phase of jerk nystagmus.

For some time, positive feedback integrator models have been dismissed as a source for congenital nystagmus (CN) because they lack linear oscillatory behavior and therefore cannot produce pendular nystagmus. Instead, authors have proposed abnormalities in the saccadic (e.g. Akman *et al.* 2005) or smooth pursuit (Jacobs and Dell'Osso 2004) systems. However, a critical look at the integrator models used by these authors reveals why they are not sufficient to encompass both jerk and pendular nystagmus; the connections in these models are symmetric (the synaptic weight from A to B is equal to the synaptic weight from B to A). In the most pared-down case, a one-dimensional integrator is used, which has only one eigenvalue which must be real.

Relaxing this constraint to include asymmetric - or more broadly *non-normal* - linear systems opens up the possibility for more flexible integrator models. Recently, there have been a number of works on the impact of non-normality on neural integration (Ganguli *et al.* 2009). While non-normal operators have a long history in the applied mathematics literature (in areas such as fluid dynamics and mathematical physics – see Trefethen and Embree 2005), their use in neuroscience is relatively new to this author's knowledge.

In this paper, we review an integrator model that exploits non-normality both for gain modification and to generate pendular CN (Barreiro *et al*. 2009). We also review some of the other recent work in non-normal integrators, not specific to CN, and provide some thoughts on how this work might be relevant to future CN research.

II. Integrator models for CN

Several authors have proposed models for CN that rely on a malfunction of the neural integrator. An early idea, based on the model that the integrator functions though positive feedback, is that an inappropriately large gain in this feedback loop can produce jerk nysgamus (Dell'Osso and Daroff 1981, Arnold and Robinson 1991). While these models give a mechanism to produce jerk nystagmus with a perturbation of a normal oculomotor system, they do not supply a mechanism for pendular nystagmus.

As an alternative to linear recurrent feedback models, Optican and Zee (1984) proposed a non-linear model that incorporates two feedback loops, one on eye position and another on eye velocity. The eye position feedback loop is positive and serves to lengthen the time constant of what would otherwise be a quickly decaying signal. The eye velocity feedback is negative and also helps to lengthen the integrator time constant; its

inclusion means that tuning of the position feedback loop can be less precise. CN arises when the eye velocity feedback loop switches sign and becomes sufficiently positive. This model can predict many waveforms of CN; however it predicts only small amplitude, low frequency pendular nystagmus, and that only close to the null position. The amplitude and frequency are limited because the oscillation depends on an interaction between eye position and a nonlinearity in the position feedback loop. Tusa *et al.* (1992) modifies this model to include a loop under voluntary control, to account for the fact that some patients can suppress their nystagmus.

Harris (1995) proposes that a co-operation between smooth pursuit and integrator malfunctions can produce CN. Later, Harris and Berry (2006) suggest that CN arises as an adaptive response to delayed sensory development, and identify integrator malfunction as being present in at least jerk nystagmus (C.M. Harris, personal communication), but do not yet present a network mechanism for producing the waveform.

In the context of modeling central myelin disorders, Das *et al.* (2000) generated pendular nystagmus by including an asymmetric proprioceptive feedback loop to a neural network model of the integrator.

Recently Anastasio and Gad (2007) proposed a linear integrator network that can produce both jerk and pendular nystagmus, as well as normal function. One characteristic that they sought to include was plasticity; a mammalian integrator generally maintains a time constant of about 20 seconds, but can be trained to become leaky or unstable. The integrator can also adjust its gain in reponse to vestibulo-ocular stimuli by a factor of two when trained with visually conflicting stimuli (Tiliket *et al.* 1994). Anastasio and Gad's

integrator included connectivity with cerebellar Purkinje cells, motivated by the observation that cerebellar lesions abolish VOR plasticity (Nagao and Kitazawa 2003). Their goal was to understand how sparse cerebellar connections can regulate the two basic characteristics of the integrator: time constant and gain. They started with a symmetric brainstem network based on Cannon *et al.*'s model of reciprocal inhibition across the midline and added asymmetric connections to cerebellar Purkinje cells. The resulting network, properly tuned, could independently regulate time constant and gain by tuning cerebellar connection weights.

Anastasio and Gad also note that mistuning can produce oscillations; in a sample survey of changes to a properly functioning network, the authors altered one brainstem-to-cerebellum connection, readjusted the cerebellar weights to regain the correct 20 second time constant, and tested the reponse of the network. In roughly one half (18 out of 40) the network became oscillatory; in the remaining networks gain was reduced.

These two observations (gain plasticity and oscillations) are intimately connected, as a full analysis in a later paper by Barreiro *et al.* (2009) would show. Cerebellar connections introduce network asymmetry to the model; while projections from the brainstem to the cerebellum are dense, projections from the cerebellum to the brainstem are sparse (Babalian and Vidal 2000). In addition, each side of the cerebellar flocculus receives input from both sides of the brainstem, but projects only to its own side of the brainstem (Büttner-Ennever 1988). So any particular Purkinje cell may receive input from many vestibular neurons from both sides of the brainstem, but will only output onto a few (in Anastasio and Gad's model, only one) ipsilateral vestibular neurons.

