



HAL
open science

Dynamic Neural Field with Local Inhibition

Nicolas P. Rougier

► **To cite this version:**

Nicolas P. Rougier. Dynamic Neural Field with Local Inhibition. *Biological Cybernetics (Modeling)*, 2006, 94 (3), pp.169-179. 10.1007/s00422-005-0034-8 . inria-00000244

HAL Id: inria-00000244

<https://inria.hal.science/inria-00000244>

Submitted on 16 Sep 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.

Dynamic Neural Field With Local Inhibition

Nicolas P. Rougier

Loria laboratory

Campus Scientifique, B.P. 239

54506 Vandœuvre-lès-Nancy Cedex, France

Fax: +(33) 383 413 079

e-mail: Nicolas.Rougier@loria.fr

Received: date / Revised version: date

Abstract A lateral-inhibition type neural field model with restricted connections is presented here and represents an experimental extension of the Continuum Neural Field Theory (CNFT) by suppression of the global inhibition. A modified CNFT equation is introduced and allows for a locally defined inhibition to spatially expand within the network and results in a global competition extending far beyond the range of local connections by virtue of diffusion of inhibition. The resulting model is able to attend to a moving stimulus in the presence of a very high level of noise, several distractors or a mixture of both.

Key words Dynamic Neural Field – CNFT – Local Inhibition – Diffusion

1 Introduction

The dynamics of pattern formation in lateral-inhibition type neural fields with global inhibition has been extensively studied in a number of works [Wilson and Cowan, 1973, Amari, 1977, Taylor, 1999, Giese, 1999] and these studies basically demonstrate that these fields are able to maintain a localized packet of neuronal activity that can for example represent the current state of an agent in a continuous space or reflect some sensory input feeding the field. Such networks most generally use excitatory recurrent collateral connections between the neurons as a function of the distance between them and global inhibition is used to ensure the uniqueness of the bubble of activity within the field. This uniqueness property is critical when it is used to either represent head direction [Zhang, 1996], place [Samsonovitch and McNaughton, 1997] or view cells [Stringer et al., 2005] and is no less critical when this kind of field is used in the context of attention where a unique localized bubble of activity is able to represent an external stimulus in spite of noise, distractors or saliency [Rougier and Vitay, 2005]. Those kinds of neural fields have been primarily inspired by the study of the cortex which has long been known for being a massively interconnected structure of elementary processing elements [Burnod, 1989] benefiting from a structural two dimensional topology ascribed in the two dimensional topology of the cortical sheet itself. Furthermore, the cortex is also known to be topographically organized and nearby neurons tend to respond to nearby areas of the input. Topographic maps form themselves

by the self-organization of afferent connections to the cortex which are driven by external input [Hubel and Wiesel, 1965, von der Malsburg, 1973, Miller et al., 1989] and the pioneer works of [Takeuchi and Amari, 1979, Amari, 1980, Kohonen, 1982] have demonstrated how such an organization can emerge from a local competition based on lateral interactions within the cortex. Resulting models generally use Winner Take All (WTA) or a k-WTA algorithms that aim at modelling lateral competition and thus implicitly use full lateral connectivity that allow for a unique winner to emerge.

However, this full connectivity, either implicit or explicit, is somehow problematic because the cortex, if richly connected, is not fully connected and it is the purpose of this paper to introduce a model that perform global competition (leading to the creation of a unique bubble of activity) by only using local excitation and inhibition without the use of any supervisor or central executive. This is made possible by using a diffusion of the inhibition throughout the network. This locality yields several advantages. First, in terms of pure computational power, it is far more quicker to have a few local interactions when computing activity within the network. Second, having real local and distributed computing make the model a real candidate for parallelization. And last, but not least, the use of diffusion makes the model scalable to virtually any size without any change in parameters. More precisely, lateral weights do not need to be adjusted for any particular size of the network since the travelling inhibition wave ultimately reaches any neurons within a map. After reminding the reader with basic equations

of the classical Continuum Neural Field Theory, proposed changes that allow for the diffusion phenomenon to emerge will be presented in details. The architecture of the model will then be introduced and experimentally compared with the fully connected version that has been experimentally and thoroughly tested in [Rougier and Vitay, 2005]. After studying the dynamic of the inhibition travelling wave and having underlined new properties implied for the model, the model will be related to known biological facts about the cortex and its elementary processing unit, the cortical column [Burnod, 1989].

2 The model

Modified CNFT equations as well as the accompanying model are presented in this in order to illustrate the principle of the diffusion of inhibition.

