



massachusetts institute of technology — artificial intelligence laboratory

Biologically Plausible Neural Circuits for Realization of Maximum Operations

Angela J. Yu, Martin A. Giese and
Tomaso A. Poggio

AI Memo 2001-022
CBCL Memo 207

September 2001

Abstract

Object recognition in the visual cortex is based on a hierarchical architecture, in which specialized brain regions along the ventral pathway extract object features of increasing levels of complexity, accompanied by greater invariance in stimulus size, position, and orientation. Recent theoretical studies postulate a non-linear pooling function such as the maximum (MAX) operation could be fundamental in achieving such invariance. In this paper, we are concerned with neurally plausible mechanisms that may be involved in realizing the MAX operation. Four canonical circuits are proposed, each based on neural mechanisms that have been previously discussed in the context of cortical processing. Through simulations and mathematical analysis, we examine the relative performance and robustness of these mechanisms. We derive experimentally verifiable predictions for each circuit and discuss their respective physiological considerations.

This report describes research done within the Center for Biological and Computational Learning in the Department of Brain and Cognitive Sciences and in the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology.

This research is sponsored by a grant from the Office of Naval Research under contract No. N00014-93-1-3085, Office of Naval Research under contract No. N00014-95-1-0600, National Science Foundation under contract No. IIS-9800032, National Science Foundation under contract No. DMS-9872936, National Science Foundation under Graduate Research Fellowship Program, Massachusetts Institute of Technology under the UROP program, and a grant to the Gatsby Computational Neuroscience Unit from the Gatsby Charitable Foundation.

1 Introduction

Neurophysiological experiments have provided evidence that object recognition in the visual cortex is accomplished via a predominantly hierarchical system, in which areas farther along the ventral pathway have larger receptive fields and are selective for increasingly complex stimulus features with increasing invariance with respect to stimulus size and position (Hubel & Wiesel, 1962; Perrett et al., 1991; Logothetis, Pauls, & Poggio, 1995; Tanaka, 1996; Pasupathy & Connor, 1999).

Many theories based on neurobiologically plausible mechanisms have been proposed to account for such invariance properties. One body of theoretical work involves flexible central mechanisms that dynamically adjust stimulus scale and position selectivity according to the input: an example is the “shifter circuit” (Anderson & Essen, 1987), which involves saliency-map driven control neurons that mediate interactions between top-down (memory) and bottom-up (retinal input) sources in order to dynamically alter synaptic strengths so that only a relevant subset of sensory information is routed through the cortex. However, although there is evidence for such dynamic modulation in a variety of visual areas (Motter, 1994; Connor, Preddie, Gallant, & Van Essen, 1997; Treue & Maunsell, 1996), it functions at a time scale too slow to account for the short latencies that are found for object recognition tasks with short stimulus presentation (Thorpe, Fize, & Marlot, 1996).

Another class of models circumvents this problem by pooling non-invariant feature detectors. The underlying idea was first postulated by Hubel and Wiesel, which accounts for complex cell invariance with respect to spatial phase shifts via linear summation of responses of rectifying, phase-sensitive neurons (Hubel & Wiesel, 1962). Fukushima’s more general, hierarchical Neocognitron network achieves invariant responses via a feedforward, hierarchical network of alternating layers of adaptive feature detectors and pooling neurons with nonlinear characteristics (Fukushima, 1980). More recently, Riesenhuber and Poggio have reproduced neurophysiological data from area IT using a hierarchical network model that combines pooling by linear summation and by MAX (or soft-max) operations (Riesenhuber & Poggio, 1999b, 1999a).

It is interesting to note that pooling by the MAX operation is computationally equivalent to “scanning” an image with a template, which is the basis for many recognition algorithms in computer vision (Riesenhuber & Poggio, 1999b, 2000). Pooling by a MAX operation as opposed to linear summation achieves high feature specificity *and* invariance simultaneously (Riesenhuber & Poggio, 1999b). For example, suppose the inputs to the system are activity levels of a population of simple cell bar-detectors that prefer the same orientation but have receptive fields in different locations. Then summing from these detectors gives the same response if the input is an oriented bar contained in any one of the receptive fields, achieving position invariance. However, the response is even stronger if multiple bars (e.g. a grating), one long bar, or background clutter are present. Taking the maximum of the responses of these feature detectors, however, solves this problem, because then the system output is solely determined by the response of its most active afferent. Thus, the MAX operation both preserves feature specificity and achieves invariance in a more robust fashion than linear summation.

In this paper, we are concerned with how the MAX operation can be implemented neurophysiologically. In addition to its potential involvement in a variety of cortical processes such as object recognition (Riesenhuber & Poggio, 1999b), motion recognition (Giese, 2000), and visual velocity estimate (Grzywacz & Yuille, 1990), the MAX operation is interesting because it is a basic nonlinear operator that can be implemented by simple, neurophysiologically plausible circuits. In section 2, we give the mathematical definition of the ideal MAX operation, and discuss the theoretical and experimental context of our work. In section 3, we isolate a small number of neural mechanisms that may be involved in realizing the MAX operation, and present four highly simplified “canonical” circuits that implement these mechanisms. In section 4, we present our main simulation results; the relevant mathematical analysis can be found in the Appendix. In section 5, we propose various neuronal implementations of the computational operations involved in each of the networks and compare the relative plausibility of these implementations. The lists of plausible mechanisms are by no means exhaustive, and their verification depends on future data from more detailed experiments, which we hope this work will inspire. A number of general and model-specific predictions are presented, along with a discussion on potential experimental frameworks that can be used to verify these predictions. Finally, section 6 gives a summary of our work, relates our efforts to previous work in neural modelling, and makes suggestions for future directions of research.

2 Definitions and Background

The ideal MAX operation is defined as a mapping from an input vector $\mathbf{x} = [x_1, x_2, \dots, x_n]$ to an output signal z , where

$$z = \max_{1 \leq i \leq n} x_i = x_{\max} \quad (1)$$

In particular, the operation should achieve the following properties:

1. **Selectivity:** The output signal z depends only on the maximum of all the input signals¹, x_{\max} , and not on the other values.
2. **Linearity:** The output signal z depends linearly on x_{\max} with a constant gain factor g , i.e. $z = gx_{\max}$.

The first property is relevant to achieving feature specificity. The second property is important for optimally recovering information about the strength of the maximal input.

Biological systems can only be expected to implement an approximation to the ideal MAX operation. Likewise, in most computation models employing the MAX operation, ideal MAX operation behavior is also not strictly necessary. For example, it has been demonstrated by Riesenhuber and Poggio that the “soft-max” operation is sufficient for obtaining good simulation results in their model of object recognition (Riesenhuber & Poggio, 1999b). In section 4, we discuss the conditions under which our models achieve good approximations to the ideal MAX operation.

Two areas of earlier work on neural modelling are closely related to the MAX operation: *winner-takes-all* (WTA) and *contrast gain control* networks.

WTA has been widely studied in the neural networks and VLSI literature (Kohonen, 1995; Lazzaro, Ryckenbusch, Mahowald, & Mead, 1989; Starzyk & Fang, 1993; Hahnloser, Douglas, Mahowald, & Hepp, 2000), and has been used to model cortical functions such as the integration of component motions (Nowlan & Sejnowski, 1995) and attentional selection (Koch & Ullman, 1985; Lee, Itti, Koch, & Braun, 1999). WTA networks select the afferent input with the largest amplitude, but its output is not required to reflect this amplitude and is often nonlinear or even binary. In general, WTA networks convey the identity of the “winner neuron” but not its precise amplitude to the downstream cortical processing areas, whereas the MAX operation communicates the amplitude and not the identity. There are exceptions among the WTA networks, however, whose output does communicate some information about the amplitude of the maximal input: Yuille and Grzywacz’s divisive feedback network (Yuille & Grzywacz, 1989) and Hahnloser’s linear threshold VLSI network (Hahnloser, 1998) both output the maximal input under certain conditions, and Lazzaro et al.’s VLSI network outputs the logarithm of the maximal input (Lazzaro et al., 1989). The main difference between these previous studies and ours is that both the mechanisms and the functional goals of these previous models are less biologically motivated. However, the underlying neural substrates of WTA and MAX operations may be similar. Indeed, some of our implementations are inspired by previously proposed models in the context of WTA.

