

A Continuous Attractor Neural Network Model of Divided Visual Attention

Dominic I. Standage and Thomas P. Trappenberg
Faculty of Computer Science
Dalhousie University
Halifax, NS, Canada
E-mail: {standage, tt}@cs.dal.ca

Raymond M. Klein
Department of Psychology
Dalhousie University
Halifax, NS, Canada
E-mail: klein@or.psychology.dal.ca

Abstract—The biologically realistic model of selective visual attention by Deco et al uses a continuous attractor neural network to simulate a saliency map in posterior parietal cortex. We test the ability of the model to explain experimental evidence on the distribution of spatial attention. The majority of evidence supports the view that attention is a unitary construct, but recent experiments provide evidence for split attentional foci. We simulate two such experiments. Our results suggest that the ability to divide attention depends on sustained endogenous signals from short term memory to the saliency map, stressing the interplay between working memory mechanisms and attention. Our results also point to a mechanism whereby inhibitory endogenous signals may play a role in dividing attention, suggesting a possible mechanism for inhibition of return.

I. INTRODUCTION

Attention is an old concept in psychology correlated with enhanced processing of objects or regions in space [1]. While attention is a multi-modal phenomenon [2], [3], the majority of research has focused on selective visual attention (SVA). The limited capacity of the visual system necessitates a mechanism to select stimuli from the visual field, and Tsotsos pointed out that attention solves the complexity problem of sensory processing [4].

A distinction can be drawn between pre-attentive and attentive visual processing [5]. Pre-attentive processing refers to bottom-up (BU) feature saliency of visual stimuli whereby items that differ from their surroundings ‘pop out’ to the viewer. Attentive processing refers to top-down (TD) influences on perception of stimuli determined by object and locational bias such as task instructions or foreknowledge of stimulus characteristics. Determining saliency, then, is both a BU and TD requirement, and computational models of SVA include maps that integrate BU salience across object features [6], TD bias [7], and the interplay of both [8], [9].

Koch and Ullman [6] provide a neural network model of SVA in which topographical feature maps are integrated by a winner-take-all (WTA) saliency map of BU stimuli. In their model, inhibiting the selected location causes a shift to the next most salient location. Wolfe [8] builds on Neisser’s pre-attentive/attentive distinction [5], integrating BU and TD saliency criteria in his Guided Search model. Treisman [7] provides a model of spatial attention to solve the Binding Problem, in which a TD saliency map determines object

features selected for further processing, and suggests parietal cortex as the biological correlate of her ‘master’ map. Deco et al [9] use inhibition to mediate BU and TD influences in an instantiation of Duncan and Humphreys’ biased competition model [10], simulating saliency in posterior parietal cortex (PP) with a Continuous Attractor Neural Network (CANN). Spatial saliency in PP interacts with BU feature maps to converge on a winning location. See [11] and [12] for a review of these and other models.

There is long-standing debate about the distribution of SVA. Many cognitive models propose a unitary focus of attention, likened to a roving spotlight over the visual field [1]. Variants of the spotlight metaphor include gradient [13], [14] and zoom lens [15] models, suggesting that attention may be a graded phenomenon, attenuated around a central focus. A large body of evidence supports such unitary models [1], [16], but several more recent experiments have provided evidence for non-contiguous allocation of SVA [17], [18], [19].

Here we study how split attention can be achieved by a dynamic implementation of a WTA map. Despite their WTA nature, CANNs are able to account for split attention when network dynamics facilitate long transition states between regimes [20] and when dominated by sustained inputs. We simulate the experiments of Müller et al [19] and Awh and Pashler [18] with a CANN model, demonstrating mechanisms by which the model can (and cannot) account for their findings, and make predictions for further study.

