

A single functional model accounts for the distinct properties of suppression in cortical area V1

M. W. Spratling

King's College London, Department of Informatics and Division of Engineering, London. UK.

Abstract

Cross-orientation suppression and surround suppression have been extensively studied in primary visual cortex (V1). These two forms of suppression have some distinct properties which has led to the suggestion that they are generated by different underlying mechanisms. Furthermore, it has been suggested that mechanisms other than intracortical inhibition may be central to both forms of suppression. A simple computational model (PC/BC), in which intracortical inhibition is fundamental, is shown to simulate the distinct properties of cross-orientation and surround suppression. The same model has previously been shown to account for a large range of V1 response properties including orientation-tuning, spatial and temporal frequency tuning, facilitation and inhibition by flankers and textured surrounds as well as a range of other experimental results on cross-orientation suppression and surround suppression. The current results thus provide additional support for the PC/BC model of V1 and for the proposal that the diverse range of response properties observed in V1 neurons have a single computational explanation. Furthermore, these results demonstrate that current neurophysiological evidence is insufficient to discount intracortical inhibition as a central mechanism underlying both forms of suppression.

1 Introduction

Inhibitory mechanisms intrinsic to primary visual cortex (V1) were originally believed to be responsible for both cross-orientation suppression (Bonds, 1989; DeAngelis et al., 1992; Morrone et al., 1982) and surround suppression (DeAngelis et al., 1994; Fitzpatrick, 2000; Knierim and van Essen, 1992). Subsequent work has demonstrated that these two forms of suppression exhibit a number of distinct properties, and hence, may result from different underlying mechanisms. A range of alternative mechanisms have been proposed for each type of suppression. Specifically, it has been suggested that cross-orientation suppression might arise from attenuation of the feedforward input due to depression of the thalamocortical synapses (Carandini et al., 2002; Freeman et al., 2002) or a reduction in feedforward drive to cortical cells caused by contrast saturation in lateral geniculate nucleus (LGN) cells (Li et al., 2006; Priebe and Ferster, 2006). It has also been suggested that surround suppression might be mediated by inhibitory mechanisms intrinsic to V1 but driven by feedback from extrastriate cortex (Angelucci et al., 2002; Bair et al., 2003; Cavanaugh et al., 2002; Sullivan and de Sa, 2006) or might be due to surround suppression in LGN (Naito et al., 2007; Ozeki et al., 2004; Webb et al., 2005).

A previous publication (Spratling, 2010) has demonstrated that a simple functional model (PC/BC), derived from the predictive coding and biased-competition theories of cortical function, can simulate a very wide range of V1 response properties including cross-orientation and surround suppression. This article extends that work by showing that the PC/BC model of V1 can also simulate the distinct behaviours exhibited by these two forms of suppression. The PC/BC model includes two mechanisms that can give rise to suppression: a mechanism of intracortical inhibition employing divisive normalisation of the inputs to a population of competing neurons; and saturation of the LGN responses to high contrast stimuli. The latter mechanism was proposed by Li et al. (2006); Priebe and Ferster (2006) to account for cross-orientation suppression. It is found that in the PC/BC model surround suppression is generated by the mechanism of cortical inhibition, while cross-orientation suppression is generated by a combination of cortical inhibition and LGN response saturation. Hence, the PC/BC model predicts that intracortical inhibition is essential for both forms of suppression, contrary to suggestions that completely separate mechanisms are required and to claims that cortical inhibition is not involved.

2 Methods

2.1 The LGN Model

The input to the PC/BC model of V1, described below, was an input image (I) pre-processed by convolution with a Laplacian-of-Gaussian (LoG) filter (l) with standard deviation equal to one. This is virtually identical to the Difference-of-Gaussians (DoG) filter that has traditionally been used to model circular receptive fields (RFs)




Figure 1: The PC/BC model of V1. The input image I is preprocessed by convolution with a circular-symmetric on-centre/off-surround kernel (to generate the input to the ON channel of the V1 model), and a circular-symmetric off-centre/on-surround kernel (to generate the input to the OFF channel of the V1 model). The prediction neurons (labelled Y) represent V1 simple cells. Responses from these neurons were recorded during most experiments. These responses were generated by convolving the outputs of the (ON and OFF channels of the) error-detecting neurons (labelled E) with (the ON and OFF channels of) a number of kernels representing V1 RFs. This convolution process effectively reproduces the same RFs at every pixel location in the image. The responses of the error-detecting neurons are influenced by divisive feedback from the prediction neurons, which is also calculated by convolving the prediction neuron outputs with the weight kernels. Complex cell responses (labelled D), which were also recorded during some experiments, were calculated by taking the maximum response of a small population of simple cells representing a single orientation at a particular spatial location.

in LGN and retina. The output from this filter was subject to a multiplicative gain (the strength of which was determined by parameter κ) followed by a saturating non-linearity, such that:

$$\mathbf{X} = \tanh \{ \kappa (I * l) \} \quad (1)$$

The positive and rectified negative responses were separated into two images \mathbf{X}_{ON} and \mathbf{X}_{OFF} simulating the outputs of cells in retina and LGN with circular-symmetric on-centre/off-surround and off-centre/on-surround RFs respectively. This pre-processing is illustrated on the left of Figure 1. Consistent with neurophysiological data (Reid and Alonso, 1995), the ON-centre model LGN neurons indirectly provided input to the ON sub-field of the model V1 simple cells, while the OFF-centre model LGN neurons indirectly provided input to the OFF sub-field of the model V1 neurons (see next section).

To explore the effects of the different mechanisms of suppression, in some experiments the suppression of the LGN responses was turned off. In this case equation 1 was replaced by:

$$\mathbf{X} = \kappa (I * l) \quad (2)$$

A value of $\kappa = 10$ was used in all experiments reported here.

2.2 The V1 Model

The PC/BC model of V1 is illustrated in Figure 1 and described by the following equations:

$$\mathbf{E}_o = \mathbf{X}_o \otimes \left(\epsilon_2 + \sum_{k=1}^p (\hat{w}_{ok} * \mathbf{Y}_k) \right) \quad (3)$$

$$\mathbf{Y}_k \leftarrow (\epsilon_1 + \mathbf{Y}_k) \otimes \sum_o (w_{ok} * \mathbf{E}_o) \quad (4)$$

Where $o \in [ON, OFF]$, \mathbf{X}_o is a 2-dimensional array, equal in size to the input image, that represents the input to the model of V1, \mathbf{E}_o is a 2-dimensional array, equal in size to the input image, that represents the error-detecting neuron responses, and \mathbf{Y}_k is a 2-dimensional array, equal in size to the input image, that represent the prediction neuron responses, w_{ok} is a 2-dimensional kernel representing the synaptic weights for a particular class (k) of neuron normalised so that sum of all the weights was equal to ψ , \hat{w}_{ok} is a 2-dimensional kernel representing the

same synaptic weights as w_{ok} but normalised so that the maximum value was equal to ψ , p is the total number of kernels, ϵ_1 , ϵ_2 , and ψ are parameters, \oslash and \otimes indicate element-wise division and multiplication respectively, \star represents cross-correlation (which is equivalent to convolution without the kernel being rotated 180°), and $*$ represents convolution (which is equivalent to cross-correlation with a kernel rotated by 180°). Parameter values $\psi = 5000$, $\epsilon_1 = 0.0001$ and $\epsilon_2 = 250$ were used in the simulations reported in this article.

Equation 4 describes the updating of the prediction neuron activations. The response of each prediction neuron is a function of its activation at the previous iteration and a weighted sum of afferent inputs from the error-detecting neurons. Equation 3 describes the calculation of the neural activity for each population of error-detecting neurons. These values are a function of the activity of the input to V1 divisively modulated by a weighted sum of the outputs of the prediction neurons in V1. The activation of the error-detecting neurons can be interpreted in two ways. Firstly, \mathbf{E} can be considered to represent the residual error between the input and the reconstruction of the input generated by the prediction neurons. The values of \mathbf{E} indicate the degree of mismatch between the top-down reconstruction of the input and the actual input (assuming ϵ_2 is sufficiently small to be negligible). When a value within \mathbf{E} is greater than $\frac{1}{\psi}$ it indicates that a particular element of the input is under-represented in the reconstruction, a value of less than $\frac{1}{\psi}$ indicates that a particular element of the input is over-represented in the reconstruction, and a value of $\frac{1}{\psi}$ indicates that the top-down reconstruction perfectly predicts the bottom-up stimulation. A second interpretation is that \mathbf{E} represents the inhibited inputs to a population of competing prediction neurons. Each prediction neuron modulates its own inputs, which helps stabilise the response of the prediction neurons, since a strongly (or weakly) active prediction neuron will suppress (magnify) its inputs, and hence, reduce (enhance) its own response. Prediction neurons that share inputs (*i.e.*, that have overlapping RFs) will also modulate each other's inputs. This generates a form of competition between the prediction neurons, such that each neuron effectively tries to block other prediction neurons from responding to the inputs which it represents. This mechanism of competition is called Divisive Input Modulation (DIM) (Spratling et al., 2009).

