Emergence of Attention within a Neural Population

Nicolas P. Rougier, Julien Vitay

► To cite this version:

Nicolas P. Rougier, Julien Vitay. Emergence of Attention within a Neural Population. Neural Networks, 2005. inria-00000143

HAL Id: inria-00000143

<https://hal.inria.fr/inria-00000143>

Submitted on 5 Jul 2005

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Emergence of Attention within a Neural Population

Nicolas P. Rougier and Julien Vitay

Loria laboratory

Campus Scientifique, B.P. 239

54506 Vandœuvre-lès-Nancy Cedex, France

Corresponding author

Dr. Nicolas P. Rougier

e-mail: Nicolas.Rougier@loria.fr

fax: (33) 3 83 41 30 79

Abstract

We present a dynamic model of attention based on the Continuum Neural Field Theory that explains attention as being an emergent property of a neural population. This model is experimentally proved to be very robust and able to track one static or moving target in the presence of very strong noise or in the presence of a lot of distractors, even more salient than the target. This attentional property is not restricted to the visual case and can be considered as a generic attentional process of any spatio-temporal continuous input.

Keywords. Attention, CNFT, dynamic neural fields, lateral interactions.

Emergence of Attention within a Neural Population

Introduction

The cortex has long been known for being a massively interconnected structure of elementary processing elements (the so-called cortical columns, see (Burnod, 1989) for further details) benefiting from a structural two dimensional topology ascribed in the two dimensional topology of the cortical sheet itself. Furthermore, along this structural topology, there exists also a topographical organization such that response properties of neurons in many sensory cortical areas are ordered such that nearby neurons tend to respond to nearby areas of the input. These topographic maps form themselves by the self-organization of afferent connections to the cortex which are driven by external input (Hubel & Wiesel, 1965; Malsburg, 1973; Miller, Keller, & Stryker, 1989).

Several theories together with their associated neural network models have demonstrated how such an organization can emerge from a local competition based on lateral interactions within the cortex (Takeuchi & Amari, 1979; Amari, 1980; Kohonen, 1982). Those models have been primarily based on predetermined lateral interactions, focusing on the learning of afferent synaptic weights. Generally, these models rely on a Winner Take All (WTA) or a k-WTA algorithm to model lateral interactions. It helps both competition and numerical simulation in term of speed. Nonetheless, a number of recent neurobiological studies (Gilbert & Wiesel, 1990) have pinpointed the importance of lateral interactions and showed that cortico-cortical connections indeed change throughout development (Katz & Callaway, 1992). Based on these studies, (Sirosh & Miikulainen, 1993, 1997; Miikulainen, Bednar, Choe, & Sirosh, 1997) have designed a self-organizing neural network model for the simultaneous and cooperative development of topographic receptive fields and lateral interactions in cortical maps that numerically demonstrates how the famous mexican hat pattern of connectivity

can develop itself through unsupervised learning.

But, if these models were able to explain to some extent some observations on the development of both afferent and lateral connections in cortical feature maps, they did not exploit the dynamic aspect of neurons as it has been originally introduced by (Wilson & Cowan, 1973; Amari, 1977). The Continuum Neural Field Theory (CNFT) has been extensively analyzed both for the one-dimensional case (Wilson & Cowan, 1973; Feldman & Cowan, 1975; Amari, 1977) and for the two-dimensional case (Taylor, 1999) where much of the analysis is extendable to higher dimensions. Those theories explain the dynamic of pattern formation for lateral-inhibition type homogeneous neural fields with general connections. They show that, in some conditions, continuous attractor neural networks are able to maintain a localised bubble of activity in direct relation with the excitation provided by stimulation.

We investigate further these theories in order to experimentally study functional properties of the CNFT and show how it is indeed tightly linked to attention defined as the capacity to attend to one stimulus in spite of noise or distractor. Attention has a long history and complex meaning in psychology. As (James, 1890) said:

Everyone knows what attention is. It is the taking possession of the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration of consciousness are of its essence. ...

In the light of the proposed experiments, we show that bottom-up (i.e. stimulus driven) attention may be seen as an emergent property of a neural population using the Continuum Neural Field Theory. From a pool of neurons spread over two maps, one input map feeding a focus map, a bubble of activity emerges within the focus map at the precise location of a stimulus presented within the input map. This could be easily interpreted as the recognition of the location of the sensory input if it was not for noise and distractors. When noise or distractors are added, the bubble of activity stay focused on the original focused stimulus and then, between “several simultaneously

possible objects”, the model is able to “attend” to the one stimulus it first focused.