A second class of neural circuits that is closely related to the MAX operation is that of contrast gain control (Reichardt, Poggio, & Hausen, 1983; Carandini & Heeger, 1994; Wilson & Humanski, 1993; Simoncelli & Heeger, 1998). Such circuits make the response of neural detectors independent of the stimulus energy or contrast, thus exhibiting some invariance in stimulus size or intensity. However, typical gain control circuits fail the linearity requirement because they renormalize all input channels regardless of input amplitude.

3 Canonical Circuits

We present four simple, deterministic “canonical” neural models that implement the MAX operation with well-studied neural principles and which allow some degree of mathematical analysis. All these circuits can be described as three layer neural networks with an input layer representing input signals x_n , a symmetrically-connected hidden layer that transforms the input signals into output signals y_n in a nonlinear fashion, and an output unit that simply sums the hidden layer activities: $z = \sum_i y_i$. In biophysiological terms, the inputs correspond to output signals from earlier stages of sensory information processing, and if these earlier feature detectors have similar response curves in all feature dimensions but for a small subset, then the output signal would reflect invariance in this subset of feature dimensions.

The main neural principles explored in our models are:

¹Sometimes also referred to as *input amplitude* in this paper.

- feedforward vs. recurrent processing
- synaptic (divisive) vs. neuronal (linear threshold) inhibition
- cellular vs. network level of interactions

To help visualize the networks, we provide schematic diagrams of a generic feedforward system and a feedback system (Figure 1). Readers should note that the circles in these diagrams represent computational units and not necessarily biological neurons, and the lines denote computational interactions rather than dendritic or axonic processes. Instances of apparent violations of Dale’s Law can be resolved via inhibitory interneurons (Li, 2000), as is apparently done in the real cortex (White, 1989; Gilbert, 1992; Rockland & Lund, 1983).

3.1 Divisive Feed-Forward (FFN) Circuit

The first circuit is a feedforward network that normalizes each individual input by a sum of signals that are derived from the inputs by transformation through a strongly nonlinearly increasing function. Similar architectures have been used to model WTA behavior (Grossberg, 1973; Koch & Ullman, 1985; Fukai & Tanaka, 1997), gain control in the fly visual system (Reichardt et al., 1983), and the modulatory effects of attention on orientation filters in human vision (Lee et al., 1999). The normalization can be modelled by divisive shunting inhibition (Naka & Rushton, 1966), which had been previously proposed as the neural basis of multiplicative “gating” and logical “AND” operations (Torre & Poggio, 1978; Koch & Poggio, 1989).

The feedforward network “dynamics”² modelled is given by:

$$y_n = \frac{x_n f(x_n)}{c + \sum_m f(x_m)} \quad (2)$$

$$z = \sum_n y_n \quad (3)$$

where c is a small positive constant that ensures the normalization factor remains bounded. $f(x)$ is a positive and strongly nonlinearly increasing function, whose precise form is non-crucial for its functionality³. If $f(x)$ is sufficiently convex and therefore sufficiently exaggerates the difference between the maximal input x_M and the other inputs, then $f(x_M)$ dominates the sum, and $z \approx y_M \approx x_M$. Note that this circuit can be thought of as gating the individual input signals with the output of a binary winner-takes-all network.

In terms of neurophysiological plausibility of this circuit, two factors stand out: nonlinearity of $f(x)$, and the matching of $f(x)$ in the numerator and denominator of equation 2. $f(x)$ can be provided by the initial segment of the current-and-firing-rate relationship of neurons (Koch & Poggio, 1989; Carandini & Ferster, 2000; Ferster & Miller, 2000) or the relationship between current and voltage at electrical synapses (Furshpan & Potter, 1959; Koch & Poggio, 1989). Approximate matching of $f(x)$ is necessary for fulfilling the linearity property of the MAX operation, and can be achieved either by providing a reproducible relationship between two different neural nonlinearities ($f(x)$ and $xf(x)$) or by using one nonlinearity to generate both quantities at the expense of an extra multiplication operation. Multiplication can be realized either by interneurons with shunting inhibition⁴, or by more complex subnetworks of groups of linear threshold neurons (Salinas & Abbott, 1996). Cascaded multiplicative operations can also be realized by nonlinear interactions between synapses on a single synaptic tree (Koch, Poggio, & Torre, 1983; Mel & Ruderman, 1998).

3.2 Divisive Feedback (DFB) Circuit

The second circuit is based on recurrent divisive normalization signals, which are derived from the hidden layer activities rather than from the input signals. Circuits of this type have been used to model gain control in the fly

²As with all of the other models we present, this feedforward network does not model synaptic delay, and so the dynamics is trivial in the sense that the output immediately and fully reflects the system’s processing of the input.

³A polynomial form of $f(x) = x^q$ was used in our simulations, where $q > 1$.

⁴The physiological basis and functional relevance of shunting inhibition have been debated for a long time and are still topics of current research (Ferster & Jagadeesh, 1992; Borg-Graham, Monier, & Fregnac, 1998; Anderson, Carandini, & D, 2000; Koch & Poggio, 1989; Holt & Koch, 1997).

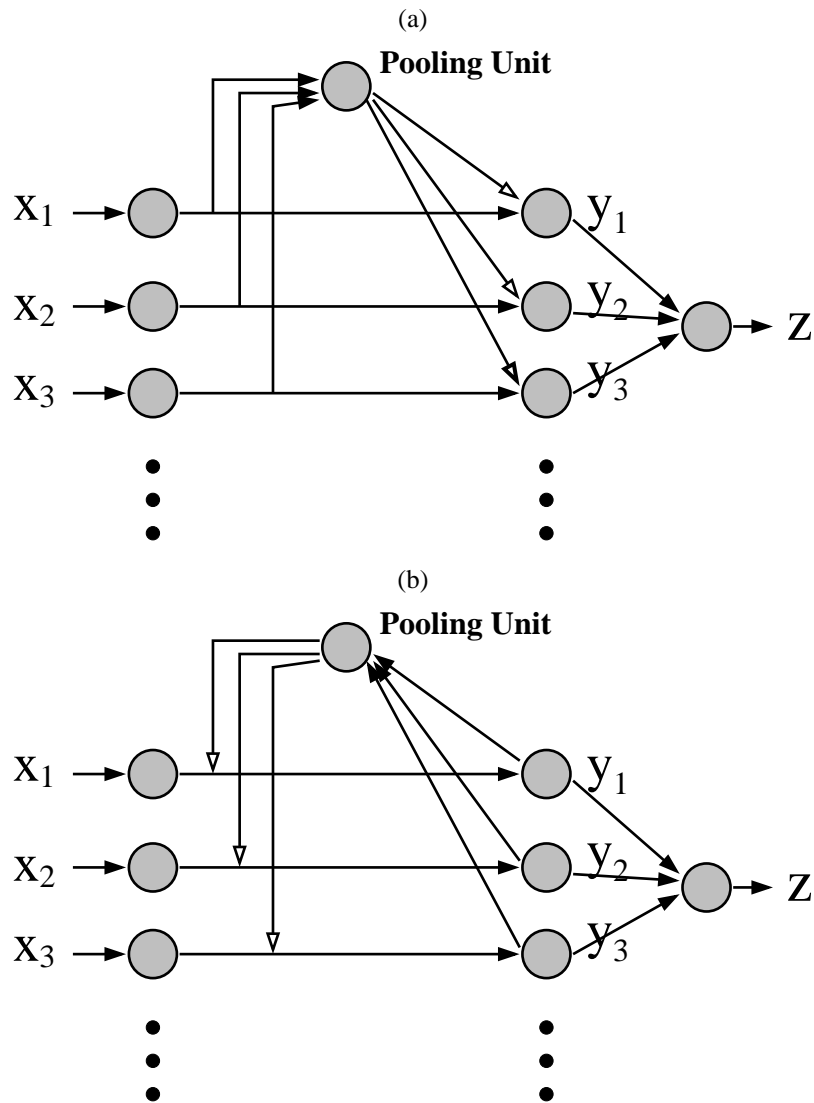


Figure 1: **FFN and FBN Schematics:** (a) Schematic diagram of the feedforward network (FFN). (b) Schematic diagram of the feedback networks (FBN). The solid arrows denote excitatory inputs, while open arrows denote inhibitory inputs. Circles and lines represent computational units and interactions rather than explicit delineation of neurons and their processes. The excitatory and inhibitory operations may be either additive or multiplicative.

