Ecovision – Work Package 6 – Deliverable 6.1

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

Computational Neuroscience Algorithm for generating Task-Optimized Spatial Representations of Neural Activities

Karl Pauwels and Marc M. Van Hulle

1 Introduction

The model discussed in this report has been developed on the basis of neurophysiological data obtained in an attention study performed by Vanduffel *et. al.* [4]. In a first section, a short discussion of these experiments and the obtained results is given. Next, the data that was generated by these experiments and that served as the basis for the model are explained. The model itself and the results obtained by training it are given in a subsequent section. We conclude with the future directions that will be taken within this project.

2 Physiological Experiments

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 [1] was used to register activation levels in these areas for awake monkeys, performing a task which involved featural attention. The main advantage of this technique is that the neural activity can be measured with very high spatial resolution. A disadvantage is the low temporal resolution; for one animal only two, temporally aggregated, activity images can be obtained. This does however allow for one control condition, which in this case involved a task with approximately similar visual stimulation but involving spatial attention.

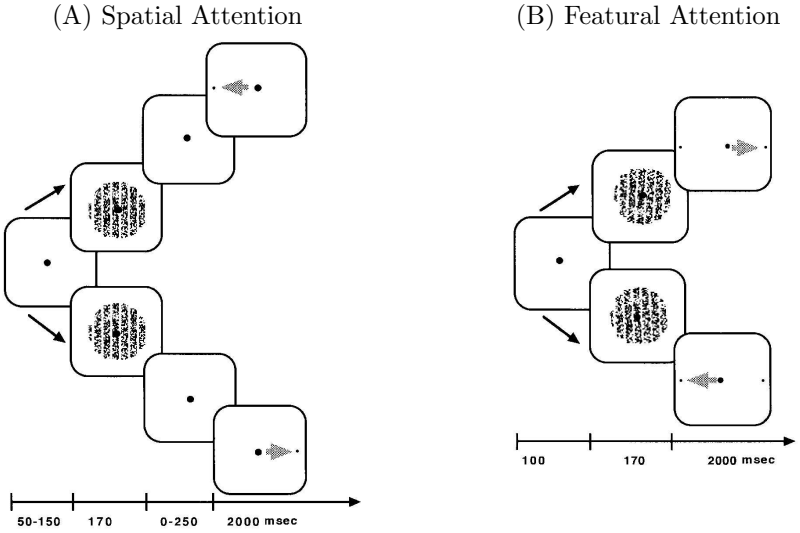

Figure 1: Task descriptions

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 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 behavioral 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 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. In Fig. 2(B) the difference in activation for both conditions is shown for layer 4Cα. It is clear from this figure that there is a region with lower activation (coded in red) during featural attention, surrounding the representation of the stimulus. At the location of the stimulus no significant differences were observed. 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.

(A) (B)

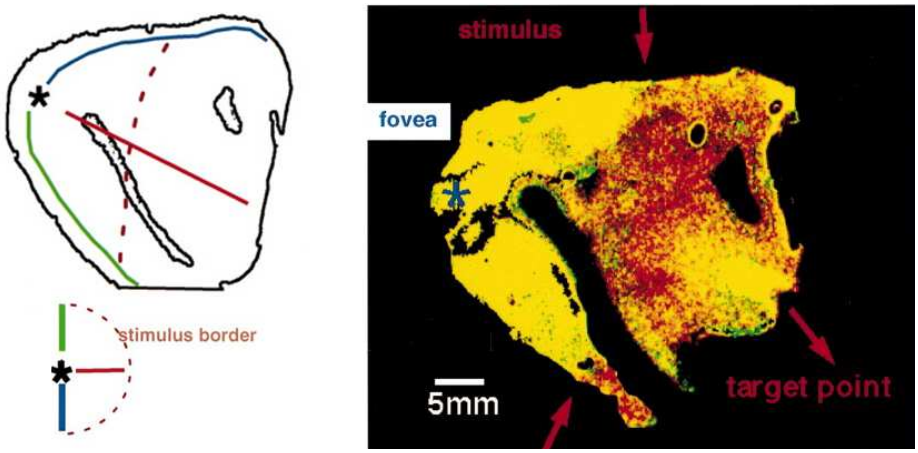

Figure 2: (A) Schematic illustration of the location of the stimulus representation in layer 4Cα. (B) Reassembled layer 4Cα showing a ring of suppression surrounding the stimulus. Color scale: red = lower featural attention-related DG uptake; yellow = no differential DG uptake