The model

Some related works (Hamker & Gross, 1997; Backer & Mertsching, 2002) have already used dynamic neural fields in the framework of attentional control and showed for example how they can be used for vision. We would like to propose a more systematic study by considering the most simple model (where a single map is laterally connected) and experimentally describe how and why attention naturally emerges from this model.

Continuum Neural Field Theory

We will use the notations introduced by (Amari, 1977) where a neural position is labelled by the vector $\mathbf{x}$ which represents a two-component quantity designing a position on a manifold M in bijection with $[-0.5, 0.5]^2$. The membrane potential of a neuron at the point $\mathbf{x}$ and time t is denoted by $u(\mathbf{x}, t)$. It is assumed that there is lateral connection weight function $w(\mathbf{x} - \mathbf{x}')$ which is in our case a difference of Gaussian function (DoG) as a function of the distance $|\mathbf{x} - \mathbf{x}'|$. There exists also an afferent connection weight function $s(\mathbf{x}, \mathbf{y})$ from the position $\mathbf{y}$ in the manifold M' to the point $\mathbf{x}$ in M . The membrane potential $u(\mathbf{x}, t)$ satisfies the following equation (1):

$$\begin{aligned} \tau \frac{\partial u(\mathbf{x}, t)}{\partial t} = & -u(\mathbf{x}, t) + \int_M w_M(\mathbf{x} - \mathbf{x}') f[u(\mathbf{x}', t)] d\mathbf{x}' \\ & + \int_{M'} s(\mathbf{x}, \mathbf{y}) I(\mathbf{y}, t) d\mathbf{y} + h \end{aligned} \quad (1)$$

where f represents the mean firing rate as some function of the membrane potential u of the relevant cell, $I(\mathbf{y}, t)$ is the output from position $\mathbf{y}$ at time t in M' and h is the neuron threshold. w_M is given by the equation (2).

$$w_M(\mathbf{x} - \mathbf{x}') = A e^{\frac{|\mathbf{x} - \mathbf{x}'|^2}{a^2}} - B e^{\frac{|\mathbf{x} - \mathbf{x}'|^2}{b^2}} \text{ with } A, B, a, b \in \mathbb{R}^{*+} \quad (2)$$

Furthermore, we use a Gaussian function for afferent connections as in equation (3).

$$s(\mathbf{x}, \mathbf{y}) = C e^{\frac{|\mathbf{x} - \mathbf{y}|^2}{c^2}} \text{ with } C, c \in \mathbb{R}^{*+} \quad (3)$$

Finally, and depending on the nature of the manifold M we consider (respectively a plane or a sphere surface), we can respectively use the Euclidean distance or the curve distance (which is defined as the shortest length of the geodesic between two points).

Discretization

In order to be able to perform numerical simulations using neural network models, we have to discretize these equations. We denote by n the discretization level which represents the regular segmentation of the interval $[-.5, .5]$ into n segments of size $1/n$. A manifold M can consequently be discretized as a set of $n \times n$ units and previous neural position $\mathbf{x}$ can be denoted $\mathbf{x}_{ij}$ with $i, j \in [0, n - 1]^2$. The corresponding neuronal position is now given by equation (4)

$$\mathbf{x}_{ij} = \left(\frac{i}{n} - 0.5, \frac{j}{n} - 0.5 \right) \quad (4)$$

and equation (1) becomes:

$$\begin{aligned} \tau \frac{du(\mathbf{x}_{ij}, t)}{dt} = & -u(\mathbf{x}_{ij}, t) + \sum_{k,l} w_M(\mathbf{x}_{ij} - \mathbf{x}'_{kl}) f[u(\mathbf{x}'_{kl}, t)] d\mathbf{x}'_{kl} \\ & + \sum_{k,l} s(\mathbf{x}_{ij}, \mathbf{y}_{kl}) I(\mathbf{y}_{kl}, t) d\mathbf{y}_{kl} + h \end{aligned} \quad (5)$$

Furthermore, In order to avoid any side effects due to the lack of connectivity along the edges of a map, we project the manifold M onto a torus in order to use a curve distance d that is defined by equation (6).

$$\begin{aligned} |\mathbf{x}_{ij} - \mathbf{x}'_{kl}| = & \min \left(\left(\frac{i-k}{n} \right)^2, \left(1 - \frac{i-k}{n} \right)^2 \right) \\ & + \min \left(\left(\frac{j-l}{n} \right)^2, \left(1 - \frac{j-l}{n} \right)^2 \right) \end{aligned} \quad (6)$$

One can observe on Figure 1 the impact of projecting the map onto a torus surface using the curve distance versus projecting onto a plane using the Euclidean distance.