(Reichardt et al., 1983) and human (Wilson & Humanski, 1993; Carandini & Heeger, 1994; Simoncelli & Heeger, 1998) visual systems, and in WTA networks (Yuille & Grzywacz, 1989).

The recurrent network dynamics is given by the following equations:

$$\tau \dot{y}_n(t) = -y_n(t) + \frac{x_n f(y_n(t))}{c + \sum_m f(y_m(t))} \quad (4)$$

$$z(t) = \sum_n y_n(t) \quad (5)$$

$$(6)$$

where $f(x)$ and c are as for the feedforward circuit, and τ is the time constant of the dynamical system⁵.

The biophysiological considerations for this model are similar to those for the feedforward version, except the multiplication of $\frac{x_n}{c + \sum_m f(y_m(t))}$ by $f(y_n(t))$ is more difficult to implement via synaptic mechanisms within a single neuron. Thus, network mechanisms such as those described by Salinas and Abbott (1996) provide more plausible explanations for this model. Another difference between the feedback and feedforward normalization model is that the recurrent dynamics requires some time to reach equilibrium, resulting in comparatively longer input-output latencies in the feedback model.

3.3 Linear Threshold (LIN) Circuit

This recurrent network uses additive feedback through a half-wave rectifying nonlinearity⁶, instead of divisive feedback. Intracellular recordings show that the half-wave rectification model provides a good fit for intracellular recording data (Carandini & Ferster, 2000). Below are the network dynamics equations:

$$\tau \dot{y}_n(t) = -y_n(t) - w \sum_m [y_m(t)]_+ + x_n \quad (7)$$

$$z(t) = w \sum_n [y_n(t)]_+$$

where $w > 0$ represents the inhibitory strength. Linear threshold networks have a *homogeneity property* (e.g. Hahnloser et al., 2000): rescaling of the input by a factor leads to the rescaling of the responses of the active units in the hidden layer and the output by the same factor.

The architecture of this network is similar to that of the divisive feedback circuit, where the pooling action could be mediated by either a dedicated pooling neuron or dendritic trees. However, since the inhibitory signal matches the network output, this feedback can be provided by the system output rather than an extra unit processing the hidden-layer activities. Networks similar to this model have been proposed to simulate the inhibitory interactions in the limulus retina (Hartline & Ratliff, 1957), orientation tuning in the visual cortex (Ben-Yishai, Lev Bar-Or, & Sompolinsky, 1995), and gain fields in the parietal cortex (Salinas & Abbott, 1996). Linear threshold models of this type also play an important role in analog VLSI circuits for the realization of winner-takes-all (WTA) behavior (Lazzaro et al., 1989; Hahnloser et al., 2000). In particular, Hahnloser et al.'s linear threshold model is very similar to our model, except for the absence of a final summing unit and the exact cancellation of self-inhibition by self-excitation. We will see in the simulation results section, however, the summing unit significantly improves the networks' performance. The issue of the relative strengths of self-inhibition and self-excitation is addressed in our work elsewhere (Giese, Yu, & Poggio, 2001).

⁵For the divisive feedback circuit, we used $f(x) = e^{qx}$.

⁶For simplicity, we assumed the rectification threshold to be zero.

3.4 Spiking Feedback (SPK) Circuit

This last circuit uses deterministic, spiking neurons to realize the nonlinearity in the network. A simple, leaky integrate-and-fire model is used for each hidden layer unit, whose output inhibits all the other units symmetrically:

$$\tau \dot{m}_n(t) = -m_n(t) - w \sum_{m \neq n} y_m(t) + x_n \quad (8)$$

$$y_n(t) = \sum_k \delta(t - t_k) \quad (9)$$

$$z(t) = \sum_n y_n(t) \quad (10)$$

The membrane potential is reset to zero after a spike at time t_k :

$$m_n(t_k^-) = \theta \quad (11)$$

$$m_n(t_k^+) = 0 \quad (12)$$

where $m_n(t)$ is the membrane potential of the n^{th} hidden neuron at time t , $y_n(t)$ is the corresponding spike output signal, τ is the time constant of the membrane, and w is the inhibitory strength.

In the limit of high firing rates, the dynamics of this model is similar to that of the linear threshold model, except there is no self-inhibition, and the inhibitory nonlinearity is explicitly modeled by the nonlinear membrane potential-firing rate relationship. In the regime of high firing rate, this model can also be thought of as a spiking implementation of the Hahnloser (1998) linear threshold model with fixed time constant and instantaneous inhibition. Typically, the membrane potentials of the most strongly stimulated neurons reach the spiking threshold earlier than the others and inhibit them from ever reaching the threshold θ .

4 Results

In the following we present a number of results from our simulation studies. The relevant mathematical analysis are to be found in the Appendix. We are interested in how well the networks achieve the linearity and selectivity properties of the MAX operation. We also analyze the responses of hidden layer and output layer units in the face of noise and their dependence on network parameters. In all of our simulations, unless otherwise noted, we used 81 input units and 81 corresponding hidden units, plus one output unit. Similarly, assume the input amplitude to be 1 unless otherwise noted.

4.1 Linearity and Selectivity

One crucial feature of selectivity is that the output of the system, z , should be minimally influenced by the non-maximal inputs. Regardless of the relative strengths of the input signals, the output should always only depend on the maximum of the inputs. As the networks were all symmetrically connected, the ordering of the inputs makes no difference, so we simulated the networks for different classes of inputs with identical amplitude α . The input units were indexed $n = -40, \dots, 40$

- Gaussian: variance $\sigma = 10$, mean $\mu = 0$
- Ramp: $x_n = \alpha * (n/80 + 1/2)$
- Uniform inputs but for one winner: $x_0 = \alpha, x_n = 0.90 * \alpha, n \neq 0$
- Uniform inputs but for two identical “winners”: $x(-40) = x(40) = \alpha, x(n) = 0.90 * \alpha, n \neq -40, 40$
- Random, iid-distributed inputs drawn from uniform distribution: inputs normalized so that the maximum is α .

Figure 2(a) is a graphical illustration of what the inputs look like. The random inputs are not depicted to avoid clutter. Figure 2(b) shows the hidden-layer activities of the divisive feedforward network to these same inputs. The network output, z , as shown in the legend, is very similar for the different types of inputs, indicating the *selectivity* property

of the ideal MAX operation is well-approximated (this was also true for all the other network models under a large regime of parameter space). Moreover, although the outputs do not always achieve unit-gain, they still appear to satisfy the *linearity* property well, as can be seen from the approximately linear response to systematic variation of input amplitude from 0 to 4.5. The network response to Gaussian and ramp inputs is an interesting illustration of how the hidden-layer interactions attenuate the input signals nonlinearly: the Gaussian peak is narrower and the ramp has been transformed into an exponential. The “uniform” inputs represent a worst-case scenario (as we shall see in the Appendix), while the “2 winner” scenario examines whether any model breaks down when multiple, identical maximal inputs are present. It is reassuring that the models behave reasonably under both conditions. We also used “random” inputs to ensure that the networks do not behave especially nicely for smoothly varying or regular inputs. The “random” inputs were independently generated for each network for each input amplitude. The results are shown in Figure 2(c)-(f). Note that the scale of the firing rate in the last panel is necessarily arbitrary, but it is not an issue as we are in any case only concerned with *linearity*. In this case, we simply summed up the number of spikes for each unit over the duration of the simulation.

The fact that the networks responded similarly to the different types of inputs is encouraging support for fulfillment of the *selectivity* property. However, to test *selectivity* more systematically, we used uniform inputs and varied the relative relationship between the magnitude of the non-maximal inputs and the maximal one. Uniform inputs were chosen because they are easy to vary systematically, and they provide a worst-case scenario for some of the network models (see Appendix).