The RF of a simple cell in primary visual cortex can be accurately modelled by a 2-dimensional Gabor function (Daugman, 1980, 1988; Jones and Palmer, 1987; Lee, 1996; Marcelja, 1980). Hence, the Gabor function was used to define the weights of each kernel w_{ok} . A definition of a Gabor function of the form proposed by Lee (1996) was used, which includes a term to remove the D.C. response of the filter:

$$g(\sigma, \gamma, \lambda, \phi, \theta) = \exp \left\{ -\frac{x'^2 + (y'/\gamma)^2}{2\sigma^2} \right\} \left[\cos \left\{ \frac{2\pi y'}{\lambda} + \phi \right\} - \cos(\phi) \exp \left\{ -\left(\frac{\pi\sigma}{\lambda} \right)^2 \right\} \right]$$

Where $\sigma = 4$ (pixels) was a constant that defined the standard deviation of the Gaussian envelope (which determines the spatial extent of the RF), $\gamma = \frac{1}{\sqrt{2}}$ was a constant that defined the aspect ratio of Gaussian envelope (which determines the ellipticity of the RF), $\lambda = 6$ (pixels) was a constant that defined the wavelength of the sinusoid, ϕ was the phase of the sinusoid, and $x' = x \cos(\theta) + y \sin(\theta)$ and $y' = -x \sin(\theta) + y \cos(\theta)$ where θ defined the orientation of the RF. A family of 32 Gabor functions (Fig. 2a) with eight orientations ($\theta = 0$ to 157.5° in steps of 22.5°) and four phases ($\phi = 0^\circ, 90^\circ, 180^\circ$, and 270°) were used to define the RFs of the neurons in the model. The weights were separated into distinct ON and OFF channels which represented the positive and negative parts of the Gabor function using separate sets of non-negative weights (Fig. 2b). The cross-correlation and convolution performed in Equations 3 and 4 mean that neurons with these RFs are reproduced at every pixel location in the image, and consequently, that the size of the population of V1 cells simulated varies with image size. For an $a \times b$ pixel image, the model simulates the response of $32ab$ prediction neurons.

The responses of complex cells (\mathbf{D}) were simulated by making the complex cell activation equal to the maximum activity of all prediction neurons (which model simple cells) with the same orientation preference in a local neighbourhood. Specifically, the maximum activation was found across the four phases and in a 3-by-3 pixel spatial neighbourhood. This model of complex cells is similar to that employed in the HMAX model (Riesenhuber and Poggio, 1999) which is in turn an idealised version of the hierarchical model proposed by Hubel and Wiesel (1962). There is some neurophysiological support for this model (Gawne and Martin, 2002; Kouh and Poggio, 2008; Lampl et al., 2004).

Complex cells in the model do not provide feedback to either the error-detecting neurons or the simple cells/prediction neurons. It would not be appropriate for the \mathbf{D} population to contribute to the top-down input to the \mathbf{E} population. From \mathbf{Y} to \mathbf{D} the sensory input is recoded to deliberately remove information about the exact location and phase of the stimulus. The response of the complex cells would therefore not contribute accurate information to the top-down reconstruction of the input that is required by the error-detecting neurons. Reciprocal connections from the \mathbf{D} population to the \mathbf{Y} population may be of value, particularly if feedback connections from subsequent processing stages target the complex cells. Connections from \mathbf{D} to \mathbf{Y} could therefore transmit that information to modulate the predictions of the causes of the sensory input represented by the prediction neurons. However, since this article only considers a single cortical region in isolation these connections have not been




Figure 2: The synaptic weights used in the PC/BC model of V1. (a) A family of 32 Gabor functions (eight orientation and four phases) used to define the RFs of the neurons in the model. (b) The actual synaptic weights of the model neurons were created by separating the positive and negative parts of the Gabor function into separate (non-negative) ON and OFF weights (shown for the bottom-right Gabor function only). Each Gabor kernel is 21×21 pixels, and hence, each prediction neuron in the model receives $21 \times 21 \times 2 = 882$ synaptic weights.

included here. Hence, the existing model of simple cell response properties is unaffected by the addition of the complex cell population and all the previous results for this model (Spratling, 2010) are still valid.

To disentangle the effects of competition between the neurons in the PC/BC model of V1 from the effects of suppression caused by the saturation of the neurons in the LGN model, experiments were also performed with a linear model of V1 simple cells. In this case Equations 3 and 4 were replaced by:

$$\mathbf{Y}_k = \epsilon_1 \sum_o w_{ok} \star \mathbf{X}_o \quad (5)$$

This represents the output produced by a set of linear filters when applied to the image. This linear response is constant over time assuming that the input image does not change.

2.3 Experimental Procedure

At the start of each simulation the prediction neuron responses (\mathbf{Y}) were initialised to zero. Then equations 3 and 4 (or equation 5) were applied recursively for a number of iterations (t). During this time the values of \mathbf{Y} , for a particular simple cell and/or the values of \mathbf{D} for a particular complex cell were recorded. Results were presented either by showing the value of response as a function of time t , or by averaging the response over t and presenting the mean evoked response as a function of changing a property of the stimulus.

As for typical physiological experiments, the stimulus parameters other than the one being varied during the experiment, were matched to the preferred parameters of the neuron under test (*e.g.*, the stimulus was centred over the RF, at the recorded neuron’s preferred orientation, spatial frequency, temporal frequency, *etc.*). Furthermore, the range of grey-scale values in the input image I were set equal to the fractional Michelson contrast used for the presentation of stimuli in the corresponding physiological experiment, if this value was reported. In each experiment where a simple cell response was recorded, this cell had identical tuning properties to those shown in (Spratling, 2010, Fig. 5). In each experiment the recorded neuron had an orientation preference of 0° . A grating at 0° is shown as vertical in the icons used in the figures in section 3 (however, note that a grating at 0° is shown as horizontal in (Spratling, 2010)).

In some experiments the image remained constant throughout the recording time, in other experiments the image changed (“transitioned”) to a second image at a particular iteration. These “constant” images presented throughout the recording, or before and after the transition, were either static sinusoidal gratings or drifting sinusoidal gratings. For a static image the elements of I , and hence, the values of \mathbf{X}_{ON} and \mathbf{X}_{OFF} remained constant during the presentation of the static image. For a drifting grating the input image was changed, and new \mathbf{X}_{ON} and \mathbf{X}_{OFF} values were calculated, at each iteration of the PC/BC algorithm. The amount the input image changed between consecutive iterations reflected the drift rate of the stimulus. This was measured in terms of cycles per iteration, where the number of cycles refers to the phase shift between sinusoids in consecutive images.

In some experiments it was necessary to calculate the latency between a transition in the input image and the changed response evoked in the recorded neuron. The same method was used as in the corresponding neurophysiological experiments (Bair et al., 2003, 2002; Smith et al., 2006), namely, the latency was measured as the time




Figure 3: Graphs of results are shown in the format illustrated. Each experiment was simulated three times; once using both mechanisms of suppression, once using only intracortical inhibition, and once using only LGN saturation. Results for these three simulations are shown in plots arranged as shown above. The main simulation results, using both forms of suppression, are surrounded by a thick line to help them stand-out. In some cases data from a neurophysiological experiment are shown above the corresponding simulation results. Many figures contain multiple panels, in which case each panel is arranged as shown above.

at which the change in response first reached 5% of the maximum change. These times were linearly interpolated between iterations of the PC/BC algorithm.

2.4 Code

Software, written in MATLAB, which implements the PC/BC model described above and performs the experiments described below is available at http://www.corinet.org/mike/Code/v1_suppression_mechanisms.zip.

3 Results

The following sections present simulations of a number of experiments performed to assess the characteristics of suppression mechanisms affecting V1 response properties. These experiments examine features of cross-orientation and surround suppression in terms of response latencies, the gain of the contrast-response function, the effects of high temporal frequency suppressive stimuli, and the effects of one suppressive stimulus on another. Each set of simulations were performed using the full model in which both mechanisms of suppression (LGN saturation and V1 inhibition) operate. These results simulate the corresponding neurophysiological experiments. Simulations were also performed using versions of the model in which only one or other of the two suppression mechanisms operated. These additional simulations help to determine the contributions of each form of suppression to the overall result. The results from these different simulations are presented in plots arranged in the layout shown in Fig. 3.

3.1 Latency of Response

Bair et al. (2003) and Smith et al. (2006) performed experiments in V1 to assess the dynamics of suppression. They recorded the responses of cells during a change in the identity of the presented stimulus and compared this to the response when the stimulus remained constant. Figure 4 shows results recorded from a prediction neuron in the PC/BC model to the same stimulus transitions, and table 1 summarises the latencies found in both the neurophysiological experiments and with the model. Note that in the model there are no transmission delays between retina and V1, and hence, model V1 neurons can respond with very little latency. However, there is a strong, and statistically significant, correlation between the latencies recorded in V1 and those recorded in the

transition	mean latency in V1 (ms)		latency in model (iterations)	
	cross-orientation suppression	surround suppression	cross-orientation suppression	surround suppression
onset	50.0	52	2.36	2.61
offset	30.1	35	0.05	0.13
suppression	42.5	61	0.20	2.90
release	40.9	60	0.28	4.53

Table 1: The latency of response to different stimulus transitions. Data on the left show the average latency recorded in V1 for transitions in cross-orientation suppression stimuli (taken from [Smith et al., 2006](#), Fig. 3), and for surround suppression stimuli (taken from [Bair et al., 2003](#), Fig. 4). Data on the right show the latency generated by the model for the same stimulus transitions.

model ($r = 0.914$, $p < 0.002$). The simulation results presented in this section were recorded from a model simple cell using static stimuli. However, very similar results were generated by model complex cells in response to drifting gratings.

When the stimulus presented to the RF of a cortical neuron changes from one that is poor at driving that neuron, to one that is good at generating a response (an “onset” transition) the latency of the change in response is significantly longer than for the reverse transition (an “offset” transition) where the stimulus changes from good to poor ([Bair et al., 2002](#); [Smith et al., 2006](#)). The model shows the same asymmetry in the latency of onset and offset responses as shown in Figs. 4a and c. [Bair et al. \(2002\)](#) proposed that the longer latency for the onset response might be due to the time taken for the cell to reach its firing threshold after stimulus onset. Although the PC/BC model does not include a firing threshold, the longer latency response to stimulus onset has a similar cause: the speed at which a prediction neuron’s response changes is proportional to the current activity of that neuron, due to the feedforward drive to the neuron being multiplicatively modulated by the current response (see equation 4). Hence, when the model neuron has low activation (*e.g.*, before stimulus onset), it takes time for its response to increase, whereas when it has high activation (*e.g.*, before stimulus offset) its response can change quickly. Hence, as shown in the insets to Figs. 4a and c, removing the saturation of the model LGN neurons has little effect on the behaviour (upper insets), whereas removing intracortical competition (which replaces equation 4 with one in which response is a linear function of the input) eliminates the onset delay (lower insets).