Architecture

The model we designed is made of two maps *input* and *focus*, each of them being of size $n \times n$ units. Map *input* corresponds to an entry that is feeding the *focus*

map as illustrated on Figure 2 while *focus* map represents a cortical layer whose units possess localized receptive fields on the surface of the input. In other words, each unit $\mathbf{x}_{ij}$ of map *focus* receives its input from the *input* map using equation 3 which corresponds to a localized receptive field, being more or less broad depending on constant c . The *input* map does not have any lateral interaction nor feedback while each unit in the *focus* map is laterally connected using a difference of Gaussian (see Appendix A for implementation details).

This architecture, as simple as it stands, implements the most rudimentary form of attention that allows a model to focus on one static or moving stimulus without being distracted by noise or distractors, even more salient ones. We will now experimentally demonstrate this attentional property.

Asynchronous evaluation

As we stated before, the CNFT relies on a continuous evaluation of both lateral and afferent connections that result in one or more localized bubble of activities, depending on some initial conditions and profile of lateral interactions. In the following experiments, we are primarily interested in having a single bubble of activity representing the position of an input stimulus. The problem in such a framework is that two stimuli of equal intensity and width may be presented within the input map, with no noise or distractor. Furthermore, if we suppose that the CNFT map starts from zero activity, the question is then, where the localized bubble of activity will emerge ? If we use a discretized synchronous evaluation of units within the CNFT map, and depending on the relative position of the two stimuli, the answer is *nowhere or in the middle* while if we use asynchronous evaluation, the answer is *on one of the two stimuli*.

Synchronous evaluation refers to a well known algorithm used in the neural networks community where evaluation of activity of a unit u at time t is performed using stored information at time $t - 1$. Using such an algorithm for lateral interaction evaluations is a source of problem in the example cited above because in this case, two

bubbles compete to emerge while trying to inhibit each other. None of these bubbles has an advantage on the other since we considered noiseless input. This results in an oscillatory symmetric behavior where each of the two bubbles starts to emerge and is immediately inhibited by the other one. Once inhibition is weak enough, the two bubbles will re-emerge and will be immediately re-inhibited, etc. The reason for this behaviour is a lack of dissymmetry in the network that should be normally provided by non-uniform noise, giving the necessary dissymmetry to the network. We have experimentally tested this hypothesis and showed that even a very small amount of noise is able to break the symmetry.

Another solution is the asynchronous evaluation where evaluation synchronicity is broken using a random evaluation order. In this case, at each time step, a unit is randomly chosen and evaluated using information available at this time. A computational step corresponds in this case to n successive evaluations.

Experiments and results

As we stated before, the goal of the model is to implement a very basic attentional apparatus (embedded in a single map) and to propose that attention may be thought as an emergent property of a neural population. Consequently, we define a target as a spatially localized stimulus onto an *input* map that is feeding the *focus* map which realizes the attentional function. In order to realize such a function, the *focus* map should then be able to remain focused on the target in spite of noise, distractors or movement of target.

Encoding

Mean input activity $S_{r,\theta,W,I}$ follows a bell-shaped profile with height proportional to contrast. A stimulus $s_{r,\theta,W,I}$ is then characterized by the tuple (r, θ, W, I) corresponding to a Gaussian profile whose center is localized at

$(r\sin\theta, r\cos\theta)$ of width W and intensity I given by equation (7).

$$s_{r,\theta,W,I}(x,y) = I e^{\frac{(x-x_c)^2}{W^2}} e^{\frac{(y-y_c)^2}{W^2}} \text{ with } (x_c, y_c) = (r\sin\theta, r\cos\theta) \quad (7)$$

Using such a symmetric function about both x -axis and y -axis yields an interesting decoding property given by equation (8)

$$\forall s/\forall x, s(x) = s(-x) \Rightarrow \forall x_c, x_c = \frac{\int_{-\infty}^{\infty} x s(x - x_c) dx}{\int_{-\infty}^{\infty} s(x - x_c) dx} \quad (8)$$