II. METHODS

In simulations, we used a fully connected recurrent rate model with N nodes. The average state u_i of a node with index i is given by

$$\tau \frac{du_i(t)}{dt} = -u_i(t) + \sum_j w_{ij} r_j(t) \Delta x + I_i^{\text{ext}}(t), \quad (1)$$

where τ is a time constant, I_i^{ext} is external input to the network, $\Delta x = 2\pi/N$ is a scale factor, and r_i is a normalized square of u_i given by

$$g(u_i) = \frac{u_i^2}{1 + \frac{1}{2} \sum_j u_j^2 \Delta x}. \quad (2)$$

We use this normalization through divisive normalization (shunting inhibition) to force more biologically realistic smooth (Gaussian) bubbles [21].

Weight matrix \mathbf{w} is determined by Gaussian function

$$w_{ij} = A_w e^{-d^2/2\sigma_w^2} - C \quad (3)$$

for all connections to node i from node j (for all $i, j \in N$), where d is given by

$$d = \min(|i - j|\Delta x, 2\pi - |i - j|\Delta x). \quad (4)$$

C is an inhibition constant describing the activity dependent inhibition of an inhibitory pool of neurons, and A_w is a scale factor.

External input I_i^{ext} is Gaussian shaped around input location j , determined by

$$I_i^{\text{ext}} = e^{-d^2/2\sigma_{ext}^2} \quad (5)$$

where d is given by Equation 4.

In all simulations, $N = 100$, $C \in \{0.1, 0.3\}$, $A_w = 10$, $t = 1$, $\tau = 10$, and $\sigma_w, \sigma_{ext} \in \{0.2, 0.4, 0.8\}$.

We classify our inputs along exogenous (exo) and endogenous (endo) dimensions. Exo inputs refer to neural responses to stimuli, here representing visual cues. Endo inputs refer to voluntary control of attention, here representing task instructions in behavioural studies.

Simulations are run with transient and sustained inputs. Under both paradigms, we equate network activity with SVA. Because transient inputs elicit WTA behaviour in CANN models, we start by demonstrating one-bubble attractor states as models of a unitary attentional focus. Transient input stimuli are the norm in biological networks, as evidenced by high firing rates at stimulus onset followed by lower rates when stimuli are sustained in experimental settings. In the exo case, this initial burst of activity serves as input to higher cortical areas such as PP. Sustained firing after transient stimulation is a property of highly specialized neural assemblies [22], and as such is the exception among biological networks, not the norm. We interpret sustained endo inputs as STM representations of task instructions in PFC.

In simulations with sustained inputs, we model exo inputs as transient when stimuli in behavioural experiments are stable, as described above, and model these stimuli as sustained when the features of spatially static stimuli are changing. We interpret these changes as providing continual ‘refreshment’ of neural representations due to novelty effects [23]. Endo inputs are always sustained in this set of simulations.

We use a Gaussian shaped input profile to approximate typical tuning curves of neurons, so their firing profiles are well approximated by smooth curves. In the case of transient input, the specific shape of localized input is unimportant because the network dynamic dominates after cessation of input. With sustained input, a Gaussian input profile leads to a good approximation of a Gaussian output profile, achieving the biological realism of our input profile described above.

We compare our CANN model of SVA to one with no lateral interaction, modelling the latter by simply adding together its Gaussian inputs.

III. SIMULATIONS

Müller et al [19] provided evidence for sustained division of visual attention by recording steady state visual evoked potentials (SSVEP) while subjects attended to two of four contiguous horizontal locations. The SSVEP is the electrophysiological response in visual cortex to a rapidly flickering stimulus, and has been shown to increase in amplitude when attention is paid to the location of the stimulus [19]. They found that target detection rates (TDR) and SSVEPs were lower at the location between separated targets in a detection task. Additionally, they showed that split locations were attended just as well as adjacent locations in their experiment. We model this experiment in Simulations 1 and 2.

