Ecovision – Work Package 6 – Deliverable 6.2

Partner Bel: Laboratorium voor Neuro- en Psychofysiologie,
KU Leuven Medical School, Belgium

CNS-solution for the Interaction of Attention and Context

Karl Pauwels and Marc M. Van Hulle

1 Introduction

The model introduced in Deliverable 6.1[1] is extended in order to fully explain the experimental results obtained by Vanduffel et al. [2000]. Contrary to the model introduced earlier, this novel model is not restricted to area V1 but includes regions from the thalamus as well. In this way, novel insight is provided with respect to the interaction between contextual effects and attention in the earliest stages of visual processing, in particular the emergence of suppressed activity surrounding the center of interest (CoI).

Before explaining the model architecture and results in full detail, a short overview will be given of the experimental results obtained.

2 Experimental Findings

The experiment was concerned with the investigation of attention-dependent modulations in the early stages of the macaque visual system. A modified double-label deoxyglucose (2DG) procedure [Geesaman et al., 1997] was used to register activation levels in these areas for awake monkeys, performing a task which involved featural attention.

2.1 Task Description

The monkeys were trained to identify the orientation of a large circular square-wave grating (Fig. 1(B)). After 100 ms of fixating, the stimulus, tilted either to the left or to the right, appeared for 170 ms. Immediately afterwards, the

[1] CNS-algorithm for generating Task-Optimised Spatial Representations of Neural Activities.

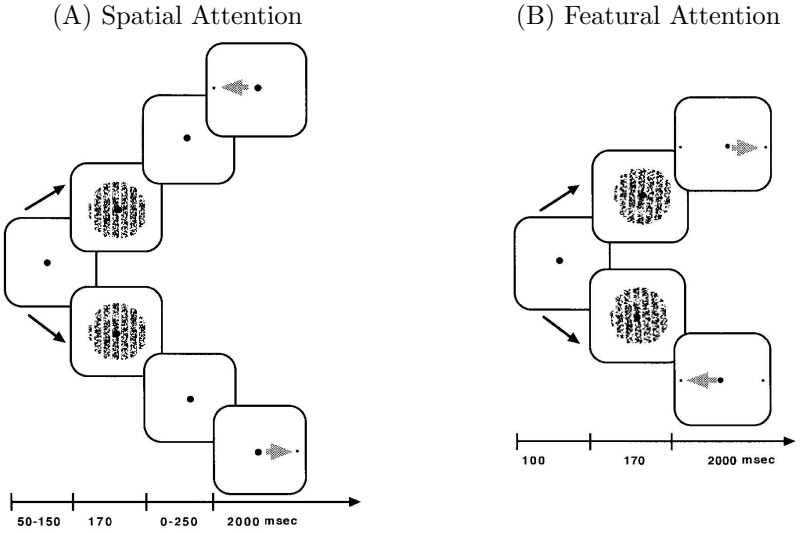

Figure 1: Task descriptions.

animal had to make a saccade, in the direction of the orientation of the grating, to one of two target points. In the spatial attention control task (Fig. 1(A)) the grating carried no behavioural significance and the monkey had to make a saccade to a single target point that appeared to the left or the right of the fixation point. For both conditions, high resolution images were recorded using 2DG.

2.2 Results

By comparing metabolic activity for the two conditions, attentional effects were observed in areas as early as the lateral geniculate nucleus (LGN) and the magnocellular-recipient layers 4Cα and 4B of the striate cortex. In these areas, attention manifests itself as a retinotopically specific band of suppressed activity, peripheral to the representation of the stimulus. Fig. 2 shows the concentration of radioactive 2DG sampled over trajectories covering different parts and different layers of striate cortex in one of the four monkeys participating in the experiment. This figure illustrates the basic finding of the study: differences in 2DG concentration were observed between the two attention conditions in certain layers and in certain parts of V1. Further analysis revealed that 1) these differences occurred mainly in the magnocellular input layer 4Cα and 2) they were restricted to an annular region surrounding the representation of the grating. No significant attention effects were observed at the location of the stimulus and outside the suppressive ring. Fig. 3 shows the distribution of the differential 2DG uptake of the two conditions in lateral geniculate nucleus (top row) and V1 layer 4Cα (bottom row) for the four different animals as a function of eccentricity, along with the radius of the grating used in each animal. This figure indicates that the ring of suppression changed in diameter with that of the grating both in LGN and V1, confirming that the suppression surrounded the stimulus representation. The strength of the suppression increased with

