

Divisive inhibition in recurrent networks

Frances S Chance and L F Abbott

Volen Center for Complex Systems and Department of Biology, Brandeis University, Waltham, MA 02454-9110, USA

Received 8 July 1999, in final form 21 February 2000

Abstract. Models of visual cortex suggest that response selectivity can arise from recurrent networks operating at high gain. However, such networks have a number of problematic features: (i) they operate perilously close to a point of instability, (ii) small changes in synaptic strength can dramatically modify the degree of amplification, and (iii) they respond slowly to rapidly changing stimuli. Divisive inhibition, acting through interneurons that are themselves divisively inhibited, can solve these problems without degrading the selectivity of a recurrent network.

1. Introduction

The extensive recurrent connectivity of cortical circuitry raises interesting questions about the relative importance of recurrent and feedforward inputs in shaping the selectivity of cortical neurons. Models of primary visual cortex have suggested that intracortical connections generate the orientation tuning (Ben-Yishai *et al* 1995, Somers *et al* 1995, Sompolinsky and Shapley 1997, Carandini and Ringach 1997) and direction selectivity (Douglas *et al* 1995, Suarez *et al* 1995, Maex and Orban 1996) of simple cells, and the spatial-phase insensitivity of complex cells (Chance *et al* 1999). These models rely on strong amplification of specific response selectivities by recurrent circuitry, which requires the networks to operate at high gain. At high gain, the amplification produced by a recurrent network is extremely sensitive to the strength of its synapses. Small changes in the strengths of the recurrent connections can dramatically change the gain or even result in the network becoming unstable. These observations are particularly worrisome because synapses in biological networks are highly plastic. Thus, while recurrent models can account for basic selectivity properties of cortical neurons, they may not have the inherent stability required of a viable cortical model.

Another problematic feature of recurrent networks operating at high gain is the sluggishness of their responses to rapidly changing inputs, a feature known as critical slowing down. The dynamics of network models of the type we are considering is controlled by a time constant that we call τ_r , which is a fixed parameter. The physical significance of τ_r is not entirely clear (for a discussion see Ermentrout 1998). Many consider it to be related to the membrane time constant of the neurons being modelled. Another more tenable view is that τ_r is related to the time constant of synaptic transmission in the network (Treves 1993). These interpretations support values of τ_r anywhere in the range from one to tens of milliseconds. The actual response time of a recurrent network that is selectively amplifying with a given gain, however, is not τ_r , but τ_r times the gain. At high gain, this slowing can severely reduce the response of the network to transient or rapidly oscillating inputs. Recurrent models typically use a small time constant to compensate for the critical slowing down that occurs at high gain

(Ben-Yishai *et al* 1995, Chance *et al* 1999). For larger τ_r values, the roll-off at high frequency in these models is unrealistically severe. This is worrisome because the values used are at the lower limit of the acceptable range for this parameter.

In this paper, we study a form of inhibition that can solve these problems. Inhibition is often incorporated into recurrent network models in a subtractive fashion to stabilize networks with excessively strong excitation (Ben-Yishai *et al* 1995, Douglas *et al* 1995). However, stabilizing a network in this fashion does not resolve the issues discussed above. Stability is not improved unless changes in the recurrent excitatory synaptic strengths are precisely matched by modifications of inhibitory synapses. Furthermore, network responses are still severely slowed at high gain.

Subtractive inhibition in a network model simulates the effect of current flowing out of a postsynaptic neuron when its inhibitory synapses are activated. However, inhibitory synapses, like all other synapses, are not merely current sources, they change the conductances of their postsynaptic targets. The effect of these conductance changes can be represented in a network model using divisive inhibition. Inhibition that acts divisively, rather than subtractively, has been implemented in feedforward models of primary visual cortex to account for a number of response features (Heeger 1992, 1993, Carandini and Heeger 1994). In addition, Battaglia and Treves (1998) have shown that divisive inhibition can both stabilize and speed up processing in an autoassociative memory model. Here we show that divisive inhibition in a model with recurrently amplified feedforward responses can stabilize network activity for arbitrarily large excitatory coupling, and can, in certain cases, eliminate the critical slowing down of a high-gain network without modifying its selectivity.