The results are shown in Figure 3. For the FFN model, y_{\max} decreases sigmoidally with the magnitude of the non-maximal inputs, while z approximates x_{\max} well both when the non-maximal inputs are small and large, but dips at an intermediate value. This dipper is larger and affects a larger range of magnitudes of non-maximal inputs when q is smaller. A similar effect is seen in the DFB model, except for sufficiently small q , the network breaks down completely and z varies linearly with the magnitude of the non-maximal inputs. The LIN model performs extremely well except when all the inputs are approximately equal, in which case all the units are active and the activity pattern oscillates rather than settling into a stable equilibrium. The SPK model behaves very similarly to the LIN model, breaking down when all the inputs are approximately equal.

4.2 Strength of Lateral Inhibition

In each of the models, there is a parameter that controls the strength of lateral inhibition: q for the divisive models and w for the linear threshold and spiking models. In one series of simulations, we varied the strength of inhibition to examine its effects on *linearity* (figures not provided, although the values of q and w in figure 2 were chosen to demonstrate some of the phenomena described below). In general, stronger inhibition lead to better realization of perfect linearity, while weaker inhibition resulted in poorer performance. In the case of the feedforward network, the approximation to perfect linearity is poorer for smaller amplitudes than for larger amplitudes. A more rigorous mathematical treatment of this relationship is given in the Appendix. For the divisive feedback network (DFB), the linear relationship breaks down for smaller input amplitudes. The point at which this breakdown occurs decreases with larger inhibitory strength q , for reasons similar to those for the gradual deterioration of performance in the FFN model. For the linear network, different inputs still lead to linear output, but at slightly different slopes; this effect is can also be explained by a mathematical analysis presented elsewhere (Giese et al., 2001). The input-output relationship of the spiking network is also roughly linear. The non-zero x-intercept results from the neurons firing only when the membrane potential rises above the threshold.

In addition, we examine how each circuit’s output varies explicitly as a function of q or w . As shown in Figure 4, the input amplitude is recovered almost perfectly by each circuit for a large range of values for q and w . It is especially interesting that, in general, even though the hidden layer activity depends strongly on the strength of lateral inhibition, the summed output signal is less affected. For instance, the hidden-layer activity levels differ greatly as a function of q for the divisive feedforward network, and yet the sum of these hidden units output still recovers the input amplitude extremely well (Figure 4). Note also that the non-monotonic relationship between activity level and w in the case of the spiking network correlated closely with the number of active units (units that emitted spikes from time to time rather than those who were completely suppressed at all times). It will be demonstrated in the mathematical Appendix why robustness is helped by the final summing operation and also how it depends on q or w .

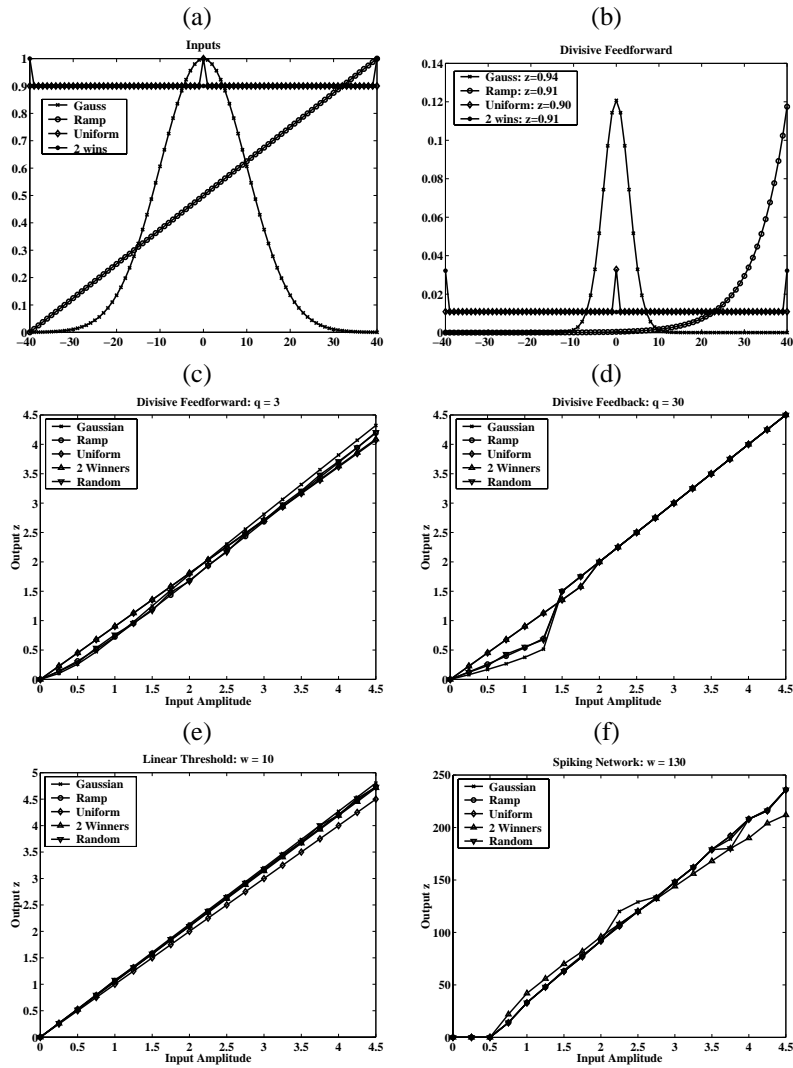


Figure 2: **Different input types:** Panel (a) shows the different input types that were tested, where the maximal input was 1 in each case. (b) shows the corresponding hidden-layer activities in the FFN (divisive feedforward network) network. $q = 10$. The remaining panels show the relationship between z and x_{\max} . (c) FFN: $q = 3$, (d) DFB (divisive feedback network): $q = 30$, (e) LIN (linear threshold network): $w = 10$, (f) SPK (spiking network): $w = 130$, output refers to firing rate.

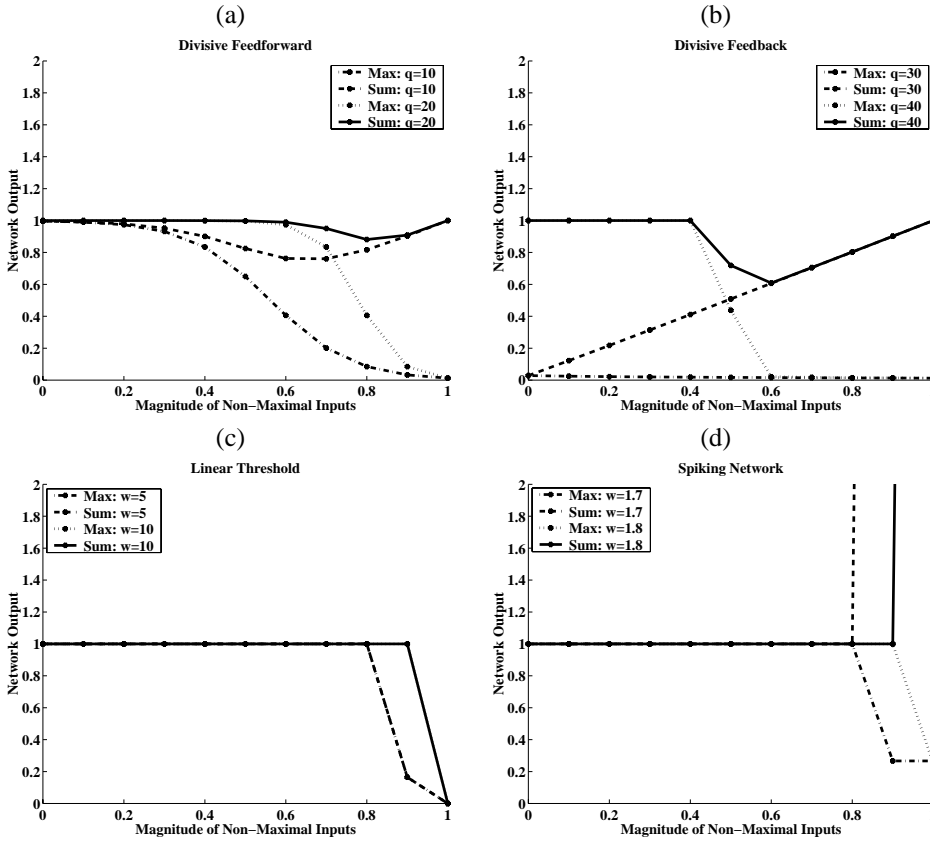


Figure 3: **Dependence on non-maximal input:** For each network, the ratio of non-maximal inputs to maximal input was varied, while the input amplitude was kept at 1 in all cases. “Sum” refers to the output signal z , while “Max” refers to the amplitude of the most strongly activated hidden layer unit. (a) FFN, (b) DFB, (c) LIN, (d) SPK. The firing rate of the spiking network has been normalized to 1 in order to be comparable to the other networks.