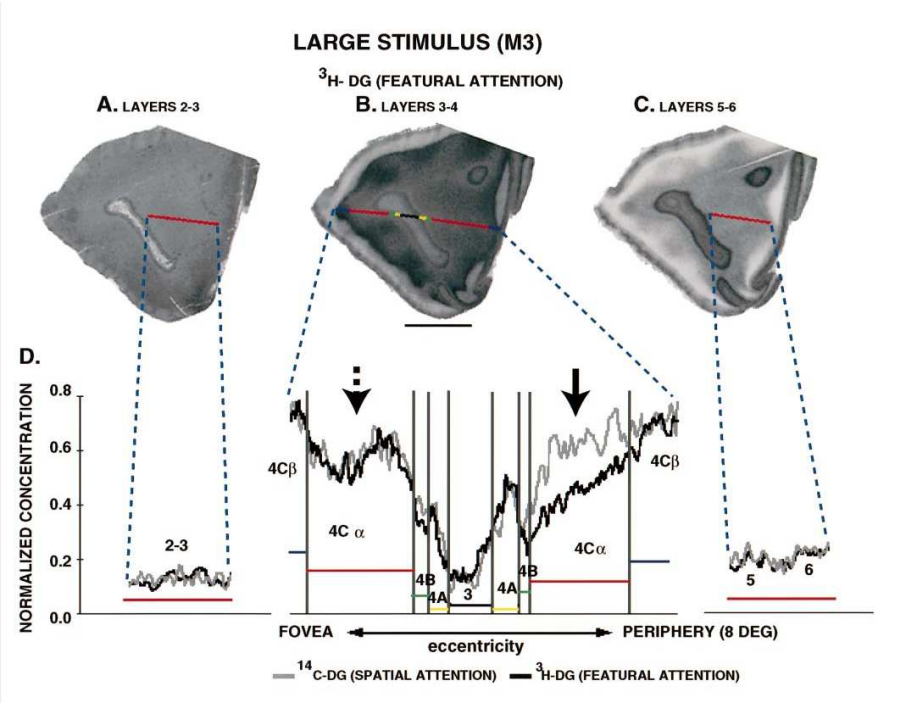

Figure 2: Plots of normalised [^{14}C]DG and [^{3}H]DG concentrations as a function of eccentricity in different layers of flattened area V1. *(A-C)* The ^{3}H signals (related to the featural-attention task, i.e. the attention-to-the-grating condition) from single V1 sections through layers 2-3, 3-4 and 5-6 respectively. In the three sections, the fovea is represented towards the left, with more peripheral visual field representations towards the right. The upper visual field is represented in the lower portion of the section. *(D)* Plots of normalised [^{3}H]DG and [^{14}C]DG concentrations as measured along the lines indicated in *(A-C)* (along the representation of the horizontal meridian from foveal to more peripheral visual field representations). Note the suppressed [^{3}H]DG concentration in more peripheral (solid black arrow in *D*), but not foveal (dotted black arrow in *D*) visual field representations of layer 4Cα.