2. Analytic results

We present both analytic and simulation results for recurrent networks with divisive inhibition. The analytic work is limited to the study of linear networks. The simulations presented in the next section are not restricted in this way and include the effects of a rectification nonlinearity. Each of the N neurons in either the linear or nonlinear models, labelled by an index i , receives feedforward input I_i , and recurrent input $\sum_j W_{ij}r_j$, where W_{ij} is the weight of the recurrent synapse from neuron j to neuron i . The firing rate of neuron i , r_i , is then determined, in the linear case, by the standard rate-model equation

$$\tau_r \frac{dr_i}{dt} = -r_i + I_i + \sum_{j=1}^N W_{ij}r_j. \quad (1)$$

To analyse equation (1), we assume that $W_{ij} = W_{ji}$ and write the firing rates in terms of a complete set of eigenvectors, $\sum_j W_{ij}\xi_j^\mu = \lambda_\mu \xi_i^\mu$ for $\mu = 1, 2, \dots, N$, where λ_μ are the eigenvalues. Writing $r_i = \sum_\mu a_\mu \xi_i^\mu$, the coefficients a_μ obey the uncoupled equations

$$\frac{\tau_r}{1 - \lambda_\mu} \frac{da_\mu}{dt} = -a_\mu + \frac{h_\mu}{1 - \lambda_\mu} \quad (2)$$

where $h_\mu = \sum_i I_i \xi_i^\mu$. For constant h_μ and $\lambda_\mu < 1$, a_μ approaches the value $h_\mu/(1 - \lambda_\mu)$ exponentially with time constant $\tau_r/(1 - \lambda_\mu)$. If the largest, or principal eigenvalue, which we label λ_1 , is just less than one, the coefficient a_1 will take a steady-state value $a_1 = h_1/(1 - \lambda_1)$ that is much larger than the input coefficient h_1 . Thus, the component of the input proportional to the principal eigenvector is selectively amplified by the network. If $\lambda_1 > 1$, the network is unstable and a_1 will grow without bound. The gain, which for a linear network is the factor multiplying h_1 in the expression for the steady-state value of a_1 , is $1/(1 - \lambda_1)$. High gain

thus requires $\lambda_1 \approx 1$, while stability requires $\lambda_1 < 1$. This is the stability problem. The approach of a_1 to its steady-state value is governed by the effective time constant of equation (2), $\tau_r/(1 - \lambda_1)$, rather than by τ_r . As recurrent amplification increases, the effective time constant grows proportional to the gain, which produces critical slowing down.

Both the stability and slowing-down problems arise from the fact that the recurrent excitation grows as the firing rates in the network increase. This growth can be suppressed through divisive inhibition. We introduce a single inhibitory neuron firing at rate R that divisively inhibits the recurrent input to every neuron in the network. The divisive inhibition does not have to come from just one neuron, but we consider this case for simplicity. With this form of divisive inhibition, the firing rates of the excitatory neurons of the network are given by

$$\tau_r \frac{dr_i}{dt} = -r_i + I_i + \frac{\sum_j W_{ij} r_j}{R}. \quad (3)$$

Note that we only divide the recurrent input, not the feedforward input, by R . This corresponds to a situation where the recurrent synaptic inputs are located at the ends of a dendritic tree along which inhibitory inputs shunt current flow, and feedforward inputs are located nearer to the soma and are unaffected by the shunting. Equation (3) is ill-defined if $R = 0$. In our simulations, we eliminate this problem by adding a constant B so that the last term in equation (3) becomes $\sum_j W_{ij} r_j / (R + B)$. However, this is not essential because the equations we use do not allow R to go to zero. As long as R and one or more of the network rates r_i are initialized to a value greater than zero, the denominator in the last term of equation (3) will never vanish. The analytic results are simpler if we leave the constant B out for the time being. Written in terms of the coefficients in an expansion of r_i in terms of eigenvectors of W , equation (3) is equivalent to

$$\tau_r \frac{da_\mu}{dt} = -a_\mu + h_\mu + \frac{\lambda_\mu a_\mu}{R}. \quad (4)$$