[Smith et al. \(2006\)](#) found that when a mask was added to the preferred stimulus of a recorded neuron (a “suppression” transition), the latency with which cross-orientation suppression was generated was similar to the latency of the change in response for the reverse transition where the mask was removed from the stimulus (a “release” transition). Hence, unlike onset and offset, suppression and release did not show any asymmetry in the latency with which they acted. The model shows the similar behaviour as shown in Figs. 4e and g.

In both V1 and the model, the latency for suppression was shorter than the latency for onset. This is difficult to reconcile with the idea that cross-orientation suppression is caused by intracortical inhibition from cells tuned to the orientation of the mask. If this were the case then suppression would be expected to occur with a latency longer than the onset time of the cells generating the suppression. [Smith et al. \(2006\)](#) therefore concluded that cross-orientation suppression is a result of suppression in the feedforward drive to the recorded neuron or is due to a fast mechanism of intracortical inhibition that is yet to be identified. The model proposes that cross-orientation suppression results from both intracortical inhibition and saturation in LGN. The LGN saturation is fast acting (since it affects the feedforward drive to the cortex) and gives rise to the rapid changes in response that occur when the mask is added or removed from the stimulus, as can be seen when the model is executed without intracortical inhibition (bottom insets to Figs. 4e and g.). The intracortical inhibition is much slower acting and the suppression it generates is delayed after mask onset by more than the onset delay for the neurons generating the suppression (compare the top insets for Figs. 4a and e.).

The asymmetry between onset and offset latency in V1 is also present when the stimulus driving the RF changes in the presence of a surround ([Bair et al., 2003](#)). The model shows a similar pattern of latency (see Figs. 4b and d) which is due to the same mechanism described above: a model neuron is slower to respond to a change in the input when its activation is low compared to when its activation is high.

[Bair et al. \(2003\)](#) found that when the central grating remained constant but the orientation of the surround grating changed from perpendicular to iso-oriented (causing greater suppression), the latency with which this suppression was generated was greater than the onset latency. This is consistent with the suppression being due to intracortical inhibition, since such inhibition should only occur after the response onset of those neurons activated by the surround. The model proposes that surround suppression results from intracortical inhibition and the onset of surround suppression does occur with a latency greater than the onset latency (Fig. 4f). Without intracortical





Onset



(a)

(b)





Offset



(c)

(d)





Suppression



(e)

(f)

Release



(g)

(h)

Figure 4: (previous page) The effect on the response of a model prediction neuron of a number of stimulus transitions compared to the response to an unchanging stimulus. The icons at the beginning and end of each trace illustrate the stimulus or stimuli that gave rise to that response: thin blue lines show the response to an unchanging stimulus, thick red lines show the response to the stimulus transition. For recordings where a transition occurred, time 0 (marked by the dotted vertical line) was the last iteration in which the initial stimulus was presented to the model, hence, time 1 was the first iteration after the transition. Each row shows a different stimulus transition condition: onset (first row), offset (second row), suppression (third row), and release (fourth row). The left column shows effects of stimulus transitions designed to test the dynamics of cross-orientation suppression. The right column shows effects of stimulus transitions designed to test the dynamics of surround suppression. The plots above each simulation result show data from corresponding neurophysiological experiments. For the left column the neurophysiological data is for a typical complex cell in anaesthetised primate V1 (adapted from [Smith et al., 2006](#), Fig. 2). For the right column the neurophysiological data is for a (different) typical complex cell in anaesthetised primate V1 (adapted from [Bair et al., 2003](#), Fig. 3). The smaller plots to the right of each simulation result show the response of the model to the same stimuli when there is no saturation to the LGN response (top) and when there is no competition between cortical neurons (bottom).

inhibition, see bottom inset to Fig. 4f, the response is very similar in both conditions, since the feedforward input received by the neuron is almost identical in both conditions, and it is only the competition in the full model that differentiates the response.

When the stimulus transition was reversed, so that the surround changed from iso-oriented to perpendicular (reducing suppression) the latency of the change in response in V1 was approximately the same as the latency for the onset of suppression ([Bair et al., 2003](#)). If surround suppression is mediated by intracortical inhibition, then it might be expected that the latency of surround suppression would be longer than the latency of the release from suppression. This is because the former should be due to the onset of RF driven responses in the cells generating the suppression and the latter should be due to the offset of RF driven responses in the cells generating the suppression. Hence, suppression and release would be expected to show the same asymmetry as onset and offset. To explain this unpredicted result [Bair et al. \(2003\)](#) proposed that the inhibitory inter-neurons transmitting the suppression become active quickly in response to the onset of input from the surround, but take a longer time to become inactive when the surround becomes inactive, effectively cancelling out the onset/offset asynchrony. However, the current model generates release from suppression with a long latency, without the need for additional mechanisms. The feedforward input received by the recorded neuron is almost identical in both surround orientations. However, it is marginally weaker for the perpendicular surround¹. The small reduction in feedforward input has the effect of delaying the increase in response that is generated by the reduction in inhibition from the surround.

Figure 5a shows the time course of the response of a model neuron to the onset of the preferred stimulus in comparison with that due to the onset of a cross-orientation stimulus. As with the equivalent experiment performed in V1 ([Li et al., 2006](#); [Smith et al., 2006](#)), suppression is evident from the start of the response (and reaches 5% of its maximum value after 3.1 iterations). The rapid onset of suppression is due to the saturation in the model LGN, since when the V1 model is replaced with a linear model (bottom inset), in which the activity represents the feedforward stimulation, it can be seen that there is rapid suppression. However, it can also be seen that the inhibition in the V1 model also gives rise to suppression, since when the LGN model is replaced by one that does not saturate (top inset), the mask stimulus produces a suppressed response compared to the preferred stimulus in isolation, however, this intracortical suppression is slightly further delayed after stimulus onset (reaching 5% of its maximum value after 3.3 iterations).

Figure 5b shows the time course of the response of a model neuron to the onset of the preferred stimulus within a perpendicular surround in comparison with the response due to the onset of the preferred stimulus within

¹The dimensions of the stimuli used in these simulations were defined using the same technique used to determine the dimensions of the stimuli used in the corresponding physiological experiments ([Bair et al., 2003](#)). The diameter of the central grating was made equal to the diameter of the smallest, high-contrast, optimally-oriented circular grating that generated 95% of the peak response from the recorded neuron. This diameter was 11 pixels (see [Spratling, 2010](#), Fig. 5b). The inner diameter of the surround was made equal to the smallest inner diameter of an optimally oriented annular grating that produced a response indistinguishable from zero. This diameter was 15 pixels (see [Spratling, 2010](#), Fig. 5b). There was thus a small uniformly gray gap, of width 2 pixels, between the centre and surround. However, the region of the image from which a prediction neuron receives connections with non-zero synaptic weights has a diameter of 21 pixels. The weak inputs at the edge of the recorded prediction neuron's RF do not generate a response when the surround stimulus is presented in the absence of the central stimulus, as in these conditions prediction neurons other than the recorded one are much more strongly activated by the annular stimulus and suppress any response from the recorded neuron. However, in the presence of a central grating, the weak peripheral inputs to the recorded prediction neuron can have an effect on its response. These peripheral inputs are more strongly stimulated by an iso-oriented, rather than a perpendicular, surround.

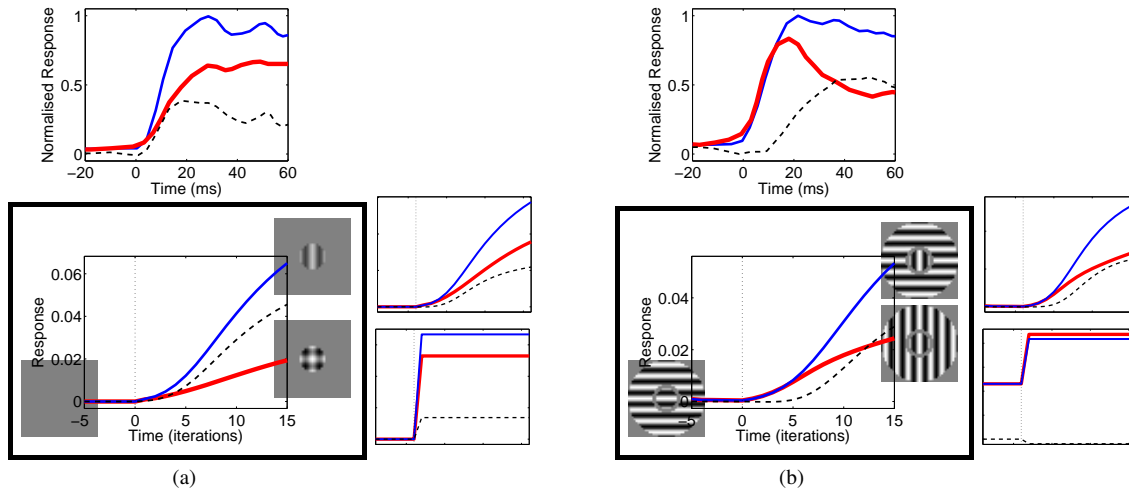


Figure 5: The response of a model prediction neuron to the onset of its preferred stimulus in the presence and absence of a second suppressive stimulus. For (a) the suppressive stimulus is a superimposed orthogonal mask, for (b) the suppressive stimulus is an iso-oriented surround. The icons at the beginning and end of each trace illustrate the stimuli that gave rise to that response: thin blue lines show response to the onset of the preferred stimulus, thick red lines show response to the onset of preferred stimulus together with the suppressive stimulus. The dashed lines represent the difference between the two responses. Time 0 (marked by the dotted vertical line) was the last iteration in which the initial stimulus was presented to the model, hence, time 1 was the first iteration after the transition. The plots above each simulation result show data from corresponding neurophysiological experiments. The neurophysiological data shows the average response of a population of complex cells in anaesthetised primate V1 (adapted from [Smith et al., 2006](#), Fig. 4). Similar neurophysiological data to that shown in (a) was also presented in [Li et al. \(2006, Fig. 2d\)](#). The smaller plots to the right of each simulation result show the response of the model when there is no saturation to the LGN response (top) and when there is no competition between cortical neurons (bottom).

a suppressive iso-oriented surround. As with the equivalent experiment performed in V1 ([Smith et al., 2006](#)), suppression is only evident after a delay from the start of the response (it reaches 5% of its maximum value after 5.7 iterations). This suppression is due entirely to the intracortical inhibition in the model. Removing this inhibition (bottom inset) shows that no suppression occurs (on the contrary, the iso-oriented surround provides marginally more feedforward excitation to the recorded neuron). In contrast, removing the LGN saturation from the model (top inset), has little effect on the time course of the suppression generated, however latency is reduced slightly so that the suppression reaches 5% of its maximum value after 4.1 iterations.