We can write a general system of linear differential equations as

$$\frac{d\vec{\mathbf{v}}}{dt} = \mathbf{A}\vec{\mathbf{v}}, \quad \vec{\mathbf{v}}(0) = \vec{\mathbf{b}}. \tag{3}$$

Assuming that the eigenvalues are distinct, the behavior of a linear system of differential equations is governed by its eigenvalues and eigenvectors:

$$\vec{\mathbf{v}}(t) = \sum_i \vec{\mathbf{e}}_\mathbf{i} \frac{\langle \vec{\mathbf{f}}_\mathbf{i}, \vec{\mathbf{b}} \rangle}{\langle \vec{\mathbf{f}}_\mathbf{i}, \vec{\mathbf{e}}_\mathbf{i} \rangle} e^{\lambda_i t} \tag{4}$$

where the left (**f**) and right (**e**) eigenvectors are defined by

$$\begin{aligned} \mathbf{A}\vec{\mathbf{e}}_i &= \lambda_i \vec{\mathbf{e}}_i \\ \mathbf{A}^T \vec{\mathbf{f}}_i &= \lambda_i \vec{\mathbf{f}}_i \end{aligned} \tag{5}$$

where $(\bullet)^T$ denotes the transpose and $\langle \bullet, \bullet \rangle$ denotes the dot product. If there are repeated eigenvalues, a modified expression holds.

A real n x n matrix is *normal* if it commutes with its adjoint:

$$\mathbf{A}\mathbf{A}^T = \mathbf{A}^T\mathbf{A}. \tag{6}$$

In the case of a normal matrix, we are guaranteed to have n true eigenvalues (both real and complex) and the eigenvectors are guaranteed to be orthogonal to one another (and as a consequence, the right and left eigenvectors coincide; $\vec{\mathbf{e}}_i = \vec{\mathbf{f}}_i$).

Nonnormal matrices do not satisfy Eq. 6, and while their eigenvectors have a property called bi-orthogonality ($\mathbf{e_i}$ is orthongonal to $\mathbf{f_j}$), they do not have to be orthogonal to one another. In fact, a system could have two eigenvectors very close

together; in this case, the scalar product $\left\langle \vec{\mathbf{f}}_1, \vec{\mathbf{e}}_1 \right\rangle$ - the denominator of one of the terms in Eq. 4 - must be close to zero.

An example is the following:

$$\mathbf{M} = \begin{bmatrix} -0.05 & 10 \\ 0 & -0.06 \end{bmatrix}. \tag{7}$$

The eigenvalue and eigenvectors are given by

$$\begin{aligned} &\lambda_1 = -0.05, \lambda_2 = -0.06 \\ &\vec{\mathbf{e}}_1 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \vec{\mathbf{f}}_1 = \begin{bmatrix} 0.001 \\ 1 \end{bmatrix}, \vec{\mathbf{e}}_2 = \begin{bmatrix} -1 \\ 0.001 \end{bmatrix}, \vec{\mathbf{f}}_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \end{aligned} \tag{8}$$

If we solve the corresponding differential equation with the initial condition

$$\vec{\mathbf{b}} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \tag{9}$$

and look at the first component of the result, we get the blue curve in Figure 1. What we observe is that although the system is linear, and all the eigenvalues are negative, and the initial condition is small, there is a dramatic transient growth: a “gain”.

Figure (1) here

One important aspect of this observation is that it requires the two eigenvalues to be close together; if we try the same exercise with the matrix

$$\mathbf{M}_2 = \begin{bmatrix} -0.05 & 10 \\ 0 & -1.05 \end{bmatrix}, \tag{10}$$

we will find the gain to be substantially reduced.

This same mechanism is operating in the model of Anastasio and Gad (2007). Anastasio and Gad originally considered relatively large networks (20-400 vestibular neurons on each side of the midline). Barreiro *et al.* (2009) considered a smaller version that retained the essential properties:

$$\mathbf{A} = \begin{bmatrix} \mathbf{T} & -\rho_1\vec{\mathbf{u}}_1 & -\rho_2\vec{\mathbf{u}}_2 \\ \vec{\mathbf{w}}_1^T & -1 & 0 \\ \vec{\mathbf{w}}_2^T & 0 & -1 \end{bmatrix} \tag{11}$$

where **T** is a 6 x 6, negative definite symmetric matrix and $\vec{\mathbf{u}}_1, \vec{\mathbf{u}}_2, \vec{\mathbf{w}}_1, \vec{\mathbf{w}}_2$ are 6 x 1 column vectors; **u** vectors represent cerebellum-to-brainstem weights, and **w** vectors represent brainstem-to-cerebellum weights.

The model is tuned by varying the parameters ρ_1 and ρ_2, which represent the weights of connections from the cerebellum. It is straightforward to see that this matrix is non-normal, as the matrices $\mathbf{A}^T\mathbf{A}$ and $\mathbf{A}\mathbf{A}^T$ will only be equal under very special circumstances; this would require, for example, that $\mathbf{w}_2^T\mathbf{w}_2 = \rho_2^2$. Because only two parameters are allowed to change, this system admits an analysis by viewing the cerebellar connections as perturbations to an underlying system in which these connections are set to zero. The characteristic polynomial of the matrix can be written as a polynomial in ρ_1, ρ_2 as follows:

$$\det(\mathbf{A} - \lambda\mathbf{I}) = D(\lambda) + P_1(\lambda)\rho_1 + P_2(\lambda)\rho_2 + Q(\lambda)\rho_1\rho_2 \tag{12}$$

where D, P_1, P_2, and Q are polynomials in the variable λ.