The steady-state values of these coefficients are determined by setting $da_\mu/dt = 0$, giving

$$a_\mu = \frac{h_\mu R}{R - \lambda_\mu}. \quad (5)$$

Because the principal ($\mu = 1$) mode is the first mode to become unstable as the gain of the network is increased, and it has the longest effective time constant, we design the inhibition to stabilize and speed up this particular mode. The most obvious way to attempt to do this is to let the excitatory neurons of the network drive the inhibitory neuron through synaptic weights that are proportional to the principal eigenvector. In this case,

$$\tau_R \frac{dR}{dt} = -R + G \sum_{i=1}^N \xi_i^1 r_i \quad (6)$$

where the parameter G determines the strength of the excitatory drive to the inhibitory neuron. However, it is easy to see that this form of inhibition does not satisfy the requirements needed to produce a fast and stable high-gain network. The steady-state value of R from equation (6) is $R = G \sum_i \xi_i^1 r_i = G a_1$. The steady-state value of the coefficient of the principal mode, a_1 , is determined by substituting this into equation (5) (for the case $\mu = 1$), which gives $a_1 = h_1 + \lambda_1 / G$. This is finite for all values of λ_1 , and thus the network is stable, but it has a gain of only one. The effect of the recurrent connections is simply to add a constant, input-independent term to the steady-state value of a_1 , which produces a background firing rate. Thus, divisive inhibition has stabilized the network, but at the expense of completely eliminating selective amplification.

This problem can be resolved if we modify equation (6) so that the feedforward input to the network divisively inhibits the excitatory input to the inhibitory neuron. This could arise if the feedforward input also was relayed to the inhibitory neuron through interneurons that shunt the excitatory input. We assume that input i contributes to the divisive inhibition with a strength proportional to ξ_i^1 , and that the effects of the different inputs add linearly. The resulting equation for R is

$$\tau_R \frac{dR}{dt} = -R + \frac{G \sum_i \xi_i^1 r_i}{\sum_i \xi_i^1 I_i + A} \quad (7)$$

where A is a constant, similar to the term B , introduced to prevent a divergence when the feedforward input is zero. A and B represent the sums of resting conductances and synaptic conductances due to any constant background activity. The steady-state value of R , in this case, is $R = Ga_1/(h_1 + A)$. Substituting this into equation (5), we find

$$a_\mu = \frac{h_\mu G a_1}{G a_1 - \lambda_\mu (h_1 + A)}. \quad (8)$$

The steady-state value of a_1 is determined by solving this equation for $\mu = 1$, which gives

$$a_1 = \left(1 + \frac{\lambda_1}{G}\right) h_1 + \frac{\lambda_1 A}{G}. \quad (9)$$

The gain of the network is thus $1 + \lambda_1/G$ and, if A and G are both small, the network can exhibit high gain without an excessively high background firing rate.

We can show that this form of divisive inhibition eliminates critical slowing down by first assuming that critical slowing down does occur, and then showing that this assumption is inconsistent. In other words, we argue by contradiction. If critical slowing down did occur, R would reach its steady-state value long before a_1 did. We could then substitute the steady-state value of R into equation (4) to find that the dynamics of a_1 are governed by

$$\tau_r \frac{da_1}{dt} = -a_1 + h_1 + \frac{\lambda_1 (h_1 + A)}{G}. \quad (10)$$

However, equation (10) indicates that the assumption that critical slowing down occurs is self-inconsistent, because the time constant in this equation is τ_r not $\tau_r/(1 - \lambda_1)$. Thus, critical slowing cannot occur.