Translated in the discrete case and considering a discretized manifold M_n (in bijection with $[-.5, .5]^2$) whose value at position $\mathbf{x}_{i,j}$ is given by $a(i, j)$, we can get an approximation of (x_c, y_c) with equation (9).

$$(\hat{x}_c, \hat{y}_c) = \left(\frac{\sum_{i,j} \frac{i}{n} a(i, j)}{\sum_{i,j} a(i, j)} - 0.5, \frac{\sum_{i,j} \frac{j}{n} a(i, j)}{\sum_{i,j} a(i, j)} - 0.5 \right) \quad (9)$$

Furthermore, noise is added at each neural position and is assumed to be independent. It follows a zero-mean Gaussian distribution whose variance is fixed at different levels (see Figure 3). Finally, values are clipped in the range $[0, 1]$ implying that addition of noise results in a non zero-mean signal.

Static stimulus

There exist several models using population codes focusing on noise clean-up such as (Douglas, Koch, Mahowald, Martin, & Suarez, 1995; Deneve, Latham, & Pouget, 1999) or more general types of computation such as sensorimotor transformations, feature extraction in sensory systems, motion perception or multisensory integration (Giese, 1999; Wu, Nakahara, & Amari, 2001; Zhang, 1996; Deneve, Latham, & Pouget, 2001; Stringer, Rolls, & Trappenberg, 2004). (Deneve et al., 1999) were able to show through analysis and simulations that it is indeed possible to implement an ideal observer using biological plausible model of cortical circuitry and it comes as no surprise that this model relies heavily on lateral interactions. The model we designed also relies heavily on lateral interactions, as dictated by the CNFT, and fall into the more general case of *recurrent network whose activity relaxes to a*

smooth curve peaking at a position that depends on the encoded variable that was analyzed as being a good implementation of a Maximum Likelihood approximator (Deneve et al., 1999).

Our experimental approach is different since we do not consider an experiment to be a sum of isolated trials but rather consider the temporal nature of stimuli succession. Consequently, there is not such thing as a “reset” of the activity in the model between each trials. The experimental protocol is the following:

1. A single stimulus (without noise or distractor) is clamped to the *input* map.
2. Noise or distractors are added
3. 10 steps of computation are performed within *focus* map.
4. Position of stimulus is recorded and we re-iterate steps 1 to 4.

There is also an initialization procedure where we let the model first converge (equivalent to 3 steps of computation) on the single stimulus present within the *input* map.

As stated before, we use a stimulus with a bell-shaped profile located at a fixed position (x_c, y_c) and we use different levels of Gaussian noise and different numbers of distractors. As illustrated on Figure 4, the model is able to quite accurately track the stimulus position in spite of an important level of noise or an important number of distractors. In the case of distractors, it is important to understand that it is not possible to decide what is the position of the target based on one trial since distractors have the exact same profile as the stimulus (see Figure 5). The only “solution” to the problem is to perform an attentional process where attention is focused on the same “stimulus”, the only one having an observable spatio-temporal continuity.

Moving stimulus

Using the same protocol as in static experiments, we tested the model against a moving target evolving around a circular path and we keep track of the decoded position of the activity bubble within the *focus* map. One can see in figure 6 the

resulting path decoded from the bubble of activity in the *focus* map. The speed of the moving target is a critical parameter on these experiments since it is directly related to the apparent spatial continuity of the target which is observable (or not) by the model. For example, in presented results, θ angle was increased every ten steps of computation by an amount of 3 degrees. These 10 steps of computation correspond roughly to the time needed for a bubble of activity to *move* from one position to another near one. If the new position is too far from the previous one (undersampling), the bubble of activity cannot *move* toward it and simply vanishes to let another bubble of activity emerge some place else. In such a case, the attentional property cannot be guaranteed, i.e. the new bubble can emerge at the new position of the target but it can also emerge at the position of a distractor. Nonetheless, when the sampling is performed in such a way that the continuity of the movement of the stimulus is observable by the model, the bubble of activity is able to *move* to the new neighborhood position because the competition is biased toward this new position that is both fed by input and some lateral excitation.