Figure 4: **Dependence on inhibitory strength:** Relationship between output and input amplitude as a function of inhibitory strength, represented by the parameter q or w . For each value of q or w , each network was simulated to convergence. The dashed line refers to y_{\max} and the solid line corresponds to z .

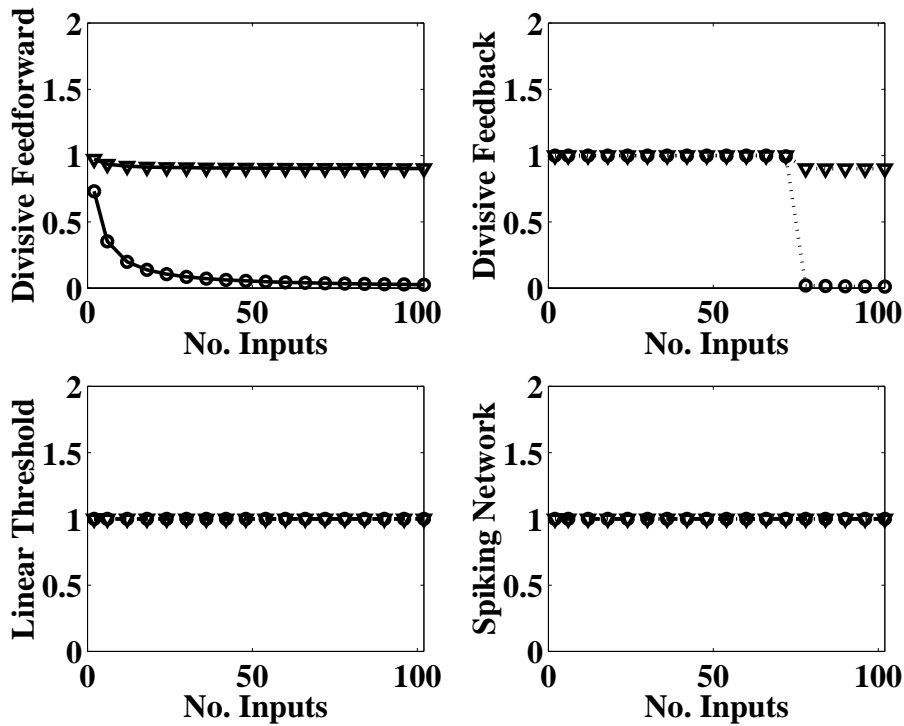


Figure 5: **Number of inputs:** The dependence of network performance on the number of inputs, ranging from 2 to 100. The inputs were uniform (magnitude = 0.9) but for one winner (magnitude = 1). FFN: $q = 10$, DFB: $q = 55$, LIN: $w = 10$, SPK: $w = 10$. The circles refer to y_{\max} and the triangles refer to z . Note the firing rate of the spiking network has been normalized to 1 to be comparable to the other models.

4.3 Number of Inputs

In some of the previous models proposed for the realization of the MAX operation or WTA, the performance of the models were dependent upon the number of inputs. As the number of inputs increased, the system tended to become increasingly inaccurate in reproducing the maximal input (e.g. Yuille & Grzywacz, 1989). To investigate this issue, we simulated our networks for different number of input units (2 to 100) and obtained the results shown in Figure 5. Uniform inputs (but for one winner) were used. For the FFN model the amplitude of the hidden-layer behavior quickly deteriorated in a graded fashion as the number of inputs increased, but z decreased only slightly initially and then stabilized. For the DFB model, up to a certain number of inputs, there was only one active unit in the hidden layer, and its activity level (as well as z) reflected x_{\max} . For larger number of inputs, however, all the hidden-units became active, albeit at a lower level, and z dropped slightly but consistently, indicating the system switched to a different attractor state (Giese et al., 2001). With a smaller q , this change of state occurs with a smaller number of inputs. In the LIN and SPK models, the number of inputs did not affect the network behaviour, neither in the hidden-layer nor in the output unit: the active set of the hidden-layer consisted of a single unit, whose activity reflected the input amplitude exactly. This last result is consistent with the mathematical result that the linear threshold model and the spiking model (in the limit of large firing rate) are monostable (see Section 4.5 and Giese et al., 2001). Suppose the system has N inputs, all but one of them is smaller than the middle one, and the system has been simulated to convergence such that only the middle one is active. Then adding another unit of sub-maximal magnitude would be the same thing as having started out with $N + 1$ units to begin with, if the system were monostable. In this case, adding another sub-maximal unit should not change the identity or amplitude of the winning hidden unit, since the winning unit would suppress the new unit so that its activity is below the threshold.

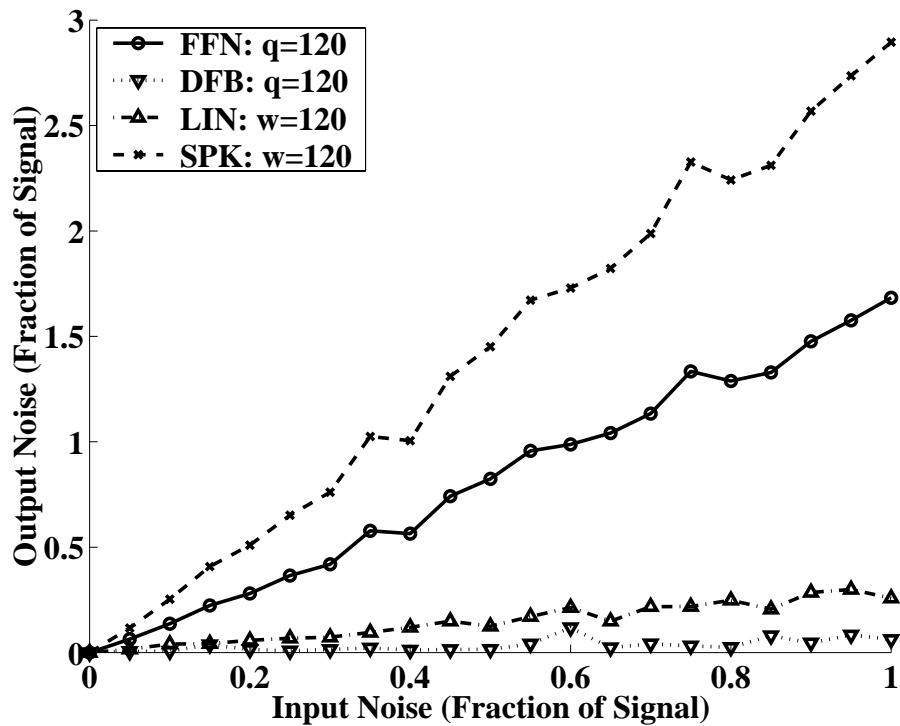


Figure 6: **Noise sensitivity:** Standard deviation of output for each network over 50 trials, normalized by the *expected* or *noiseless* output, as a function of the amount of uncorrelated noise (measured in terms of standard deviation from the *noiseless* input signal) added independently to each unit in each iteration. Variance of Gaussian input was 10 in all cases, and $q = w = 120$.