The unassuming equation

$$\det(\mathbf{A} - \lambda\mathbf{I}) = 0 \tag{13}$$

turns out to give a lot of information about the qualitative behavior of this linear system, which can be viewed in a nice graphical way. For each fixed value of λ, that equation describes a curve (or pair of curves) in the $(\rho 1, \rho 2)$ plane. For example, an integrator with a 20 second time constant must have an eigenvalue of $\lambda = -0.05$; input that projects onto the corresponding eigenvector will be maintained with the correct decay rate (see Eq. 4). Therefore properly functioning integrators should live on, or near (if we tolerate some variation in the time constant), the curve in Eq. 12 with parameter $\lambda = -0.05$.

If I want to know where I have a double eigenvalue of value λ, I look for the simultaneous solution of Eq. 12 and its derivative with respect to λ;

$$\begin{aligned} D(\lambda) + P_1(\lambda)\rho_1 + P_2(\lambda)\rho_2 + Q(\lambda)\rho_1\rho_2 &= 0 \\ D'(\lambda) + P_1'(\lambda)\rho_1 + P_2'(\lambda)\rho_2 + Q'(\lambda)\rho_1\rho_2 &= 0 \end{aligned} \tag{14}$$

This yields a curve (or set of curves) in the $(\rho 1, \rho 2)$ plane. We refer to this as the "envelope" because it is the classic envelope of a family of curves (here, the constant eigenvalue curves).

A third curve which will be useful is the location of purely imaginary eigenvalues. Because our matrix **A** is real, complex eigenvalues must occur in complex conjugate

pairs (example). They are purely imaginary if the following equations are satisfied simultaneously for some real ω:

$$\begin{aligned} \mathrm{Re}\left(\det\left(\mathbf{A} - i\omega\mathbf{I}\right)\right) &= 0 \\ \mathrm{Im}\left(\det\left(\mathbf{A} - i\omega\mathbf{I}\right)\right) &= 0 \end{aligned} \tag{15}$$

On one side of this curve, the eigenvalues have negative real part (stable); on the other side they have positive real part (unstable); therefore we refer to this curve as the "Hopf curve" in reference to a dynamical bifurcation in which, among other things, an equilibrium point changes stability.

Also of interest are points where a real eigenvalue switches dominance with a complex pair. We refer to this as the "dominance curve". We don't have a nice way to compute it analytically but we can compute it numerically and have shown it on the following diagrams.

Because we are only changing two parameters, we can look at these curves in the two-dimensional plane, as in Figure 2:

(Figure (2) here)

Each of the curves that we described is important because there is a qualitative change in the spectrum as the curve is crossed; in crossing the envelope, a pair of real eigenvalues must come together and split off into the complex plane (or a complex pair must merge onto the real axis, in the other direction); in crossing the Hopf curve a pair of complex eigenvalues must change stability. In the open regions bounded by a set of curves, the

relative positions of the eigenvalues can be mapped onto the complex plane as in the insets in Figure 2: each inset shows a cartoon of the complex plane and the "x"s shows the locations of the three dominant eigenvalues. The information preserved in each region is the relative ordering of the real parts of the eigenvalues, and whether each one is in the left- or right-hand plane.

Three sets of parameters are identified with colored dots; systems with these parameters generated the traces shown in Figure 3. Note the intersection of the eigenvalue line (red dashed) with the envelope curve (blue dotted); this is the location of a double eigenvalue of the appropriate value. As the parameters are varied to approach this point, the gain increases. If the network strays to the left on this parameter plane, it will become unstable (showing the slow phase of jerk nystagmus); if it strays to the right it becomes oscillatory.

(Figure (3) here)

Recall that Anastasio and Gad perturbed their network, readjusted their parameters so that the dominant eigenvalue was reset to the correct value, and evaluated the network again. They found either reduced gain or oscillations, each for roughly half of the perturbations. This can be readily understood from the figures we have seen: one of their perturbations – adding or removing a vestibular-to-cerebellar connection – is equivalent to a subtle shift of the curves we have described under a fixed (ρ_1, ρ_2) value pair. The curves might just be distorted, without topological change, or there could be a

modest topological change. In the former case, the new point would be found further down on the dashed line, at a more distant point from the double eigenvalue point – hence a decreased gain. In the latter case, the topology of the graph could be shifted so that the dashed line makes a short excursion across the Hopf curve, representing a regime in which we have, simultaneously an integrating eigenvalue as well as a pair of oscillating eigenvalues. Such a phase plane might look like the following:

(Figure (4) here)

As in the previous parameter space diagram (Figure 2), the Hopf curve and eigenvalue line intersect the envelope at nearby points. However, instead of continuing in different directions, these curves have a second intersection in the lower left quadrant of the diagram; this intersection represents a point where there is an integrating eigenvalue, as well as a superimposed, non-decaying oscillation.

(Figure (5) here)

If a controller attempts to increase gain in this network, the response will be very different than that displayed in Figure 3. The colored dots in Figure 4 mark parameter settings that generated the traces. As the parameter settings approach the intersection point with the Hopf curve, a complex pair will approach the integrating eigenvalue from the left-hand side. At first this will result in a transient, quickly decaying oscillation, and

eventually a steady oscillation riding on top of the otherwise normal integrator response. The system cannot increase gain arbitrarily without activating this oscillation.

