

Computational differences between asymmetrical and symmetrical networks

Zhaoping Li and Peter Dayan

Gatsby Computational Neuroscience Unit, University College, 17 Queen Square, London WC1N 3AR, UK

Received 13 August 1998

Abstract. Symmetrically connected recurrent networks have recently been used as models of a host of neural computations. However, biological neural networks have asymmetrical connections, at the very least because of the separation between excitatory and inhibitory neurons in the brain. We study characteristic differences between asymmetrical networks and their symmetrical counterparts in cases for which they act as selective amplifiers for particular classes of input patterns. We show that the dramatically different dynamical behaviours to which they have access, often make the asymmetrical networks computationally superior. We illustrate our results in networks that selectively amplify oriented bars and smooth contours in visual inputs.

1. Introduction

A large class of nonlinear recurrent networks, including those studied by Grossberg (1988), the Hopfield net (Hopfield 1982, 1984), and those suggested in many more recent proposals for the head direction system (Zhang 1996), orientation tuning in primary visual cortex (Ben-Yishai *et al* 1995, Carandini and Ringach 1997, Mundel *et al* 1997, Pouget *et al*), eye position (Seung 1996), and spatial location in the hippocampus (Samsonovich and McNaughton 1997) make a key simplifying assumption that the connections between the neurons are symmetric (we call these S systems, for short), i.e. the synapses between any two interacting neurons have identical signs and strengths. Analysis is relatively straightforward in this case, since there is a Lyapunov (or energy) function (Cohen and Grossberg 1983, Hopfield 1982, 1984) that guarantees the convergence of the state of the network to an equilibrium point. However, the assumption of symmetry is broadly false in the brain. Networks in the brain are almost never symmetrical, if for no other reason than the separation between excitation and inhibition, notorious in the form of Dale's law. In fact, it has never been completely clear whether ignoring the polarities of cells is simplification or over-simplification. Networks with excitatory and inhibitory cells (EI systems, for short) have certainly long been studied (e.g. Ermentrout and Cowan 1979b), for instance from the perspective of pattern generation in invertebrates (e.g. Stein *et al* 1997) and oscillations in the thalamus (e.g. Destexhe *et al* 1993, Golomb *et al* 1996) and the olfactory system (e.g. Li and Hopfield 1989, Li 1995). Further, since the discovery of 40 Hz oscillations (or at least synchronization) amongst cells in primary visual cortex of anaesthetized cats (Gray *et al* 1989, Eckhorn *et al* 1988), oscillatory models of V1 involving separate excitatory and inhibitory cells have also been popular, mainly from the perspective of how the oscillations can be created and sustained and how they can be used for feature linking or binding (e.g. von der Malsburg 1981, 1988, Sompolinsky *et al* 1990, Sporns *et al* 1991,




Figure 1. Three effects that are observed and desired for the mapping between visual input and output and which constrain recurrent network interactions. The strengths of all the input bars are the same; the strengths of the output bars are proportional to the displayed widths of the bars, but normalized separately for each figure (which hides the comparative suppression of the texture).

König and Schillen 1991, Schillen and König 1991, König *et al* 1992, Murata and Shimizu 1993). However, the full scope for computing with dynamically stable behaviours such as limit cycles is not yet clear, and Lyapunov functions, which could render analysis tractable, do not exist for EI systems except in a few special cases (Li 1995, Seung *et al* 1998).

A main inspiration for our work is Li's nonlinear EI system that models how the primary visual cortex performs input contour enhancement and pre-attentive region segmentation (Li 1997, 1998). Figure 1 shows two key phenomena that are exhibited by orientation-tuned cells in area V1 of visual cortex (Knierim and van Essen 1992, Kapadia *et al* 1995) in response to the presentation of small edge segments that can be isolated, or parts of smooth contours or texture regions. First, the activities of cells whose inputs form parts of smooth contours that could be connected are *boosted* over those representing isolated edge segments. Second, the activities of cells in the centres of extended texture regions are comparatively *suppressed*. A third, which is computationally desirable, is that unlike the case of hallucinations (Ermentrout and Cowan 1979a), non-homogeneous spatial patterns of response should *not* spontaneously form in the central regions of uniform texture. These three phenomena tend to work against each other. A uniform texture is just an array of smooth contours, and so enhancing contours whilst suppressing textures requires both excitation between the contour segments and inhibition between segments of different contours. This competition between contour enhancement and texture suppression tends to lead to spontaneous pattern formation (Cowan 1982)—i.e. the more that smooth contours are amplified, the more likely it is that, given a texture, random fluctuations in activity favouring some contours over others will grow unstably. Indeed, studies by Braun *et al* (1994) had suggested that an S-system model of the cortex cannot stably perform contour enhancement unless mechanisms for which there is no neurobiological support are used. Li (1997, 1998) showed empirically that an EI system built using just the Wilson–Cowan equations (1972, 1973) can comfortably exhibit the three phenomena, and she used this model to address an extensive body of neurobiological and psychophysical data. This poses a question, which we now answer, as to what are some of the critical computational differences between EI and S systems.

The computational underpinning for contour enhancement and texture suppression is the operation of *selective amplification*—magnifying the response of the system to input patterns that form smooth contours and weakening responses to those that form homogeneous textures. Selective amplification also underlies the way that many recurrent networks for

orientation tuning work—selectively amplifying any component of the input that is well tuned in orientation space and rejecting other aspects of the input as noise (Suarez *et al* 1995, Ben-Yishai *et al* 1995, Pouget *et al* 1998). Therefore, in this paper we study the computational properties of a family of EI systems and their S-system counterparts as selective amplifiers. We show that EI systems can take advantage of non-trivial dynamical behaviour through delayed inhibitory feedback (i.e. giving limit cycles) in order to achieve much higher selective amplification factors than S systems. Crudely, the reason is that over the course of a limit cycle, units are sometimes above and sometimes below the activity (or firing) threshold. **Above** threshold, the favoured input patterns can be substantially amplified, even to the extent of leading to a tendency towards spontaneous pattern formation. However, **below** threshold, in response to homogeneous inputs, these tendencies are corrected.

In section 2, we describe the essentials of the EI systems and their symmetric counterparts. In section 3 we analyse the behaviour of what is about the simplest possible network, which has just two pairs of units. In section 4 we consider the more challenging problem of a network of units that collectively represent an angle variable such as the orientation of a bar of light. In section 5 we consider Li's (1997, 1998) original contour and region network that motivated our study.