2.1 Continuum Neural Field Theory

Notations introduced by [Amari, 1977] are used and a neural position is labelled by the vector \mathbf{x} which represents a two-component quantity designating 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 as a function of the distance $|\mathbf{x} - \mathbf{x}'|$. There also exists an afferent connection weight function $s(\mathbf{x}_{M'}, \mathbf{x}_M)$

from the position $\mathbf{x}_{M'}$ in the manifold M' to the point \mathbf{x}_M in M . The membrane potential $u(\mathbf{x}, t)$ satisfies the following equation (1):

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

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' , h is the neuron threshold and α is a scaling term.

$w_M(\mathbf{x} - \mathbf{x}')$ is given by equation (2).

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

and afferent connections are given by equation (3).

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

Furthermore, activity of a neuron is bound between 0 and 1 using following conditions:

$$\begin{aligned} \text{if } u(\mathbf{x}, t) > 1, u(\mathbf{x}, t) &= 1 \\ \text{if } u(\mathbf{x}, t) < 0, u(\mathbf{x}, t) &= 0 \end{aligned} \quad (4)$$

2.2 Modified CNFT equation

As stated before, the goal of the proposed model is to implement the Continuum Neural Field Theory using a restricted and local set of connections and to ensure in the same time the uniqueness of the bubble of activity. The

major problem and the main question is how to make two distant neurons from the same map to reciprocally influence themselves knowing there is no direct connections between them ? The natural answer, that has been used and is at the essence of the vast majority of artificial neural networks, is to use information relays throughout the network. Those information relays, which are neurons themselves, are able to convey the required information under given or learned circumstances. However, if we consider the most generic form of the formal neuron equation:

$$u(\mathbf{x}) = \Phi\left(\sum_i w_i x_i\right) \quad (5)$$

where $u(\mathbf{x})$ denotes the membrane potential of a neuron \mathbf{x} receiving inputs from a set of neurons x_i with weights w_i . It appears quite evidently that the multiplicative nature of the terms of the sum prevents any neuron with a null activity to influence anything. This null activity is generally obtained thanks to the bounding of activity between 0 and some strictly positive value. In the end, only those “activated” neurons are able to communicate some information to their neighbours. In the present case, the problem is a bit different because there is a need to convey inhibition information to some distant neurons. If equation (1) is used for computing activation and clamp neuron activity between 0 and 1, then no inhibition can travel around: a null activated neuron is unable to influence its neighbours. The solution is then first to bound the neuron activity between a strictly negative value that represent the inhibited state and a strictly positive one that represent the excited state. Consequently, a first proposition is to replace conditions

(4) with the following conditions:

$$\begin{aligned} &\text{if } u(\mathbf{x}, t) > 1, u(\mathbf{x}, t) = 1 \\ &\text{if } u(\mathbf{x}, t) < -1, u(\mathbf{x}, t) = -1 \end{aligned} \tag{6}$$

These new conditions immediately bring some problems within equation (1) because an inhibited neuron (neuron with a strictly negative mean firing rate) having an inhibitory connection (strictly negative weight) with another neuron is now able to positively influence this last one. Intuitively, this does not make much sense and it should be fixed by simply preventing an inhibited neuron to propagate its activity through negative connections and equation (1) now becomes¹:

$$\begin{aligned} \tau \frac{\partial u(\mathbf{x}, t)}{\partial t} = & -u(\mathbf{x}, t) + h \\ & + \frac{1}{\alpha} \int_M w_M^+(\mathbf{x} - \mathbf{x}') f[u(\mathbf{x}', t)] d\mathbf{x}' \\ & + \frac{1}{\alpha} \int_M w_M^-(\mathbf{x} - \mathbf{x}') f^+[u(\mathbf{x}', t)] d\mathbf{x}' \\ & + \frac{1}{\alpha} \int_{M'} s(\mathbf{x}, \mathbf{y}) I(\mathbf{y}, t) d\mathbf{y} \end{aligned} \tag{7}$$

The biological signification of the proposed model of neuron is that it possesses an activation function that is bound between -1 to 1 while it is supposed to represent the mean firing rate. This negative mean firing rate does not make any sense and consequently the proposed model of neuron cannot claim any biological plausibility. However, there exists an equivalence of this model if each neuron is replaced by a system of two neurons, one being excitatory and the other being inhibitory. The condition for the strict

¹ For any real function f , $f^+ = \max(0, f)$, $f^- = \min(0, f)$

equivalence of that system is explained in detailed in annex B. Finally, this system may be related to cortical columns [Burnod, 1989] that can be described as a set of a several neurons that possesses some coherent activity between each other.

2.3 Architecture

The model itself is made of two neural 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 1 while *focus* map represents a cortical layer whose units possess very 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 correspond 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 and locally connected using a difference of Gaussian as illustrated on figure 2 (see section C for implementation details). As explained in [Rougier and Vitay, 2005], this architecture implements a rudimentary form of attention that allows the model to focus on one static or moving stimulus without being distracted by noise or distractors, even more salient ones.