3.2 Gain of Contrast-Response Function

Another distinction between cross-orientation suppression and surround suppression is manifest in the contrast-response functions generated by each form of suppression ([Sengpiel et al., 1998](#)). If the mean evoked response is plotted as a function of the contrast of the recorded neuron’s preferred grating, then cross-orientation suppression (generated by a fixed contrast orthogonal mask) gives rise to a rightward shift in this contrast-response function; it produces contrast gain ([Sengpiel et al., 1998](#)). Whereas surround suppression (generated by a fixed contrast iso-oriented surround) reduces the magnitude of the response by a constant scale factor; it produces response gain ([Cavanaugh et al., 2002; Sengpiel et al., 1998](#)). Consistent with these neurophysiological observations, the PC/BC model also produces contrast gain for cross-orientation suppression and response gain for surround suppression, see Fig. 6. The simulation results presented in this section were recorded from a model simple cell using static or slowly drifting gratings. Results for a model complex cell in response to fast drifting gratings are presented in the bottom row of Fig. 7.

For cross-orientation suppression, when the test grating is in phase with the RF of the recorded neuron and is presented at high contrast, the addition of the mask reduces the feedforward drive to the recorded model V1 neuron. This can be seen when the intracortical inhibition is removed from the model and so LGN saturation operates in isolation, see the solid lines in the bottom inset to Fig. 6a. At locations where the mask decreases the contrast of the input, the LGN response is reduced. Whereas, at locations where the mask increases the

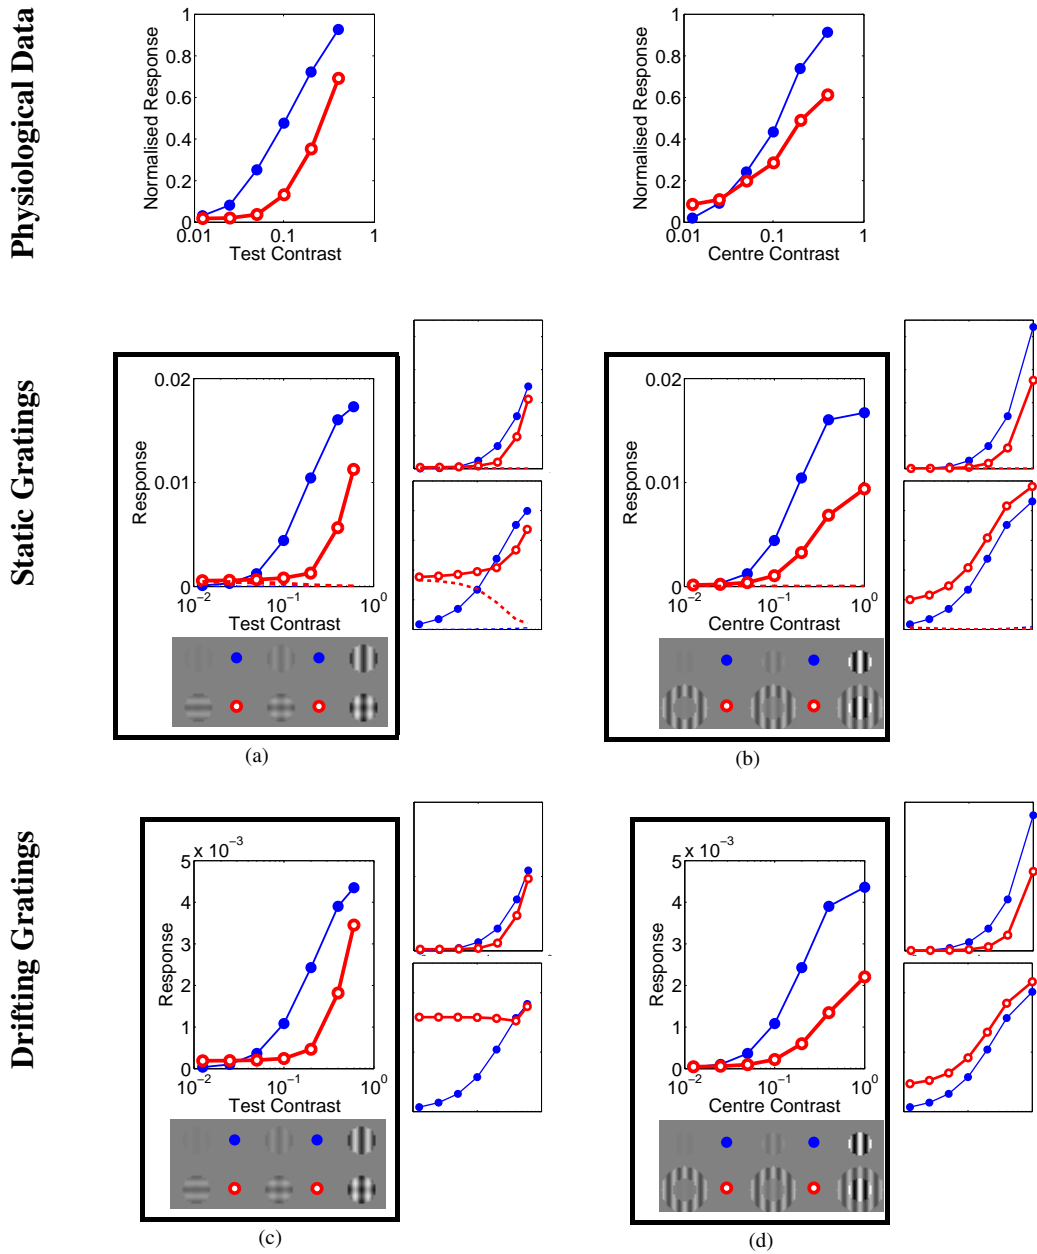


Figure 6: Contrast response functions for (left column) cross-orientation suppression, and (right column) surround suppression. In each plot solid markers show the response generated by the preferred grating presented in isolation at various contrasts. The open markers show the response for various contrasts of the preferred grating in the presence of (left) a superimposed orthogonal mask grating at 40% contrast, (right) an iso-oriented annular surround grating at 40% contrast. (a) and (b) show responses (averaged over 10 iterations) to static gratings: solid lines for when the preferred grating is in phase with the recorded neuron’s RF, dashed lines for when the preferred grating is 180° out of phase with the recorded neuron’s RF. (c) and (d) show responses (averaged over 180 iterations) to gratings drifting at $\frac{1}{90}$ cycles/iteration. The smaller plots to the right of each graph show the response of the model when there is no saturation to the LGN response (top) and when there is no competition between cortical neurons (bottom). The plots above the simulation results show data from corresponding neurophysiological experiments. For cross-orientation suppression (left column) the neurophysiological data is the average response of a population of 48 cells in anaesthetised cat V1 (adapted from Sengpiel et al., 1998, Fig. 3). For surround suppression (right column) the neurophysiological data is the average response of a population of 25 cells in anaesthetised cat V1 (adapted from Sengpiel et al., 1998, Fig. 5a).

contrast of the input, the LGN response increases. However, this increase is reduced due to saturation, and hence, does not balance the decreases. Therefore, the overall response, and hence, the feedforward drive delivered to the recorded V1 neuron, is reduced. This behaviour is consistent with the explanation for cross-orientation suppression proposed by [Priebe and Ferster \(2006\)](#). However, LGN saturation alone is insufficient to account for the full range of neurophysiological data. Firstly, when the test grating is presented at low contrast the addition of the mask increases the feedforward drive to the recorded model V1 neuron. At locations where the mask decreases the contrast of the input, the LGN response which is already very weak can only be reduced slightly until it reaches zero. Whereas, at locations where the mask increases the contrast of the input, the LGN response increases markedly. The net result is a strong increase in the feedforward drive delivered to the recorded V1 neuron. A second reason why LGN saturation does not provide a complete account of cross-orientation suppression is that when the test grating is out-of-phase with the RF of the recorded neuron, the input from the LGN cells is weak at all contrasts. The addition of the mask increases the feedforward drive to the recorded model V1 neuron (see the dashed lines in the bottom inset to [Fig. 6a](#)) for the same reason that it increases drive for a low contrast, in-phase, test grating. The increase in input from LGN caused by a mask when the test grating is out-of-phase balances any reduction in the input from LGN caused when the test grating is in-phase and at high contrast. Hence, if the average response is measured to drifting gratings (as is typically the case for neurophysiological experiments), LGN suppression has no net contribution to cross-orientation suppression (see the bottom inset to [Fig. 6c](#)). Intracortical inhibition, of the type proposed by the PC/BC model, suppresses response to the mask when the test grating is at low contrast or is out-of-phase with the RF (see top insets to [Fig. 6a](#) and [c](#)). This enables the full model to show contrast gain of the contrast-response function, as illustrated in [Fig. 6c](#).

For surround suppression, the presence of the iso-oriented surround gives rise to an approximately constant increase in feedforward drive to the recorded model V1 neuron at all contrasts. This can be seen when the intracortical inhibition is removed from the model, see bottom inset to [Fig. 6b](#). This is due to the edges of the recorded neuron's RF overlapping with the surround, and hence, receiving weak input from the surround. The suppression in the full model is thus due entirely to intracortical inhibition. When this mechanism operates in isolation (top inset to [Fig. 6b](#)), it can be seen that there is strong suppression at all contrasts. This is due to neurons that are activated by the surround competing with the recorded neuron to represent the stimulus. Performing this experiment with drifting gratings produces the same pattern of results ([Fig. 6d](#)), however, the absolute strength of the response is significantly reduced due to the drifting central grating only matching the recorded neuron's RF part of the time, and hence, producing a weaker temporally averaged response.