3 Available Data

The source images, from which the difference image in Fig. 2(B) was constructed, were used to construct our attention model. These images are shown in Fig. 3(A) for the case of spatial attention and in Fig. 3(B) for featural attention. It is apparent that these images are very noisy and contain many artifacts. In the figures, air bubbles and missing pieces are clearly visible. Furthermore, due to the complicated process of reconstructing layer 4Cα, misalignment errors may be present as well.

(A) Spatial Attention (B) Featural Attention

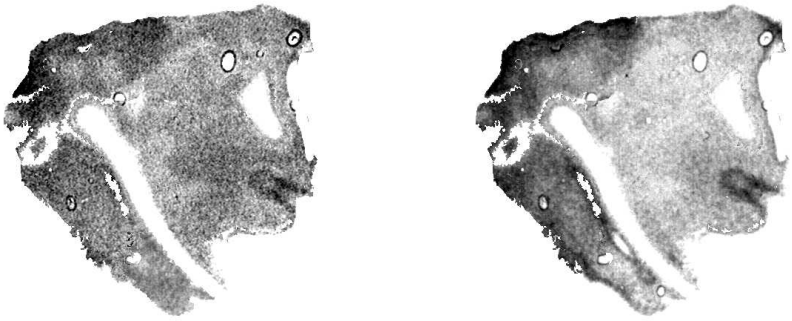

Figure 3: Available 2DG data

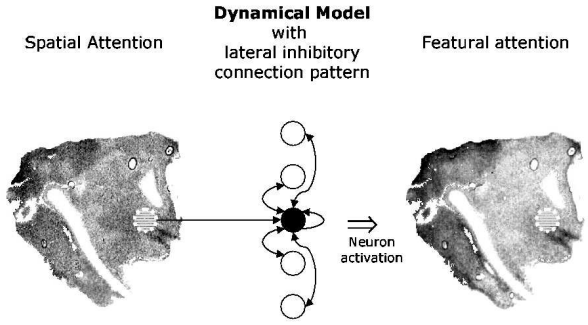

Figure 4: Model Overview

4 Model

In the following subsections the proposed model for attentional modulation in early-visual processing is explained in detail.

4.1 Assumptions

The model is built on the following assumptions. The task is to detect very small differences in the orientation of the grating. It is assumed that, in order to do this, the monkey uses an attentional mechanism that allows it to attribute increased processing to the feature 'orientation', which results in an increased sensitivity for orientation differences. The attentional mechanism we propose, temporarily changes the inputs received from lateral inhibitory neurons by changing their effective connectivity. In this way, activation levels are changed in order to suppress activity in non-attended regions.

4.2 Model Neurons

Since the available data only contains the neuron activation levels during two different task conditions, we need to restrict our model to this part of the cortex. Consequently we do not include external stimulation and conceptualise the model as an array of neurons (with constant neural density) corresponding to layer 4Cα of monkey primary visual cortex. Since visual input is not part of the model, issues related to the cortical magnification factor [2] need not be considered here.

4.3 Attentional Modulation

A mechanism is proposed to model the attentional effects observed when attending the feature 'orientation'. The activation observed while performing the spatial attention task is considered as 'passive' activation, caused by observing the stimulus (left-hand side of Fig. 4). Consequently, we assume this activation remains present during the dynamical evolution of the proposed modulatory

mechanism and serves as a constant input to the layer of neurons that comprise the model. The neurons themselves are fully interconnected and specific weights are assigned to these connections to dynamically alter the neurons activation levels (center of Fig. 4). Goal of this modulation is to transform the 'passive' activation into the observed featural attention activation, including the ring of suppression surrounding the representation of the stimulus (right-hand side of Fig. 4).