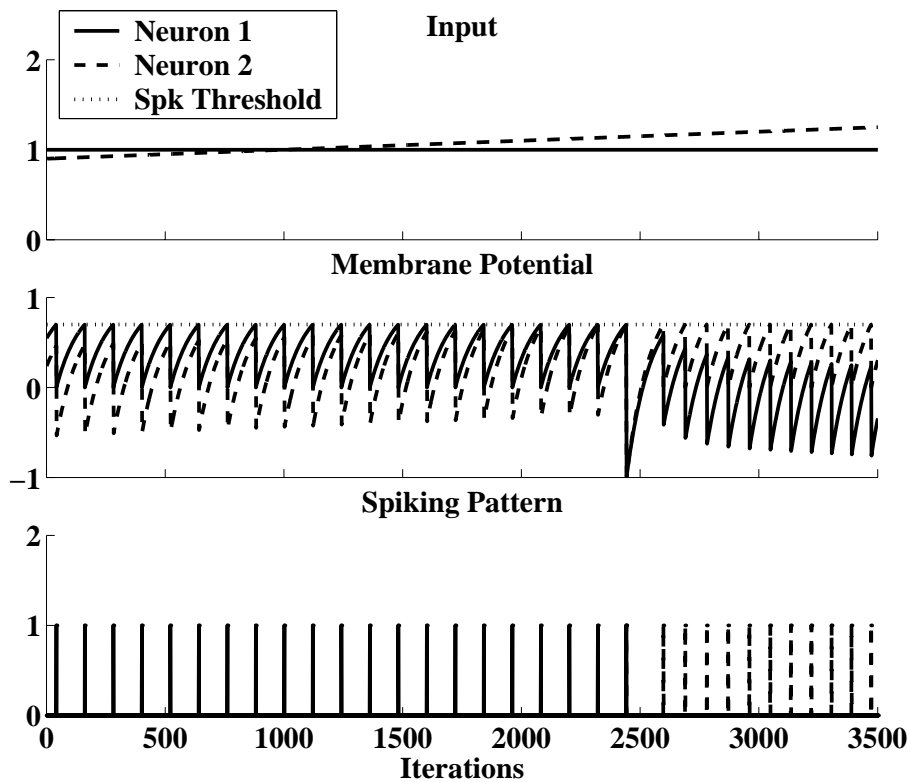


Figure 7: **Hysteresis:** Hysteresis in the spiking model. Initially, higher input to neuron 2 induces it to fire regularly and suppresses the membrane potential of neuron 1 and prevents neuron 1 from reaching the spiking threshold (the initial segment of the simulation has been cut off in order to show the latter segment more clearly). As the input to neuron 1 slowly but steadily increased, while the input to neuron 2 was kept constant, the excitatory input to neuron 1 eventually overcomes the suppression from neuron 2 and neuron 1 starts firing instead. However, this switch does not take place until long after the relative strength of inputs have swapped, indicating hysteresis.

4.4 Noise Sensitivity

In order to test the robustness of the networks, normally distributed, random noise was added independently to the input at each iteration. Figure 4.4 shows the result: Noise from 0% to 100% of the input signal was added to the inputs. In accordance with previous work on recurrent networks, the feedback action performs noise-reduction (Yuille & Grzywacz, 1989) on uncorrelated noise in the input: the recurrent mean-firing-rate models are more resistant to noise than the feedforward model. In particular, suppose the acceptable level of output noise is no more than 10% (for the purpose of downstream decoding), then the feedforward network can tolerate no more than 10% noise in the input, while the linear feedback network can tolerate up to 40%, and the divisive feedback network can tolerate up to as much as 60%. Note that the noise behaviour of the spiking model is not directly comparable to the other models, as the input noise of the spiking model was measured in terms of membrane potential, while the output noise was measured in terms of firing rate; for all the other models, both the input and output noise were measured in terms of mean firing rate. Overall, it is comforting that despite the presence of high-gain amplifiers in the systems, output noise increases linearly as a function of input noise with acceptably small slope.

Prediction / Model	FFN	DFB	LIN	SPK
Shunting/gating	cellular implementation: yes network implementation: no	cellular implementation: yes network implementation: no	no	no
Homogeneity property	no	no	yes	yes
Sparse hidden-layer activation	no	yes	no	yes
Hysteresis	no	yes	yes	yes

Table 1: **Predictions from the different neural mechanisms:** **FFN:** divisive feed-forward, **DFB:** divisive feedback, **LIN:** linear threshold, **SPK:** spiking mechanism

4.5 Stability and Hysteresis

It can be shown that a Lyapunov function exists for the recurrent models, implying the networks converge to stable fixed points⁷. However, the Lyapunov functions have multiple minima, raising the possibility of hysteresis. To investigate this issue, simple two-hidden-unit versions of the different networks were used: while the input signal of one hidden unit was kept constant, the input signal of the other neuron was gradually increased from being smaller to larger than the first one. Hysteresis was found for both models in the appropriate parameter regimes. The feedforward network is monostable and exhibits no hysteresis. The SPK model is shown empirically to exhibit hysteresis (see Figure 7(b)).

5 Predictions

As we are interested in how the MAX operation is implemented in the nervous system, we derive from our simulation and theoretical results a number of predictions which may help guide future experimental research.

The first task is to determine where the MAX operation is actually performed in the brain. It is postulated to be a fundamental operation involved at many stages of the visual processing system, and possibly may also be found other modalities (Riesenhuber & Poggio, 1999a). There is already evidence MAX-like mechanisms exist in area V1 (Sakai & Tanaka, 1997), where the linear response of simple cells and the larger receptive field and nonlinear response of complex cells make them potential candidates for the input and output units in our models, respectively. Preliminary data from intracellular recordings in area V1 of anesthetized cats indicate that some of the complex cells's behaviors approximate the MAX operation when 2 spots or optimally oriented bars are presented alone or simultaneously at different separations and contrast levels (Lampl, Riesenhuber, Poggio, & Ferster, 2001). There is also support for the existence of MAX operation in the inferotemporal cortex (Sato, 1989).

Another very important question is whether the MAX operation is implemented at a cellular or networks level. That is, are the computational operations mediated by dendritic/synaptic mechanisms or network dynamics? For the divisive feedforward and feedback networks, synaptic shunting inhibition is the mostly likely candidate, while multiplication achieved by dynamic interactions of linear threshold neurons is also plausible (Salinas & Abbott, 1996). Although the topic of shunting inhibition is still under heated debate, the latest evidence indicate that not only do conductance changes contribute to shunting inhibition, but that they are more prominent in simple cells than in complex cells (Anderson et al., 2000), in accordance with our hypothesis that pooling from simple cells operates in a MAX-like fashion in the visual cortex. For the linear threshold and spiking models, network interactions are the prime candidates. To investigate the potential synaptic/dendritic mechanisms, injection of currents or monitoring of excitatory and inhibitory synaptic currents in conjunction with patch-clamping can be used to tease apart the possibilities (Aksay, Gamkrelidze, Seung, Baker, & Tank, 2001).

⁷The DFB model has the following Lyapunov function: $L(\mathbf{y}) = -\log(c + \mathbf{1}^T \mathbf{y}) + \sum_n (1/x_n) \int_0^{\log f(y_n)} f^{-1}(e^\gamma) d\gamma$. A Lyapunov function has also been shown to exist for the LIN model (Grossberg, 1988; Hahnloser, 1998). The same Lyapunov function also applies to the SPK model in the limit of high firing rate.

When a neuron or a group of neurons have reliably been confirmed to be involved in a MAX-like operation via the postulated three-layer architecture, we would then need a set of model-specific predictions to test whether these neurons implement the MAX operation via a network structure similar to any of our models. Some of these model-specific predictions are summarized in Table 1, touching upon anatomical, neurophysiological, and computational aspects of the individual models that we have discussed in this paper.

The active involvement of shunting inhibition or gating would suggest divisive interactions at the synaptic level, lending more plausibility to the cellular implementation of the divisive feedforward and feedback models. The homogeneity property was discussed in section 3.3, and is expected to be found for the *active* hidden units and the output unit of the linear threshold and spiking models. This can be investigated by manipulating the sensory input and monitoring the input-output relationship. Sparse hidden-layer activation is expected to be found in models where MAX-operation-like behaviour is achieved only when there is a single significantly active unit in the hidden unit, as in the case of the DFB and SPK models (see Figure 4 for simulation results and (Giese et al., 2001) for mathematical analysis). Hysteresis was discussed in the previous section. If it can be induced in neurons presumably involved in the MAX operation, then there must recurrent interactions involved.