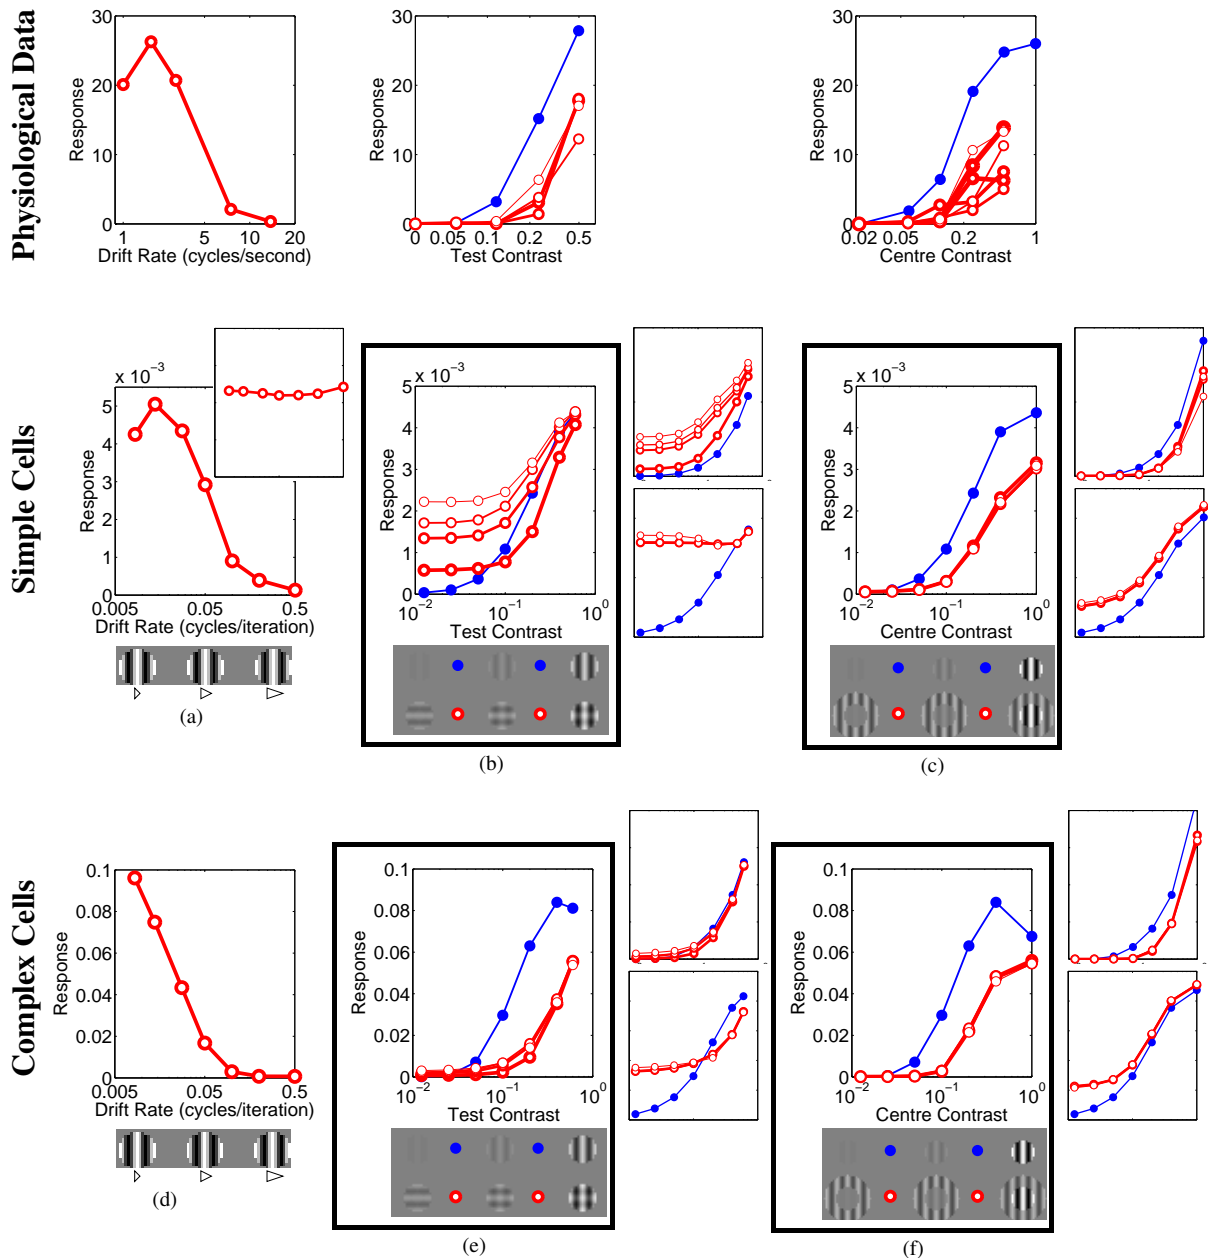
3.3 High Temporal Frequency Suppressive Stimuli

A grating presented at high temporal frequency generates a weak response in individual V1 cells ([Freeman et al., 2002](#)). However, high temporal frequency gratings produce strong suppression in experiments on both cross-orientation and surround suppression ([Durand et al., 2007](#); [Freeman et al., 2002](#)). In light of these results, it has been claimed that cross-orientation suppression can not be due to intracortical inhibition, since if it were it should disappear when the mask is presented at temporal frequencies in excess of those that produce a response in neurons tuned to the orientation of the mask ([Carandini et al., 2002](#); [Freeman et al., 2002](#)). Alternatively, it has been claimed that surround suppression could be entirely due to intracortical inhibition if those neurons generating the suppression respond to higher temporal frequencies than is typical for neurons measured in V1 ([Durand et al., 2007](#)). The current model demonstrates that neither of these claims are necessarily true since cortical inhibition could remain strong at high temporal frequencies due to it being generated by many weakly active cells rather than a few strongly active ones.

Increasing the temporal frequency of a drifting grating reduced the response of a neuron in the model in a similar manner to that observed for neurons in V1. This was true for both a model simple cell ([Fig. 7a](#)) and for a model complex cell ([Fig. 7d](#)). In the model this effect is due to a fast moving grating only matching the recorded prediction neuron's RF part of the time, and hence, producing a weaker temporally-averaged response. A fast moving grating also activates many other prediction neurons (since the stimulus matches different neuron's RFs at different times), and hence, there is increased competition further suppressing the recorded neuron's response. In effect the response to the stimulus become distributed across many prediction neurons and the total activity across the local population of prediction neurons remained approximately constant with temporal frequency (inset to [Fig. 7a](#)).

Strong suppression resulting from high temporal frequency gratings has been demonstrated in V1 by measuring contrast response functions, like those presented in [Section 3.2](#), using suppressive gratings drifting at various temporal frequencies ([Durand et al., 2007](#); [Freeman et al., 2002](#)). Simulations of these experiments are shown in [Fig. 7](#).

For a model simple cell, cross-orientation suppression becomes gradually less effective as the temporal fre-



quency of the mask increases (Fig. 7b). The failure of suppression is due to the strong response of the recorded prediction neuron to the mask when the preferred grating is out-of-phase with the neuron's RF. At low mask frequencies this response is suppressed by intracortical inhibition (see Fig. 6c). However, intracortical inhibition takes several iterations to become significant (see section 3.1), and hence, at high mask frequencies intracortical suppression does not have time to act. In contrast, for a model complex cell the preferred grating is never out-of-phase with the neuron's RF, and hence, cross-orientation suppression remains strong at all mask temporal frequencies (Fig. 7e). Essentially, the model complex cell responds only to model simple cells for which the preferred grating is in-phase. For such cells, LGN response saturation is effective at reducing the feedforward activation received by the simple cell at high contrast (bottom inset to Fig. 6a), and in turn, the complex cell (bottom inset to Fig. 7e). At low contrasts, the intracortical inhibition received from the local population of prediction neurons is sufficiently strong to suppress activation to the mask.

For a model simple cell, surround suppression with a high temporal frequency surround (Fig. 7c) is weaker than for a low temporal frequency surround (Fig. 6d). For the result presented in Fig. 6d the centre and surround drifted at the same rate, and hence, were always in-phase. For the result presented in Fig. 7c the surround drifted at a different rate than the centre and was, hence, out-of-phase with the centre part of the time. Out-of-phase surrounds generate much weaker suppression in the model (Spratling, 2010, Fig. 10f) and in cortex (Xu et al.,

Figure 7: (previous page) Effects of stimulus temporal frequency. Left column shows response as a function of grating temporal frequency. The inset to (a) shows the response summed over all prediction neurons within ± 5 pixels of the neuron recorded in the main figure. Middle column shows contrast response functions for cross-orientation suppression with a superimposed orthogonal mask grating at 40% contrast. Right column shows contrast response functions for surround suppression with a iso-oriented annular surround grating at 50% contrast. (a), (b) and (c) show responses for model simple cells. (d), (e) and (f) show responses for model complex cells. In the middle and right columns, solid markers show the response generated by the preferred grating presented in isolation at various contrasts. The open markers show the response for various contrasts of the preferred grating in the presence of the suppressive grating. The thickness of each line corresponds to the temporal frequency of the suppressive grating: 0.05, (thickest) 0.1, 0.2, 0.5 (thinnest) cycles/iteration in the simulations. In several plots the curves overlap for suppressive gratings at different frequencies. The preferred grating was shown at a temporal frequency of $\frac{1}{90}$ cycles/iteration and responses were averaged over 180 iterations. The smaller plots to the right of each graph in (b), (c), (e), and (f) show the response of the model when there is no saturation to the LGN response (top) and when there is no competition between cortical neurons (bottom). The plots above the simulation results show data from corresponding neurophysiological experiments. For temporal frequency tuning (left column) the neurophysiological data is for a typical cell in anaesthetised cat V1 (adapted from [Freeman et al., 2002](#), Fig. 3c). For cross-orientation suppression (middle column) the neurophysiological data is for a typical cell in anaesthetised cat V1 (adapted from [Freeman et al., 2002](#), Fig. 4). Suppression was measured for masks with temporal frequencies of 3Hz, 6Hz, 12Hz, and 24Hz. For surround suppression (right column) the neurophysiological data is for a typical simple cell in anaesthetised cat V1 (adapted from [Durand et al., 2007](#), Fig. 3). Suppression was measured for masks with temporal frequencies of 1Hz, 2Hz, 4Hz, 8Hz, 16Hz, and 21Hz. In both cases the thickness of the line, with open markers, used to draw the suppressed contrast response function decreases with increasing frequency of the suppressive grating.