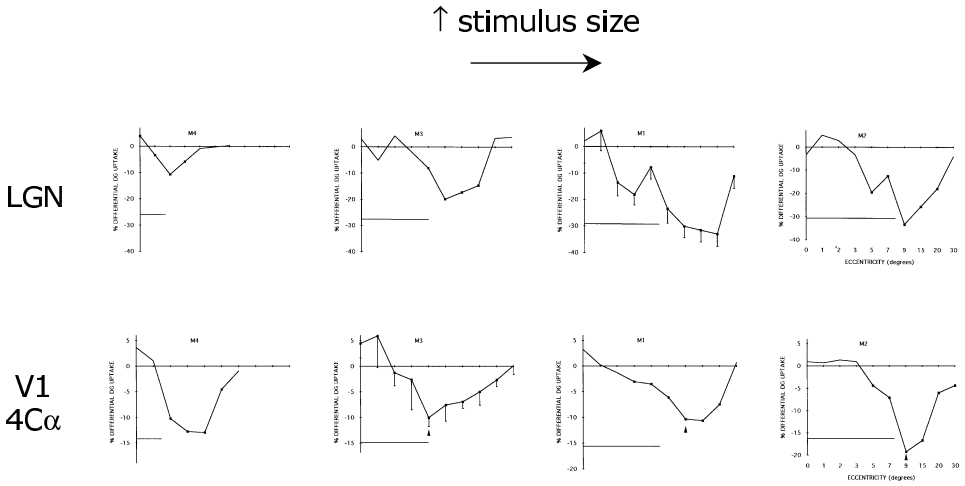

Figure 3: Normalised differential DG uptake (using the spatial attention task as baseline) in LGN (top row) and V1 4Cα (bottom row) as a function of eccentricity is plotted for different stimulus sizes. Stimulus size increases from left to right. The standard deviations are shown for monkey M1 in LGN and for monkey M3 in V1. The radius of the gratings which were presented in each experiment is indicated by the lines at the bottom of each panel.

stimulus diameter in the magnocellular layers of the LGN. This effect was less pronounced in V1 4Cα. In addition to these results, another subcortical change was observed in the visual thalamus: metabolic activity of the reticular thalamic nucleus (RTN) increased in the attention-to-the-grating condition relative to the attention-away condition. This increase strengthened with increasing stimulus size [Vanduffel et al., 2000]. Due to the absence of precise data on the retinotopy of RTN, the location of this increased activity remains unclear.

These results are indicative of an early selection/filtering gating mechanism. By suppressing irrelevant visual information outside the focus of attention, an increased signal-to-noise ratio is obtained for the processing of the attended feature in (less retinotopically-organised) extrastriate areas.

3 Computational Modelling

In the following subsections the proposed corticothalamic model for attentional modulation in early-visual processing is explained in full detail. First, a number of neurophysiological constraints are explained that helped shape the context in which the model was developed. Next, two existing models which focus on attention effects in early visual processing are described and their limitations are exposed. The neurons, architecture, and training procedure of the proposed model are explained subsequently and the results obtained are compared to the experimental results discussed in Section 2.2. Finally, the underlying mechanism, responsible for the model's functionality, is clarified.

3.1 Neurophysiological Constraints

A recent overview of the synaptic circuitry of the most relevant regions that are contained in the proposed model, is given by Guillery and Sherman [2002]. Fig. 4 gives a schematic overview of these different areas and how they are related to each other. The representation of the stimulus in the retina causes excitation of LGN relay cells and interneurons, the latter cause inhibition in the relay cells. From the relay cells, the activation is simultaneously transported to V1 layer 4 and the RTN (the abbreviation TRN is used in Fig. 4). Output layer 6 of V1 generates excitatory feedback in the RTN and LGN relay cells. Finally and most importantly, RTN cells are known to provide strong inhibition in LGN relay cells.

It is well-known that all connections between LGN, RTN and V1 are topographical [Mitrofanis and Guillery, 1993]. The receptive field size in the RTN is much larger than both LGN relay cells and cells in V1. Consequently, the connections to and from the RTN have a much larger spread than those between V1 and LGN [Bourassa and Deschênes, 1995, Yen et al., 1985].

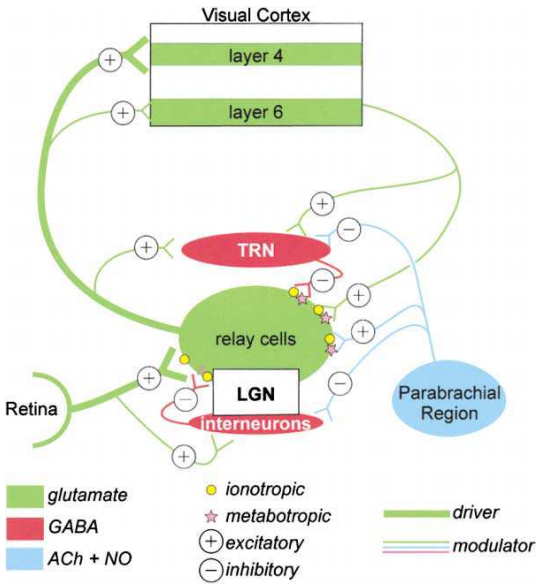

Figure 4: Schematic overview of the synaptic circuitry in cat LGN (reproduced from Guillery and Sherman [2002]).

3.2 Existing Models

Two popular models that explain some of the attention effects, observed in early visual processing, are discussed in more detail.