2. Excitatory–inhibitory and symmetric networks

Consider a simple, but biologically significant, EI system in which excitatory and inhibitory cells come in pairs and, as is true neurobiologically, there are no ‘long-range’ connections from the inhibitory cells (Li 1997, 1998)

$$\dot{x}_i = -x_i + \sum_j J_{ij}g(x_j) - h(y_i) + I_i \quad (1)$$

$$\tau_y \dot{y}_i = -y_i + \sum_j W_{ij}g(x_j). \quad (2)$$

Here, x_i are the principal excitatory cells, which receive external or sensory input I_i , and generate the network outputs through activation functions $g(x_i)$; y_i are the inhibitory interneurons (taken, for simplicity, as having no external input) which inhibit the principal neurons through their activation function $h(y_i)$; τ_y is the time-constant for the inhibitory cells; and J_{ij} and W_{ij} are the output connections of the excitatory cells. For analytical convenience, we choose $g(x)$ as a threshold linear function

$$g(x) = [x - T]_+ = \begin{cases} x - T & \text{if } x > T \\ 0 & \text{otherwise} \end{cases}$$

and $h(y) = y - T_y$. However, the results are generally similar if $h(y)$ is also threshold linear. Note that $g(x)$ is the *only* nonlinearity in the system. All cells can additionally receive input noise. Note that neither of the Lyapunov theories of Li (1995) nor Seung *et al* (1998) applies to this case.

In the limit that the inhibitory cells are made infinitely fast ($\tau_y = 0$), they can be treated as if they are constantly at equilibrium

$$y_i = \sum_j W_{ij}g(x_j) \quad (3)$$

leaving the excitatory cells to interact directly with each other

$$\begin{aligned}\dot{x}_i &= -x_i + \sum_j J_{ij}g(x_j) - h \left(\sum_j W_{ij}g(x_j) \right) + I_i \\ &= -x_i + \sum_j (J_{ij} - W_{ij})g(x_j) + I_i + T_y.\end{aligned}\quad (4)$$

In this reduced system, the effective neural connections $J_{ij} - W_{ij}$ between any two cells x can be either excitatory or inhibitory, as in many abstract neural network models. We call this reduced system in equation (4) the *counterpart* of the original EI system. The two systems have the same fixed points; that is, $\dot{x} = \dot{y} = \mathbf{0}$ for the EI system and $\dot{x} = \mathbf{0}$ for the reduced system (with $\tau_y = 0$) happen at the same values of $\bar{x}$ (and $\bar{y}$). Since there are many ways of setting J_{ij} and W_{ij} in the EI system whilst keeping constant the effective weight in its reduced system, $J_{ij} - W_{ij}$, and the dynamics in the EI system take place in a space of higher dimensionality, one may intuitively expect the EI system to have a broader computational range. In cases for which the connection weights are symmetric ($J_{ij} = J_{ji}$, $W_{ij} = W_{ji}$), the reduced system is an S system. In such cases, however, the EI network is asymmetric because of the asymmetrical interactions between the two units in each pair. We study the differences between the behaviour of the full system in equations (1) and (2) (with $\tau_y = 1$) and the behaviour of the S system in equation (4) (with $\tau_y = 0$).

The response of either system to given inputs is governed by the location and linear stability of their fixed points. Note that the input–output sensitivity of both systems at a fixed point $\bar{x}$ is given by

$$d\bar{x} = (\mathbb{I} - \mathbf{J}\mathbf{D}_g + \mathbf{W}\mathbf{D}_g)^{-1} dI$$

where $\mathbb{I}$ is the identity matrix, $\mathbf{J}$ and $\mathbf{W}$ are the connection matrices, and the diagonal matrix $[\mathbf{D}_g]_{ii} = g'(\bar{x}_i)\dagger$. Although the locations of the fixed points are the same for the EI and S systems, the dynamical behaviour of the systems about those fixed points are quite different, and this is what leads to their differing computational power.

To analyse the stability of the fixed points, consider, for simplicity, the case that the matrices $\mathbf{J}\mathbf{D}_g$ and $\mathbf{W}\mathbf{D}_g$ commute. This means that they have a common set of eigenvectors, say with eigenvalues $\lambda_k^{\mathbf{J}}$ and $\lambda_k^{\mathbf{W}}$, respectively, for $k = 1, \dots, N$ where N is the dimension of x . The local deviations $\Delta x \equiv x - \bar{x}$ near the fixed points along each of the N eigenvectors of $\mathbf{J}\mathbf{D}_g$ and $\mathbf{W}\mathbf{D}_g$ will grow in time $\Delta x_k(t) = \Delta x_k(0)\mathbf{e}^{\gamma_k t}$ if the real parts of the following values are positive:

$$\begin{aligned}\gamma_k &= \gamma_k^{\text{EI}} \equiv -1 + \frac{1}{2}\lambda_k^{\mathbf{J}} \pm \left(\frac{1}{4}(\lambda_k^{\mathbf{J}})^2 - \lambda_k^{\mathbf{W}} \right)^{1/2} && \text{for the EI system} \\ \gamma_k &= \gamma_k^{\text{S}} \equiv -1 - \lambda_k^{\mathbf{W}} + \lambda_k^{\mathbf{J}} && \text{for the S system.}\end{aligned}$$

For the case of real $\lambda_k^{\mathbf{J}}$ and $\lambda_k^{\mathbf{W}}$, the fixed point is *less stable* in the EI system than in the reduced system. That is, an unstable fixed point in the reduced system, $\gamma_k^{\text{S}} > 0$, leads to an unstable fixed point in the EI system, $\gamma_k^{\text{EI}} > 0$, since $\frac{1}{4}(\lambda_k^{\mathbf{J}})^2 - \lambda_k^{\mathbf{W}} > (-1 + \frac{1}{2}\lambda_k^{\mathbf{J}})^2$. However, if γ_k^{EI} is complex ($\lambda_k^{\mathbf{W}} > \frac{1}{4}(\lambda_k^{\mathbf{J}})^2$), i.e. if the EI system exhibits (possibly unstable) oscillatory dynamics around the fixed point, then the reduced system is stable: $\gamma_k^{\text{S}} < -1 + \lambda_k^{\mathbf{J}} - \frac{1}{4}(\lambda_k^{\mathbf{J}})^2 < 0$. In the general case, for which $\mathbf{J}\mathbf{D}_g$ and $\mathbf{W}\mathbf{D}_g$ do not commute or when $\lambda_k^{\mathbf{J}}$ and $\lambda_k^{\mathbf{W}}$ are not real, the conclusion that the fixed point in the EI system is *less stable* than that in the reduced system is merely a conjecture. However, this conjecture is consistent with results from singular perturbation theory (e.g. Khalil 1996) that when the reduced system is stable, the original system in equations (1) and (2) is also stable as $\tau_y \rightarrow 0$ but may be unstable for larger τ_y .

$\dagger$ We ignore subtleties such as the non-differentiability of $g(x)$ at $x = T$ that do not materially affect the results.