Awh and Pashler [18] used a partial report procedure to test subjects’ ability to divide spatial attention. Subjects viewed a 5x5 array of alpha-numeric characters containing 23 letters and 2 digits. They fixated at a central location before the presentation of two cues indicating the probable location of the digits. The character array was then provided and the subjects’ task was to identify the digits. Digits appeared between the cued locations on twenty percent of trials and performance at the cued and intervening locations was compared. If SVA can be divided, subjects should perform better at the cued locations than in between. Subjects’ ability to divide SVA was found to depend on the presence of a subsequent noise mask, but the removal of array noise alone was sufficient to significantly reduce division of attention, regardless of subsequent masking. We model this work in Simulation 3.

A. Simulation 1

Adopting Müller’s terminology, we refer to the locations of stimuli as 1, 2, 3 and 4, where 1 is the left-most location and 4 is the right-most location. A 1+2 trial refers to trials in which subjects were instructed to direct their attention to locations 1 and 2, a 2+4 trial refers to instructions to attend to locations 2 and 4, and so forth for other combinations of the four locations. We equate activity in the CANN with spatial attention.

The network was given exo and endo inputs for 300 iterations of dt/τ , followed by 300 iterations without either source of input. These times were sufficient for the network to stabilize under both dynamic regimes (both during and after input). The input profile, final network state, and the network state over time are shown in Figure 1A, B, and C respectively.

In 1+2 and 3+4 (adjacent) trials, bubbles corresponding to the attended locations merge into a single bubble, predicting a unitary focus of attention maximal between the two locations. In 1+3 and 2+4 (split) trials, one bubble wins at the expense of the other, predicting that only one of the locations should be attended. These results conflict with Müller’s findings, and are shown in Figure 1.

B. Simulation 2

Network configuration and the shape and location of inputs is identical to Simulation 1. Exo and endo inputs are sustained

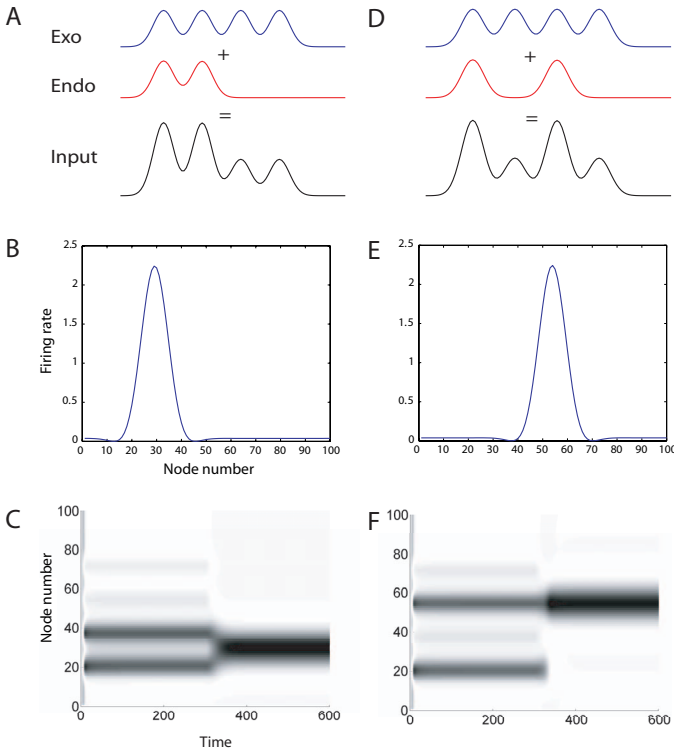


Fig. 1. (A) Nodes 20, 37, 54 and 71 correspond to locations 1, 2, 3 and 4 respectively. Exo inputs are applied to all locations. Endo inputs are applied to locations 1 and 2 only. Combined exo and endo input activity shown at bottom. Gaussian width factors $\sigma_w = 0.4$, $\sigma_{ext} = 0.2$. (B) Stable bubble following transient input. The bubble is centred on node 29, reflecting the merge between locations 1 and 2. (C) Network activity over time. Input is stopped after 200 iterations, followed by transition to a one bubble (merged) state by approximately 250 iterations. (D) All parameters are identical to A except endo inputs are applied to locations 1 and 3. (E) Stable bubble following transient input, centred on node 54 (location 3). (F) Network states over time. Location 3 wins by approximately 250 ms.

simultaneously for 500 iterations, sufficient for the network to stabilize.