3.2.1 Inhibition of Interneurons

A first important model involving the RTN was introduced by Steriade et al. [1986] and centres around the inhibitory interneurons, present in the LGN. According to Steriade et al. [1986], RTN neurons inhibit the inhibitory interneurons in LGN more strongly than they inhibit relay neurons, effectively causing a disinhibition of the latter. A similar idea was also used in a model by LaBerge et al. [1992] of the pulvinar system. An overview of the circuit is shown in Fig. 5. Combined with lateral inhibition within RTN neurons this model can account for a ring of suppression surrounding the stimulus. Through excitatory connections from LGN relay cells to RTN, increased activation will be present at the stimulus representation. Lateral inhibition within RTN will cause a ring of suppression surrounding this representation and, through the inhibition of interneurons, this will faithfully project back to the LGN. In this model, attention can manifest itself as increased activation of the respective RTN neurons. This will result in all the desired effects: in the presence of a stimulus, stronger excitation at the stimulus representation and stronger suppression surrounding it, and in the absence of a stimulus increased baseline activity (and suppressive surround).

A crucial assumption of this model is the stronger influence of RTN neurons on interneurons than on relay cells. This is however contradicted by recent

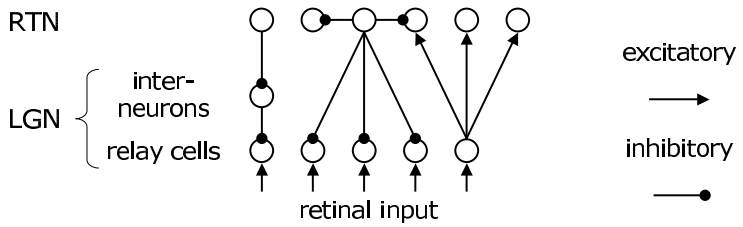

Figure 5: Model introduced by Steriade et al. [1986], which explains attentional effects through the inhibition of LGN interneurons by RTN neurons.

physiological evidence [Wang et al., 2001]. As can also be seen in Fig. 4, these connections are no longer considered significant.

3.2.2 Different Firing Modes of LGN Relay Cells

A different model was proposed by Crick [1984], which exploits the two different firing modes in which LGN relay cells are known to operate: burst and tonic mode. In tonic mode (the standard operating mode of the cells) the response varies linearly with stimulus intensity. When the cells are mildly inhibited, the response gradually becomes ohmic. However, in the case of extreme inhibition, burst mode is initiated. In this mode, the cells exhibit a nonlinear response with a high signal-to-noise ratio. In the absence of a stimulus, spontaneous activity is almost reduced to zero whereas the presence of a stimulus results in strong bursts of activity. It has been found that inhibitory input from the RTN can induce mode switching from tonic to burst mode [Wang et al., 2001].

Similar to the previous model, an attentional input is assumed to excite RTN neurons corresponding to the stimulus representation. Due to the one-to-many inhibitory mapping from RTN to LGN, LGN neurons at the stimulus representation will become strongly inhibited and neurons surrounding the representation will become weakly inhibited. This could result in a switch from tonic to burst mode at the representation, from tonic to an ohmic response (or burst mode) in the surrounding region and no change outside this region. The result will be a higher signal-to-noise ratio at the stimulus and a reduction in spontaneous activity surrounding the stimulus due to either ohmic response or burst mode.

A significant problem with this model is the fact that when the same task is performed without a stimulus, an increase in baseline activity should be observed during attention [O'Connor et al., 2002]. Since burst mode reduces spontaneous activity [Sherman, 1996], this model is unable to account for this observation.

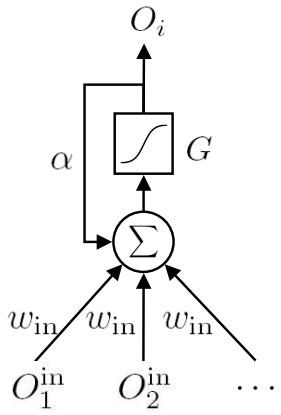

Figure 6: Schematic diagram of the computations performed by a model neuron.

For this reason, in the proposed model, LGN relay cells are assumed to be operating in tonic mode (with is linear response properties).

3.3 Proposed Information 'Flow'

The novel model introduced here can generate the desired effects with a global, uniform attention signal. In this way, the attention signal that enters the model is the simplest one conceivable. Since the experiment is concerned with featural attention, this uniformity assumption is necessary. The main mechanism involves a diffusion of stimulus driven relay cell activity to RTN (both directly and via V1) and a subsequent injection of inhibition from RTN to LGN relay cells in regions of the LGN surrounding the stimulus representation. A relatively similar hypothesis was proposed by Montero [1999]. The model is now explained in full detail.