Equation (9) shows that the a_1 mode remains finite for all values of λ_1 . If we substitute the steady-state value of a_1 into equation (5), we can determine the steady-state coefficients for the other eigenvectors of the system,

$$a_\mu = \frac{h_\mu (\lambda_1 (h_1 + A) + h_1 G)}{(\lambda_1 - \lambda_\mu) (h_1 + A) + h_1 G}. \quad (11)$$

This is finite for arbitrarily large λ_1 provided that $\lambda_\mu < \lambda_1$. It is instructive to consider the case $h_1 = 0$, for which equation (11) reduces to $a_\mu = h_\mu / (1 - \lambda_\mu / \lambda_1)$. This means that the modes with $\mu \neq 1$ are slightly more amplified than they were without divisive inhibition. However, the steady-state value for any mode that has $\lambda_\mu = \lambda_1$ diverges as $h_1 \rightarrow 0$. Equal eigenvalues are called degenerate, and this case is an important one because one of the models we consider in the simulation studies has degenerate principal eigenvectors.

3. Simulation results

The previous section discussed the behaviour of linear recurrent networks, but real neurons are inherently nonlinear if for no other reason than that they always have positive firing rates.

Our simulation results are based on two models of neurons in the primary visual cortex. These neurons are often characterized by their responses to grating images in which the luminance varies sinusoidally as a function of position. The responses of neurons in the primary visual cortex are sensitive to a number of attributes of such sinusoidal gratings. Of interest here are selectivity for the orientation of the grating and for its spatial phase, which specifies where the light and dark regions are located on the viewing screen. Neurons are divided into simple and complex classes on the basis of their spatial-phase selectivity. Both simple and complex cells display orientation selectivity, but only simple cells are strongly selective for spatial phase. Complex cells exhibit approximately spatial-phase invariant responses. These selectivities are most easily seen by plotting orientation and spatial-phase tuning curves, which are responses plotted as a function of either the orientation or the spatial phase of the visual stimulus. The two models we study both use recurrent amplification to generate realistic tuning curves, but they do so in opposite ways. The orientation model (based on the model of Ben-Yishai *et al* 1995) uses recurrent coupling to increase selectivity to orientation, allowing a weakly tuned input to generate a more sharply tuned response. The complex cell model (based on the model of Chance *et al* 1999) amplifies the phase-invariant component of the response, allowing an input that is tuned to spatial phase to generate a response that is relatively insensitive to spatial phase.

In these models, the firing rates are determined by an equation similar to equation (3) of the previous section,

$$\tau_r \frac{dr_i}{dt} = -r_i + \left[I_i + \frac{\sum_j W_{ij} r_j}{R + B} \right]_+ \quad (12)$$

A rectification nonlinearity has been introduced through the notation $[]_+$, which is defined as $[x]_+ = x$ if $x \geq 0$ and $[x]_+ = 0$ if $x < 0$. Also, for the simulations, we have added a constant B to the denominator of the interaction term in this equation to keep it well defined as $R \rightarrow 0$. In the simulations, we use the value $B = 1$. The overall magnitude of the firing rates in the models we consider is determined by the magnitude of the feedforward input. Instead of adjusting this to obtain realistic rates in Hz, we compute the firing rates in arbitrary units and scale them to the maximum firing rate obtained in each case. This allows us to plot our results in a way that is applicable to inputs of any magnitude.

In both the models we consider, R is determined by

$$\tau_R \frac{dR}{dt} = -R + \frac{G \sum_j r_j}{\sum_j I_j + A} \quad (13)$$

No rectification is needed in this equation because the excitatory firing rates and inputs we consider are always positive. Equation (13) corresponds to equation (7) with $\xi_i^1 = 1$. We will discuss the rationale for this choice when we analyse each of the models. The constant A in equation (13) assures that the last term on the right-hand side remains well-defined if the input goes to zero, and we use the value $A = 0.01$. In the simulations, $G = 0.1$, and we set $\tau_r = \tau_R$, using a value of either 1 or 5 ms for both of these time constants. We set $G = 0$ when we wish to compare our results with those of the original models without divisive inhibition.