Fig. 1 The model is made of two maps of $n \times n$ units each ($n=30$ on figure). The *input* map receives its input from an external moving stimulus that evolves along a circular trajectory whose center corresponds to the center of the map and radius to the equivalent of 10 units. On the displayed figure, the *focus* map has settled itself on a pattern of activity that is representing the actual input.

3 Experiments and results

There exist several models using population codes focusing on noise clean-up such as [Douglas et al., 1995, Deneve et al., 1999] or more general types of computation such as sensorimotor transformations, feature extraction in sensory systems, motion perception or multisensory integration [Giese, 1999, Wu et al., 2001, Zhang, 1996, Deneve et al., 2001, Stringer et al., 2004]. For example, [Deneve et al., 1999] were able to show through analysis and simulations that it is possible to implement an ideal observer using biological plau-




Fig. 2 Lateral connectivity pattern is a local 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. Lateral weights have been drawn for unit at position $(.4,.4)$. Furthermore, the manifold M is projected onto a torus and the distance used is the curve distance, defined as the shortest length of the geodesic between two points. This projection coupled with the curve distance prevent side effects along the border of the map where otherwise, there could be a lack of connection. This side effect is a pure modelling artefact of the finite nature of the map that implies possessing borders. If we consider a real brain, there is no such direct notion as “borders” and the toric projection is thus not required.

sible model of cortical circuitry. The presented model falls into the 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].

However, the experimental protocol that has beenj used is different since an experiment is not considered as being 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 is clamped to the *input* map.
2. Noise and/or distractors are added
3. 10 steps of computation are performed.
4. Position of stimulus is recorded.
5. Steps 1 to 5 are re-iterated.

There is an initialization procedure where the model is allowed to converge on the single stimulus present within the *input* map (this is equivalent to 3 steps of computation). Stimulus and distractors are defined as spatially localized gaussian-shaped activity profiles clamped in the *input* map and a zero mean Gaussian noise with a fixed variance can be added. The *input* map that is feeding the *focus* map with inpouts and this *focus* map is actually realizing the attentional function. In order to realize such a function, the *focus* map must be able to remained focused on the target in spite of noise,

distractors or movement of target. Finally, the performance of the network is measured as the distance between the position of the original stimulus in the “ideal” *input* map and the position of the encoded stimulus in the *focus* map.

3.1 Stimulus Encoding

Mean input activity $S_{r,\theta,W,I}$ of the stimulus follows a bell-shaped profile with height proportional to contrast. A stimulus $s_{r,\theta,W,I}$ is 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 (8).

$$s_{r,\theta,W,I}(x, y) = I e^{-\frac{(x-x_c)^2}{W^2}} e^{-\frac{(y-y_c)^2}{W^2}} \text{ with } (x_c, y_c) = (r \sin \theta, r \cos \theta) \quad (8)$$

Using such a symmetric function about both x -axis and y -axis yields an interesting decoding property given by equation (9)

$$\forall s/\forall x, s(x) = s(-x) \Rightarrow \forall x_c, x_c = \frac{\int_{-\infty}^{\infty} x s(x - x_c) dx}{\int_{-\infty}^{\infty} s(x - x_c) dx} \quad (9)$$

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 $a(i, j)$, a fairly good approximation of (x_c, y_c) is given by equation (10).

$$(\hat{x}_c, \hat{y}_c) = \left(\frac{\sum_{i,j} \frac{i}{n} a(i, j)}{\sum_{i,j} a(i, j)} - 0.5, \frac{\sum_{i,j} \frac{j}{n} a(i, j)}{\sum_{i,j} a(i, j)} - 0.5 \right) \quad (10)$$

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. Finally, input values are clipped in the range $[0, 1]$ such that addition of noise results in a non zero-mean signal.

3.2 Results

Figure 3 present results concerning performances in the presence of noise only or in the presence of distractors only. In the two cases, the model is able to accurately track the moving target quite accurately (maximum theoretical error is 1). But while overall performances are better in the presence of noise compared to the fully connected model, performances in the presence of an increasing number of distractors are first better and then very rapidly degrade. Figures for 10 or 25 distractors are not represented because in these situations, the model is no more able to track the target and performances are too much degraded (answer is equivalent to a random one). The fully connected model in comparison is able to track the target even in the presence of 10 or 25 distractors. It illustrates the fundamental difference between the two models: in one case, the action of inhibition is instantaneous (fully connected model) while in the other case (locally connected model), the propagation of inhibition takes some time and this is enough for a group of strong distractors to disturb the model.