The constant spatial attention activation of neuron i, A_i also serves as its initial featural attention activation F_i at iteration 0:

$$F_i^{(0)} = A_i \ . \tag{1}$$

The connection weight from neuron j to neuron i is depicted by w_{ij}. Since the neurons are fully and recurrently connected, the network is a dynamical system, the final activation F_i^* of which needs to be determined iteratively. Based on the featural attention activation at iteration k, the activation at the next iteration is determined as follows:

$$F_i^{(k+1)} = \sum_{j=1}^{N} f(w_{ij} F_j^{(k)} + A_j) \ , \tag{2}$$

where

$$f(x) = \begin{cases} 0 & x \leq 0 \\ x & 0 < x \leq 1 \\ 1 & x > 1 \end{cases} \ . \tag{3}$$

The weights w_{ij} are determined using a learning algorithm that tries to match the final activations F_i to the featural attention activations observed in the 2DG-study (see Fig. 3(B)). Specific constraints are enforced on the shape of the weights interconnection pattern. These constraints are the subject of the next section.

4.4 Connection Pattern

Our model of V1 consists of a two dimensional array of fully interconnected neurons. A weight can be assigned to each connection to control the excitatory or inhibitory effect a neuron's activation has on other neurons. Obviously, such a massively interconnected lattice has far too many degrees of freedom. To prevent the network from overfitting the data, the connection pattern needs to be constrained. Since we are mainly interested in explaining the effects using center-surround interaction patterns, the connection weights are sampled from a continuous Mexican-hat-shaped kernel. In Fig. 5 an example kernel is shown. With increasing distance (in lattice coordinates), the connection strengths change from an excitatory center, to an inhibitory surround until they finally saturate at zero. Furthermore, the kernel is constructed in such a way that the average incoming effect equals zero and the self-connection strength equals one.

A variety of different kernels have been investigated. First of all, classical Laplacian or Gabor kernels are not suitable since they only have one adjustable parameter. Since our goal is to control the ranges of the center and surround parts

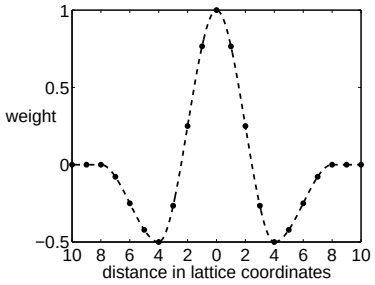

Figure 5: Example Mexican-hat kernel

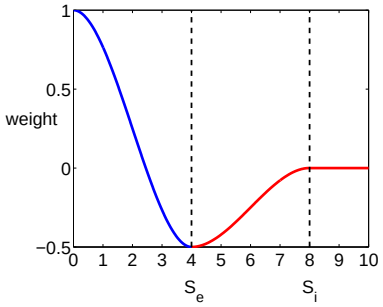

Figure 6: One dimensional, continuous, splines kernel

separately and for each neuron individually, two different parameters should be available to accomplish this. The most obvious choice are Difference Of Gaussians (DOG). They have however proved to be unstable during training, mainly because of a singularity that originates when the standard deviations of the excitatory and inhibitory component are nearly identical. To solve this problem, a novel kernel has been developed that consists of two cubic splines, parameterised in such a way to demonstrate all of the desirable properties mentioned above.

The shape of the kernels is determined by the parameters S_e and S_i, which are related to the span of the excitatory, respectively inhibitory region. The first spline is defined as:

$$y_e = 2\left(\frac{S_e + S_i}{S_e^3 S_i}\right) x_e^3 - 3\left(\frac{S_e + S_i}{S_e^2 S_i}\right) x_e^2 + 1 \ , \tag{4}$$

where $0 \leq |x_e| \leq S_e$. The second spline equals:

$$y_i = \frac{S_e}{S_i \left(S_i - S_e\right)^3} \left[S_i^2 \left(3S_e - S_i\right) - 6S_e S_i x_i + 3\left(S_e + S_i\right) x_i^2 - 2x_i^3\right] \ , \tag{5}$$

where $S_e < |x_i| \leq S_i$. The right side of an example splines kernel with parameters $S_e = 4$ and $S_i = 8$ is shown in Fig. 6. In the two-dimensional case, x-values are replaced by the distance from the center $r = \left(x_1^2 + x_2^2\right)^{1/2}$. To compensate for discretisation errors which may cause the weights not to sum to zero, a normalisation is performed by rescaling the inhibitory part of the kernel. An example of a discrete, two dimensional kernel is shown in Fig. 7.