3.3.1 Model Neurons

Classical artificial neurons were used in the model. A schematic diagram is shown in Fig. 6. The model is dynamical in nature and operates in discrete time steps. The output of unit i at time $t+1$, $O_i(t+1)$, is a weighted summation of the outputs of the neuron's fan-in at the previous time step and a decayed version of the current neuron's previous output, which is sent through a nonlinear squashing function $G(\cdot)$. This computation can be summarised by the following equation:

$$O_i(t+1) = G\left[\alpha O_i(t) + \sum_{j \in f_{\text{in}}} w_{\text{in}} O_j^{\text{in}}(t) + \ldots\right] , \tag{1}$$

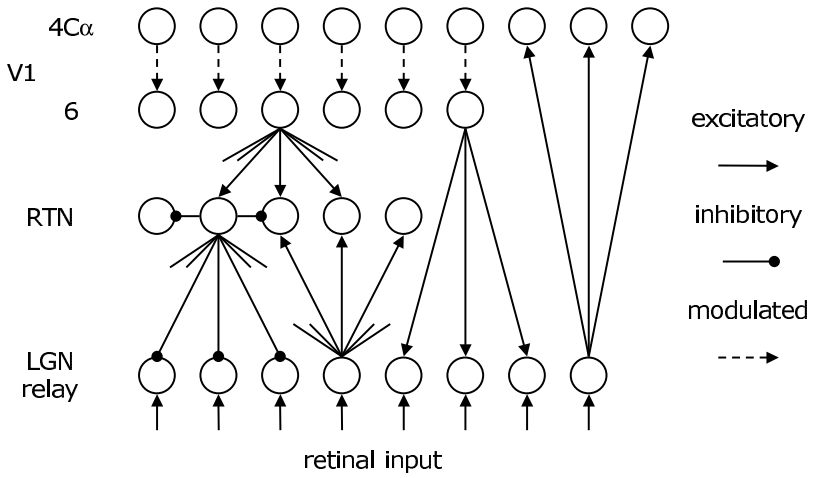

Figure 7: Thalamocortical model architecture.

where α is the proportion of neuron decay, w_{in} is the connection weight between the two layers under consideration (the connection weights are assumed constant for all neurons withing the same layer), f_{in} is the set of all neurons that connect to the current neuron and O_j^{in} is the output of the j-th incoming neuron. The nonlinear transfer function $G(\cdot)$ is a sigmoid, parameterised so as to operate on the interval $[0, 1]$:

$$G[x] = \frac{1}{1 + e^{-10x+5}} \ . \tag{2}$$

3.3.2 Model Architecture

For simplicity, a one-dimensional stimulus is considered and consequently, all layers in the network are one-dimensional chains of neurons. All layers are connected topographically. A schematic diagram of the complete model is shown in Fig. 7. The model consists of four layers, corresponding to LGN relay, RTN and area V1 layers 4Cα and 6. As in Fig. 5, excitatory connections are depicted by arrows, inhibitory by circles. The connections on which the attentional mechanism is assumed to operate are shown dashed.

The retinal input (an on/off representation of the stimulus was used) arrives in the LGN relay cells. From the LGN, this activation is transfered upwards through excitatory connections with RTN and V1 4Cα. Note that the spread of connections to and from the RTN is larger. By means of attentionally-modulated connections, activation flows further towards output layer 6 of V1. There, it is fed back to the thalamus through excitatory connections between layer 6 and RTN and LGN relay cells. In the RTN, lateral inhibition is active and finally widespread inhibition is injected back into LGN relay cells.

The attentional mechanism is located on the connections between layers V1 4Cα and 6, in agreement with a large body of anatomical results indicating that feedback connections end outside layer 4. In the presence of featural attention, the weights of all connections between these two layers are multiplied with a common factor (or a common input is provided to these target neurons). This indirectly strengthens the feedback of activation towards the RTN and the LGN.

In the actual model, activation does not flow sequentially through the different layers but instead all updates are performed simultaneously. In this way, the recurrent model behaves truly dynamically and a number of iterations (typically around 10) need to be performed before a steady-state is obtained.