Conclusion

A dynamic model of attention has been described using the Continuum Neural Field Theory that explains attention as being an emergent property of a neural population. Using distributed and iterative computation, this model has been proved very robust and to be able to track one static or moving target in the presence of noise with very high intensity or in the presence of a lot of distractors, possibly more salient than the target. The main hypothesis concerning target stimulus is that it possesses a spatio-temporal continuity that should be observable by the model, i.e. if the movement of the target stimulus is too fast, then the model can possibly loose its focus. Nonetheless, this hypothesis makes perfect sense when considering *real world* robotic applications. We have been able to succesfully implement this simple model on a robot watching perfectly identical targets and it revealed itself able to focus on the

first presented target and to remain focused on it, even when other targets were added or removed from the perceived scene or when any of them were moved (including the target). Nevertheless, and as model stands, one can object that this model is not able to switch attention between available stimulus. The reason is that we wanted to introduce one of the most simple model able to exhibit some kind of early attention. We have now extended the basic model as to implement attentional switch between relevant object and successfully implemented it on a real robot (Vitay, Rougier, & Alexandre, 2005). The robot revealed itself able to scan successively different identical and moving targets without ever focusing twice on the same target.

Finally, attention as it has been introduced in this work and implemented in the model is not restricted to visual attention. Provided there exists some map with some coherent bubbles of activity, a focus map can be used to attend to one or the other bubble. This may shed a new light on prefrontal cortex and working memory where it would become highly dynamic.

References

- Amari, S. (1977). Dynamic of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27, 77–78.
- Amari, S. (1980). Topographic organization of nerve fields. *Bulletin of Mathematical Biology*, 42, 339–364.
- Backer, G., & Mertsching, B. (2002). Using neural field dynamics in the context of attentional control. In *Icann 2002* (pp. 1237–1242).
- Burnod, Y. (1989). *An adaptive neural network: the cerebral cortex*. Masson.
- Deneve, S., Latham, P., & Pouget, A. (1999). Reading population codes: a neural implementation of ideal observers. *Nature Neurosciences*, 2, 740–745.
- Deneve, S., Latham, P., & Pouget, A. (2001). Efficient computation and cue integration with noisy population codes. *Nature Neuroscience*, 4(8), 826–831.
- Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A., & Suarez, H. H. (1995). Recurrent excitation in neocortical circuits. *Science*, 269, 981–985.
- Feldman, J., & Cowan, J. (1975). Large-scale activity in neural nets. i. theory with applications to motoneuron pool responses. *Biological Cybernetics*, 17, 29–38.
- Giese, M. (1999). *Dynamic neural field theory for motion perception*. Kluwer Academic Publishers.
- Gilbert, C. D., & Wiesel, T. (1990). Lateral interactions in visual cortex. In C. S. H. L. Press (Ed.), *Cold spring harbor symposia on quantitative biology* (Vol. LV, pp. 663–677).
- Hamker, F. H., & Gross, H. M. (1997). Object selection with dynamic neural maps. In *Icann 1997* (pp. 919–924).
- Hubel, D., & Wiesel, T. (1965). Receptive fields and functional architecture in two non-striate visual areas (18 and 19) of the cat. *Journal of Neurophysiology*, 28, 229–289.

- James, W. (1890). *The principles of psychology*. New York: Holt.
- Katz, L. C., & Callaway, E. M. (1992). Development of local circuits in mammalian visual cortex. *Annual Review of Neurosciences*, 15, 31–56.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43, 59–69.
- Malsburg, C. von der. (1973). Self-organization of orientation-sensitive cells in the striate cortex. *Kybernetic*, 15, 85–100.
- Miikulainen, R., Bednar, J. A., Choe, Y., & Sirosh, J. (1997). Self-organization, plasticity, and low-level visual phenomena in a laterally connected map model of the primary visual cortex. *Psychology of Learning and Motivation*, 36, 257–308.
- Miller, K. D., Keller, J. B., & Stryker, M. P. (1989). Ocular dominance column development: analysis and simulation. *Science*, 245, 605–615.
- Sirosh, J., & Miikulainen, R. (1993). How lateral interaction develops in a self-organizing feature map. In *Proceedings of the ieee international conference on neural networks*.
- Sirosh, J., & Miikulainen, R. (1997). Topographic receptive fields and patterned lateral interaction in a self-organizing model of the primary visual cortex. *Neural Computation*, 577–594.
- Stringer, S. M., Rolls, E. T., & Trappenberg, T. P. (2004). Self-organising continuous attractor networks with multiple activity packets, and the representation of space. *Neural Networks*, 17, 5–27.
- Takeuchi, A., & Amari, S. (1979). Formation of topographic maps and columnar microstructures. *Biological Cybernetics*, 35, 63–72.
- Taylor, J. G. (1999). Neural bubble dynamics in two dimensions: foundations. *Biological Cybernetics*, 80, 5167–5174.
- Vitay, J., Rougier, N. P., & Alexandre, F. (2005). A distributed model of spatial visual

attention. In S. Wermter & G. Palm (Eds.), *Neural learning for intelligent robotics*. Springer-Verlag.