2005), and this gives rise to the apparent reduction in suppression with high frequency surrounds. However, surround suppression is still strong, this is due to strong suppression when the surround and centre are in-phase and due to some suppression being generated when the surround and centre are out-of-phase. This latter effect results from the high temporal frequency surround weakly activating many prediction neurons which can provide a similar degree of suppression as the few neurons strongly activated by a slowly drifting surround (see inset to Fig. 7a). For a model complex cell the effect of surround suppression is similar to that observed for model simple cells, since the complex cell is just taking the maximum response from a local population of simple cells tuned to the same orientation.

A limitation of the model is that, surround suppression in simple cells, and both forms of suppression in complex cells, fail to show any reduction with increasing temporal frequency. Such a reduction in suppression at very high temporal frequencies is observed in V1 ([Durand et al., 2007](#); [Freeman et al., 2002](#)) and may be due to a reduction in LGN response at high frequency, which has not been modelled here.

3.4 Suppression of Suppressive Stimuli

Psychophysical correlates of surround suppression and cross-orientation suppression have been observed in humans ([Petrov et al., 2005](#)). This work also revealed differences between surround suppression in foveal and peripheral vision. Here we consider only those experiments performed using stimuli presented in the periphery, as is the case for the vast majority of neurophysiological experiments. In these experiments [Petrov et al. \(2005\)](#) showed differences in cross-orientation and surround suppression by employing a second grating designed to suppress the response to another suppressive grating but not the target grating.

In one experiment, a fixed contrast central grating was presented together with an iso-oriented surround with the same fixed contrast. The strength of suppression was then measured for variable contrasts of an orthogonal surround grating additively superimposed on the iso-oriented surround. It was found that increasing the contrast of the orthogonal surround reduced suppression of the central target ([Petrov et al., 2005](#)). This suggests that iso-oriented surround suppression becomes less effective in the presence of a perpendicular surround. [Petrov et al. \(2005\)](#) interpreted this result as showing that cross-orientation suppression occurred in the surround before surround suppression could take place. A similar result is generated by the PC/BC model, see Fig. 8a, and consistent neurophysiological data has been obtained in cat V1 ([Walker et al., 2002](#)). In the model, the high contrast orthogonal surround suppressed (via cross-orientation suppression) the response to the iso-oriented surround which in turn reduced the surround suppression generated by the iso-oriented surround.




Figure 8: Effects of suppression on suppressive stimuli. (a) Response as a function of the contrast of an orthogonal surround grating superimposed upon an iso-oriented surround grating in the presence of an optimally oriented centre. The contrast of the centre and the iso-oriented surround were fixed at 20%. The plot above this simulation result shows data from a corresponding neurophysiological experiment. The neurophysiological data shows the response of a typical simple cell in anaesthetised cat V1 (adapted from Walker et al., 2002, Fig. 2). (b) Response as a function of the contrast of a perpendicular mask grating superimposed upon an optimally oriented grating in the presence of an orthogonal surround. The contrast of the optimally-oriented centre and surround were fixed at 20%. (c) Response as a function of the contrast of an orthogonal surround grating to a fixed contrast central plaid stimulus. The contrast of the optimally-oriented centre and superimposed orthogonal mask were fixed at 20%. (d) Response as a function of the contrast of an iso-oriented surround grating superimposed upon an orthogonal surround grating in the presence of an optimally oriented centre. The contrast of the centre and the orthogonal surround were fixed at 20%. In each plot, the horizontal line indicates the response to the central grating in isolation. The smaller plots to the right of each graph show the response of the model when there is no saturation to the LGN response (top) and when there is no competition between cortical neurons (bottom).

In a second experiment, a fixed contrast optimally-oriented central grating was presented together with an orthogonal surround with the same fixed contrast. The strength of suppression was then measured for variable contrasts of a perpendicular mask grating additively superimposed on the central test grating. It was found that increasing the contrast of the mask grating resulted in stronger suppression of the target, *i.e.*, greater cross-orientation suppression (Petrov et al., 2005). This suggests that the surround, which was oriented to generate the maximum suppression in the mask, was ineffective. Petrov et al. (2005) interpreted this result as showing that cross-orientation suppression occurred in the centre before surround suppression could take place. A similar result is generated by the PC/BC model, see Fig. 8b. In the model, a high contrast mask reduces, via LGN saturation, the feedforward input to the recorded neuron in the model V1. Furthermore, at high mask contrasts the intracortical inhibition generated by the mask becomes more effective.

Note that the second experiment performed by Petrov et al. (2005) is not exactly the converse of the first. In the first experiment the contrast is varied for the grating that should reduce suppression, while in the second experiment the grating that should reduce suppression has a fixed contrast. The model was therefore tested using two additional experiments which were the converse of the two previous experiments.

In the first additional experiment, a fixed contrast optimally-oriented central grating was presented together with a superimposed perpendicular mask grating with the same fixed contrast. The strength of suppression was then measured for variable contrasts of an orthogonal surround grating. It was found that increasing the contrast of the orthogonal surround resulted in slightly stronger suppression of the target, see Fig. 8c. In the model, a high contrast orthogonal surround provides a small additional feedforward input to neurons representing the orientation of the mask. This enables the neurons representing the mask to more effectively suppress, via intracortical inhibition, the response of the recorded neuron.

In the second additional experiment, a fixed contrast optimally-oriented central grating was presented together with an orthogonal surround with the same fixed contrast. The strength of suppression was then measured for variable contrasts of an iso-oriented surround grating. It was found that increasing the contrast of the iso-oriented surround resulted in stronger suppression of the target, see Fig. 8d. In the model, a high contrast iso-oriented surround suppresses, via intracortical inhibition, the response of the recorded neuron.

The simulations presented in this section were recorded from model simple cells using static stimuli. However, very similar results are generated by model complex cells in response to drifting gratings.

4 Discussion

The results presented above together with those presented in a previous publication (Spratling, 2010) show that the PC/BC model is able to accurately simulate many aspects of cross-orientation and surround suppression, as well as other tuning properties of cells in primary visual cortex. PC/BC thus provides a comprehensive model of the behaviour (*i.e.*, the neural response properties) of cells in cortical area V1.

All models simulate behaviour. However, models (of V1 or any other kind of system) can be categorised depending on what other aspects of the modelled system they do or do not account for (Table 2; Dayan, 2001; Seriès et al., 2003). Descriptive models characterise the behaviour of the system being modelled. They provide a summary of “what” the system does. However, these models usually fail to provide any information about the mechanisms by which that system operates or the purpose of the computations it is performing. Structural models consider how the underlying mechanisms (*e.g.*, biophysical and neural processes) give rise to the observed behaviour. They address “how” the system operates, however, these models usually fail to provide any computational explanation for the modelled behaviour. Structural models provide a reductionist explanation of the observed behaviour, however, the models of the underlying mechanisms may themselves be descriptive models of lower-level processes. Functional models start from a consideration of what computation the system must be performing. They address “why” the system operates in the way it does.

In neuroscience there is a need for functional models (Carandini et al., 2005; Olshausen and Field, 2005). PC/BC represents such a model. It proposes that V1 response properties are a result of the cortex performing efficient coding using an over-complete set of neural representations (Olshausen and Field, 2005). This is achieved by minimising the error between the observed sensory input and the expectations stored in the synaptic weights of V1 cells. The model can be related to the biased-competition and predictive coding theories of cortical function (Spratling, 2008a,b), and in turn to more general theories of hierarchical perceptual inference. It thus provides a functional explanation for V1 response properties and relates V1 neurophysiology to a wider, computationally principled, framework for understanding cortical function.

Functional models do not necessarily consider the mechanisms used by the system being modelled. For example, much of the field of artificial intelligence (AI) is concerned with developing computationally principled models of human intelligence without being concerned with the mechanisms used in the human brain. However,

Type of Model	Characteristics Modelled			Examples
	computations	behaviours	mechanisms	
descriptive	✗	✓	✗	Busse et al. (2009); Cavanaugh et al. (2002); Sceniak et al. (1999)
structural	✗	✓	✓	Adorján et al. (1999); Carandini et al. (2002); Dagoi and Sur (2000); Priebe and Ferster (2006); Schwabe et al. (2006); Somers et al. (1995); Stetter et al. (2000)
functional	✓	✓	✗	Ben-Shahar and Zucker (2004); Olshausen and Field (1996); Schwartz and Simoncelli (2001)

Table 2: Types of model. Models can be classified according to the level(s) of analysis at which they successfully described the system being modelled. Some examples are listed of models of V1 which fall into each of these categories.

in contrast to AI, neuroscience is concerned with how the brain performs the information processing that underlies cognition, not with any other possible ways in which this could be done. Neuroscientists must, therefore, attempt to synthesise theories that are consistent across levels of analysis (Bechtel, 2006; Mareschal et al., 2007). Hence, functional models that simulate neural behaviour in a biologically plausible manner are most relevant to neuroscience.

Many aspect of PC/BC are biologically plausible. For example, the algorithm employs only non-negative firing rates and non-negative synaptic weights. It is also possible to define unsupervised learning rules that require only information local to each synapse in order to independently learn the reciprocal feedforward and feedback weights required by the model (Spratling, sub). The drive to the error-detecting neurons is divisively modulated. Such divisive modulation is compatible with mechanisms of shunting inhibition or synaptic depression (Carandini, 2004; Mitchell and Silver, 2003; Rothman et al., 2009). In the full PC/BC model (Spratling, 2008a, 2010), cortical feedback connections have a modulatory effect on activity, whereas, cortical feedforward connections drive neural response. This is consistent with the asymmetry in the functional roles of cortical forward and feedback connections (Anderson and Martin, 2009; Crick and Koch, 1998; Friston, 2003; Friston and Büchel, 2000; Sherman and Guillery, 1998; Spratling, 2002).