In the network we have just described, non-normality comes from the anatomical asymmetry of the brainstem-cerebellar conections. In fact, just about feedback loop that makes an excursion outside the brainstem would be likely to add network asymmetry because it generally won't originate and terminate on the same cell, for example in the proprioceptive feedback loop of Das *et al*. (2000). However, our current understanding of the oculomotor integrator – at least in the goldfish – confirms that there are asymmetric structures in the brainstem itself. Aksay *et al.* (2007) propose that there are two separate excitatory populations, one that governs each side of the visual plane, each with an associated inhibitory population that inhibits the opposite side, so that only one population is active at a time. Because inhibitory cells on the right (left) side projects to excitatory cells on the left (right) side but there is clearly not a returning inhibitory connection, the system seems asymmetric. However, because of nonlinear cut-offs in the system, the system appears to perform as though the right side excitatory population drives the right side inhibitory population, which suppresses the left side excitatory population; but without a feedback connection (M.S. Goldman, personal communication). In other words, the three populations may act like a chain of three (normal) networks, with the identity of the active networks shifting as the eyes move from left to right, rather than an asymmetric feedback network.

III. Other non-normal integrator models

The linear integrator models that we have considered remember through recurrent feedback. Generally, such models require precise tuning of the dominant eigenvalue of the system to a number close to zero (or even *two* eigenvalues, as in Barreiro *et al.* (2009) in high gain regimes). As such the models must be very finely tuned; synaptic weights in a recurrent network model such as Cannon *et al.* (1983) must be tuned to within a fraction of a percent. This has motivated criticism that these models are simply not biologically plausible. In contrast, some interesting recent work has identified the ability of "functionally feedforward" networks to do integration. These networks can exhibit enhanced robustness to mistuning and noise when compared to recurrent networks; these properties are ultimately due to their non-normal linear dynamics.

An equivalent definition of an *n* x *n* normal matrix (defined in Eq. 6) is that it has a complete set of orthogonal eigenvectors; in other words it can be written as follows:

$$\mathbf{A} = \mathbf{V}\mathbf{D}\mathbf{V}^T \tag{16}$$

where the columns of **V** are orthogonal and **D** is diagonal. In the differential equation $\frac{d\vec{\mathbf{x}}}{dt} = \mathbf{A}\vec{\mathbf{x}}$, each eigenvector acts on its own with no interaction with the others, which we can check by looking at the evolution of $\vec{\mathbf{v}}_i^T\vec{\mathbf{x}}$, the coefficient of projection of **x** onto the *i*th eigenvector. This is illustrated in Figure 6a.

A nonnormal matrix cannot be diagonalized in this way; it will look more like

$$\mathbf{A} = \mathbf{U}\mathbf{D}\mathbf{V}^T \tag{17}$$

where **U** and **V** are different (this is equivalent to having "different left and right eigenvectors") and it is not obvious how to disentangle the contribution of one

eigenvector with another. However, a nonnormal matrix can be treated with a Schur decomposition, where instead of a diagonal matrix we have an lower (or upper)-triangular matrix.

$$\mathbf{A} = \mathbf{U}\mathbf{T}\mathbf{U}^T \tag{18}$$

Now the evolution of **x** can be cleanly separated into *n* distinct modes, although modes can interact through the matrix **T**.

(**Figure (6) here)**

For example, consider the matrix

$$\mathbf{T} = \begin{bmatrix} -1 & 0 & 0 & 0 \\ 1 & -1 & 0 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & 1 & -1 \end{bmatrix} \tag{19}$$

where there are ones along the off-diagonal. This is a network in which the output of cell 1 feeds into cell 2, the output of cell 2 feeds into cell 3, and so forth, illustrated in Figure 6b. Goldman (2009) identifies this as a "pure" feedforward network, in which activity propagates from one pattern to the next, in contrast with the feedback network of Figure 6a. Goldman (2009) analyzes this network (with arbitrary size) and finds that this network can integrate a pulse for a time proportional to the size of the network. This is possible despite the fact that the eigenvalues of this system are all far from zero (in (19), all eigenvalues of **T** are -1).

We can easily imagine transforming by some orthogonal matrix **U** to yield a dense matrix in which all the cells interact superficially; however the essential dynamics remain the same. Mathematically, this is a matrix which has one true eigenvector and *n-1* generalized eigenvectors which behave like the following:

$$\mathbf{A}\vec{\mathbf{v}}_{i-1} = \lambda\vec{\mathbf{v}}_i \tag{20}$$

Goldman's pure feed-forward integrator has several advantages for integration. The degenerate structure of its eigenspace allows it to lengthen the time an input is stored without fine-tuning the dominant eigenvalue of the integrator. In addition, because information is not stored in a single eigenvector, each neuron in the integrator has a complex time-varying response. This opens the possibility to match experiments which show time-varying persistent activity during working memory tasks (for example Batuev (1994)).

One might notice that the presence of a degenerate eigenspace, itself, requires some careful tuning, because the network must have multiple identical eigenvalues. However, Goldman shows that the pure feed-forward network described above is very robust to mistuning when compared with a feedback network (see Figure 7 in Goldman 2009); a perfect feedforward structure does not seem to be necessary. Among other reasons, the feedforward network cannot sustain exponentially growing activity because its eigenvalues are far from the boundary of stability. Another attractive feature of feedforward architecture is that it can be created by architectural constraints (for example, the fact that nucleus A projects to nucleus B but not vice-versa), rather than mediation of the parameters in all-to-all circuitry.