With sustained inputs, our model replicates Müller’s findings for both adjacent and split trials, as shown in Figure 2A. Results show that network activity conforms to the shape of the inputs, though areas receiving only exo inputs are partially suppressed by lateral inhibition from more excitatory areas receiving both exo and endo inputs. Despite this sharpening of the activity profile, the network activity under sustained inputs resembles a simple addition of Gaussian curves (AOG).

To test if an AOG provides a model of CANN behaviour under sustained inputs, we measured the reduction or ‘dip’ in activity between bubbles at different distances between inputs, comparing it to peak activity in the bubbles. Correspondingly, we measured the height of the midpoint between two Gaussian curves as a function of the distance between them, comparing it to their maxima. The width of CANN bubbles is matched to the width of Gaussians in AOG. The shape of these curves is similar, but the AOG is shifted. Unfortunately, without a means to relate distance to physical parameters, this shift does not help to distinguish one curve from the other. Both

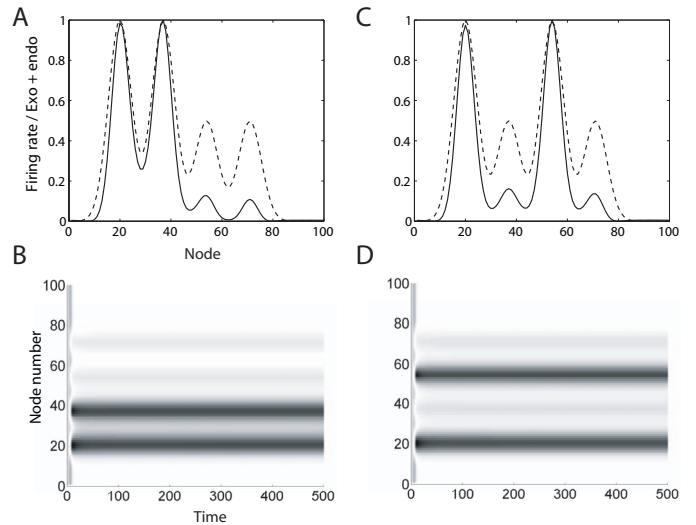


Fig. 2. (A) Input profile as described in Figure 1A. (B) Network activity conforms to the sustained input profile, where lateral inhibition sharpens the areas between peaks. For clarity, input (dashed) and output (solid) are normalized to 1 in the figure. (C) Network activity over time for sustained 1+2 trial. (D) 1+3 input profile as described in Figure 1D. (E) Activity profile conforms to 1+3 sustained input. (F) Network states over time.

models predict that SVA should be easier to divide as the distance between attended locations is increased, and it would be difficult to choose one model over the other if behavioural data were to validate this prediction. Results are shown in Figure 3A. As models of brain networks, the CANN suggests a lateral interaction profile, while the addition of Gaussians suggests a purely feed-forward topographic organization.

Because the width of Gaussian inputs effects the curves in Figure 3A, we plot these curves for different values of σ_{ext} in Figure 3B. This figure shows curves for the AOG case. The same effect was observed for CANN bubbles (results not shown). We equate the overlap between Gaussians with the overlap between neural representations driving the input signals, and predict that the more stimuli overlap in space (or feature space) the harder it is to divide SVA.

C. Simulation 3

Network configuration is identical to Simulations 1 and 2 except for $\sigma_w = 0.8$ and $C = 0.3$. The number and location of inputs are adjusted to reflect experimental conditions (described in Figure 4). Inputs are centred on nodes 10, 20, 30, 40 and 50, hereafter referred to as locations 1, 2, 3, 4 and 5. Exo inputs at locations 1, 3 and 5 represent letters (array noise) in Awh and Pashlers’ study. Exo inputs at locations 2 and 4 represent digits. Endo inputs at locations 2 and 4 represent task instructions to attend to these locations. In these simulations, we model Awh and Pashlers’ experiments only as far as the presentation of the character array. The effect of noise masks and subsequent identification of target digits presumably involve STM and object recognition processing not included in our model.