Models have been further tested using noise **and** distractors (figure 4). Once again, results are comparable and of the same magnitude. Worst error cases for the locally connected model come from cases where 5 distractors

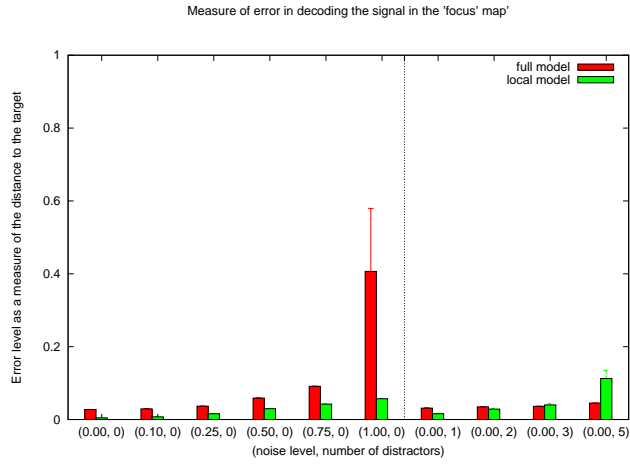


Fig. 3 Every 10 epochs, the position s of a moving target has been decoded in both *input* (s_I) and *focus* (s_F) map. Distance $|s_I - s_F|$ has been used as measures of error that is reported here (each plotted figures is an average over 1200 trials). Zero-mean Gaussian noise with various intensities has been added to the stimulus or some distractors (up to five) were added. Performances are drawn for a fully connected model and the locally connected model.

are simultaneously presents and this was already the worst case when there was no noise.

When distractors are present, 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. The only “solution” to the problem is to perform an attentional process where attention is focused on the same “stimulus” throughout time because this is the only one having an observable spatio-temporal continuity. In this sense, 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

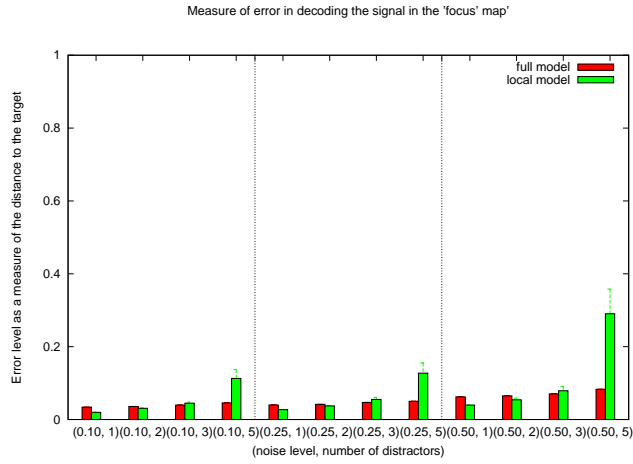


Fig. 4 Every 10 epochs, the position s of a moving target has been decoded in both *input* (s_I) and *focus* (s_F) map. Distance $|s_I - s_F|$ has been used as measures of error that is reported here (each plotted figures is an average over 1200 trials). Zero-mean Gaussian noise with various intensities has been added to the stimulus mixed with one to five. Performances are drawn for a fully connected model and the locally connected model.

observable (or not) by the model. 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 stim-

ulus 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.

4 Dynamic of the network

As explained previously, the uniqueness property of the bubble is ensured by the diffusion of the inhibition activity. This is clearly illustrated on figure 5 where the inhibition travelling wave is clearly shown from epoch 1 through epoch 40. One critical parameter of the model is the constant h (neuron threshold) that needs to be strictly positive. This corresponds to a spontaneous activity of the neuron when no input is present. This is rather counter-intuitive because the question that immediately arise concerns the behavior of the network when no input is present. In fact, in this case the network converges very quickly towards a fully inhibited state. This is a direct consequence of the local pattern of connectivity and the modified CNFT equations. Each neuron having a spontaneous activity, it tends to excite very local neighbours and inhibit more distant ones. Hence, and because of asynchronous evaluation (see Appendix A), some localized packets of excited neurons emerge while other “interneurons” activity goes below 0. But, and since those localized packets are not fed by any external input, they are very sensitive and cannot resist inhibition coming from those “interneurons” that have been previously inhibited (activity below 0). This is a subtle feedback mechanism where “winning” neurons induce their own “fate” by

