On practical neural field parameters adjustment

Frédéric Alexandre, Jérémy Fix, Axel Hutt, Nicolas P. Rougier, Thierry Viéville

ON PRACTICAL NEURAL FIELD PARAMETERS ADJUSTMENT

Frédéric Alexandre, Jéremy Fix, Axel Hutt, Nicolas Rougier, Thierry Viéville
INRIA Cortex http://cortex.loria.fr

ABSTRACT

Revisiting CNFT calculation maps in both the continuous and discrete temporal cases, we propose a set of results allowing to choose the right set of parameters in order to both (i) guaranty the stability of the calculation and (ii) tune the shape of the outputs map.

With such parameters it appears that large sampling steps can be used, speeding up overall calculation. Furthermore, we report experimenting the fact that rectification is the only required non-linearity and formalize the use of this simplified but efficient mechanism.

The outcome is shared as an open-source plug-in module to be used in existing simulation software.

KEY WORDS

Neural field. Bump. Parameter adjustment.

1 Introduction

Using the CNFT and its extensions. The Continuum Neural Field Theory is mainly concerned with the functional modeling of neural structures where information is considered to be encoded at the level of the population rather than at the level of single neurons. Such models first appeared in the 50s, but the theory really took off in the 70s with the works of Wilson and Cowan [1] and Amari [2]. At the level of a single neuron, the model that is used is a mean frequency model where the electrical activity of a neuron is approximated by a single potential. However, there also exists several spiking neuron models that represent both a finer and more accurate model of a real biological neuron. In the framework of the CNFT, these models allow to bypass the inherent time discretization implicitly brought by mean frequency models.

The dynamics of pattern formation in lateral-inhibition type neural fields with global or local inhibition has been extensively studied in a number of works where it has been demonstrated that these kinds of 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. The characteristic behavior of such fields is this formation of a very localized packets of neural activity that tend to represent some consistent information which is present at the level of the input.

These models 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 [3]. They exhibit so-called bump patterns, which have been observed in the prefrontal cortex and are involved in working memory tasks [4].

However, the exact shape of these output bumps is quite difficult to predict since it is a generally a non linear consequence of both the lateral connectivity pattern and the input pattern. Thus the question remains on how to control the shape, first in a purely mechanical way (i.e. finding the proper mathematics behind), second in a more biological plausible way. This second step, which is beyond the scope of this paper, requires to complete the first step, which is going to provide some description of the space of possible parameters.

The parameter adjustment problem. More precisely, here, we want to study which set of lateral weights allow to achieve a specific function, i.e. obtaining one unique bump in the output. At a very pragmatical level, however, the parameters of neural field are always adjusted empirically (even if authors are unobtrusive about it), although constructive results exist and should be used.

In [5], for instance, homogeneous stationary solutions (i.e independent of the spatial variable) and bump stationary solutions (i.e. localized areas of high activity) in two-dimensional neural field models composed of excitatory and inhibitory neurons have been studied, in the continuous case after [6, 7], providing an explicit solution in terms of Bessel functions and Hankel transform. Basic properties of bounded neural fields (well-posedness, stability of their solutions in the homogeneous or locally homogeneous case) have also been studied, accounting for partial or global synchrony among the cortical columns composing the field [8]. These results are based on the compactness of the related operators yielding stability results related to contracting mappings.

These results however are not easy to implement in practical codes. A caveat limits their use. In order to obtain analytical results they consider formulations (using e.g. Heaviside non-linearity and/or unlimited precision assumptions) which do not correspond to numerical implementations.

The goal of this work is thus to contribute to the understanding of the effect of some important parameters on the global behavior of a neural field and to use this knowledge for its on-line control. It also provides results directly usable at the implementation level to choose the right set of parameters to (i) guaranty the stability of the calculation and (ii) tune the shape of the map's output.

The former problem is solved here in the linear and non-linear case, while the latter is solved here only in the linear case, solving the non-linear case being in progress.