We provide an exo input to location 3 for 200 iterations, simulating subjects’ fixation between the target locations. This

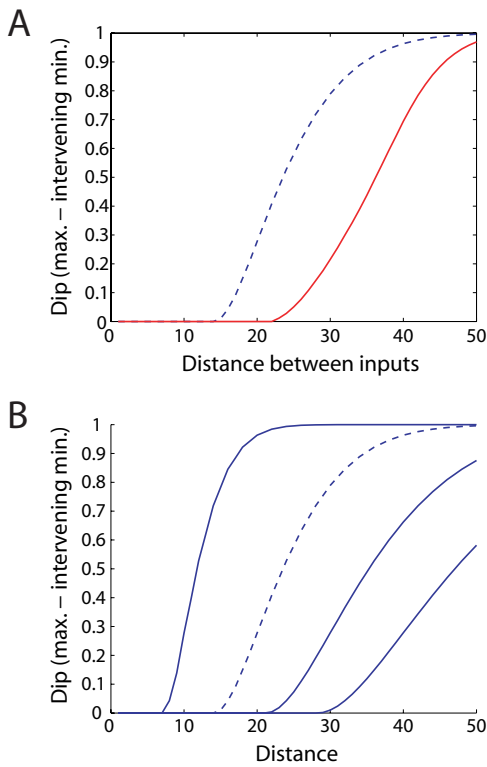


Fig. 3. (A) Dip between CANN bubbles (max. - intervening min.) plotted against the distance between peaks (solid line). Dip between peaks of an AOG plotted against distance between them (dashed line). (B) AOG distance vs. dip (as in (A)). From left to right, $\sigma_i = 5, 10, 15, 20$.

input is followed by 200 iterations of the same exo input, and endo inputs at locations 2 and 4. We then provide simultaneous exo inputs to all locations (the full character array) and endo inputs to locations 2 and 4 (the attended locations) for 100 iterations. These numbers of iterations are sufficient to model the duration of input screens in Awh and Pashlers’ study.

Our results replicate those of Awh and Pashlers’ Experiment 1. The sustained endo inputs dominate the network, facilitating the two bubbles shown in Figure 4A.

In Simulation 3b, we model Awh and Pashlers’ Experiments 4 and 4a, simulating removal of non-target characters by removing all exogenous input except at locations 2 and 4. Our results do not coincide with those Awh and Pashler. Without exogenous stimuli at locations 1, 3 and 5, bubbles are just as cleanly differentiated, predicting divided attention where Awh and Pashler found greatly reduced division of attention in Experiment 4, and unitary attention in Experiment 4a.

IV. DISCUSSION

The parameter dependence of our model raises several issues. We use a normalised square gain function to create smooth output. In Simulation 1, bubbles at attended locations are sustained until cessation of input, after which a single winner dominates. In adjacent trials, the bubbles at input locations merge, peaking between the two centres of input. This activity profile suggests that attention should be maximal