Ganguli *et al.* (2008) identify another attractive characteristic of functionally feed-forward networks: robustness to noise. They examine the information carrying capacity of a general discrete time network, in which the dynamics are given by acting on the network states with a linear transformation followed by a sigmoidal nonlinearity. They contrast normal and non-normal networks and find that non-normal matrices, and especially the pure feedforward network (which they refer to as a "delay line") can maintain a signal-to-noise ratio (SNR) for a time proportional to the number of states of the network; in contrast a network with normal linear dynamics maintains this ratio for a period of time which does *not* increase as the number of network states increase.

While these other nonnormal network models have computational benefits for integration, it remains to be seen whether such a model – say a oculomotor integrator model relying primarily on feed-forward structure to produce persistent input – can be adapted to produce CN waveforms. One natural idea is to use the network of Anastasio and Gad (2007) and Barreiro *et al.* (2009), replacing the vestibular connections with a feedforward chain. Because the spectrum no longer determines the interesting feature of the response (maintenance of input for a time longer than the intrinsic time scale of individual cells), a different method of analysis will be needed. Another possibility is to construct a network with two feedforward chains; one that performs integration and another that oscillates.

Acknowledgements

A.B. would like to thank Mark Goldman for useful discussions, and the Nystagmus Network and her fellow participants for a very stimulating workshop.

References

Akman, O.E., D.S. Broomhead, R.V. Abadi, and R.A. Clement (2005). "Eye movement instabilities and nystagmus can be predicted by a nonlinear dynamics model of the saccadic system," *Journal of Mathematical Biology* **51**, pp. 661-694.

Aksay, E., R. Barker, H.S. Seung, and D.W. Tank (2000). "Anatomy and Discharge Properties of Pre-Motor Neurons in the Goldfish Medulla That Have Eye-Position Signals During Fixations," *Journal of Neurophysiology* **84**, pp. 1035-1049.

Aksay, E., I. Olasagasti, B.D. Mensh, R. Baker, M.S. Goldman, and D.W. Tank (2007). "Function dissection of circuitry in a neural integrator," *Nature Neuroscience* **10**, pp. 494-504.

Anastasio, T.J. and Y.P. Gad (2007). "Sparse cerebellar innervation can morph the dynamics of a model oculomotor neural integrator," *Journal of Computational Neuroscience* **22**, pp. 239-254.

Arnold, D.B. and D.A. Robinson (1991). "A learning network of the neural integrator of the oculomotor system," *Biological Cybernetics* 64, pp. 447-454.

Babalian, A.L. and P.P. Vidal (2000). "Floccular modulation of vestibuloocular pathways and cerebellum-related plasticity: An in vitro whole brain study," *Journal of Neurophysiology* **84**, pp. 2514–2528.

Barreiro, A.K., J.C. Bronski and T.J. Anastasio (2009). "Bifurcation theory explains waveform variability in a congenital eye movement disorder," *Journal of Computational Neuroscience* **26**, pp. 321-329.

Batuev, A.S. (1994). Two neuronal systems involved in short-term spatial memory in monkeys. *Acta Neurobiologiae Experimentalis.* (Wars.) **54**, 335–344.

Büttner-Ennever, J.A. (1988). *Neuroanatomy of the oculomotor system*. Amsterdam: Elsevier.

Cannon, S.C., D.A. Robinson and S. Shamma (1983). "A proposed neural network for the integrator of the oculomotor system," *Biological Cybernetics* **49**, pp. 127-136.

Das, V.E., P. Oruganti, P.D. Kramer and R.J. Leigh (2000). "Experimental tests of a neural-network model for ocular oscillations caused by a disease of central myelin," *Experimental Brain Research* **133**, pp. 189-197.

Dell'Osso, L.F. and R.B. Daroff (1981). "Clinical disorders of oculomotor movement," In: Zuber, B.L., ed. *Models of oculomotor behavior and control*, CRC Press, pp. 233-256.

Ganguli, S., D. Huh and H. Sompolinsky (2008). "Memory traces in dynamical systems," *Proceedings of the National Academy of Sciences* **105**, pp. 18970-18975.

Ganguli, S. and P. Latham (2009). "Feedforward to the Past: The Relation between Neuronal Connectivity, Amplification, and Short-Term Memory," *Neuron* **61**, pp. 499-501.

Goldman, M.S. (2009). "Memory without Feedback in a Neural Network," *Neuron* **61**, pp. 621-634.

Harris, C.M. (1995). "Problems modeling congenital nystagmus: towards a new model," In: Findlay J.M. et al. eds. *Eye Movement Research: Processes, Mechanisms and Applications*. Amsterdam: Elsevier, pp. 239-253.

Harris, C. M. and D.L. Berry (2006). "A distal model of congenital nystagmus as nonlinear adaptive oscillations," *Nonlinear Dynamics* **44**, pp. 367–380.

Jacobs, J.B. and L.F. Dell'Osso (2004). "Congenital nystagmus: Hypothesis for its genesis and complex waveforms within a behavioral ocular motor system model," *Journal of Vision* **4**, pp. 604-625.

Nagao, S. and H. Kitazawa (2003). "Effects of reversible shutdown of the monkey flocculus on the retention of adaptation of the horizontal vestibulo-ocular reflex," *Neuroscience* **118**, pp. 563-570.