3.1. Complex cell model

The complex cell model used here is a version of a previously published model (Chance *et al* 1999) that is simplified so that we only consider spatial-phase selectivity. The feedforward input to neuron i in equation (12), for this model, is

$$I_i = c [\cos(\Phi - \phi_i)]_+ \quad (14)$$




Figure 1. Model complex cell tuning curves with and without divisive inhibition. The top traces show the response of a model neuron as a function of the spatial phase of the grating stimulus. The bottom traces indicate the feedforward input to the model neuron resulting from visual stimuli at different spatial phases. On the left is the original model without divisive inhibition. On the right is the model with divisive inhibition included. The excitatory recurrent synaptic strength was adjusted so that the gain of the network is 20 in both cases.

where Φ is the spatial phase of the grating stimulus, ϕ_i is a parameter characterizing the spatial-phase preference of the input to neuron i , and the factor c represents the contrast of the visual stimulus.

Spatial-phase independent responses arise in this model because the principal mode that is highly amplified is invariant to spatial phase. This is achieved by making the recurrent weights independent of the preferred spatial phase parameters, setting

$$W_{ij} = \frac{g}{(N-1)}. \quad (15)$$

For these synaptic weights, the principal eigenmode has equal firing rates for all the neurons. This means that $\xi_i^1 = 1$ is the principal eigenvector, and equation (13) is completely equivalent to equation (7) for this model.

The lower panels in figure 1 show that the feedforward input to a neuron of this network model is tuned as a function of spatial phase, as determined by equation (14). The upper panels indicate that this spatial phase tuning is greatly reduced by the recurrent circuitry, producing the characteristic weak spatial-phase tuning of a complex cell. The left column of this figure shows the model without divisive inhibition, and the right panel shows what happens when divisive inhibition on both excitatory and inhibitory neurons is included. The introduction of divisive inhibition has no visible effect on the response selectivity of the model, so the model is still viable when modified in this way. We study the visible effects of divisive inhibition in a later section.

3.2. Orientation model

The orientation model is also a simplified version of a previously published model (Ben-Yishai *et al* 1995). In this case, the feedforward input to neuron i of the network is

$$I_i = c [1 - \epsilon + \epsilon \cos(2(\Theta - \theta_i))]_+ \quad (16)$$

where Θ represents the orientation of the stimulus, θ_i is the preferred orientation of neuron i , and $\epsilon = 0.1$ is a parameter that determines the tuning of the input. Again, c represents the




Figure 2. Orientation model tuning curves with and without divisive inhibition. The top traces depict the responses of a neuron as a function of the orientation angle of the visual stimulus. The bottom traces show the feedforward input resulting from visual stimuli of different orientations. On the left is the original model without divisive inhibition. On the right is the model with divisive inhibition included. The recurrent synaptic strength was adjusted so that the network operates at a gain of 20 in both cases.

contrast of the visual stimulus. The recurrent weights are chosen to amplify orientation tuning and are given by

$$W_{ij} = \frac{J_2 \cos(2(\theta_i - \theta_j))}{N} \quad (17)$$

where J_2 is a parameter that controls that amount of recurrent amplification. This matrix has two degenerate principal eigenvectors, $\cos(2\theta_i)$ and $\sin(2\theta_i)$. Note that, despite the fact that $\xi_i^1 = 1$ is not a principal eigenvector in this model, we use equation (13) to generate the inhibition for the orientation model. Thus, for this model, equation (13) is not equivalent to equation (7). We do this because the principal eigenvectors in the orientation model are degenerate. This degeneracy arises because the model is invariant to translations of the orientation angle variable, but the principal eigenvectors are not. In the analytic results, we showed that the divisive inhibition of equation (7) does not stabilize degenerate eigenvectors. It is important that the divisive inhibition does not destroy the translation invariance of the orientation model. Using the unweighted sums in equation (13) assures that translation invariance is not broken and, as we will see, the nonlinearity of the orientation model allows this form of divisive inhibition to stabilize both of the degenerate principal eigenvectors of the orientation model. This is because the nonlinearity in the model mixes the degenerate principal modes with the constant mode used to drive the inhibition. Similarly, the modulated and constant components of the input are linked through equation (16).