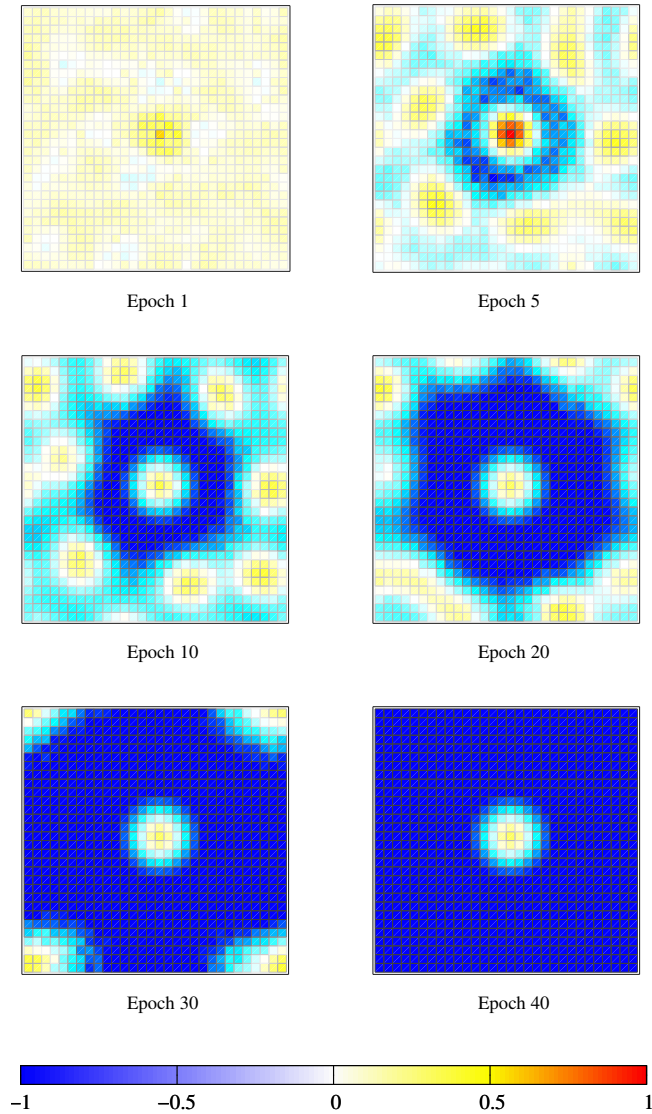


Fig. 5 Input is a bell-shaped stimulus centered on $(0,0)$. At epoch 1, all units of the focus map get some activation because of the positive baseline threshold that corresponds to a spontaneous activity. Four epochs later, neurons receiving input from the stimulus get a stronger activation and start to inhibit their direct neighbours. In the same time, some isolated packs of neurons appears because of lateral inhibition and excitation. At epoch 10, all direct neighbours have been inhibited while previous isolated packs of neurons get sharper and tend to resist propagating inhibition. At epoch 20, one can see that packs of neurons near the inhibition wave frontier have disappeared, favoring the creation of new pack of neurons at farther distance. Finally on epoch 40, the propagation of inhibition has reached all neurons that are not strong enough to resist and becomes inhibited. The only remaining pack of activated neuron is the one corresponding to the input stimulus.

winning the competition in the first place and get inhibited in turn by neurons they have inhibited. This is clearly illustrated on the figure 5 where the low-range neuron curve is made of several cycles of excitation/inhibition.

Another important aspect of the network is the hysteresis property that defines a system whose response depends not only on its current state, but also upon its past history. This is best illustrated when three stimuli are presented within the focus map. If the *focus* map was previously in a null state (all activity set to zero), three bubbles of activity are able to simultaneously form themselves without inhibiting each other. This can be explained by considering the inhibition propagation delay that makes the inhibition wave (expanding from one bubble) to hit other bubbles too late. In the meantime those bubble have reach the point where they cannot be inhibited anymore because they're too strong activated enough. Now, considering the case where a bubble was already present within both *input* and *focus* map, the addition of two stimuli within the *input* map will not induce the formation of two new bubble within the *focus* map because corresponding location are currently fully inhibited and burst of activity provided by new inputs are not strong enough to overcome this inhibition.

5 Conclusion

A lateral-inhibition type neural field model with restricted connections has been presented and represents an experimental extension of the Continuum Neural Field Theory. We proposed to slightly modify original CNFT equa-