The behaviour of model prediction neurons is consistent with the physiology of cortical pyramidal cells in V1. The connectivity of these neurons in the model is also consistent with the anatomy of cortical feedforward and feedback connections (Barbas and Rempel-Clower, 1997; Barone et al., 2000; Felleman and Van Essen, 1991; Johnson and Burkhalter, 1997). Specifically, cortical feedforward connections originate from pyramidal cells in layers II and III and feedback connections terminate outside layer IV. This suggests that the prediction neurons correspond to pyramidal cells in the superficial layers of cortex. Similarly, the error-detecting neurons in the model are driven by feedforward connections from LGN, and hence, have centre-surround RFs. Non-orientation selective centre-surround cells are common in cortical layer IV of V1 in some species (although rare in others) (Van Hooser, 2007). The physiology would thus suggest that error-detecting neurons correspond to a sub-population of cells in cortical layer IV. This is also consistent with cortical anatomy, since cortical feedforward connections predominantly target layer IV.

Other aspects of the model are less easily reconciled with cortical anatomy and physiology. For example, a mechanism that would allow the inputs to the predictions neurons to multiplicatively modulate the firing rates of those neurons (equation 4) is unknown. The model also proposes a one-to-one connectivity pattern between the error-detecting neurons and the neurons from which they receive their afferent inputs. This is clearly implausible. Also, the model proposes that there should be an asymmetry in the targets for lateral inhibitory and excitatory connections, with the former targeting error-detecting neurons only, and the latter targeting prediction neurons only. In contrast, cortical neurons receive both intracortical inhibition and excitation. In common with many other neural models, PC/BC does not include inhibitory neurons, but rather allows excitatory neurons (*i.e.*, the prediction neurons) to directly inhibit other excitatory neurons (*i.e.*, the error-detecting neurons). In a more anatomically accurate implementation these connections from the prediction neurons to the error-detecting neurons would need to be made via inhibitory interneurons. Such connections would be consistent with cortical anatomy, and the identification of model neurons with cortical populations made above, as inhibitory pathways from superficial layers to layer 4 mirror the strong excitatory pathway for layer 4 to the superficial layers (Binzegger et al., 2004; Thomson and Lamy, 2007).

The PC/BC model depends on a specific form of intracortical inhibition. This mechanism of suppression

needs to be combined with LGN response saturation to account for cross-orientation suppression. A number of other mechanisms have previously been suggested to account for suppression: thalamocortical synaptic depression (Carandini et al., 2002; Freeman et al., 2002); surround suppression in LGN (Naito et al., 2007; Ozeki et al., 2004; Webb et al., 2005); and feedback from extrastriate cortex (Angelucci et al., 2002; Bair et al., 2003; Cavanaugh et al., 2002; Sullivan and de Sa, 2006).

A reduction in synaptic efficiency, *i.e.*, synaptic depression, caused by adaptation to the stimulus has been shown to be too slow to account for the dynamics of cross-orientation suppression (Li et al., 2006). Consistent with this result, the current model does not propose that this form of synaptic depression has a significant role in generating suppression in V1. However, depression of excitatory synapses provides one potential mechanism via which inhibitory inputs can modulate neuronal gain (Rothman et al., 2009). Hence, one potential mechanism for implementing the specific form of divisive inhibition proposed by the PC/BC model might involve synaptic depression. However, the PC/BC model suggests that the mechanism driving the suppression is the response of simple cells within V1 rather than the response of cells in LGN.

The current model employs a bank of linear filters as a simplistic model of LGN, and hence, it does not simulate surround suppression in the LGN. The current model, therefore, also does not propose that surround suppression in LGN is important for generating suppression in V1. However, suppression does occur in LGN (Jones et al., 2000) and might be responsible for V1 surround suppression caused by uniform surrounds (Webb et al., 2005), which the current model fails to simulate. LGN suppression could be modelled by using neurons with Gaussian RFs competing via Divisive Input Modulation. In which case the LGN would be modelled in the same way as the PC/BC model of V1 and would be the first stage in a PC/BC hierarchy (Spratling, 2008a).

The version of the PC/BC model implemented here does not include feedback connections from extrastriate cortical regions, and hence, does not propose that top-down predictions from other parts of the cortex have a significant role in generating suppression in the experiments simulated here. However, all the simulated neurophysiological experiments were performed on anaesthetised animals and anaesthesia may block the effects of cortical feedback (Lamme and Spekreijse, 2000; Lamme et al., 1998). The PC/BC model does predict that cortical feedback will have a role in generating suppression in other experiments. For example, suppression of the response to a non-attended stimulus in a selective attention experiment (Spratling, 2008a), or suppression of the response to any stimulus when competition is biased by top-down signals in favour of a competing interpretation. PC/BC defines the influence that cortical feedback connections have (Spratling, 2008a, 2010), and hence, these influences could easily be incorporated into the model. Such feedback driven suppression might be necessary to simulate surround-suppression caused by surrounds very far from the classical RF of the recorded neuron (Angelucci et al., 2002; Levitt and Lund, 2002; Sceniak et al., 2001). Cortical feedback connections directly exciting prediction neurons, which in turn induce suppression in other prediction neurons, might conceivably account for both suppression and facilitation generated by the far surround (Ichida et al., 2007).

The success of the PC/BC model does not exclude the possibility that other mechanisms are actually responsible for suppression in primary visual cortex. However, it does demonstrate that intracortical inhibition, driven by activity intrinsic to V1, remains a viable mechanism for explaining cortical suppression and that the same mechanisms can underlie both cross-orientation and surround suppression.

Acknowledgements

Thanks to two anonymous referees for helpful comments on an earlier draft of this article. This work was funded by the Engineering and Physical Sciences Research Council grant number EP/D062225/1.

References

- Adorján, P., Levitt, J. B., Lund, J. S., and Obermayer, K. (1999). A model for the intracortical origin of orientation preference and tuning in macaque striate cortex. *Vis. Neurosci.*, 16:303–18.
- Anderson, J. C. and Martin, K. A. C. (2009). The synaptic connections between cortical areas V1 and V2 in macaque monkey. *J. Neurosci.*, 29(36):11283–93.
- Angelucci, A., Levitt, J., Walton, E. J. S., Hupé, J.-M., Bullier, J., and Lund, J. S. (2002). Circuits for local and global signal integration in primary visual cortex. *J. Neurosci.*, 22(19):8633–46.
- Bair, W., Cavanaugh, J. R., and Movshon, J. A. (2003). Time course and time-distance relationships for surround suppression in macaque V1 neurons. *J. Neurosci.*, 23(20):7690–701.
- Bair, W., Cavanaugh, J. R., Smith, M. A., and Movshon, J. A. (2002). The timing of response onset and offset in macaque visual neurons. *J. Neurosci.*, 22(8):3189–205.

- Barbas, H. and Rempel-Clower, N. (1997). Cortical structure predicts the pattern of corticocortical connections. *Cerebral Cortex*, 7:635–46.
- Barone, P., Batardiere, A., Knoblauch, K., and Kennedy, H. (2000). Laminar distribution of neurons in extrastriate areas projecting to visual areas V1 and V4 correlates with the hierarchical rank and indicates the operation of a distance rule. *J. Neurosci.*, 20:3263–81.
- Bechtel, W. (2006). Reducing psychology while maintaining its autonomy via mechanistic explanation. In Schouten, M. and de Jong, H. L., editors, *The Matter of the Mind: Philosophical Essays on Psychology, Neuroscience and Reduction*, chapter 8. Blackwell, Oxford, UK.
- Ben-Shahar, O. and Zucker, S. W. (2004). Geometrical computations explain projection patterns of long range horizontal connections in visual cortex. *Neural Comput.*, 16(3):445–76.
- Binzegger, T., Douglas, R. J., and Martin, K. A. C. (2004). A quantitative map of the circuit of cat primary visual cortex. *J. Neurosci.*, 24(39):8441–53.
- Bonds, A. B. (1989). Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Vis. Neurosci.*, 2:41–55.
- Busse, L., Wade, A. R., and Carandini, M. (2009). Representation of concurrent stimuli by population activity in visual cortex. *Neuron*, 64(6):931–42.
- Carandini, M. (2004). Receptive fields and suppressive fields in the early visual system. In Gazzaniga, M. S., editor, *The Cognitive Neurosciences, 3rd edition*, pages 313–26. MIT Press, Cambridge, MA.
- Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., Gallant, J. L., and Rust, N. C. (2005). Do We Know What the Early Visual System Does? *J. Neurosci.*, 25(46):10577–97.
- Carandini, M., Heeger, D. J., and Senn, W. (2002). A synaptic explanation of suppression in visual cortex. *J. Neurosci.*, 22(22):10053–65.
- Cavanaugh, J. R., Bair, W., and Movshon, J. A. (2002). Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. *J. Neurophysiol.*, 88(5):2530–46.
- Crick, F. and Koch, C. (1998). Constraints on cortical and thalamic projections: the no-strong-loops hypothesis. *Nature*, 391:245–50.
- Dagoi, V. and Sur, M. (2000). Dynamic properties of recurrent inhibition in primary visual cortex: contrast and orientation dependence of contextual effects. *J. Neurophysiol.*, 83:1019–30.
- Daugman, J. G. (1980). Two-dimensional spectral analysis of cortical receptive field profiles. *Vision Res.*, 20:847–56.
- Daugman, J. G. (1988). Complete discrete 2-D Gabor transformations by neural networks for image analysis and compression. *IEEE Trans. Acoust.*, 36(7):1169–79.
- Dayan, P. (2001). Levels of analysis in neural modelling. In *Encyclopedia of Cognitive Science*. MacMillan Press, London, UK.
- DeAngelis, G. C., Freeman, R. D., and Ohzawa, I. (1994). Length and width tuning of neurons in the cat’s primary visual cortex. *J. Neurophysiol.*, 71(1):347–74.
- DeAngelis, G. C., Robson, J. G., Ohzawa, I., and Freeman, R. D. (1992). Organization of suppression in receptive fields of neurons in cat visual cortex. *J. Neurophysiol.*, 68(1):144–63.
- Durand, S., Freeman, T. C. B., and Carandini, M. (2007). Temporal properties of surround suppression in cat primary visual cortex. *Vis. Neurosci.*, 24:679–90.
- Felleman, D. J. and Van Essen, D. C. (1991). Distributed hierarchical processing in primate cerebral cortex. *Cerebral Cortex*, 1:1–47.
- Fitzpatrick, D. (2000). Seeing beyond the receptive field in primary visual cortex. *Curr. Opin. Neurobiol.*, 10:438–43.
- Freeman, T. C. B., Durand, S., Kiper, D. C., and Carandini, M. (2002). Suppression without inhibition in visual cortex. *Neuron*, 35(4):759–71.
- Friston, K. J. (2003). Learning and inference in the brain. *Neural Netw.*, 16(9):1325–52.
- Friston, K. J. and Büchel, C. (2000). Attentional modulations of effective connectivity from V2 to V5/MT in humans. *Proc. Natl. Acad. Sci. U.S.A.*, 97(13):7591–6.
- Gawne, T. J. and Martin, J. M. (2002). Responses of primate visual cortical V4 neurons to simultaneously presented stimuli. *J. Neurophysiol.*, 88:1128–35.
- Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *J. Physiol. (Lond.)*, 160:106–54.
- Ichida, J. M., Schwabe, L., Bressloff, P. C., and Angelucci, A. (2007). Response facilitation from the ‘suppressive’ receptive field surround of macaque V1 neurons. *J. Neurophysiol.*, 98:2168–81.
- Johnson, R. R. and Burkhalter, A. (1997). A polysynaptic feedback circuit in rat visual cortex. *J. Neurosci.*, 17(18):7129–40.
- Jones, H. E., Andolina, I. M., Oakely, N. M., Murphy, P. C., and Sillito, A. M. (2000). Spatial summation in lateral