Wilson, H. R., & Cowan, J. D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, *13*, 55–80.

Wu, S., Nakahara, H., & Amari, S. (2001). Population coding with correlation and an unfaithful model. *Neural Computation*, *13*, 775–797.

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience*, *16*, 2112–2126.

Appendix A

A

Using equations (2), (3) and equation (5), simulation parameters are

$$n = 30, \tau = .75$$

$$A = \frac{1.4}{\alpha}, a = \frac{5}{n}, B = \frac{0.65}{\alpha}, b = \frac{17}{n}, C = \frac{1}{\alpha}, c = 0.1 \text{ with } \alpha = 13$$

Appendix B

B

Figures B1 and B2 are two screenshots from simulations displaying focus profile in the presence of noise or distractors. Demonstration movies can be downloaded from <http://www.loria.fr/~rougier/research/demos.html>

Author Note

This work is supported by the MirrorBot european project and the Robéa CNRS French initiative

Figure Captions

Figure 1. Lateral connectivity pattern is a simple difference of Gaussian function (DoG) between a sharp positive Gaussian function and a wider negative one with different intensity and same center. The profile of the DoG is the same for every unit in a map and drives the global activity profile of the whole map. The distance used (Euclidean or curve) depends on the type of projection of the manifold M . On both (a) and (b), lateral weights have been drawn for unit at position $(-.3, -.3)$. On (a) the projection has been made onto a plane and the Euclidean distance has been used whereas on (b), the projection has been made onto a sphere surface and the curve distance has been used.

Figure 2. The model is made of two maps of $n \times n$ units each ($n = 30$ on figure). The “input” map receives its inputs from an external moving stimulus that evolves along a circular trajectory and whose center corresponds more or less to the center of the map. The “focus” map receives its inputs from the “input” one, using a one-to-one connection pattern. On the example displayed, the “focus” map has settled itself on a pattern of activity that is representing the actual input.

Figure 3. Input is a bell-shaped curve centered around (x_c, y_c) representing an external stimulus. Noise is assumed to be independent and to follow a zero-mean Gaussian distribution whose variance has been set to different values: (a) noiseless input (b) variance is .1 (c) variance is .25 (d) variance is .5 (e) variance is .75 (f) variance is 1.0. All input values are clipped within interval $[0, 1]$ implying that a variance of 1 is not equivalent to a signal-noise ratio of 1.

Figure 4. Every 10 steps of computation, the position s of a static target has been decoded in both *input* (s_I) and *focus* (s_F) map. Distances $|s - s_I|$ and $|s - s_F|$ have been used as measures of error and are reported here (each plotted figures is an average over 1000 trials). On figure (a), a zero-mean Gaussian noise with various intensities

has been added to the stimulus. Clearly, the *focus* map is able to accurately extract the original position of the stimulus. On figure (b), zero to 25 distractors (with same width and intensity as the original stimulus) were added in *input* and *focus* map is also able to accurately extract the original position. Error within input map (with presence of noise and distractor) have been plotted as an element of comparison.

Figure 5. Figure (a) represents a moving noiseless stimulus without any distractors. Figure (b) represents a moving noiseless input with 10 distractors. Without considering the spatio-temporal nature of the stimulus, it is not possible to decide where is the target in Figure (b).

Figure 6. A moving target is evolving around a circular path within *input* map and the position of the bubble of activity is decoded within *focus* map at each time step. Figures present the interpolated path (a line is drawn between two successive position) for different intensity of noise ((a) 0, (b) 0.1, (c) 0.25, (d) 0.5, (e) 0.75 and (f) 1). Even with a noise of intensity one, the model is able to track the moving target along its circular path.

Figure B1. Screenshot from the simulation showing an input with a level of .5. The bubble of activity within the focus map is still focused on the original stimulus.

Figure B2. Screenshot from the simulation showing an input with 10 distractors added. The bubble of activity within the focus map is still focused on the original stimulus.

(a)

(b)

(c)

(d)

(e)

(f)

(a)

(b)

(a)

(b)

(a)

(b)

(c)

(d)

(e)

(e)