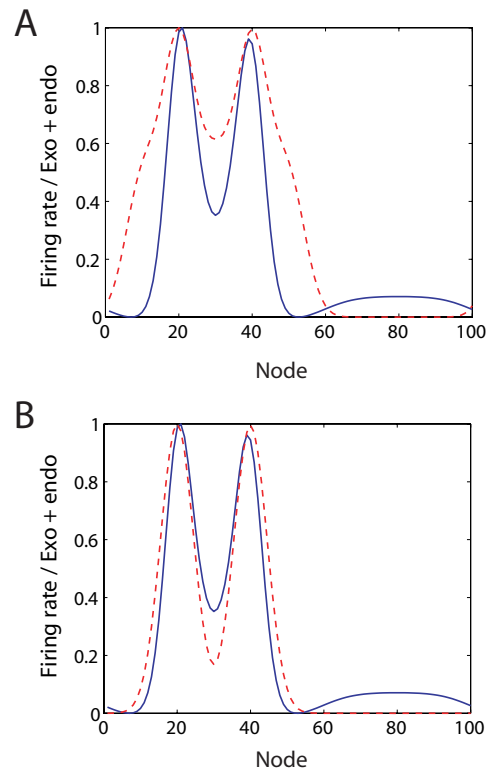


Fig. 4. (A) Nodes 10, 20, 30, 40 and 50 are referred to as locations 1, 2, 3, 4 and 5. Locations 1, 3 and 5 represent locations of letters. Locations 2 and 4 represent locations of digits. Exo input is applied to location 3 for 200 ms (fixation) followed by 200 ms of exo input (location 3) and endo inputs at locations 2 and 4. Exo inputs are then applied to all locations, while endo inputs are applied to locations 2 and 4 for 100 ms. Sustained endo inputs dominate, predicting divided attention in Awh and Pashlers’ Experiment 1. (B) Exo and endo inputs at locations 2 and 4 only (no array noise). The CANN incorrectly predicts sharply divided attention in Awh and Pashlers’ Experiments 4/4a.

between the two adjacent attended locations, tailing off on either side per the gradient model. Our choice of parameters leads to a steep curve, such that the instructed locations are poorly attended. This effect can be manipulated by adjustments to σ_w and σ_{ext} , but is more readily overcome by the use of a stiffer gain function (eg Sigmoid). In this case, output becomes flatter and sharper, predicting less of an attentive gradient.

Our choice of gain function also effects Simulation 2. Our modelling of adjacent trials reveals a dip in activity between foci, suggesting that the subjects of Müller et al may have divided their attention between adjacent locations. Because Müller et al didn’t test subjects’ attention between adjacent stimuli, this effect doesn’t conflict with their results, but it further exposes the model’s parameter dependence. Flatter bubbles would cancel the effect. The model predicts that attention should be easier to divide as foci become more distant (within a reasonable visual area). With stiffer gain, the model still predicts divided SVA under sustained input, but in all-or-none fashion. As such, the model no longer resembles the AOG model shown in Figure 3.

A second parameter that effects our simulations is the

strength of connectivity in the network. By reducing the ratio of input strength to connection strength from 1 to 0.1, connectivity dominates sustained input and the WTA nature of the model re-emerges. In this case, model behaviour under sustained input resembles behaviour under transient input and a flatter bubble may again be desirable.

Despite the difficulties raised by these parameter dependencies, they suggest an encouraging symbiosis between the model and behavioural and physiological experiments. Results of Simulation 1 predict that in adjacent trials, attention should be maximal between instructed locations in Müller’s experiments, whereas Simulation 2 predicts divided attention in adjacent trials. Results of Simulation 2 predict attention should be easier to divide in a 1+4 trial than in 1+3 and 2+4 trials. These predictions could readily be tested with a probe stimulus in the former case, and by adding a 1+4 test in the latter. Their results may help tune the parameters.

Regardless of configuration issues, results of Simulations 2 and 3 show that the CANN model is able to account for divided SVA under sustained inputs. As such, we believe that divided SVA may be possible for as long as endo and/or exo signals are provided to PP, and that differences in behavioural findings may reflect differences in experimental conditions rather than subjects’ ability to divide their attention. This conjecture echoes that of Schneider [24] that different experimental paradigms may facilitate measurements of different attention-related phenomena.

Figure 3A shows that the more our inputs overlap, the more readily a single bubble dominates. By increasing the spatial distance between locations, we reduce the overlap and facilitate multiple bubbles. By extending this principle to feature space, we conjecture that the overlap between integrated object representations should increase as a function of their number of common features. For example, neural representations of a red circle and a red square should overlap more than a red circle and a blue square, and attention should be more difficult to divide between them.