- geniculate nucleus and visual cortex. *Exp. Brain Res.*, 135(2):279–84.
- Jones, J. P. and Palmer, L. A. (1987). An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *J. Neurophysiol.*, 58(6):1233–58.
- Knierim, J. J. and van Essen, D. C. (1992). Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J. Neurophysiol.*, 67(4):961–980.
- Kouh, M. and Poggio, T. (2008). A canonical neural circuit for cortical nonlinear operations. *Neural Comput.*, 20(6):1427–51.
- Lamme, V. A. F. and Spekreijse, H. (2000). Modulations of primary visual cortex activity representing attentive and conscious scene perception. *Frontiers in Bioscience*, 5:232–43.
- Lamme, V. A. F., Zipser, K., and Spekreijse, H. (1998). Figure-ground activity in primary visual cortex is suppressed by anesthesia. *Proc. Natl. Acad. Sci. U.S.A.*, 96(6):3263–8.
- Lampl, I., Ferster, D., Poggio, T., and Riesenhuber, M. (2004). Intracellular measurements of spatial integration and the MAX operation in complex cells of the cat primary visual cortex. *J. Neurophysiol.*, 92:2704–13.
- Lee, T. S. (1996). Image representation using 2D Gabor wavelets. *IEEE Trans. Pattern Anal. Mach. Intell.*, 18(10):959–71.
- Levitt, J. B. and Lund, J. S. (2002). The spatial extent over which neurons in macaque striate cortex pool visual signals. *Vis. Neurosci.*, 19:439–52.
- Li, B., Thompson, J. K., Duong, T., Peterson, M. R., and Freeman, R. D. (2006). Origins of cross-orientation suppression in the visual cortex. *J. Neurophysiol.*, 96:1755–64.
- Marcelja, S. (1980). Mathematical description of the responses of simple cortical cells. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.*, 70:1297–1300.
- Mareschal, D., Johnson, M. H., Siros, S., Spratling, M. W., Thomas, M. S. C., and Westermann, G. (2007). *Neuroconstructivism: How the Brain Constructs Cognition*. Oxford University Press, Oxford, UK.
- Mitchell, S. J. and Silver, R. A. (2003). Shunting inhibition modulates neuronal gain during synaptic excitation. *Neuron*, 38(3):433–45.
- Morrone, M. C., Burr, D. C., and Maffei, L. (1982). Functional implications of cross-orientation inhibition of cortical visual cells. i. neurophysiological evidence. *Proc. R. Soc. Lond., B, Biol. Sci.*, 216(1204):335–54.
- Naito, T., Sadakane, O., Okamoto, M., and Sato, H. (2007). Orientation tuning of surround suppression in lateral geniculate nucleus and primary visual cortex of cat. *Neuroscience*, 149(4):962–75.
- Olshausen, B. A. and Field, D. J. (1996). Natural image statistics and efficient coding. *Network*, 7(2):333–9.
- Olshausen, B. A. and Field, D. J. (2005). How close are we to understanding V1? *Neural Comput.*, 17:1665–99.
- Ozeki, H., Sadakane, O., Akasaki, T., Naito, T., Shimegi, S., and Sato, H. (2004). Relationship between excitation and inhibition underlying size tuning and contextual response modulation in the cat primary visual cortex. *J. Neurosci.*, 24:1428–38.
- Petrov, Y., Carandini, M., and McKee, S. P. (2005). Two distinct mechanisms of suppression in human vision. *J. Neurosci.*, 25:8704–7.
- Priebe, N. J. and Ferster, D. (2006). The mechanism underlying cross-orientation suppression in cat visual cortex. *Nat. Neurosci.*, 9(4):552–61.
- Reid, R. C. and Alonso, J. M. (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378:281–4.
- Riesenhuber, M. and Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nat. Neurosci.*, 2(11):1019–25.
- Rothman, J., Cathala, L., Steuber, V., and Silver, R. A. (2009). Synaptic depression enables neuronal gain control. *Nature*, 457:1015–8.
- Sceniak, M. P., Hawken, M. J., and Shapley, R. M. (2001). Visual spatial characterization of macaque V1 neurons. *J. Neurophysiol.*, 85:1873–87.
- Sceniak, M. P., Ringach, D. L., Hawken, M. J., and Shapley, R. (1999). Contrast’s effect on spatial summation by macaque V1 neurons. *Nat. Neurosci.*, 2(8).
- Schwabe, L., Obermayer, K., Angelucci, A., and Bressloff, P. C. (2006). The role of feedback in shaping the extra-classical receptive field of cortical neurons: A recurrent network model. *J. Neurosci.*, 26(36):9117–29.
- Schwartz, O. and Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nat. Neurosci.*, 4:819–25.
- Sengpiel, F., Baddeley, R. J., Freeman, T. C. B., Harrad, R., and Blakemore, C. (1998). Different mechanisms underlie three inhibitory phenomena in cat area 17. *Vision Res.*, 38(14):2067–80.
- Seriès, P., Lorceau, J., and Frégnac, Y. (2003). The ‘silent’ surround of V1 receptive fields: theory and experiments. *J. Physiol. Paris*, 97(4-6):453–74.
- Sherman, S. M. and Guillery, R. W. (1998). On the actions that one nerve cell can have on another: distinguishing “drivers” from “modulators”. *Proc. Natl. Acad. Sci. U.S.A.*, 95:7121–6.

- Smith, M. A., Bair, W., and Movshon, J. A. (2006). Dynamics of suppression in macaque primary visual cortex. *J. Neurosci.*, 26(18):4826–34.
- Somers, D. C., Nelson, S. B., and Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *J. Neurosci.*, 15:5448–65.
- Spratling, M. W. (2002). Cortical region interactions and the functional role of apical dendrites. *Behav. Cogn. Neurosci. Rev.*, 1(3):219–28.
- Spratling, M. W. (2008a). Predictive coding as a model of biased competition in visual selective attention. *Vision Res.*, 48(12):1391–408.
- Spratling, M. W. (2008b). Reconciling predictive coding and biased competition models of cortical function. *Front. Comput. Neurosci.*, 2(4):1–8.
- Spratling, M. W. (2010). Predictive coding as a model of response properties in cortical area V1. *J. Neurosci.*, 30(9):3531–43.
- Spratling, M. W. (sub.). Unsupervised learning of generative and discriminative weights encoding elementary image components in a predictive coding model of cortical function. *submitted*.
- Spratling, M. W., De Meyer, K., and Kompass, R. (2009). Unsupervised learning of overlapping image components using divisive input modulation. *Comput. Intell. Neurosci.*, 2009(381457):1–19.
- Stetter, M., Bartsch, H., and Obermayer, K. (2000). A mean-field model for orientation tuning, contrast saturation, and contextual effects in the primary visual cortex. *Biol. Cybern.*, 82(4):291–304.
- Sullivan, T. J. and de Sa, V. R. (2006). A model of surround suppression through cortical feedback. *Neural Networks*, 19(5):564–72.
- Thomson, A. M. and Lamy, C. (2007). functional maps of neocortical local circuitry. *Frontiers in Neuroscience*, 5(0):12.
- Van Hooser, S. D. (2007). Similarity and diversity in visual cortex: Is there a unifying theory of cortical computation? *Neuroscientist*, 13:639–56.
- Walker, G. A., Ohzawa, I., and Freeman, R. D. (2002). Disinhibition outside receptive fields in the visual cortex. *J. Neurosci.*, 22(13):5659–68.
- Webb, B. S., Dhruv, N. T., Solomon, S. G., Tailby, C., and Lennie, P. (2005). Early and late mechanisms of surround suppression in striate cortex of macaque. *J. Neurosci.*, 25:11666–75.
- Xu, W.-F., Shen, Z.-M., and Li, C.-Y. (2005). Spatial phase sensitivity of V1 neurons in alert monkey. *Cereb. Cortex*, 15:1697–702.