3. The two-point system

A particularly simple case to consider has just two neurons (for the S system; two pairs of neurons for the EI system) and weights

$$\mathbf{J} = \begin{pmatrix} j_0 & j \\ j & j_0 \end{pmatrix} \quad \mathbf{W} = \begin{pmatrix} w_0 & w \\ w & w_0 \end{pmatrix}$$

The idea is that each node coarsely models a group of neurons, and the interactions between neurons within a group (j_0 and w_0) are qualitatively different from interactions between neurons between groups (j and w). The form of selective amplification here is that symmetric or ambiguous inputs $\mathbf{I}^a = I(1, 1)$ should be suppressed compared with asymmetric inputs $\mathbf{I}^b = I(1, 0)$ (and, equivalently, $I(0, 1)$). In particular, given $\mathbf{I}^a$, the system should not spontaneously generate a response with x_1 significantly different from x_2 . In terms of figure 1, $\mathbf{I}^a$ is analogous to the uniform texture and $\mathbf{I}^b$ to the isolated contour. Define the fixed points to be $\bar{x}_1^a = \bar{x}_2^a > T$ under $\mathbf{I}^a$ and $\bar{x}_1^b > T > \bar{x}_2^b$ under $\mathbf{I}^b$, where T is the threshold of the excitatory neurons. These relationships will be true across a wide range of input levels I .

We quantify the selective amplification of the networks by the ratio

$$R = \frac{d\bar{x}_1^b/dI}{d\bar{x}_1^a/dI} \quad (5)$$

where the terms $\bar{x}$ are averages or maxima over the outputs of the network. This compares the gains of the system to the input for $\mathbf{I}^a$ and $\mathbf{I}^b$. Large values imply high selective amplification. To be computationally useful, the S systems must converge to finite fixed points, in which case $\tilde{x} = \bar{x}$ and

$$R^S = \frac{1 + ((w_0 + w) - (j_0 + j))}{1 + (w_0 - j_0)} = 1 + \frac{(w - j)}{1 + (w_0 - j_0)}. \quad (6)$$

If an EI system undergoes limit cycles, then the location of its fixed points may only be poorly related to its actual output. We will therefore use the maximum or mean of the output of the network over a limit cycle as $\tilde{x}$. We will show that EI systems can stably sustain larger values of R than S systems.

Consider the S system. Since $\bar{x}_2^b$ is below threshold ($\bar{x}_2^b < T$) in response to the selective input $\mathbf{I}^b$, the stability of the fixed point, determined by $\bar{x}_1^b$ alone, is governed by the sign of

$$\gamma^S = -(1 + w_0 - j_0). \quad (7)$$

The stability of the response to the unselective input $\mathbf{I}^a$ is governed by

$$\gamma_{\pm}^S \equiv -(1 + (w_0 \pm w) - (j_0 \pm j)) \quad (8)$$

for the two *modes* of deviation $\Delta x_{\pm} \equiv (x_1 - \bar{x}_1^a) \pm (x_2 - \bar{x}_2^a)$ around fixed point a .

We derive constraints on the maximum value of the selectivity ratio R^S of the S system from constraints on γ^S and $\gamma_{\pm}^S$. First, since we only consider cases when the input-output relationship $d\bar{x}/dI$ of the fixed points ($d\bar{x}_1^a/dI = -(1/\gamma_+^S)$ and $d\bar{x}_1^b/dI = -(1/\gamma^S)$) is well defined, we have to have that $\gamma^S < 0$ and $\gamma_+^S < 0$. Second, in response to $\mathbf{I}^a$, we require that the Δx_- mode does *not* grow, as otherwise symmetry between x_1 and x_2 would spontaneously break. Given the existence of stable $\bar{x}^b$ under $\mathbf{I}^b$, dynamic system theory dictates that the Δx_- mode becomes unstable when two additional, stable, and uneven fixed points $\bar{x}_1^a \neq \bar{x}_2^a$ for (the even) input $\mathbf{I}^a$ appear. Hence the motion trajectory of the system will approach one of these stable uneven fixed points from the unstable even fixed point. Avoiding this requires that $\gamma_-^S < 0$. From equation (8), this means that $w - j < 1 + w_0 - j_0$, and therefore that

$$R^S < 2.$$



The symmetry preserving network(A) I^a (B) I^b **The symmetry breaking network**(C) I^a (D) I^b 

Figure 2. Phase portraits for the S system in the two-point case. (A, B) Evolution in response to $I^a \propto (1, 1)$ and $I^b \propto (1, 0)$ for parameters for which the response to I^a is stably symmetric. (C, D) Evolution in response to I^a and I^b for parameters for which the symmetric response to I^a is unstable, inducing two extra equilibrium points. The dotted lines show the thresholds T for $g(x)$.

Figure 2 shows phase portraits and the equilibrium points of the S system under input I^a and I^b for the two different parameter regions.

As we have described, the EI system has exactly the same fixed points as the S system, but there are parameters for which the fixed points can be stable for the S system but unstable for the EI system. The stability around the symmetric fixed point under I^a is governed by

$$\gamma_{\pm}^{\text{EI}} = -1 + \frac{1}{2}(j_0 \pm j) \pm \sqrt{\frac{1}{4}(j_0 \pm j)^2 - (w_0 \pm w)}$$


while that of the asymmetric fixed point under I^a (if it exists) or I^b is controlled by

$$\gamma^{\text{EI}} = -1 + \frac{1}{2}j_0 \pm \sqrt{\frac{1}{4}j_0^2 - w_0}.$$


Consequently, when there are three fixed points under I^a , all of them can be unstable in the EI system, and the motion trajectory cannot converge to any of them. In this case, when both the Δx_+ and Δx_- modes around the symmetric fixed point $\bar{x}_1^a = \bar{x}_2^a$ are unstable, the global dynamics can constrain the motion trajectory to a limit cycle around the fixed points. If $x_1^a \approx x_2^a$ on this limit cycle, then the EI system will not break symmetry, while potentially giving a high selective amplification ratio $R^{\text{EI}} \gg 2$. Figure 3 demonstrates the performance of

Response to $I^a = I(1, 1)$


(A)



(B)

**Response to $I^b = I(1, 0)$**

(C)



(D)




Figure 3. Projections of the response of the EI system. (A, B) Evolution of response to I^a . Plots of (A) x_1 versus y_1 and (B) $g(x_1) - g(x_2)$ (solid) and $g(x_1) + g(x_2)$ (dotted) versus time show that the $x_1 = x_2$ mode dominates and the growth of $x_1 - x_2$ when both units are above threshold (the downward ‘blips’ in the lower curve in (B) are strongly suppressed when x_1 and x_2 are both below threshold. (C, D) Evolution of the response to I^b . Here, the response of x_1 always dominates that of x_2 over oscillations. The difference between $g(x_1) + g(x_2)$ and $g(x_1) - g(x_2)$ is too small to be evident on the figure. Note the difference in scales between (A, B) and (C, D). Here $j_0 = 2.1$, $j = 0.4$, $w_0 = 1.11$ and $w = 0.9$.

the EI system in this regime. Figure 3(A,B) shows various aspects of the response to input I^a which should be comparatively suppressed. The system oscillates in such a way that x_1 and x_2 tend to be extremely similar (including being synchronized). Figure 3(C,D) shows the same aspects of the response to I^b , which should be amplified. Again the network oscillates, and, although $g(x_2)$ is not driven completely to zero (it peaks at 15), it is very strongly dominated by $g(x_1)$, and further, the overall response is much stronger than in figure 3(A,B). Note also the difference in the oscillation period—the frequency is much lower in response to I^b than I^a .