Our focus has been on stable attractor states in this paper. Transitions between dynamic regimes tend to be rapid, and given the large number of parameters that effect the model, stable states provide a better foundation for our simulations. As we show in [20], parameter adjustments effect transitions between regimes. A possible explanation of the findings of Müller et al is that divided SVA on split trials corresponds to the transition between two-bubble and one-bubble states in a WTA model. If so, subjects should be able to divide attention for the duration of these meta stable states. Given more time, subjects may have been unable to divide their attention in Müller’s experiments.

Awh and Pashlers’ experiments are difficult to explain with the CANN model. Awh and Pashler found that division of SVA was greatly reduced following removal of noise surrounding target stimuli. We believe the model can account for this result in two ways. Firstly, Awh and Pashlers’ methods of probing subjects’ attention may have interfered with its spatial allocation. Secondly, if the overlap between inputs is sufficiently

large to create a single winning bubble under sustained inputs, inhibitory inputs may be able to divide the bubble, accounting for results of all experiments in their study. One interpretation of such negative input is that subjects were ‘unattending’ to the noise. We believe that purposefully ‘unattending’ intervening locations is a possible strategy for dividing SVA, rather than purposefully attending to locations of interest. In Awh and Pashlers’ experiments, it’s possible that subjects were able to ignore invalid locations only when there was something to be ignored. Removal of noise may have removed their means of doing so. The use of negative inputs further suggests the possibility that such activity could be involved in inhibition of return (IOR), whereby the currently selected location in a saliency map receives or produces inhibitory activity, allowing visual search to move to other salient locations.

Because we interpret endo inputs as subjects’ representations of task instructions in WM, our model predicts that interference with STM should abolish split attention in both Müller’s and Awh and Pashlers’ experimental conditions. This prediction could be tested in a dual task paradigm. By equating task instructions in behavioural studies with STM representations in WM, and by modelling these representations as sustained endo inputs to the CANN, we revisit the relationship between WM and attention. Attention has often been cited as the primary constraint on WM capacity [25], but here we view WM representations as the driving force behind attention.

V. CONCLUSIONS AND FUTURE WORK

The model of SVA by Deco et al [9] implements a saliency map in PP with a CANN network. This instantiation of biased competition [10] integrates BU and TD influences in a biologically realistic computational architecture. Our simulations test this promising model on its ability to explain behavioural and physiological evidence on the spatial distribution of SVA.

Our results demonstrate that CANNs provide a model of spatial attention in PP capable of explaining divergent experimental findings. With transient inputs, the model’s WTA nature predicts a unitary attentional focus. With sustained inputs, the model accounts for divided SVA. As such, our predictions depend on the nature of exo and endo signals in attentive phenomena. Here, the use of sustained inputs replicates the findings of Müller et al [19]. The findings of Awh and Pashler [18] are difficult to reconcile with the CANN model, but they suggest a role for inhibitory mechanisms that may be modelled by negative inputs. The possibility that negative inputs provide a mechanism to simulate IOR in the model warrants further research.

The interplay between WM and SVA is paramount to our model. If divided attention is facilitated by STM representations providing endo inputs to PP, then disruption of STM should abolish divided attention. We believe further research in this area would improve our understanding of the relationship between WM and attention.

ACKNOWLEDGMENT

This work was supported in part by the NSERC grant RGPIN 249885-03.