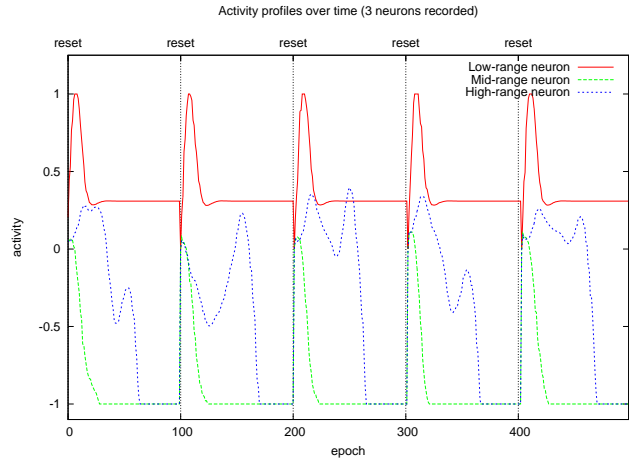


Fig. 6 Input is a bell-shaped stimulus centered on $(0,0)$. Three neuron activities have been recorded during 500 epochs. The low-range neuron is a neuron whose position within the *focus* map corresponds exactly to the center of the stimulus. The mid-range neuron is at a mid-distance position from the stimulus while the high-range neuron is at maximal distance from the stimulus (bottom-left position on the map). A reset of the network (clearing all computed values) is performed every 100 epochs in order to get several settling phases. The activity profile of the low-range neuron is quite characteristic of neurons which remain activated after the settling phase. There is first a burst of activity peaking at 1 until a decrease of activation because of lateral interactions. It is remarkable to see the brief undershoot period preceding the stable state. This undershoot is generally characteristic of neurons of type I. Mid-range neuron activity profile is best characterized by a very short period of activity increase until a more or less regular decrease until it reaches its final and fully inhibited state. High-range neuron activity profile is quite different because of its distant position from the center of the stimulus. More precisely and depending on local neighbourhood, its activity is able to go up and down several times (2 at most on the figure) until the inhibition wave reaches it. This is not possible for the chosen mid-range neuron which is nearer from the center of the stimulus and which is struck by the inhibition wave very quickly.

tions in order to take into account an “active” inhibitory state resulting in the new ability for the model to perform a global competition using only local connectivity, thanks to the diffusion of inhibition. This model is easily scalable since weights do not need to be changed when the size is changed, the only difference being in the time for the inhibition wave to travel and reach any neuron within a given map. The model has been experimentally proved to be 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 several distractors with same intensity as the target. The main hypothesis, that is true for any such lateral-inhibition type neural field model, is that the target possesses some spatio-temporal continuity that is observable by the model.

The modeling framework that has been used in this paper falls into the general category of distributed asynchronous numerical processing where no global supervisor or central executive is allowed. This is tightly linked to the way that most people think the biological brain is actually processing information (even though there are theories that claims there is some central executive somewhere within the brain). The model proposed here clearly demonstrates that at least an early form of attention, i.e. a global cognitive function, is calculable and can emerge from local computations and interactions only.

References

- [Amari, 1977] Amari, S. (1977). *Biological Cybernetics* 27, 77–78.
- [Amari, 1980] Amari, S. (1980). *Bulletin of Mathematical Biology* 42, 339–364.
- [Burnod, 1989] Burnod, Y. (1989). *An adaptive neural network: the cerebral cortex*. : Masson.
- [Deneve et al., 1999] Deneve, S., Latham, P., and Pouget, A. (1999). *Nature Neurosciences* 2, 740–745.
- [Deneve et al., 2001] Deneve, S., Latham, P., and Pouget, A. (2001). *Nature Neuroscience* 4 (8), 826–831.
- [Douglas et al., 1995] Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A., and Suarez, H. H. (1995). *Science* 269, 981–985.
- [Giese, 1999] Giese, M. (1999). *Dynamic Neural Field Theory for Motion Perception*. : Kluwer Academic Publishers.
- [Hubel and Wiesel, 1965] Hubel, D. and Wiesel, T. (1965). *Journal of Neurophysiology* 28, 229–289.
- [Kohonen, 1982] Kohonen, T. (1982). *Biological Cybernetics* 43, 59–69.
- [Miller et al., 1989] Miller, K. D., Keller, J. B., and Stryker, M. P. (1989). *Science* 245, 605–615.
- [Rougier and Vitay, 2005] Rougier, N. P. and Vitay, J. (2005). *Neural Networks*. In press.
- [Samsonovitch and McNaughton, 1997] Samsonovitch, A. and McNaughton, B. (1997). *Journal of Neuroscience* 17, 5900–5920.
- [Stringer et al., 2004] Stringer, S., Rolls, E., and Trappenberg, T. (2004). *Neural Networks* 17 (1), 5–27.
- [Stringer et al., 2005] Stringer, S., Rolls, E., and Trappenberg, T. (2005). *Neurobiol Learn Mem* 83 (1), 79–92.

[Takeuchi and Amari, 1979] Takeuchi, A. and Amari, S. (1979). Biological Cybernetics 35, 63–72.

[Taylor, 1999] Taylor, J. G. (1999). Biological Cybernetics 80, 5167–5174.