The phase-space plot in figure 4 (which expands on that in figure 2(C)) illustrates the pertinent difference between the EI and S systems in response to the symmetric input pattern I^a . When $\mathbf{J}$ and $\mathbf{W}$ are strong enough to provide substantial amplification of I^b , the S system can only roll down the local energy landscape

$$E(x) = -\frac{1}{2} \sum_{ij} (J_{ij} - W_{ij}) \Delta x_i \Delta x_j + \frac{1}{2} \sum_i \Delta x_i^2 + E(\bar{x}^a)$$




Figure 4. Phase-space plot of the motion trajectory of the S system under input $I^a \propto (1, 1)$. Amplifying sufficiently the asymmetric inputs $I^b \propto (1, 0), (0, 1)$ leads to the creation of two energy wells (marked by ✱) which are the two asymmetric fixed points under input I^a . This makes the symmetric fixed point (marked by ◊) unstable, and actually a saddle point in the energy landscape that diverts all motion trajectories towards the energy wells. There is no energy landscape in the EI system. Its fixed points (also marked by ✱) can be unstable and unapproachable. This makes the motion trajectory oscillate (into the y dimensions) around the fixed point ◊, whilst preserving $x_1 \approx x_2$ and thus not breaking symmetry.

(when $h(y)$ is linear) for $x = \bar{x}^a + \Delta x$, from the point $\bar{x}^a$ (◊), which is a saddle point, since the Hessian $\partial^2 E / \partial \Delta x_i \partial \Delta x_j$ has eigenvalues $-\gamma_+^S > 0$ and $-\gamma_-^S < 0$, to one of the two stable fixed points, and thereby break the input symmetry. However, the EI system can resort to global limit cycles (on which $x_1(t) \approx x_2(t)$) between unstable fixed points, and so maintain symmetry. The conditions under which this happens in the EI system are:

- (a) At the symmetric fixed point under input I^a , while the Δx_- mode is guaranteed to be unstable because, by design, it is unstable in the S system, the Δx_+ mode should be unstable and oscillatory such that the Δx_- mode does not dominate the motion trajectory and break the overall symmetry of the system.
- (b) Under the ambiguous input I^a , the asymmetric fixed points (whose existence is guaranteed from the S system) should be unstable, to ensure that the motion trajectory will not converge to them.
- (c) Perturbations in the direction of $\Delta x_1 = -\Delta x_2$ about the limit cycle defined by $x_1 = x_2$ should shrink under the global dynamics, as otherwise the overall behaviour will be asymmetric.

The last condition is particularly interesting since it can be that the Δx_- mode is locally *more* unstable (at the symmetric fixed point) than the Δx_+ mode, since the Δx_- mode is more strongly suppressed when the motion trajectory enters the subthreshold region $x_1 < T$ and $x_2 < T$ (because of the location of its fixed point). As we can see in figure 3(A,B), this acts to suppress any overall growth in the Δx_- mode. Since the asymmetric fixed point under I^b is just as unstable as that under I^a , the EI system responds to asymmetric input I^b also by a stable limit cycle around the asymmetric fixed point.

Using the mean responses of the system during a cycle to define $\tilde{x}$, the selective amplification ratio in figure 3 is $R^{EI} = 97$, which is significantly higher than the $R^S = 2$ available from the S system. One can analyse the three conditions theoretically (though there appears to be no closed-form solution to the third constraint), and then choose parameters for which the selectivity ratio is greatest. For instance, figure 5 shows the range of achievable




Figure 5. Selectivity ratio R as a function of w_0 and w for $j_0 = 2.1$ and $j = 0.4$ for the EI system. The ratio is based on the maximal responses of the network during a limit cycle; results for the mean response are similar. The ratio is shown as 0 for values of w_0 and w for which one or more of the conditions is violated. The largest value shown is $R^{\text{EI}} = 103$, which is significantly greater than the maximum value $R^{\text{S}} = 2$ for the S system.

ratios as a function of w_0 and w for $j_0 = 2.1$, $j = 0.4$. The steep peak comes from the region around $w_0 \sim j_0 - 1$. To reiterate, the S system is not appropriately stable for these parameters. Clearly, very high selectivity ratios are achievable. Note that this analysis says nothing about the transient behaviour of the system as a function of the initial conditions. This and the oscillation frequency are also under ready control of the parameters.

This simple example shows that the EI system is superior to the S system, at least for the computation of the selective amplification of particular input patterns without hallucinations or other gross distortions of the input. If, however, spontaneous symmetry breaking is desirable for some particular computation, the EI system can easily achieve this too. The EI system has extra degrees of freedom over the counterpart S system in that $\mathbf{J}$ and $\mathbf{W}$ can both be specified (subject to a given difference $\mathbf{J} - \mathbf{W}$) rather than only the difference itself. In fact, it can be shown in this two-point case that the EI system can reproduce qualitatively all behaviours of the S system, i.e. with the same fixed points and the same linear stability (in *all* modes). The one exception to this is that if the demands of the computation require that the ambiguous input I^a be comparatively amplified and the input I^b be comparatively suppressed using an overly strong self-inhibition term w_0 , then the EI system has to be designed to respond to I^a with oscillations along which $x_1(t) \simeq x_2(t)$.

4. The orientation system

One recent application of symmetric recurrent networks has been to the generation of orientation tuning in primary visual cortex. Here, neural units i have preferred orientations $\theta_i = (i - N/2)\pi/N$ for $i = 1, \dots, N$ (the angles range in $[-\pi/2, \pi/2]$ since direction is ignored). Under ideal and noiseless conditions, an underlying orientation θ^* generates input

to individual units of $I_i = f(\theta_i - \theta^*)$, where $f(\cdot)$ is the input tuning function, which is usually unimodal and centred around zero. In reality, of course, the input is corrupted by noise of various sorts. The network should take noisy (and perhaps weakly tuned) input and selectively amplify the component $f(\theta_i - \theta^*)$ that represents θ^* in the input. Based on the analysis above, we can expect that if an S network amplifies a tuned input enough, then it will break input symmetry given an untuned input and thus hallucinate a tuned response. However, an EI system can maintain untuned and suppressed responses to untuned inputs to reach far higher amplification ratios. We study an abstract version of this problem, and do not attempt to match the exact tuning widths or neuronal oscillation frequencies recorded in experiments.