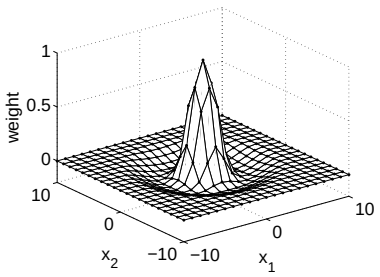

Figure 7: Two dimensional, discrete, splines kernel

4.5 Learning Rule

A novel learning rule has been developed for selecting the sizes of the excitatory and inhibitory regions of the kernels for each neuron. The goal of the learning process is to generate a mapping from the activation observed during spatial attention to that observed during featural attention. Since the model neuron activations evolve dynamically over time, it is not possible to use a conventional gradient descent scheme. For this reason we locally approximate the error gradient numerically. A good approximation is obtained by reiterating the model after small parameter updates and evaluating the error. The advantage of this procedure is that it avoids the cumbersome training procedures, often used with recurrent neural networks, and that it enables training a very complicated model using a straightforward procedure. The cost of this simplicity is a longer training duration, which is not really a concern here. A procedure like this has additional advantages since it can, in principle, even model the dynamical behavior of the network. In order to accomplish this, data with higher temporal resolution is however required.

5 Results

Since the full model suffers from issues relating to border effects and is difficult to train, we have performed a proof of concept simulation using a simplified version of the model and artificial data. Synthetic, one-dimensional data has been created which closely resembles a cross-section of the observed 2DG-data. This dataset is shown in Fig. 8. The stimulus is represented on the left-hand side of the figure. Since, in accordance with the 2DG data, there is less visual stimulation outside the representation, the activation is slightly lower at those locations. In the featural dataset, activation is lowered with respect to the spatial activation in a limited region (ring) outside the stimulus representation. The only simplifications that were made to the model are the removal of the dynamical aspects and the nonlinear transfer function. In other words, for each neuron individually, the splines excitatory and inhibitory parameters are determined so as to generate the functional mapping from the spatial to the featural attention dataset. The parameters were initialised randomly and the best (in terms of minimal deviation from the featural attention activation) results of different trails were retained. These results are shown in Fig. 9. Since

(A) Spatial Attention

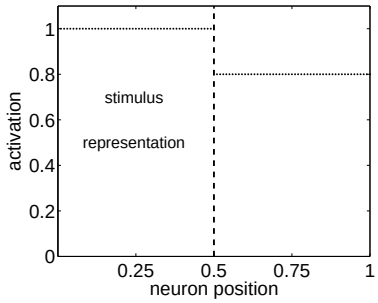

(B) Featural Attention

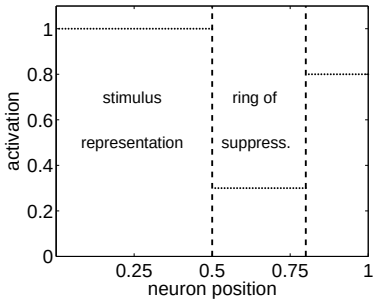

Figure 8: Synthetic datasets

(A) Splines Parameter Values

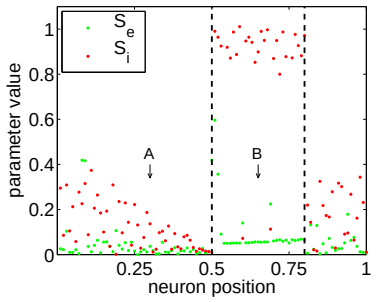

(B) Model Activation

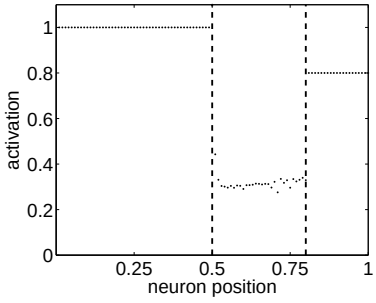

Figure 9: Featural attention activation (B) obtained by the model with trained excitatory and inhibitory parameter values shown in (A)