The response tuning curves for a neuron of the orientation model are plotted as a function of orientation angle in the upper panels of figure 2. Comparison with the weakly tuned input shown in the lower panels indicates that orientation tuning has been greatly enhanced by the recurrent coupling. The left column shows results in the model without divisive inhibition, while the right panel has the type of divisive inhibition we are analysing included. As before, the recurrent synaptic strengths were adjusted so that the network has the same gain in the two cases. The tuning curves of the model with and without divisive inhibition are identical.

3.3. Effects of divisive inhibition

Divisive inhibition does not modify the selectivity of the models we have discussed, but does improve their stability and temporal response properties. The results for stability are shown in figure 3, and the temporal response properties of the complex cell model are illustrated in figure 4.




Figure 3. Gain of the model networks as a function of overall recurrent synaptic strength. Left panel: the complex cell model without divisive inhibition is unstable for $g > g_{\max}$. With divisive inhibition, the gain of the network increases approximately linearly for $g > g_{\max}$. Right panel: the orientation model without divisive inhibition is unstable for $J_2 > 4$. With divisive inhibition included, gain grows approximately linearly for $J_2 > 4$. In both panels, gain of 20 is indicated by a dashed line. In the complex cell network without divisive inhibition, this is equivalent to $g/g_{\max} = 0.95$, and $g/g_{\max} = 2.81$ for the network with divisive inhibition, where g_{\max} is the maximum stable value of g in the network without divisive inhibition. In the orientation model, gain of 20 corresponds to $J_2 = 2.45$ and $J_2 = 2.764$ in the cases without and with divisive inhibition, respectively.




Figure 4. Responses of the complex cell model to a 2 Hz counterphase grating when $\tau_r = \tau_R = 1$ ms (top) or $\tau_r = \tau_R = 5$ ms (bottom). The left traces show the behaviour of the model without divisive inhibition, and the right traces with divisive inhibition. The excitatory synaptic strength was chosen in both cases so that the gain is 20. 100% maximum response indicates the maximum firing rate when the smaller time constant is used.

In the complex cell model, the amplified mode is one in which all firing rates are equal, so the gain of the network is defined as the sum of the output firing rates divided by the sum of the inputs. The gain of the network grows as a function of the principal eigenvalue, which is directly proportional to the overall strength of the excitatory recurrent connections, g . As shown in figure 3, the gain of the network without divisive inhibition grows without bound as g approaches a critical value, g_{\max} . When divisive inhibition on excitatory and inhibitory neurons is included, the gain of the network grows approximately linearly for $g > g_{\max}$. In the simulations of figure 1, the model network operated at a gain of 20, indicated in figure 3 by the dashed line. To operate at this gain without divisive inhibition, the network must be set close to the point of instability. With divisive inhibition, a stronger excitatory synaptic strength g is required to achieve the same gain, but there is no critical value of g above which the network becomes unstable. Furthermore, small changes in the value of g have only a modest effect on the gain of the network.

Similar behaviour is seen in the orientation model. Here the gain is calculated as the amplitude of the cosine that best fits the population activity profile divided by $c\epsilon$, the corresponding amplitude of the input. The overall strength of the recurrent connections is determined by the parameter J_2 . In the orientation model without divisive inhibition, the network is unstable for $J_2 > 4$, but gain grows approximately linearly when $J_2 > 4$ in the model with divisive inhibition.

To study the speeding up of the network by divisive inhibition, we examine the complex cell model responses to a counterphase grating (a grating in which the contrast is modulated sinusoidally in time). Complex cell responses to counterphase gratings oscillate at twice the frequency of the stimulus. The responses of the model with and without divisive inhibition with a time constant of either 1 or 5 ms are illustrated in figure 4. Increasing the time constant has a large effect on the network without divisive inhibition, as seen in the left panels of figure 4. For a time constant of 1 ms, the response is strongly modulated by the oscillating input, but for a larger time constant of 5 ms, the modulation amplitude is severely reduced. With divisive inhibition (right panels of figure 4), the degree of modulation is essentially the same with both time constants.