Consider first a simple EI system for orientation tuning, patterned after the cosine S-system network of Ben-Yishai *et al* (1995). In the simplest case, the connection matrices $\mathbf{J}$ and $\mathbf{W}$ are the Töplitz:

$$\begin{aligned} J_{ij} &= \frac{1}{N} (A + B \cos(2(\theta_i - \theta_j))) \\ W_{ij} &= \frac{1}{N} C. \end{aligned} \quad (9)$$

This is a handy form for the weights, since the net input sums in equations (1) and (2) are functions of just the zeroth- and second-order Fourier transforms of the thresholded input. Making W_{ij} a constant is solely for analytical convenience—we have also simulated systems with cosine-tuned connections from the excitatory cells to the inhibitory cells. For simplicity, assume an input of the form

$$I_i = a + b \cos(2\theta_i) \quad (10)$$

generated by an underlying orientation $\theta^* = 0$. In this case, the fixed point of the network is known to take the form

$$\bar{x}_i = \alpha + \beta \cos(2\theta_i) \quad (11)$$

where α and β are determined by a and b . Also, for $T = 1$,

$$\begin{aligned} g(\bar{x}_i) &= [\alpha - 1 + \beta \cos(2\theta_i)]_+ \\ &= \beta [\cos(2\theta_i) - \cos(2\theta_c)]_+ \end{aligned} \quad (12)$$

where θ_c is a cut-off. Note that the form in equation (12) is only valid for $\theta_c < \pi/2$.

In the same way that we designed the two-point system to amplify contentful patterns such as $I^b \propto (1, 0)$ selectively compared with the featureless pattern $I^a \propto (1, 1)$, we would like the orientation network to amplify patterns for which $b \gg 0$ in equation (10) selectively over those for which $b \sim 0$. In fact, in this case, we can also expect the network to filter out any higher spatial frequencies in the input that come from noise (see figure 6(A)), although, as Pouget *et al* (1998) discuss, the statistical optimality of this depends on the actual noise process perturbing the inputs.

Ben-Yishai *et al* (1995) analysed in some detail the behaviour of the S-system version of this network, which has weights $J_{ij} - W_{ij} = \frac{1}{N} [(A - C) + B \cos(2(\theta_i - \theta_j))]$. These authors were particularly interested in a regime they called the *marginal phase*, in which even in the absence of tuned input $b = 0$, the network spontaneously forms a pattern of the form $x_i = \alpha + \beta \cos(2(\theta_i - \phi))$, for arbitrary ϕ . In terms of our analysis of the two-point system, this is exactly the case that the symmetric fixed point is unstable for the S system, leading to symmetry breaking. However, this behaviour has the unfortunate consequence that the network is forced to hallucinate that the input contains some particular angle (ϕ) even when none is presented. It is this behaviour that we seek to avoid.

It can be shown that, above threshold, the gain $\partial\beta/\partial b$ of the fixed point is

$$\frac{\partial\beta}{\partial b} = \frac{1}{1 - B(2\theta_c - \frac{1}{2}\sin(4\theta_c))/2\pi}$$

which increases[†] with θ_c . For the S system, we require that the response to a flat input $b = 0$ is stably flat. Making the solution $g(x) = \alpha - 1$ stable against fluctuations of the form of $\cos(2\theta)$ requires $B < 2$. This implies that $\partial\beta/\partial b < 2$ for $\theta_c = \pi/4$. The gain $\partial\alpha/\partial a$ for the flat mode is

$$\frac{\partial\alpha}{\partial a} = \frac{1}{1 - (A - C)}$$

and so, if we impose the extra condition that $J_{ii} - W_{ii} > 0$, in accordance with neurobiological expectations that the weights have the shape of a Mexican hat with net excitation at the centre, then we require that $A - C > -B$ and so that

$$\frac{1}{1 - (A - C)} > \frac{1}{3}.$$

Hence, the amplification ratio $R^S = (\partial\beta/\partial b)/(\partial\alpha/\partial a) < 6$ for a tuning width $\theta_c = \pi/4$, or $R^S < 3.75$ for a smaller (and more biologically faithful) width $\theta_c = 30^\circ$.

The EI system will behave appropriately for large amplification ratio R if the same set of constraints as for the two-point case are satisfied. This means that in response to the untuned input, at least:

- (a) The untuned fixed point $x = \alpha$ should be unstable. The behaviour of the system about this fixed point should be oscillatory in the mode $x_i = x_j$. These conditions will be satisfied if $A^2 < 4C$ and $A > 2$.
- (b) The ring of fixed points that are not translationally invariant ($x_i = \alpha + \beta \cos(2(\theta_i - \phi))$ for arbitrary ϕ) should exist under translation-invariant input and be unstable and oscillatory.
- (c) Perturbations in the direction of $\cos(2(\theta - \phi))$ (for arbitrary ϕ) about the final limit cycle defined by $x_i = x_j$ should shrink. The conditions under which this happens are very similar to those for the two-point system, which were used to help derive figure 5.

Although we have not been able to find closed-form expressions for the satisfaction of all these conditions, we can use them to delimit sets of appropriate parameters. In particular, we may expect that, in general, large values of B should lead to large selective amplification of the tuned mode, and therefore we seek greater values of B subject to the satisfaction of the other constraints. Figure 6 shows the response of one network designed to have a very high selective amplification factor of ~ 1500 . Figure 6(A) shows noisy versions of both flat and tuned input. Figure 6(B) shows the response to tuned and flat inputs in terms of the mean over the oscillations. Figure 6(C) shows the structure of the oscillations in the thresholded activity of two units in response to a tuned input. The frequency of the oscillations is greater for the untuned than for the tuned input (not shown).

The suppression of noise for nearly flat inputs is a particular nonlinear effect in the response of the system. Figure 7(A,B) shows one measure of the response of the network as a function of the magnitude of b for different values of a . The sigmoidal shape of these curves shows the way that noise is rejected. Indeed, b has to be sufficiently large to excite the tuned mode of the network. Figure 7(B) shows the same data, but appropriately normalized, indicating that, if $r(a, b)$ is the peak response when the input is $a + b \cos(2\theta)$, then $r(a, b) \simeq ar(1, b/a)$. The scalar dependence on a was observed for the S system by Salinas and Abbott (1996). When b/a is so large that the response is away from the flat portion of the sigmoid, the response of

[†] Note that θ_c also changes with the input b , but only by a small amount when there is substantial amplification.




Figure 6. Cosine-tuned 64-unit EI system. (A) Tuned (dotted and solid lines) and untuned input (dashed line). All inputs have the same DC level $a = 10$ and the same random noise; the tuned inputs include $b = 2.5$ (dotted) and $b = 5$ (solid). (B) Mean response of the system to the inputs in (A), the solid, dotted, and dashed curves being the responses of all units to the correspondingly designated input curves in (A). The network amplifies the tuned input enormously, albeit with rather coarse tuning. Note that the response to the noisy and untuned input is almost zero at this scale (dashed curve). If a is increased to 50, then the response remains indistinguishable from the dashed line in the figure although its peak value does actually increase very slightly. (C) Temporal response of two units to the solid input from (A). The solid line shows the response of the unit tuned for 0° and the dashed line that for 36.5° . The oscillations are clear. Here $A = 6.5$, $B = 8.5$ and $C = 14.5$.

the network at the peak has the same width (θ_c) for all values of a and b , being determined just by the weights.