REFERENCES

- [1] M. Posner, C. Snyder, and B. Davidson, "Attention and detection of signals," *Journal of Experimental Psychology, General*, vol. 109, pp. 160–174, 1980.
- [2] E. Cherry, "Some experiments on the recognition of speech, with one and with two ears," *Journal of Acoustical Society of America*, vol. 25, pp. 975–979, 1953.
- [3] C. Zelano, M. Bensafi, J. Porter, J. Mainland, B. Johnson, E. Bremner, C. Telles, R. Khan, and N. Sobel, "Attentional modulation in human primary olfactory cortex," *Nature*, vol. 8, pp. 114–120, 2004.
- [4] J. K. Tsotsos, "On the relative complexity of active vs. passive visual search," *International Journal of Computer Vision*, vol. 7, pp. 127–141, 1992.
- [5] U. Neisser, *Cognitive Psychology*. New York: Appleton, Century, Crofts, 1967.
- [6] C. Koch and S. Ullman, "Shifts in selective visual attention: towards the underlying neural circuitry," *Human Neurobiology*, vol. 4, no. 4, pp. 219–227, 1985.
- [7] A. Treisman, "Feature binding, attention and object perception," *Proceedings of the Royal Society*, 1998.
- [8] J. Wolfe, "Guided search 2.0: a revised model of visual search," *Psychonomic Bulletin and Review*, vol. 1, no. 2, pp. 202–238, 1994.
- [9] G. Deco, O. Pollatos, and J. Zihl, "The time course of selective visual attention: theory and experiments," *Vision Research*, vol. 42, pp. 2925–2945, 2002.
- [10] J. Duncan and G. Humphreys, "Visual search and stimulus similarity," *Psychological Review*, vol. 96, pp. 433–458, 2002.
- [11] S. Shipp, "The brain circuitry of attention," *Trends in Cognitive Sciences*, vol. 8, no. 5, pp. 223–230, 2004.
- [12] L. Itti and C. Koch, "Computational modelling of visual attention," *Nature Reviews: Neuroscience*, vol. 2, pp. 194–203, 2001.
- [13] C. Downing and S. Pinker, *Attention and Performance*. Earlbaum, 1985, vol. 2, ch. The spatial structure of visual attention, pp. 171–187.
- [14] D. LaBerge and V. Brown, "Theory of attentional operations in shape identification," *Psychological Review*, vol. 96, pp. 101–124, 1989.
- [15] C. Eriksen and J. S. James, "Visual attention within and around the field of focal attention: A zoom lens model," *Perception and Psychophysics*, vol. 40, pp. 225–240, 1986.
- [16] P. McCormick, R. Klein, and S. Johnston, "Splitting vs. sharing focal attention: comment on castiello and umilta (1992)," *Journal of Experimental Psychology*, vol. 24, no. 1, pp. 350–357, 1998.
- [17] S. Hahn and A. Kramer, "Further evidence for the division of attention between non-contiguous locations," *Visual Cognition*, vol. 5, pp. 217–256, 1998.
- [18] E. Awh and H. Pashler, "Evidence for split attentional foci," *Journal of Experimental Psychology, Human Perception and Performance*, vol. 26, no. 2, pp. 834–846, 2000.
- [19] M. Muller, P. Malinowski, T. Gruber, and S. Hillyard, "Sustained division of the attentional spotlight," *Nature*, vol. 424, 2003.
- [20] T. P. Trappenberg and D. I. Standage, "Multi-packet regions in stabilized continuous attractor networks," *Neurocomputing, IN PRESS*, 2004.
- [21] S. Deneve, A. Pouget, and P. E. Latham, "Divisive normalization, line attractor networks and ideal observers," *Advances in Neural Processing Systems*, vol. 11, pp. 104–110, 1999.
- [22] S. Funahashi, C. Bruce, and P. Goldman-Rakic, "Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex," *Journal of Neurophysiology*, vol. 61, pp. 331–349, 1989.
- [23] C. L. Colby and M. E. Goldberg, "Space and attention in parietal cortex," *Annu. Rev. Neurosci.*, vol. 22, pp. 319–349, 1999.
- [24] W. X. Schneider, "An introduction to 'mechanisms of visual attention: a cognitive neuroscience perspective'," *Visual Cognition*, vol. 5, pp. 1–8, 1998.
- [25] N. Cowan, "The magical number 4 in short-term memory: A reconsideration of mental storage capacity," *Behavioral and Brain Sciences*, vol. 24, pp. 87–114, 2001.