Divisive inhibition speeds up the responses of the complex cell model by acting directly on the amplified mode. Although the divisive inhibition in the orientation model does not respond directly to the degenerate amplified modes, a similar speeding up occurs if we examine its response to contrast oscillations. This occurs because the divisive inhibition varies with the constant part of the feedforward input, which is directly related to the contrast of the stimulus. On the other hand, the divisive inhibition we have introduced is invariant with respect to rotations of the orientation angle. As a result, it does not affect the response to changes solely in the orientation angle. For example, the response of the orientation model to a rotating grating is equally sluggish whether or not divisive inhibition is included.

4. Discussion

We have shown that an appropriate form of divisive inhibition can stabilize and speed up a recurrent network while maintaining a constant input-independent gain. This suggests that nonlinear inhibition could have important implications for the computational potential and dynamics of recurrent networks.

Nonlinear inhibition was implemented in the firing-rate models we considered as a divisive form of shunting inhibition. Significant shunting effects will only occur if synaptic inputs produce fairly large conductance changes in their postsynaptic targets. Recently, such large conductance changes have been reported in recordings of neurons responding to visual input

(Borg-Graham *et al* 1998), although earlier studies did not see such large effects (Berman *et al* 1991, Douglas *et al* 1988, Ferster and Jagadeesh 1992, Pei *et al* 1994, Ahmed *et al* 1997). Holt and Koch (1997) have pointed out that, even with large conductance changes, shunting effects may be difficult to achieve, even in fairly realistic neuronal models. Thus, the biophysical implementation of the proposed divisive shunting effects is far from clear.

The effects of divisive inhibition discussed here may be achievable using other forms of nonlinear inhibition. To maintain stability of a recurrent network, inhibition is only required to grow with the output of the network faster than excitation. Maintaining a relatively constant gain with changes in stimulus intensity requires that the inhibition also depend on the feedforward input to the network. Inhibition that depends on both the output and the input of a recurrent network but is not necessarily strictly divisive could produce, at least, an approximately constant gain. For example, inhibitory inputs along a dendritic shaft can generate exponential attenuation of an excitatory input entering near the end of the shaft (Abbott 1991), and such a form of inhibition can fill at least some of the requirements for the inhibition needed in our model.

As an example of a more general inhibitory scheme, consider a generalization of equation (3),

$$\tau_r \frac{dr_i}{dt} = I_i - r_i + F(R) \sum_{j=1}^N W_{ij} r_j \quad (18)$$

for an arbitrary function F . Written in terms of the coefficient of the principal eigenvector of the recurrent coupling matrix, this becomes

$$\tau_r \frac{da_1}{dt} = h_1 - a_1 + F(R) \lambda_1 a_1. \quad (19)$$

This equation will not allow a_1 to grow without bound for any value of λ_1 provided that R approaches its steady-state value rapidly enough, and F , expressed as a function of this steady-state R , goes to zero as $a_1 \rightarrow \infty$.

The model we propose requires both direct feedforward inhibition and a novel form of divisive inhibition of inhibition itself, which might be termed divisive disinhibition. There is ample evidence of feedforward inhibition in cortex (see, for example, Berman *et al* 1991) and some evidence for disinhibition in primary visual cortex and other areas (Somogyi *et al* 1983, Freund and Antal 1988). However, there is no direct evidence for the precise type of circuitry required in our model.

Our proposal involves solving both the stability and the slowing down problems with a single mechanism, divisive inhibition. It is possible that, instead, separate mechanisms might solve these two problems. Divisive inhibition speeds up a network in part by providing strong inhibition when the feedforward input decreases rapidly. Speeding up might be achieved by other mechanisms that produce strong inhibition in response to rapidly decreasing stimuli. For example, transient inhibition working in a push-pull arrangement with excitation might reproduce the speeding up of the network. Thus, biophysically plausible mechanisms exist to either realize or mimic the effects achieved in the model discussed here by divisive inhibition.

Acknowledgments

We thank one of the anonymous reviewers for helpful suggestions. Research supported by the Sloan Center for Theoretical Neurobiology at Brandeis University, the National Science Foundation (IBN-9817194), and the W M Keck Foundation.