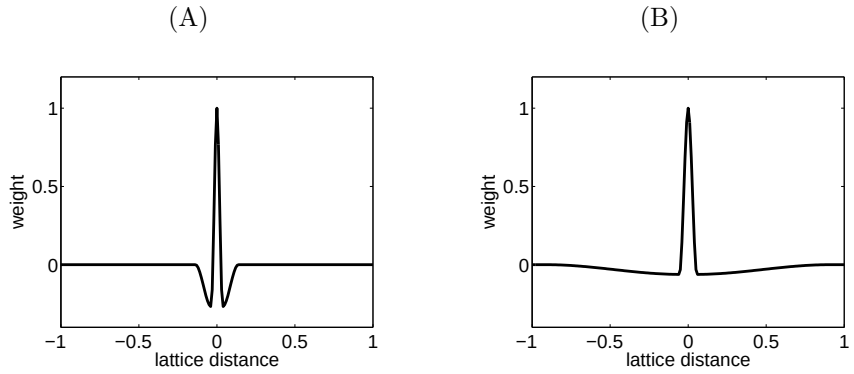

Figure 10: Representative kernel configurations in the region corresponding to the stimulus representation (A) and the suppressive region surrounding the stimulus representation (B)

the activation does not have to be changed in the areas corresponding to the stimulus representation and outside the ring of suppression, different parameter configurations are possible. For this reason, the fit is perfect and the parameter values are always relatively small. No modulatory mechanisms are required here since the neurons are already driven to saturation. The apparent noise present in the parameter estimates is due to the discrete nature of the kernel parametrisation and the consequent difficulties in training the model. The higher values on the left-hand side of the figure are caused by border effect. The area outside the stimulus representation is more interesting and clearly shows a consistent parameter configuration. The excitatory center is always relatively small and the inhibitory surround is always large. Fig. 10 shows the connection pattern for the two kernels marked with arrows in Fig. 9.

These results indicate that the observed ring of suppression may originate from long-range inhibitory connections that start from within the stimulus representation and end in a peripheral region. The connections use the activation, generated by the stimulus, to inhibit a spatially localised, surrounding, region.

6 Conclusions

A novel model architecture has been introduced to investigate attentional effects observed in early visual cortex during 2DG experiments. The model involves a new center-surround kernel and requires a novel learning rule to be trained. Using a synthetic dataset which closely resembles the observed physiological data, it has been shown that a specific spatial configuration of long-range inhibitory connections (originating at the location of the representation of the attended stimulus) can generate a 'ring of suppression' surrounding the focus of attention.

7 Further Steps

Since further 2DG experiments are not supported by the project, and since this class of experiments are extremely demanding in terms of material and personnel costs, additional data are not to be expected in the near future. As a consequence this WP is revised towards a more technical approach. Specifically, the interactions between attention and context will be further investigated. We will adopt the view that attention and context interact simultaneously [3] and will involve lateral connectivity to be in line with the suppression observed in the 2DG experiments.

References

[1] Geesaman, B.J. Born, R.T., Andersen, R.A. & Tootell, R.B. (1997) 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.

[2] Johnston, A. (1989) The Geometry of the Topographic Map in Striate Cortex. *Vision Research*, 11:1493–1500.

[3] Raizada, R. & Grossberg, S. (2001) Context-Sensitive Bindings by the Laminar Circuits of V1 and V2: A Unified Model of Perceptual Grouping, Attention, and Orientation Contrast. *Visual Cognition*, **8**(3/4/5), 341–466.

[4] Vanduffel, W. Tootell, R.B.H. & Orban, G.A. (2000) Attention-dependent Suppression of Metabolic Activity in the Early Stages of the Macaque Visual System. *Cerebral Cortex*, 10:109–126.