3.3.3 Model Training

To investigate whether the proposed architecture is capable of generating all the desired effects, a training procedure was used. This training procedure results in a single weight for each layer-to-layer connection, together with the magnitude of the weight increase between V1 layers 4Cα and 6 during attention. Due to the complexity of the dynamical model and the qualitative nature of the associated cost function, an analytical evaluation of *e.g.* the gradient of this cost function, is extremely difficult to obtain. For this reason, a simplified simulated annealing-like optimisation was opted for. Through manual experimentation, a 'good' set of initial parameters was obtained first. From this initial model an (almost random) exploration of the parameter space was performed next. The direction of the weight updates was decided upon by a numerical evaluation of the cost function's derivative.

3.4 Results

An overview of the results obtained with the trained model is shown in Fig. 8. The activation pattern for all neurons in LGN, RTN and V1 (2 layers combined) are shown in the presence (solid line) and absence (dashed line) of attention and for both a large and small stimulus. The large stimulus activated LGN relay cells at index 110 up to 190. The small stimulus was present at neurons 140 up to 160. The magnitude of spontaneous and stimulus activity were chosen so as not to saturate LGN relay cells (values of 0.4 and 0.8 were chosen respectively). Fig. 8 contains the steady-state activations for all the different conditions.

Both for large and small stimuli, a ring of suppressed activity surrounding the representation of the stimulus can be observed both in LGN and V1. The differential effects in V1 are smaller than in LGN. At the representation and outside the ring, there is no significant change in activity. In correspondence with the experimental results and the inhibitory nature of area RTN, an increased activation can be observed during attention in this layer.

When comparing the activity patterns for large and small stimulus size it becomes apparent that the effects scale with stimulus size. The suppression in LGN and V1 is smaller for the small stimulus and so is the increase in activation in RTN.

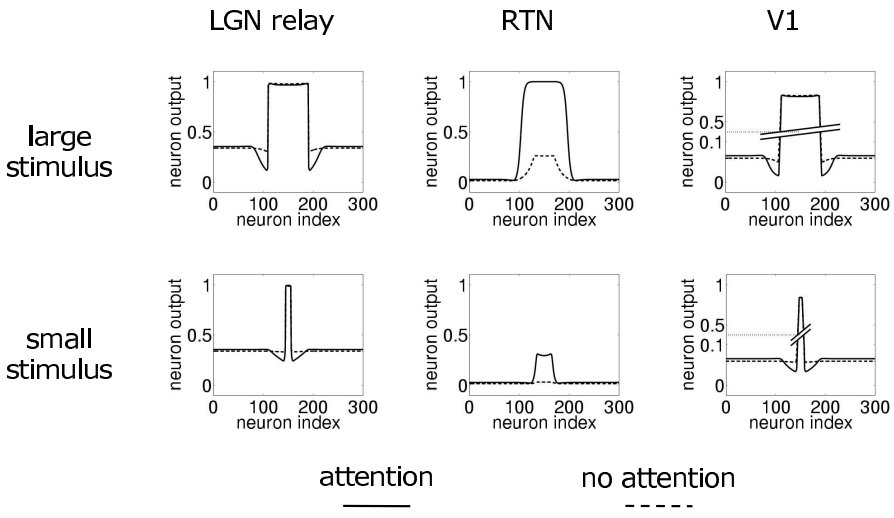

Figure 8: Neuron outputs of the different layers in the model for a large and small stimulus. Solid and dashed lines correspond respectively to outputs obtained in the presence and absence of attention.

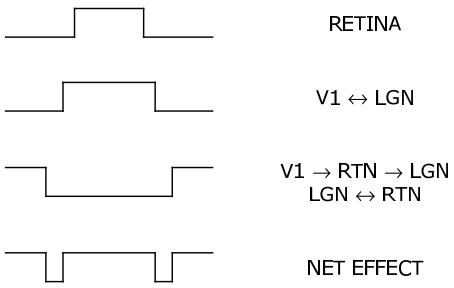

Figure 9: Origin of surround suppression. Rows illustrate the different effects that operate simultaneously on LGN relay cells in the presence of a stimulus.

In conclusion, all model results are in correspondence with the experimental results described in Section 2.2.

3.5 Underlying Mechanism

We next further elaborate on the workings of the proposed model and more precisely on the origin of the surround suppression and the effects related to the stimulus size.