References

- Abbott L F 1991 Realistic synaptic inputs for network models *Network: Comput. Neural Syst.* **2** 245–58
- Ahmed B, Allison J D, Douglas R J and Martin K A C 1997 An intracellular study of the contrast-dependence of neuronal activity in cat visual cortex *Cereb. Cortex* **7** 559–70
- Battaglia F P and Treves A 1998 Stable and rapid recurrent processing in realistic autoassociative memories *Neural Comput.* **10** 431–50
- Ben-Yishai R, Bar-Or L and Sompolinsky H 1995 Theory of orientation tuning in visual cortex *Proc. Natl Acad. Sci., USA* **92** 3844–8
- Berman N J, Douglas R J, Martin K A C and Whitteridge D 1991 Mechanisms of inhibition in cat visual cortex *J. Physiol.* **400** 697–722
- Borg-Graham L J, Monier C and Frégnac Y 1998 Visual input evokes transient and strong shunting inhibition in visual cortical neurons *Nature* **393** 369–73
- Carandini M and Heeger D J 1994 Summation and division by neurons in primate visual cortex *Science* **264** 1333–6
- Carandini M and Ringach D L 1997 Predictions of a recurrent model of orientation selectivity *Vision Res.* **37** 3061–71
- Chance F S, Nelson S B and Abbott L F 1999 Complex cells as cortically amplified simple cells *Nature Neurosci.* **2** 277–82
- Douglas R J, Koch C, Mahowald M, Martin K A C and Suarez H H 1995 Recurrent excitation in neocortical circuits. *Science* **269** 981–5
- Douglas R J, Martin K A C and Whitteridge D 1988 Selective responses of visual cortical cells do not depend on shunting inhibition *Nature* **332** 642–4
- Ermentrout B 1998 Neural networks as spatio-temporal pattern-forming systems *Rep. Prog. Phys.* **61** 353–430
- Ferster D and Jagadeesh B 1992 EPSP-IPSP interactions in cat visual cortex studied with *in vivo* whole-cell patch recording *J. Neurosci.* **12** 1262–74
- Freund T F and Antal M 1988 GABA-containing neurons in the septum control inhibitory interneurons in the hippocampus *Nature* **336** 170–3
- Heeger D J 1992 Normalization of cell responses in cat striate cortex *Visual Neurosci.* **9** 181–98
- 1993 Modeling simple-cell direction selectivity with normalized, half-squared, linear operators *J. Neurophysiol.* **70** 1885–98
- Holt G R and Koch C 1997 Shunting inhibition does not act divisively on firing rates *Neural Comput.* **9** 1001–13
- Maex R and Orban G A 1996 Model circuit of spiking neurons generating directional selectivity in simple cells *J. Neurophysiol.* **75** 1515–45
- Pei X, Vidyagarar T R, Volgushev M and Creutzfeldt O D 1994 Receptive field analysis and orientation selectivity of postsynaptic potentials of simple cells in cat visual cortex *J. Neurosci.* **14** 7130–40
- Somers D C, Nelson S B and Sur M 1995 An emergent model of orientation selectivity in cat visual cortical simple cells *J. Neurosci.* **15** 5448–65
- Somogyi P, Freund T F, Wu J Y and Smith A D 1983 The section-Golgi impregnation procedure. 2. Immunocytochemical demonstration of glutamate decarboxylase in Golgi-impregnated neurons and in their afferent synaptic boutons in the visual cortex of the cat *Neurosci.* **9** 475–90
- Sompolinsky H and Shapley R 1997 New perspectives on the mechanisms for orientation selectivity *Current Opinion Neurobiol.* **7** 514–22
- Suarez H, Koch C and Douglas R J 1995 Modeling direction selectivity of simple cells in striate visual cortex within the framework of the canonical microcircuit *J. Neurosci.* **15** 6700–19
- Treves A 1993 Mean-field analysis of neuronal spike dynamics *Network: Comput. Neural Syst.* **4** 259–84