2 Methods

The ideal linear continuous map. Let us consider a *vectorial* map (see [9] for a discussion) $\mathbf{u} : \mathcal{R}^n \rightarrow \mathcal{R}^m$ (e.g., for a 2D scalar map $n = 2$, $m = 1$) defined from:

$$\begin{cases} \tau\, \dot{\mathbf{u}}(\mathbf{p},t) &= -\mathbf{u}(\mathbf{p},t) + \int_{\mathbf{q}} \mathbf{W}_\theta(\mathbf{p}-\mathbf{q})\, \mathbf{u}(\mathbf{q},t) + \mathbf{i}(\mathbf{p}) \\ \mathbf{u}(\mathbf{p},0) &= \mathbf{i}(\mathbf{p}), \end{cases}$$

where $\mathbf{u}(\mathbf{p},t)$ is the *positive* output activity at location $\mathbf{p}$ and time t and $\mathbf{i}(\mathbf{p})$ the *positive* stationary input. The map is supposed large enough to neglect boundary conditions at this stage.

This corresponds to what is usually considered for analytical studies, up to some additional non-linearity.

The weights $\mathbf{W}_\theta$ may correspond, for instance (but not only), to a "Mexican hat" (i.e. with exponential excitatory/inhibitory connections) profile:

$$\mathbf{W}_\theta(\mathbf{q}) = \mathbf{A}_+\, e^{-\frac{|\mathbf{q}^2|}{\sigma_+^2}} - \mathbf{A}_-\, e^{-\frac{|\mathbf{q}^2|}{\sigma_-^2}},$$

thus parametrized by $\theta = (\mathbf{A}_+, \mathbf{A}_-, \sigma_+, \sigma_-)$.

Here, the goal is to tune the parameters θ in order to control the map's output.

The real synchronous non-linear discrete map. Let us also consider the more realistic non-linear discrete scheme in time and space, thus, now for $\mathbf{p} \in \{0..N\{^n$:

$$\mathbf{u}(\mathbf{p},t+1) = \rho\left(\mathbf{u}(\mathbf{p},t) + \delta\, \left(-\mathbf{u}(\mathbf{p},t) + \textstyle\sum_{\mathbf{q}} \mathbf{W}_\theta(\mathbf{p}-\mathbf{q})\, \mathbf{u}(\mathbf{q},t) + \mathbf{i}(\mathbf{p})\right)\right),$$

providing $0 < \delta < 1$, while a non-linearity $\rho(.)$ has been introduced. This writes in compact matrix form:

$$\bar{\mathbf{u}}(t+1) = \mathbf{F}\left(\bar{\mathbf{u}}(t)\right) = \rho\left(\bar{\mathbf{K}}\, \bar{\mathbf{u}}(t) + \delta\, \bar{\mathbf{i}}\right),$$

where $\bar{\mathbf{u}}$ is a N^n vector, with the $N^n \times N^n$ Toeplitz matrix:

$$\bar{\mathbf{K}} = 1 - \delta\, \left(1 - \bar{\mathbf{W}}_\theta\right).$$

This corresponds to what is implemented in usual simulation codes.

Here $\rho_i(u_j) = 0 \vee u_j = u_j \geq 0\ ?\ u_j : 0$ is a simple "rectification". This somehow strange and minimal choice results from the following experimental fact. *It has been observed that (e.g. Fig.1) with "suitable" parameters* (here obtained by "chance") *all interesting properties of neural fields (filtering, selection, dynamic input following, etc..) [10, 3] can be obtained using rectification* (instead of the usual piece-wise linear or regular sigmoid, while not Heaviside profiles).