3.5.1 Surround Suppression

To obtain the surround suppression, it is crucial that the connections between V1 and LGN have a smaller spread than the connections to and from the RTN. To see this it is helpful to look at Fig. 9 which contains a simplified illustration of how the surround suppression builds up in the network. By means of the, albeit relatively small, divergence of activity between V1 and LGN, the positive effect on LGN relay cells (second row) is slightly smoothed (and wider) as compared to the retinal stimulation (top row). The much larger spread of connections between V1, RTN and LGN has a similar but even wider and inhibitory effect on LGN relay cells (third row). Provided both the excitatory and inhibitory effects are balanced at and outside the stimulus representation, the net effect is a suppression surrounding the stimulus representation in LGN (bottom row). Note that this effect is already present in the absence of attention. Since the attentional mechanism (strengthening the connection weights between V1 layer 4Cα and 6) indirectly strengthens the connections from V1 to RTN and from V1 to LGN in a comparable manner, the excitatory and inhibitory effects are balanced and cancel each other out at the stimulus representation and outside the suppressive ring. In the region surrounding the stimulus however, the suppression is deepened.

3.5.2 Stimulus Size Effects

Another mechanism is responsible for the deeper suppression and higher RTN activation when the stimulus size increases. Crucial here is that the spread of connections to and from the RTN is larger than the stimulus width. In this way, increasing the stimulus size results in more stimulus-induced inhibition towards neurons located at the suppressive ring. With a smaller stimulus, these surrounding neurons receive inhibition from the same number of neurons but since now some of them do not receive stimulus-induced activation, the total inhibition will be smaller.

4 Conclusions

A novel thalamocortical model has been introduced to investigate attentional effects observed in early visual cortex during 2DG experiments. The model demonstrates that a uniform attention signal in V1, combined with a diffusion of stimulus activity in RTN can account for all the observed experimental results. The model architecture was validated through simulation.

References

J. Bourassa and M. Deschênes. Corticothalamic projections from the primary visual cortex in rats: a single fiber study using biocytin as anterograde tracer. *Neuroscience*, 253:253–263, 1995.

F. Crick. Function of the thalamic reticular complex: the searchlight hypothesis. *Proc Natl Acad Sci USA*, 81:4586–4590, 1984.

B.J. Geesaman, R.T. Born, R.A. Andersen, and R.B. Tootell. Maps of complex motion selectivity in the superior temporal cortex of the alert macaque monkey: a double-label 2-deoxyglucose study. *Cerebral Cortex*, 7:749–757, 1997.

R.W. Guillery and S. Murray Sherman. Thalamic relay functions and their role in corticocortical communication: Generalizations from the visual system. *Neuron*, 33:163–175, 2002.

D. LaBerge, M. Carter, and V. Brown. A network simulation of thalamic circuit operations in selective attention. *Neural Computation*, 4:318–331, 1992.

J Mitrofanis and R.W. Guillery. New views of the thalamic reticular nucleus in the adult and the developing brain. *Trends in Neurosciences*, 16:240–245, 1993.

V.M. Montero. Amblyopia decreases activation of the corticogeniculate pathway and visual thalamic reticularis in attentive rats: A 'focal attention' hypothesis. *Neuroscience*, 91:805–817, 1999.

D.H. O'Connor, M.M. Fukui, M.A. Pinsk, and S. Kastner. Attention modulates responses in the human lateral geniculate nucleus. *Nature Neuroscience*, 5 (11):1203–1209, 2002.

S.M. Sherman. Dual response modes in lateral geniculate neurons: Mechanisms and functions. *Visual Neuroscience*, 13(2):205–213, 1996.

M. Steriade, L. Domich, and G. Oakson. Reticularis thalami neurons revisited: activity changes during shifts in states of vigilance. *Journal of Neuroscience*, 6:68–81, 1986.

W. Vanduffel, R.B.H. Tootell, and G. Orban. Attention-dependent suppression of metabolic activity in the early stages of the macaque visual system. *Cerebral Cortex*, 10:109–126, February 2000.

S. Wang, M.E. Bickford, S.C. Van Horn, A. Erisir, D.W. Godwin, and S.M. Sherman. Synaptic targets of thalamic reticular nucleus terminals in the visual thalamus of the cat. *Journal of Comparative Neurology*, 440:321–341, 2001.

C.T. Yen, M. Conley, S.H.C. Hendry, and E.G. Jones. The morphology of physiologically identified gabaergic neurons in the somatic sensory part of the thalamic reticular nucleus in the cat. *Journal of Neuroscience*, 5:254–268, 1985.