Although cosine tuning is convenient for analytical purposes, it has been argued that it is too broad to model cortical responsivity (see, in particular, the statistical arguments in Pouget *et al* 1998). One side effect of this is that the tuning widths in the EI system are uncomfortably large. It is not entirely clear why they should be larger than for the S system. However, in the reasonable case that the tuning of the input is also sharper than a cosine, for instance, the Gaussian, and the tuning in the weights is also Gaussian, sharper orientation tuning can be achieved. Figure 8(B,C) shows the oscillatory output of two units in the network in response to a tuned input, indicating the sharper output tuning and the oscillations. Figure 8(D) shows the activities of all the units at three particular phases of the oscillation. Figure 8(A) shows how the mean activity of the most activated unit scales with the levels of tuned and untuned




Figure 7. Mean response of the $\theta_i = 0^\circ$ unit as a function of b for three values of a . (A) The mean responses for $a = 20$ (solid), $a = 10$ (dashed) and $a = 5$ (dotted) are indicated for different values of b . Sigmoidal behavior is prominent. (B) Rescaling b and the responses by a makes the curves lie on top of each other.

input. The network amplifies the tuned inputs dramatically more—note the logarithmic scale. The S system breaks symmetry to the untuned input ($b = 0$) for these weights. If the weights are scaled uniformly by a factor of 0.22, then the S system is appropriately stable. However, the magnification ratio is 4.2 rather than something greater than 1000.

The orientation system can be understood to a largely qualitative degree by looking at its two-point cousins. Many of the essential constraints on the system are determined by the behaviour of the system when the mode with $x_i = x_j$ dominates, in which case the complex nonlinearities induced by θ_c and its equivalents are irrelevant. Let $\tilde{J}(f)$ and $\tilde{W}(f)$ for (angular) frequency f be the Fourier transforms of $J(i - j) \equiv J_{ij}$ and $W(i - j) \equiv W_{ij}$ and define

$$\lambda(f) = \text{Re} \left\{ -1 + \frac{1}{2} \tilde{J}(f) + i \sqrt{\tilde{W}(f) - \frac{1}{4} \tilde{J}^2(f)} \right\}.$$

Then, let $f^* > 0$ be the frequency such that $\lambda(f^*) \geq \lambda(f)$ for all $f > 0$. This is the non-translation-invariant mode that is most likely to cause instabilities for translation-invariant behaviour. A two-point system that closely corresponds to the full system can be found by solving the simultaneous equations

$$\begin{aligned} j_0 + j &= \tilde{J}(0) & w_0 + w &= \tilde{W}(0) \\ j_0 - j &= \tilde{J}(f^*) & w_0 - w &= \tilde{W}(f^*). \end{aligned}$$

This design equates the $x_1 = x_2$ mode in the two-point system with the $f = 0$ mode in the orientation system and the $x_1 = -x_2$ mode with the $f = f^*$ mode. For smooth $J(i - j)$ and $W(i - j)$, f^* is often the smallest or one of the smallest non-zero spatial frequencies. It is easy to see that the two systems are exactly equivalent in the translation-invariant mode $x_i = x_j$ under translation-invariant input $I_i = I_j$ in both the linear and nonlinear regimes. A coarse sweep over the parameter space of Gaussian-tuned $\mathbf{J}$ and $\mathbf{W}$ in the EI system showed that for all cases tried, the full orientation system broke symmetry if, and only if, its two-point equivalent also broke symmetry. Quantitatively, however, the amplification ratio differs between the two systems, since there is no analogue of θ_c for the two-point system.




Figure 8. The Gaussian orientation network. (A) Mean response of the $\theta_i = 0^\circ$ unit in the network as a function of a (untuned) or b (tuned) with a log scale. (B) Activity of the $\theta_i = 0^\circ$ (solid) and $\theta_i = 30^\circ$ (dashed) units in the network over the course of the positive part of an oscillation. (C) Activity of these units in (B) over all time. (D) Activity of all the units at the three times shown as (i), (ii) and (iii) in (B), where (i) (dashed) is in the rising phase of the oscillation, (ii) (solid) is at the peak, and (iii) (dotted) is during the falling phase. Here, the input is $I_i = a + be^{-(\theta_i)^2_\pi/2\sigma^2}$, with $\sigma = 13^\circ$, and the Töplitz weights are $J_{ij} = (3 + 21e^{-(\theta_i - \theta_j)^2_\pi/2\sigma'^2})/N$, with $\sigma' = 20^\circ$ and $W_{ij} = 23.5/N$, and $(\theta)_\pi = \pi/2 - |\theta| - \pi/2$.

5. The contour-region system

The final example is the application of the EI and S systems to the task described in figure 1 of contour enhancement and texture region segmentation. In this case, the neural units represent visual stimuli in the input at particular locations and orientations. Hence the unit $x_{i\theta}$ (or the pair $(x_{i\theta}, y_{i\theta})$) corresponds to a small bar or edge located at (horizontal, vertical) image location (m_i, n_i) in a discrete (for simplicity, Manhattan) grid and oriented at $\theta = k\pi/K$ for $k = 0, 1, \dots, K-1$ for a finite K [†]. The neural connections $\mathbf{J}$ and $\mathbf{W}$ link units $i\theta$ and $j\theta'$ symmetrically and locally. The desired computation is to amplify the activity of unit $i\theta$ selectively if it is part of an isolated smooth contour in the input, and suppress it selectively if it is part of a homogeneous input region.

[†] We do not consider here how orientation tuning is achieved as in the orientation system; hence the neural circuit within given a grid point i is not the same as the orientation system we studied above.




Figure 9. The four particular visual stimulus patterns A, B, C and D discussed in the text.

In particular, consider the four input patterns, A, B, C and D shown in figure 9, when all input bars have $\theta = \pi/2$, either located at every grid point i as in pattern A or at selective locations as in patterns B, C and D. Here, wrap-around boundary conditions are employed, so the top and bottom of the plots are identified, as are the right and left.

Given that all the visible bars in the four patterns have the same input strength I , the computation performed by the network should be such that the outputs for the visible bar units be weakest for pattern A (which is homogeneous), stronger for pattern B, and even stronger still for pattern C, and also such that all visible units should have the same response levels within each example. For these simple input patterns, we can ignore all other orientations for simplicity, denote each unit simply by its location i in the image, and consider the interactions J_{ij} and W_{ij} restricted to only these units. Of course, this is not true for more complex input patterns, but will suffice to derive some constraints. The connections should be translation and rotation invariant and mirror symmetric; thus J_{ij} and W_{ij} should depend only on $i - j$ and be symmetric. Intuitively, weights J_{ij} should connect units i, j when they are more or less vertically displaced from each other locally to achieve contour enhancement, and weights W_{ij} should connect those i, j that are more or less horizontally displaced locally to achieve activity suppression.