[von der Malsburg, 1973] von der Malsburg, C. (1973). Kybernetic 15, 85–100.

[Wilson and Cowan, 1973] Wilson, H. R. and Cowan, J. D. (1973). Kybernetic 13, 55–80.

[Wu et al., 2001] Wu, S., Nakahara, H., and Amari, S. (2001). Neural Computation 13, 775–797.

[Zhang, 1996] Zhang, K. (1996). Journal of Neuroscience 16, 2112–2126.

A Appendix

In order to be able to perform numerical simulations using neural network models, we have to discretize CNFT 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 a neural position \mathbf{x} can be denoted \mathbf{x}_{ij} with $i, j \in [0, n - 1]^2$. Corresponding neuronal position is then given by equation (11)

$$\mathbf{x}_{ij} = \left(\frac{i}{n} - 0.5, \frac{j}{n} - 0.5 \right) \quad (11)$$

and equation (7) becomes:

$$\begin{aligned}
\tau \frac{du(\mathbf{x}_{ij}, t)}{dt} &= -u(\mathbf{x}_{ij}, t) + h \\
&+ \frac{1}{\alpha} \sum_{k=0}^{n-1} \sum_{l=0}^{n-1} w_M^+(\mathbf{x}_{ij} - \mathbf{x}'_{kl}) f[u(\mathbf{x}'_{kl}, t)] d\mathbf{x}'_{kl} \\
&+ \frac{1}{\alpha} \sum_{k=0}^{n-1} \sum_{l=0}^{n-1} w_M^-(\mathbf{x}_{ij} - \mathbf{x}'_{kl}) f^+[u(\mathbf{x}'_{kl}, t)] d\mathbf{x}'_{kl} \\
&+ \frac{1}{\alpha} \sum_{k=0}^{n-1} \sum_{l=0}^{n-1} s(\mathbf{x}_{ij}, \mathbf{y}_{kl}) I(\mathbf{y}_{kl}, t) d\mathbf{y}_{kl}
\end{aligned} \tag{12}$$

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 (13).

$$\begin{aligned}
|\mathbf{x}_{ij} - \mathbf{x}'_{kl}| &= \min \left(\left(\frac{i-k}{n} \right)^2, \left(1 - \frac{i-k}{n} \right)^2 \right) \\
&+ \min \left(\left(\frac{j-l}{n} \right)^2, \left(1 - \frac{j-l}{n} \right)^2 \right)
\end{aligned} \tag{13}$$

Clearly, the notion of “map edges” is a modelling artefact because in the framework of computer simulation, we have to consider finite maps. Consequently, the toric projection we’re using is a way to fix this artefact while for a real brain, there is no such notion of edges and no need of the toric projection. Finally and in order to avoid oscillatory symmetric behavior due to synchronous evaluation of the neurons, evaluation synchronicity is broken using a random evaluation order. 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^2 successive evaluations.

B Appendix

Let us consider a neuron \mathbf{x} using $u(\mathbf{x}, t)$ as the membrane potential with:

$$-1 \leq u(\mathbf{x}, t) \leq +1 \quad (14)$$

Let us also consider two neurons \mathbf{x}_+ , \mathbf{x}_- such that:

$$u(\mathbf{x}_+, t) = u(\mathbf{x}, t), 0 \leq u(\mathbf{x}_+, t) \leq +1 \quad (15)$$

$$u(\mathbf{x}_-, t) = -u(\mathbf{x}, t), 0 \leq u(\mathbf{x}_-, t) \leq +1 \quad (16)$$

By definition of \mathbf{x}_+ and \mathbf{x}_- , we have:

$$u(\mathbf{x}, t) = u(\mathbf{x}_+, t) - u(\mathbf{x}_-, t) \quad (17)$$

Let us now consider a neuron \mathbf{y} receiving some input from the neuron \mathbf{x} using a weight function $w(\mathbf{x}, \mathbf{y})$. We have immediately:

$$w(\mathbf{x}, \mathbf{y}).u(\mathbf{x}, t) = w(\mathbf{x}, \mathbf{y})u(\mathbf{x}_+, t) - w(\mathbf{x}, \mathbf{y}).u(\mathbf{x}_-, t) \quad (18)$$

that gives us the equivalent pattern of connection between a neuron \mathbf{y} and the neuron \mathbf{x} and a neuron \mathbf{y} and the system $(\mathbf{x}_+, \mathbf{x}_-)$.

Consequently, the system of the two strictly positive firing rate neurons $(\mathbf{x}_+, \mathbf{x}_-)$ is strictly equivalent to the single neuron system (\mathbf{x}) provided that the equivalent pattern of connectivity is used throughout the system.