Optican, L.M. and D.S. Zee (1984). "A hypothetical explanation of congenital nystagmus," *Biological Cybernetics* **50**, pp. 119-134.

Romo, R., C.D. Brody, A. Hernandez and L. Lemus (1999). "Neuronal correlates of parametric working memory in the prefrontal cortex," *Nature* **399**, pp. 470-473.

Shadlen, M.N. and W.T. Newsome (2001). "Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey," *Journal of Neurophysiology* **86**, pp. 1916-1936.

Tiliket, C., M. Shelhammer, D. Roberts, and D.S. Zee (1994). "Short-term vestibulo-ocular reflex adaptation in humans. I. Effect on the ocular motor velocity-to-position neural integrator," *Experimental Brain Research* **100**, pp. 316-327.

Trefethen, L.N. and M. Embree (2005). *Spectra and Psuedospectra: the Behavior of Nonnormal Matrices and Operators*, Princeton University Press.

Tusa, R.J., D.S. Brody, T.C. Hain and H.J. Simonsz (1992). “Voluntary control of congenital nystagmus,” *Clinical Vision Sciences* **7** (3), pp. 195-210.

Zee, D.S., A. Yamazaki, P.H. Butler and G. Gücer (1981). “Effects of ablation of flocculus and parafloccullus on eye movements in primate,” *Journal of Neurophysiology* **46**, pp. 878-899.

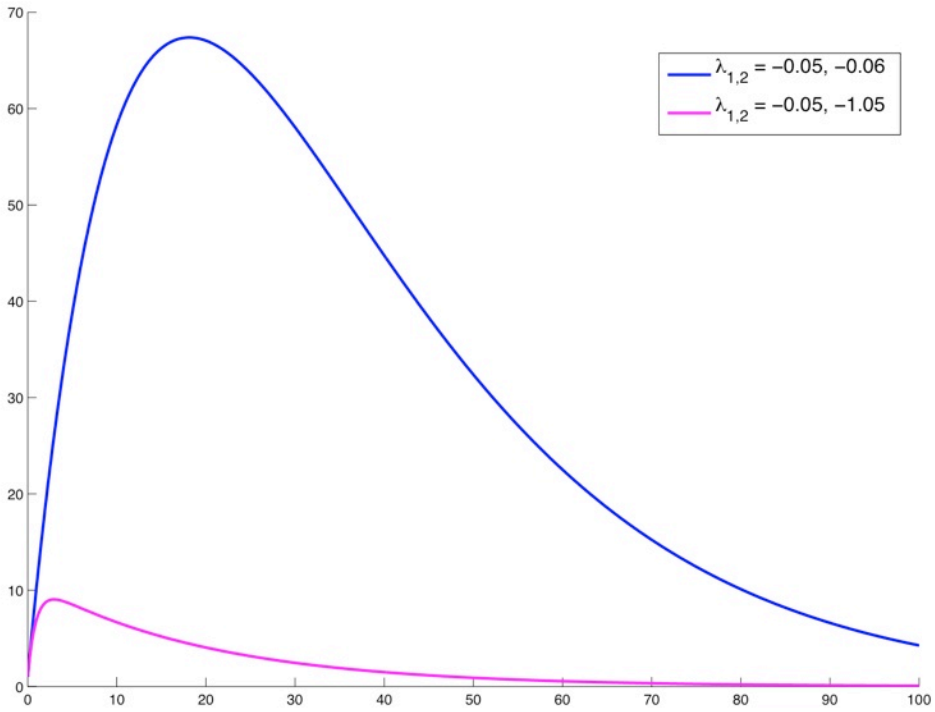

Figure 1: Response of a non-normal network to an impulsive input at time zero (blue). A transient gain drives the system response to 70 times the input strength. The magnitude of this gain depends on the proximity of two eigenvalues; a system with more widely spaced eigenvalues will show lesser gain (magenta).

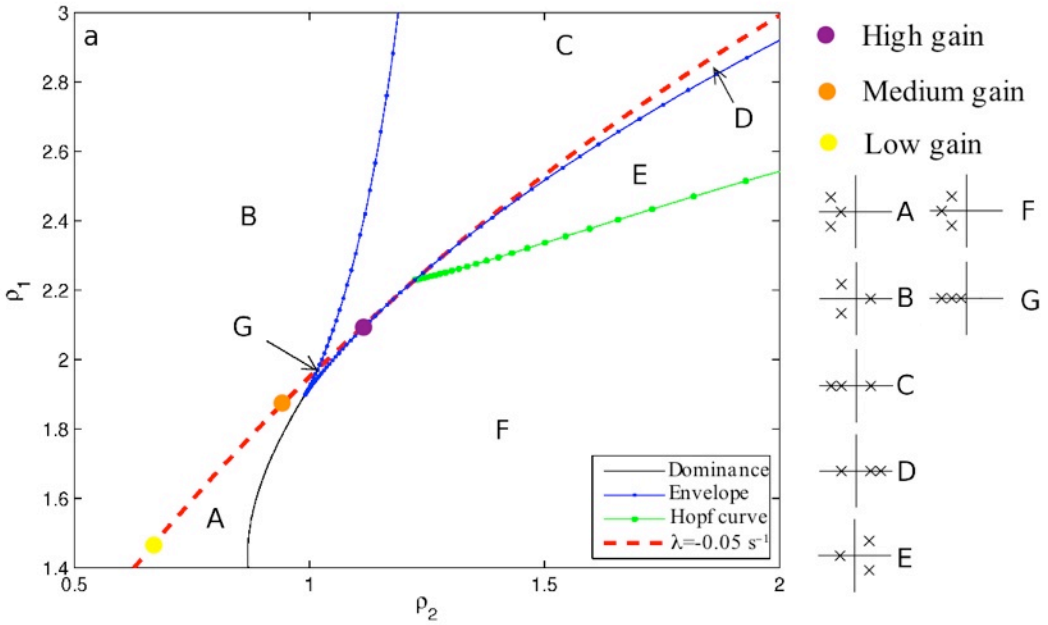

Figure 2: Parameter space of a network that performs integrator function correctly.