Define

$$J'_c \equiv \sum_{j:m_j=m_i+c} J_{ij} \quad j_0 \equiv J_{ii}$$

$$W'_c \equiv \sum_{j:m_j=m_i+c} W_{ij} \quad w_0 \equiv W_{ii}.$$

The input gains to patterns A, B and C at the fixed points will be roughly

$$g^A = \left[1 + \sum_c W'_c - \sum_c J'_c \right]^{-1}$$

$$g^B = [1 + w_0 - j_0]^{-1}$$

$$g^C = [1 + W'_0 - J'_0]^{-1}.$$
(13)

The relative amplification or suppression can be measured by ratios $g^C : g^B : g^A$. Let $\Delta \equiv (J'_0 - j_0) - (W'_0 - w_0)$. Then, the degree of contour enhancement, as measured by g^C/g^B , is

$$g^C/g^B = \frac{1 + w_0 - j_0}{1 + w_0 - j_0 - \Delta}.$$

To avoid symmetry breaking between the two straight lines in the input pattern D, we require, just as in the two-point system, that

$$W'_1 - J'_1 < (1 + W'_0 - J'_0) = (1 + w_0 - j_0 - \Delta). \quad (14)$$

In the simplest case, let all connections J_{ij} and W_{ij} connect elements displaced horizontally for no more than one grid distance, i.e. $|m_i - m_j| \leq 1$. Then, $J'_c = W'_c = 0$ for $|c| > 1$. For $g^A < g^B$, we require

$$\sum_c W'_c - \sum_c J'_c > w_0 - j_0$$

or, equivalently,

$$2(W'_1 - J'_1) > \Delta.$$

Combining this with equation (14), we get

$$\Delta < \frac{2}{3}(1 + w_0 - j_0) \quad \Rightarrow \quad g^C/g^B < 3.$$

In the EI system, however, $W'_1 - J'_1$ can be very large without breaking the symmetry between the two lines in pattern D. This is the same effect we investigated in the two-point system. This allows the use of

$$\Delta = (1 - \delta)(1 + w_0 - j_0)$$

with very small values of $\delta \ll 1/3$ and thus a large contour enhancement factor $g^C/g^B \sim 1/\delta \gg 3$. Other considerations do limit Δ , but to a lesser extent (Li 1998). This simplified analysis is based on a crude approximation to the full, complex, system. Nevertheless, it may explain the comparatively poor performance of many S systems designed for contour enhancement, such as the models of Grossberg and Mingolla (1985) and Zucker *et al* (1989), by contrast with the performance of a more biologically based EI system (Li 1998). Figure 10 demonstrates that to achieve reasonable contour enhancement, the reduced S system (using $\tau_y = 0$ and keeping all the other parameters the same) breaks symmetry and hallucinates stripes in response to a homogeneous input. As one can expect from our analysis, the neural responses in the EI system are oscillatory for the contour and line segments as well as all segments in the texture.

6. Conclusions


We have studied the dynamical behaviour of networks with symmetrical and asymmetrical connections and have shown that the extra degrees of dynamical freedom of the latter can be put to good computational use. Many applications of recurrent networks involve selective amplification—and the selective amplification factors for asymmetrical networks can greatly exceed those of symmetrical networks without their having undesirable hallucinations or grossly distorting the input signal. If, however, spontaneous pattern formation or hallucination by the network is computationally necessary, such that the system gives preferred output patterns even with ambiguous, unspecified or random noise inputs, the EI system, just like the S system, *can* be so designed, at least for the paradigmatic case of the two-point system. Oscillations are a key facet of our networks. Although there is substantial controversy surrounding the computational role and existence of sustained oscillations in cortex, there is ample evidence that oscillations of various sorts can certainly occur, which we take as hinting at the relevance of the computational regime that we have studied.

We have demonstrated the power of excitatory–inhibitory networks in three cases, the simplest having just two pairs of neurons, the next studying their application to the well studied case of the generation of orientation tuning, and, finally, in the full contour and region segmentation system of Li (1997, 1998) that inspired this work in the first place. For analytical convenience, all our analysed examples have translation symmetry in the neural connections and the preferred output patterns are translational transforms of each other. This translation

(A) Contour enhancement

Input image

Output image

**(B) Responses to homogeneous inputs**

From EI system

From reduced system




Figure 10. Demonstration of the performance of the contour-region system. (A) Input image $I_{i\theta}$ and the mean output response $g_x(x_{i\theta})$ from the EI system of Li (1998). $I_{i\theta}$ is the same for each visible bar segment, but $g_x(x_{i\theta})$ is stronger for the line and circle segments, shown in the plot as proportional to the bar thicknesses. The average response from the reduced system (taking $\tau_y = 0$) is qualitatively similar. (B) In response to a homogeneous texture input, the EI system responds faithfully with homogeneous output, while the reduced system hallucinates stripes.

symmetry is not an absolutely necessary condition to achieve selective amplification of some input patterns against others, as is confirmed by simulations of systems without translation symmetry.

We made various simplifications in order to get an approximate analytical understanding of the behaviour of the networks. In particular, the highly distilled two-point system provides much of the intuition for the behaviour of the more complex systems. It suggests a small set of conditions that must be satisfied to avoid spontaneous pattern formation. We also made the unreasonable assumption that the inhibitory neurons are linear rather than sharing the nonlinear activation function of the excitatory cells. In practice, this seems to make little difference in the behaviour of the network, even though the linear form $h(y) = y - T_y$ has the paradoxical property that inhibition turns into excitation when $y < T_y$. The analysis of the contour integration and texture segmentation system is particularly impoverished. Li (1997,

1998) imposed substantial extra conditions (e.g. that an input contour of a finite length should not grow because of excessive contextual excitation) and included extra nonlinear mechanisms (a form of contrast normalization), none of which we have studied.

A prime fact underlying asymmetrical networks is neuronal inhibition. Neurobiologically, inhibitory influences are, of course, substantially more complicated than we have suggested. In particular, inhibitory cells do have somewhat faster time constants than excitatory cells (though they are not zero), and are also not so subject to short-term plasticity effects such as spike rate adaptation (which we have completely ignored). Inhibitory influences also play out at a variety of different time scales by dint of different classes of receptor on the target cells. Nevertheless, there is ample neurobiological and theoretical reason to believe that inhibition has a critical role in shaping network dynamics, and we have suggested one computational role that can be subserved by this. In our selective amplifiers, the fact that inhibition comes from interneurons, and is therefore delayed, both *introduces* local instability at the fixed point and *removes* the global spontaneous, pattern-forming instability arising from the amplifying positive feedback.

It is not clear how the synaptic weights $\mathbf{J}$ and $\mathbf{W}$ in the EI system may be learnt. Most intuitions about learning in recurrent networks come from S systems, where we are aided by the availability of energy functions. Showing how learning algorithms can sculpt appropriate dynamical behaviour in EI systems is the next and significant challenge.

Acknowledgments

We are grateful to Boris Hasselblatt, Jean-Jacques Slotine and Carl van Vreeswijk for helpful discussions, and to three anonymous reviewers for comments on an earlier version. This work was funded in part by a grant from the Gatsby Charitable Foundation and by grants to PD from the NIMH (1R29MH55541-01), the NSF (IBN-9634339) and the Surdna Foundation.