C Appendix

Using equations (2), (3) and equation (12), simulation parameters are

$$n = 30$$

$$\tau = 0.75$$

$$h = 0.10$$

$$\alpha = 12.5$$

$$A = \frac{3.15}{\alpha}, a = \frac{2}{n}$$

$$B = \frac{0.90}{\alpha}, b = \frac{4}{n}$$

$$C = \frac{1.25}{\alpha}, c = \frac{1}{2n}$$

D Appendix

Figures 7 and 8 are two screenshots from simulations displaying two situations where input is the same but history of the network is different, leading to two different states within the *focus* map. Demonstration movies can be downloaded from <http://www.loria.fr/~rougier>. Figure 9 show a representation of interpolated paths as presented in the results section.




Fig. 7 In this experiment, the network has been first presented with a single stimulus until the *focus* map has settled on it. Then, and only then, two additional stimuli have been introduced within the *input* map. This introduction did not produce significant activity within the *focus* map since the inhibition is too strong for any other else coherent activity to emerge.




Fig. 8 In this experiment, the network has been presented with three stimuli at once and we wait for the *focus* map to settle its activity. One can observe on the figure that there are now three coherent pack of activity within the *focus* map that correspond exactly to the three stimuli.

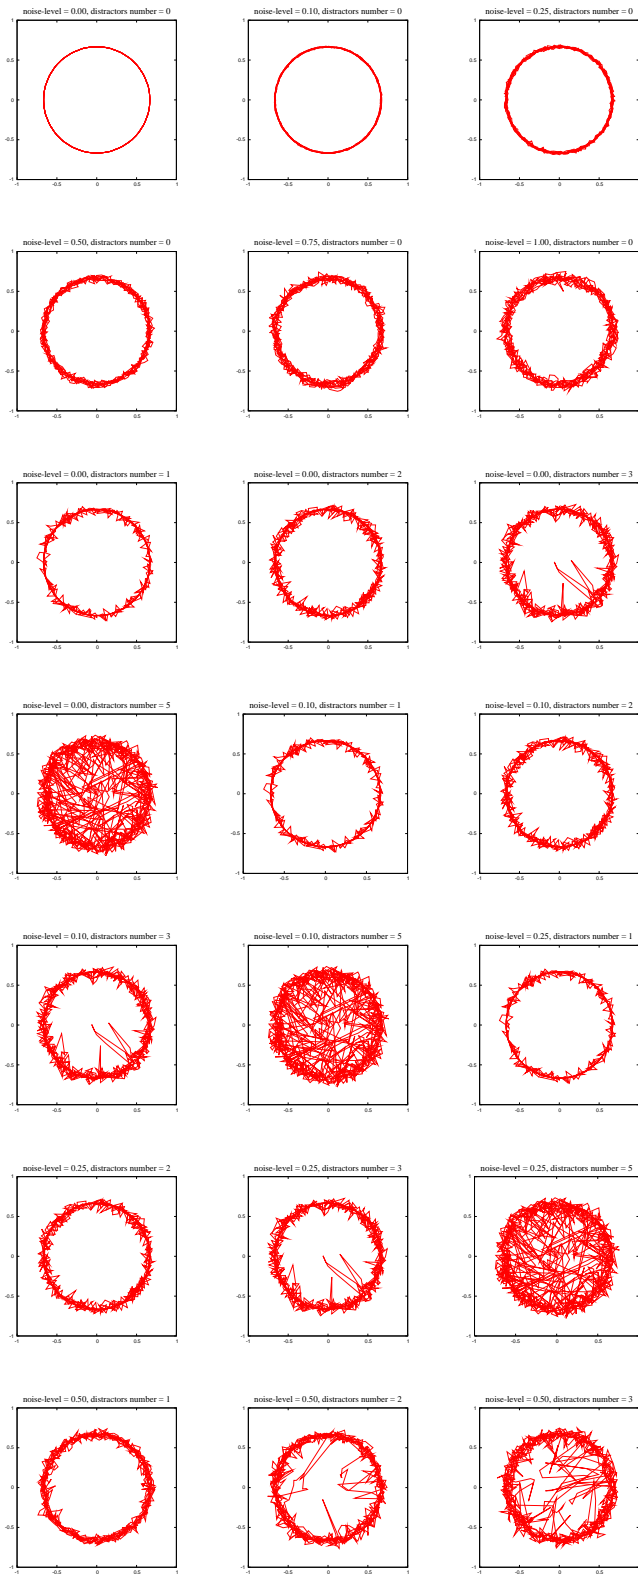


Fig. 9 Interpolated path (a line is drawn between two successive decoded positions within *focus* map) for different intensity of noise and different number of distractors.