6 Discussion

In this paper we have reviewed and analyzed a number of neural circuits that provide good approximations to the MAX operation, which has been proposed to play a significant role in various processes of the visual system. The MAX operation is interesting because it is an example of a fundamental, nonlinear computational operation that can be realized with neurophysiologically plausible mechanisms. The circuits were chosen in order to demonstrate different neural principles for the realization of this computational operation. They are simple enough to allow an understanding of the underlying parametric dependencies and some mathematical analysis. The neural mechanisms on which our models are based are also fundamental in models for *contrast gain control* and *winner-take-all*, both of which have already been extensively studied in the context of important cortical processes. Our analysis suggests that these processes and the MAX operation may share similar or overlapping neural substrates.

From the neurophysiological perspective, it appears that the divisive feedforward network and the linear threshold model (the spiking model is a variation on the latter) are more plausible than the divisive feedback network. The divisive feedforward model uses a neurophysiologically less complicated form of shunting inhibition/multiplication than the divisive feedback model; the linear threshold model uses a standard form of linear threshold feedback, which has been demonstrated experimentally in many different systems in many different contexts. Moreover, the feedforward model and the linear threshold model are both superior over the divisive feedback model in that they require smaller inhibitory gains in order to achieve equivalent proximity to the ideal MAX operation. From the computational perspective, the divisive feedback network has the advantage over the others in that it is particularly resistant to input noise. The remaining differences among the models do not lead to immediate support for one model over any other due to the lack of sufficient experimental data.

This work gives rise to a number of potential directions of future theoretical research. One obvious extension of the current work is to analyze neurophysiologically more detailed and more realistic models, which could involve more stochastic descriptions of network dynamics, or biophysically more realistic neurons. Another interesting question is how any of these models may be learned through experience or wired up during development. Mechanisms similar to those proposed by Fukushima to explain learning in the Neocognitron may be explored (Fukushima, 1980). The details of these future studies are, however, contingent upon the availability of relevant experimental data that are yet to be obtained. The simulations and mathematical results from our analysis give rise to a number of predictions, which can be used to guide future experimental investigations on the MAX operation: where it occurs in the brain, and which of the four circuits, if any, is actually implemented by the biological system. Although the experimental demonstration of the different properties predicted by the models is non-trivial, this preliminary theoretical analysis should be a helpful first step for the preparation of more detailed neurophysiological experiments.

7 Acknowledgement

We thank Peter Dayan, Richard Hahnloser, and Maximilian Riesenhuber for helpful discussions. This work was supported by the Office of Naval Research under contract No. N00014-93-1-3085, Office of Naval Research under

contract No. N00014-95-1-0600, National Science Foundation under contract No. IIS-9800032, National Science Foundation under contract No. DMS-9872936, National Science Foundation under Graduate Research Fellowship Program, Massachusetts Institute of Technology under the UROP program, and the Gatsby Charitable Foundation under a grant to the Gatsby Computational Neuroscience Unit. Additional support is provided by Honda R&D, U.S.A.

Appendix

The mathematical discussions in this section are intended to help explain and support the experimental results presented in section 4. We are mainly interested in how well the networks satisfy the selectivity and linearity properties of the MAX operation, and the relationship between the hidden layer population activities and the final output.

A Divisive Feedforward Network

Here, we only give a detailed analysis of the case where $f(x) = e^{qx}$, since it is a highly tractable model. But the insight we gain from this example generalize to other examples of $f(x)$.

Without loss of generality, first assume $x_1 \geq x_i$, for all $1 < i \leq N$. Let $a_i \geq x_1 - x_i$ and assume $c \ll 1$, then we have the following:

$$z \approx x_1 - \frac{\sum_{i>1} a_i e^{-qa_i}}{1 + \sum_{j>1} e^{-qa_j}} \quad (13)$$

Let $L = \frac{\sum_{i>1} a_i e^{-qa_i}}{1 + \sum_{j>1} e^{-qa_j}}$, then finding the worst-case performance or the lower-bound on z is equivalent to finding the maximum for L . Setting the partial derivative of L with respect to each a_i to 0, we obtain the following:

$$qa_i = \frac{\sum_{j>1} qa_j e^{qa_j}}{1 + \sum_{j>1} e^{-qa_j}} + 1 \quad (14)$$

Notice the right hand side does not depend on i . Therefore the worst case scenario is when all of the sub-maximal inputs are equal to each other. Let b denote the quantity qa_i , and b^* denote the value of b that solves Equation 14, then we have the following:

$$b^* = \frac{(N-1)be^{-b^*}}{1 + (N-1)e^{-b^*}} + 1 \quad (15)$$

After re-arranging the terms, it turns out $b^* - 1$ is equal to the Lambert W function of $(N-1)e^{-1}$. Let a denote the magnitude of all the non-maximal inputs, such that $qa = b^*$ for a given q , and let $L^*(q)$ denote the corresponding maximal value of L as a function of q , then we have the following:

$$z^* = x_1 - L^*(q) = x_1 - \frac{q^{-1}}{1 + (N-1)^{-1}e^b} \quad (16)$$

Note $L^*(q)$ is proportional to a or inversely proportional to q .

Equation 16 implies that both the *selectivity* and the *linearity* properties of the ideal MAX operation are better approximated as q becomes large. More precisely, as $q \rightarrow \infty$, $L(x, q) \rightarrow 0$, and $z \rightarrow x_{\max}$, achieving perfect selectivity and unit-gain linearity. This is the trend we see in Figure 4(b). At the other extreme, if q is small, so that 1 dominates in the denominator of the fraction in L , then L is linear in a . If then a scales with the input amplitude x_1 , as may well be the case, then z is linear in x_1 . Therefore, even though selectivity monotonically deteriorates with smaller q , good linearity can still be achieved. Notice how this agrees with simulation results shown in Figure 2(c) and 3(a). Figure 2(c) shows that the FFN model achieves good linearity with q as small as 3. Figure 3(a) shows that even with q as large as 20, the model is significantly affected by non-maximal inputs valued around 0.8. The phenomenon that the dip increases in size and expands toward smaller magnitudes of non-maximal inputs with smaller q is also explained by the the relationship between the maximal value of L and the product qa . In Figure 5, the slight dip and then the stabilization or z as a function of N is also explained by equation 16.

To analyze the exact relationship between the maximal response in the hidden layer and the final output, we examine the following form of y_1 in the worst-case scenario:

$$y_i^* = \frac{x_i}{1 + (N-1) * e^{-b^*}} \quad (17)$$

We see from this equation that the relative order of the inputs are preserved in the hidden layer: $y_i \geq y_j$ if and only if $x_i \geq x_j$, and $y_i = y_j$ if $x_i = x_j$. Also, since $z \leq x_1$ and $y_1 \leq z$, z is always a better estimate of x_1 than y_1 . Equation 17 also explains the sigmoidal decay of y_1 as a function of magnitude of non-maximal inputs in Figure 3 and its decay as a function of N in Figure 5.