References

- Ben-Yishai R, Bar-Or R L and Sompolinsky H 1995 Theory of orientation tuning in visual cortex *Proc. Natl Acad. Sci. USA* **92** 3844–8
- Braun J, Niebur E, Schuster H G and Koch C 1994 Perceptual contour completion: a model based on local, anisotropic, fast-adapting interactions between oriented filters *Soc. Neurosci. Abstr.* **20** 1665
- Carandini M and Ringach D L 1997 Predictions of a recurrent model of orientation selectivity *Vision Res.* **37** 3061–71
- Cohen M A and Grossberg S 1983 Absolute stability of global pattern formation and parallel memory storage by competitive neural networks *IEEE Trans. Systems, Man Cybern.* **13** 815–26
- Cowan J D 1982 Spontaneous symmetry breaking in large scale nervous activity *Int. J. Quantum Chem.* **22** 1059–82
- Destexhe A, McCormick D A and Sejnowski T J 1993 A model for 8–10 Hz spindling in interconnected thalamic relay and reticularis neurons *Biophys. J.* **65** 2473–7
- Eckhorn R, Bauer R, Jordan W, Brosch M, Kruse W, Munk M and Reitboeck H J 1988 Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat *Biol. Cybern.* **60** 121–30
- Ermentrout G B and Cowan J D 1979a A mathematical theory of visual hallucination patterns *Biol. Cybern.* **34** 137–50
- 1979b Temporal oscillations in neuronal nets *J. Math. Biol.* **7** 265–80
- Golomb D, Wang X J and Rinzel J 1996 Propagation of spindle waves in a thalamic slice model *J. Neurophysiol.* **75** 750–69
- Gray C M, König P, Engel A K and Singer W 1989 Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties *Nature* **338** 334–7
- Grossberg S 1988 Nonlinear neural networks: principles, mechanisms and architectures *Neural Networks* **1** 17–61
- Grossberg S and Mingolla E 1985 Neural dynamics of perceptual grouping: textures, boundaries, and emergent segmentations *Perception Psychophys.* **38** 141–71

- Hopfield J J 1982 Neural networks and systems with emergent selective computational abilities *Proc. Natl Acad. Sci. USA* **79** 2554–8
- 1984 Neurons with graded response have collective computational properties like those of two-state neurons *Proc. Natl Acad. Sci. USA* **81** 3088–92
- Kapadia M K, Ito M, Gilbert C D and Westheimer G 1995 Improvement in visual sensitivity by changes in local context: parallel studies in human observers and in V1 of alert monkeys *Neuron* **15** 843–56
- Khalil H K 1996 *Nonlinear Systems* (Englewood Cliffs, NJ: Prentice-Hall) 2nd edn
- Knierim J J and van Essen D C 1992 Neuronal responses to static texture patterns in area V1 of the alert macaque monkeys *J. Neurophysiol.* **67** 961–80
- Konig P, Janosch B and Schillen T B 1992 Stimulus-dependent assembly formation of oscillatory responses *Neural Comput.* **4** 666–81
- Konig P and Schillen T B 1991 Stimulus-dependent assembly formation of oscillatory responses: I. Synchronization *Neural Comput.* **3** 155–66
- Li Z 1995 Modeling the sensory computations of the olfactory bulb *Models of Neural Networks* ed J L van Hemmen, E Domany and K Schulten (Berlin: Springer)
- 1997 Primary cortical dynamics for visual grouping *Theoretical Aspects of Neural Computation* ed K Y M Wong, I King and D Y Yeung (Berlin: Springer)
- 1998 A neural model of contour integration in the primary visual cortex *Neural Comput.* **10** 903–40
- Li Z and Hopfield J J 1989 Modeling the olfactory bulb and its neural oscillatory processes *Biol. Cybern.* **61** 379–92
- Mundel T, Dimitrov A and Cowan J D 1997 Visual cortex circuitry and orientation tuning *Advances in Neural Information Processing Systems 9* ed M C Mozer, M I Jordan and T Petsche (Cambridge, MA: MIT Press) pp 887–93
- Murata T and Shimizu H 1993 Oscillatory binocular system and temporal segmentation of stereoscopic depth surfaces *Biol. Cybern.* **68** 381–91
- Pouget A, Zhang K C, Deneve S and Latham P E 1998 Statistically efficient estimation using population coding *Neural Comput.* **10** 373–401
- Salinas E and Abbott L F 1996 A model of multiplicative neural responses in parietal cortex *Proc. Natl Acad. Sci. USA* **93** 11956–61
- Samsonovich A and McNaughton B L 1997 Path integration and cognitive mapping in a continuous attractor neural network model *J. Neurosci.* **17** 5900–20
- Schillen T B and Konig P 1991 Stimulus-dependent assembly formation of oscillatory response: II. Desynchronization *Neural Comput.* **3** 167–78
- Seung H S 1996 How the brain keeps the eyes still *Proc. Natl Acad. Sci. USA* **93** 13339–44
- Seung H S, Richardson T J, Lagarias J C and Hopfield J J 1998 Minimax and Hamiltonian dynamics of excitatory–inhibitory networks *Advances in Neural Information Processing Systems 10* ed M I Jordan, M Kearns and S Solla (Cambridge, MA: MIT Press)
- Sompolinsky H, Golomb D and Kleinfeld D 1990 Global processing of visual stimuli in a neural network of coupled oscillators *Proc. Natl Acad. Sci. USA* **87** 7200–4
- Sporns O, Tononi G and Edelman G M 1991 Modeling perceptual grouping and figure–ground segregation by means of reentrant connections *Proc. Natl Acad. Sci. USA* **88** 129–33
- Stein P S G, Grillner S, Selverston A I and Stuart D G 1997 *Neurons, Networks, and Motor Behavior* (Cambridge, MA: MIT Press)
- Steriade M, McCormick D A and Sejnowski T J 1993 Thalamocortical oscillations in the sleeping and aroused brain *Science* **262** 679–85
- Suarez H, Koch C and Douglas R 1995 Modeling direction selectivity of simple cells in striate visual cortex within the framework of the canonical microcircuit *J. Neurosci.* **15** 6700–19
- von der Malsburg C 1981 *The Correlation Theory of Brain Function* Report 81-2, MPI Biophysical Chemistry
- 1988 Pattern recognition by labeled graph matching *Neural Networks* **1** 141–8
- Wilson H R and Cowan J D 1972 Excitatory and inhibitory interactions in localized populations of model neurons *Biophys. J.* **12** 1–24
- Wilson H R and Cowan J D 1973 A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue *Kybernetik* **13** 55–80
- Zhang K 1996 Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory *J. Neurosci.* **16** 2112–26
- Zucker S W, Dobbins A and Iverson L 1989 Two stages of curve detection suggest two styles of visual computation *Neural Comput.* **1** 68–81