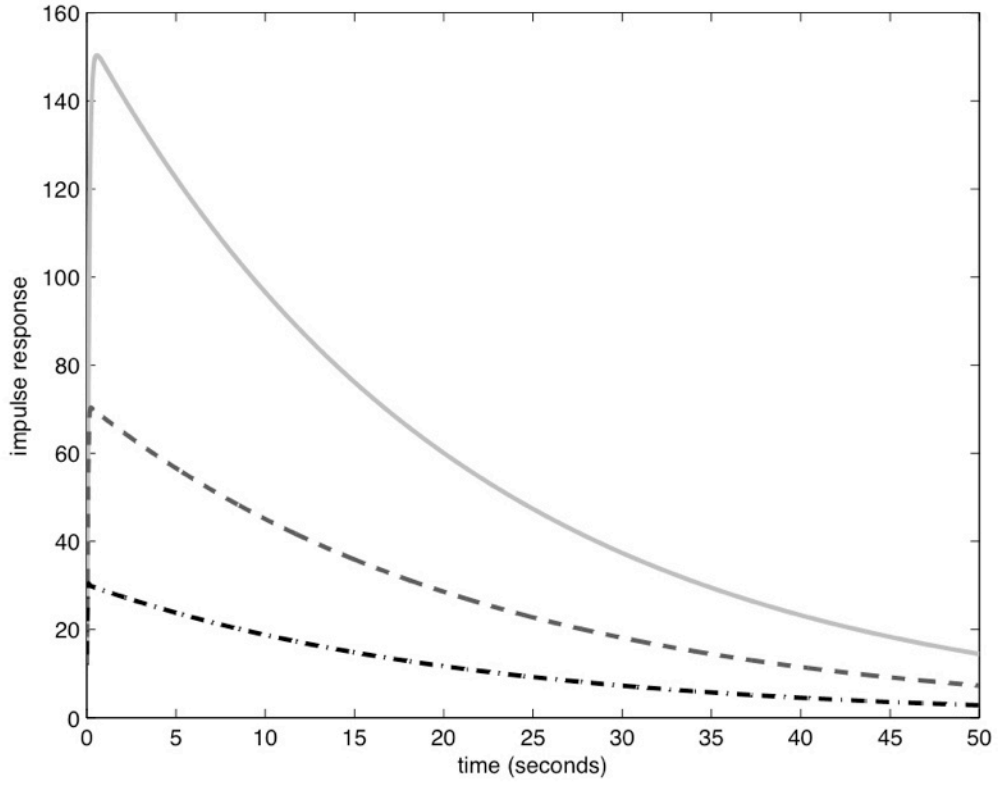

Figure 3: Response of the first network as gain is increased.

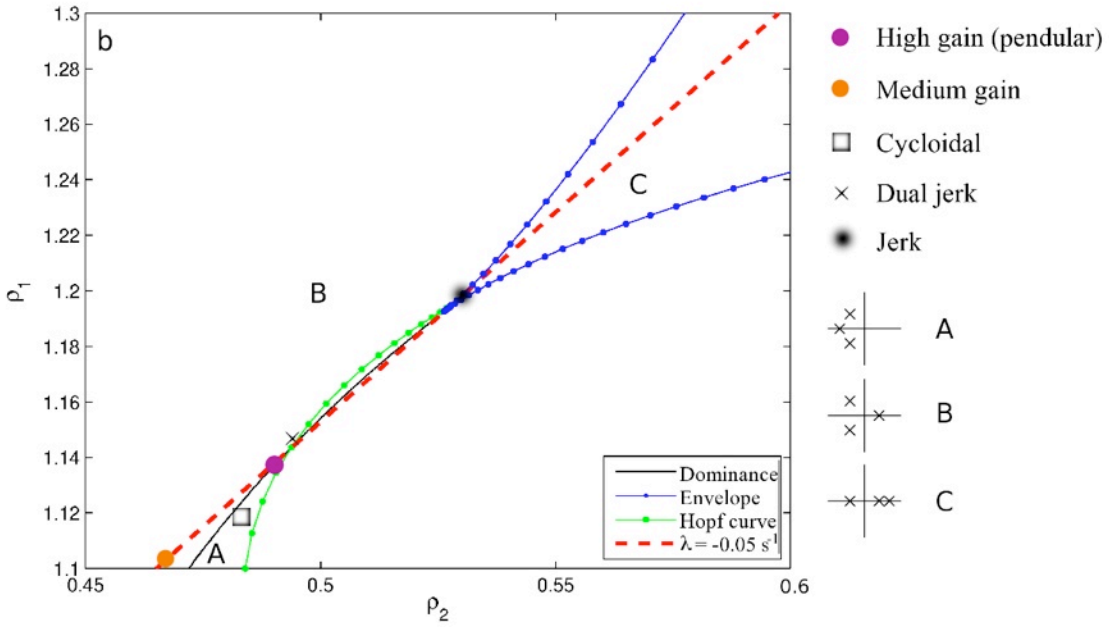

Figure 4: Parameter space of a network that displays pendular nystagmus as gain is increased. (The "low gain" location is located off the diagram).