References

- Aksay, E., Gamkrelidze, G., Seung, H. S., Baker, R., & Tank, D. W. (2001). In vivo intracellular recording and perturbation of persistent activity in a neural integrator. *Nature*, 4(2), 184-193.
- Anderson, C. H., & Essen, D. C. van. (1987). Shifter circuits: a computational strategy for dynamic aspects of visual processing. *Proceedings of the National Academy of Sciences (USA)*, 84(17), 6297-6301.
- Anderson, J. S., Carandini, M., & D, F. (2000). Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *Journal of Neurophysiology*, 84, 909-926.
- Ben-Yishai, R., Lev Bar-Or, R., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences (USA)*, 92, 3844-3848.
- Borg-Graham, L., Monier, C., & Fregnac, Y. (1998). Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature*, 393, 367-373.
- Carandini, M., & Ferster, D. (2000). Membrane potential and firing rate in cat primary visual cortex. *Journal of Neurophysiology*, 20(1), 470-484.
- Carandini, M., & Heeger, D. J. (1994). Summation and division by neurons in primate visual cortex. *Science*, 264, 1333-1336.
- Connor, C. E., Preddie, D. C., Gallant, J. L., & Van Essen, D. C. (1997). Spatial attention effects in macaque area V4. *Journal of Neuroscience*, 17(9), 3201-3214.
- Ferster, D., & Jagadeesh, B. (1992). EPSP-IPSP interactions in cat visual cortex studied with in vivo whole-cell patch recording. *Journal of Neuroscience*, 12(4), 1262-1274.
- Ferster, D., & Miller, K. D. (2000). Neural mechanisms of orientation selectivity in the visual cortex. *Annual Review of Neuroscience*, 23, 441-471.
- Fukui, T., & Tanaka, S. (1997). A simple neural network exhibiting selective activation of neural ensembles: From winner-takes-all to winner-shares-all. *Neural Computation*, 9, 77-97.
- Fukushima, K. (1980). Neocognitron: a self organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36(4), 193-202.
- Furshpan, E. J., & Potter, D. D. (1959). Transmission of giant motor synapses in the crayfish. *Journal of Physiology*, 145, 289-325.
- Giese, M. A. (2000). Neural field model for the recognition of biological motion. *Sec. Int. IcSC Symp. on Neur. Comp. NC 2000*. (in press)
- Giese, M. A., Yu, A. J., & Poggio, T. A. (2001). *Mathematical analyses for circuits that implement the maximum operation*. (Manuscript in preparation.)
- Gilbert, C. D. (1992). Horizontal integration and cortical dynamics. *Neuron*, 9(1), 1-13.
- Grossberg, S. (1973). Contour enhancement, short term memory, and constancies in reverberating neural networks. *Studies in Applied Mathematics*, 52, 213-257.
- Grossberg, S. (1988). Nonlinear neural networks: principles, mechanisms, and architectures. *Neural Networks*, 1, 17-61.
- Grzywacz, N. M., & Yuille, A. L. (1990). A model for the estimate of local image velocity by cells in the visual cortex. *Proc. R. Soc. Lond. B*, 239, 129-161.
- Hahnloser, R. (1998). On the piecewise analysis of networks of linear threshold neurons. *Neural Networks*, 11, 691-697.

- Hahnloser, R., Douglas, R. J., Mahowald, M., & Hepp, K. (2000). Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature*, 405(6789), 947-951.
- Hartline, H. K., & Ratliff, F. (1957). Spatial summation of inhibitory influences in the eye of limulus, and the mutual interaction of receptor units. *Journal of General Physiology*, 41(5), 1049-1066.
- Holt, G. R., & Koch, C. (1997). Shunting inhibition does not have a divisive effect on firing rates. *Neural Computation*, 9, 1001-1013.
- Hubel, D. H., & Wiesel, T. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology*, 160, 106-154.
- Koch, C., & Poggio, T. (1989). The synaptic veto mechanism: does it underlie direction and orientation selectivity in the visual cortex. In D. Rose & V. G. Dobson (Eds.), *Models of the visual cortex* (p. 15-34). John Wiley.
- Koch, C., Poggio, T., & Torre, V. (1983). Nonlinear interactions in the dendritic tree: localization timing and the role in information processing. *Proceedings of the National Academy of Sciences (USA)*, 80, 2799-2802.
- Koch, C., & Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. *Human Neurobiology*, 4, 219-227.
- Kohonen, T. (1995). *Self organizing maps*. Springer-Verlag, Berlin.
- Lampl, I., Riesenhuber, M., Poggio, T., & Ferster, D. (2001). The max operation in cells in the cat visual cortex. *Society of Neuroscience Abstracts*.
- Lazzaro, J., Ryckenbusch, S., Mahowald, M. A., & Mead, C. A. (1989). Winner-take-all network of O(n) complexity. In D. Touretzky (Ed.), *Advances in neural information processing systems 1* (p. 703-711). Morgan Kaufmann, San Mateo, CA.
- Lee, D. K., Itti, L., Koch, C., & Braun, J. (1999). Attention activates winner-takes-all competition among visual filters. *Nature Neuroscience*, 2, 375-381.
- Li, Z. (2000). *Computational design and nonlinear dynamics of recurrent models of the primary visual cortex* (Technical report No. GCNU TR 2000-001). 17 Queen Sq, London WC1N 3AR, U.K.
- Logothetis, N. K., Pauls, J., & Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. *Current Biology*, 5, 552-563.
- Mel, B. W., & Ruderman, K. A., D L Archie. (1998). Translation-invariant orientation tuning in visual "complex" cells could derive from intradendritic computations. *Journal of Neuroscience*, 18(11), 4325-4334.
- Motter, B. C. (1994). Neural correlates of attentive selection for color or luminance in extrastriate area V4. *Journal of Neuroscience*, 14(4), 2178-2189.
- Naka, K. I., & Rushton, W. A. H. (1966). S-potential from luminosity units in the retina of fish (cyprinidae). *Journal of Physiology*, 185, 587-599.
- Nowlan, S. J., & Sejnowski, T. J. (1995). A selection model for motion processing in area mt of primates. *Journal of Neuroscience*, 15(2), 1195-1214.
- Pasupathy, A., & Connor, C. E. (1999). Responses to contour features in macaque area v4. *Journal of Neurophysiology*, 82(5), 2490-2502.
- Perrett, D. I., Oram, M. W., Harries, M. H., Bevan, R., Hietanen, J. K., Benson, P. J., & Thomas, S. (1991). Viewer-centred and object-centred coding of heads in the macaque temporal cortex. *Experimental Brain Research*, 86(1), 159-173.
- Reichardt, W., Poggio, T., & Hausen, K. (1983). Figure-ground discrimination by relative movement in the visual system of the fly. Part II. towards the neural circuitry. *Biological Cybernetics*, 46.

- Riesenhuber, M. K., & Poggio, T. (1999a). Are cortical models really bound by the 'binding problem'? *Neuron*, *24*, 87-99.
- Riesenhuber, M. K., & Poggio, T. (1999b). Hierarchical models of object recognition in cortex. *Nature Neuroscience*, *2*, 1019-1025.
- Riesenhuber, M. K., & Poggio, T. (2000). Models of object recognition. *Nature Neuroscience*, *3*, Supp., 1199-1204.
- Rockland, K. S., & Lund, J. S. (1983). Intrinsic laminar lattice connections in primate visual cortex. *Journal of Comparative Neurology*, *216*, 303-318.
- Sakai, K., & Tanaka, S. (1997). *Soc. Neurosci. Abs.*, *23*, 453.
- Salinas, E., & Abbott, L. F. (1996). A model of multiplicative neural responses in parietal cortex. *Proceedings of the National Academy of Sciences (USA)*, *93*, 11956-11961.
- Sato, T. (1989). Interactions of visual stimuli in the receptive fields of inferiortemporal neurons in awake monkeys. *Experimental Brain Research*, *77*, 23-30.
- Simoncelli, E. P., & Heeger, D. J. (1998). A model of neural responses in visual area MT. *Vision Research*, *38*, 743-761.
- Starzyk, J. A., & Fang, X. (1993). CMOS current mode winner-take-all circuit with both excitatory and inhibitory feedback. *Electronics Letters*, *29*(10), 908-910.
- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience*, *19*, 109-139.
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, *381*, 520-522.
- Torre, V., & Poggio, T. (1978). A synaptic mechanism possibly underlying directional selectivity motion. *Proceedings of the Royal Society London B*, *202*, 409-416.
- Treue, S., & Maunsell, J. H. L. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature*, *382*, 539-541.
- White, E. L. (1989). *Cortical circuits*. Birkhauser.
- Wilson, H. R., & Humanski, R. (1993). Spatial frequency adaptation and contrast gain control. *Vision Research*, *33*, 1133-1149.
- Yuille, A. L., & Grzywacz, N. M. (1989). A winner-take-all mechanism based on presynaptic inhibition feedback. *Neural Computation*, *1*, 334-347.