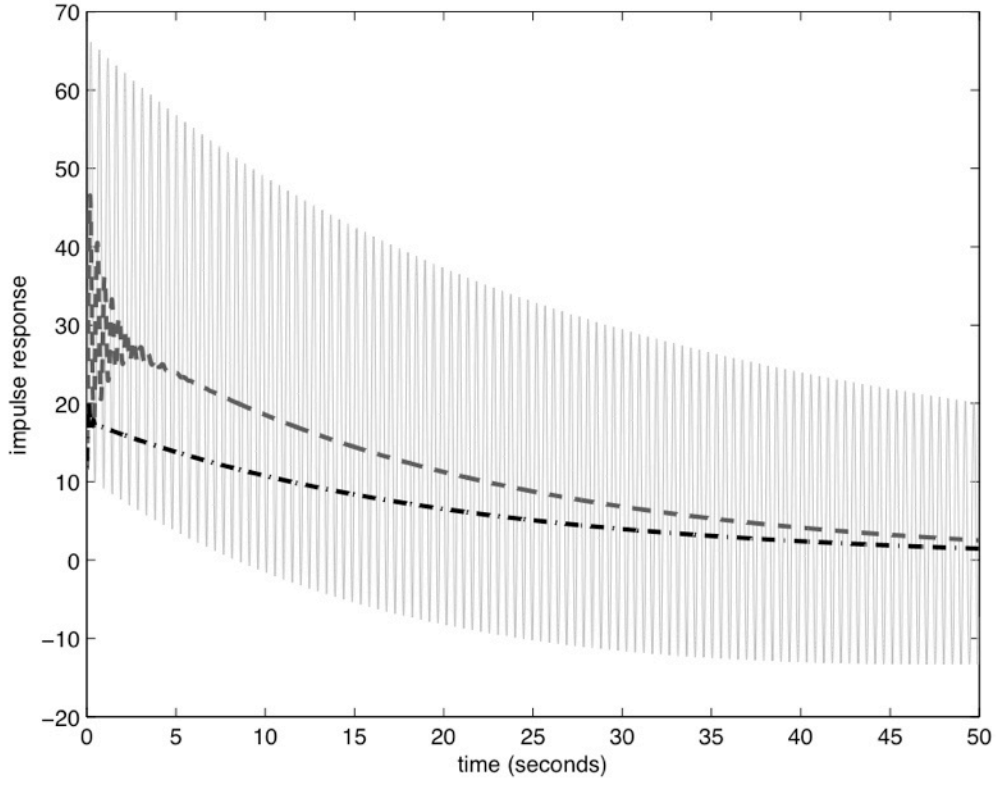

Figure 5: Response of second network to an impulse at time zero, as gain increases, at parameter choices shown in Figure 4. As the parameters ascend the 20 second time constant curve, the integrating eigenvalue exchanges dominance with a complex pair and the network response becomes oscillatory.

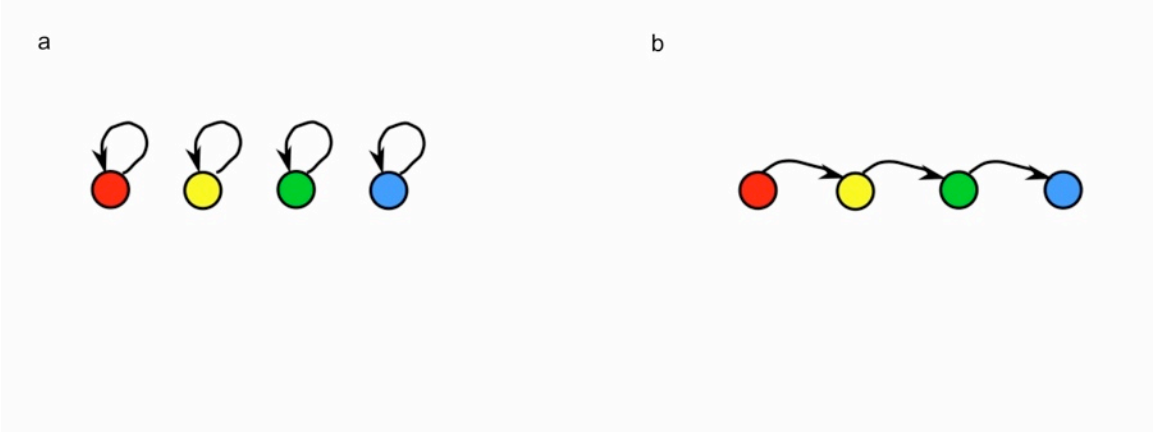

Figure 6: Schematic of dynamics in a feedback (a) vs. feedforward (b) network. In a feedback (specifically a normal) network, the vector space of possible network responses can be divided into *n* orthogonal modes, each of which acts independently of the others. In a nonnormal network, different eigenvectors interact; (b) shows a pure feedforward network.