Bump as required output In this context, the desired output profile is a bump, i.e. a *positive, decreasing, radial symmetric profile*,as observed for bubble of activities in the cortex [2, 4]. A Gaussian profile is one example of bump. In order to capture the notion of bump at a more general level, we propose to consider Gaussian enumerable linear combinations, i.e. *Gaussian series*:

$$\mathbf{b}(\mathbf{p}) = \textstyle\sum_{s \in \{s_1, \cdots\}, s>0} \mathbf{b}^*(s)\, e^{-s\,|\mathbf{p}|^2},$$

writing $s = \frac{1}{2\,\sigma^2}$ for an enumerable set of "width" s, which is of common use (e.g. [11]). Small s for flat Gaussian ($s = 0$ is the "constant part"), Large s for tight Gaussian.

The goal is thus to find the parameters allowing to generate bumps, controlling their width, amplitude and shape.

3 Results about the convergence

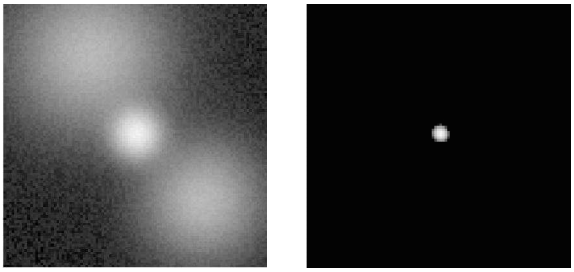

Figure 1. An example of input (left-view) output (right-view) transformation via a synchronous non-linear discrete map with rectification. The input is noisy (more than 20% of the signal) with three bumps. The output "filters" the noise, and select the highest bump even if the lower bumps have more energy (being wider). This result is obtained with $\theta = (0.0015, 0.0015, 0.45\,N, N)$, $N = 100$ and $\delta = 0.99$. Furthermore, with the same parameters, if the bump is moving, the mechanism tracks the highest bump.

Let us first discuss how to control the discrete implementation convergence, in both linear and non-linear cases.

Implementation as a contracting mapping. In order to obtain the convergence of the iterative equation toward a fix point we simply require the related mapping to be contracting, i.e.

$$\left|\tfrac{\partial \mathbf{F}}{\partial \mathbf{u}}\right| = \left|\Sigma'\, \bar{\mathbf{K}}\right| < 1$$

where $\Sigma' = \mathrm{diag}(\cdots, \rho'_i, \cdots)$ is the diagonal matrix with $\rho'_i = \rho'\left(\bar{\mathbf{K}}\, \bar{\mathbf{u}}(t) + \delta\, \bar{\mathbf{i}}\right) \in \{0, 1\}$ whether the value is negative or positive, respectively.

This allows to write pragmatic conditions of convergence, as discussed now.

This also restrains our study to convergence to a fixed point, without taking other asymptotic dynamics into account [6].

Convergence in the linear case. Assume that all output values are positive, i.e. $\rho_i = 1$, thus no rectification. This occurs typically with excitatory weights only. This corresponds to a pure linear transformation.

The convergence condition writes in this case:

$$|\bar{\mathbf{K}}| < 1 \Leftrightarrow |\bar{\mathbf{W}}_\theta| < 1 \text{ as soon as } 0 < \delta < 1,$$

i.e. in words: the system is contracting if the magnitude of the global weight matrix is lower than one. If higher than one, the system diverges.

This condition is obvious to check[1], using for instance the power method, which very efficiently calculates the highest matrix eigen-value, i.e., the magnitude of the matrix in our case.

Tuning the parameters for convergence. As a rule of thumb, we have observed that this magnitude is always higher than the weight's average value, and close to it for high magnitude.

[1]Numerical implementations used in this paper are available at `http://enas.gforge.inria.fr`.

Furthermore, we have observed for Mexican-hat profiles, that convergence decreases with $|\mathbf{A}_-|/\mathbf{A}_+|$ and increases with σ_-/σ_+.

More than that, the calculation not only tells if the system is contracting or not, but since the equation is linear, allows to calculate the maximal weight magnitude which guarantees the convergence. In other words, given a set of parameters θ the previous algorithm tells how to *rescale* the parameters in order to guaranty the convergence.

For Mexican-hat profiles, in this linear case, it simply means dividing $|\mathbf{A}_-|$ and $|\mathbf{A}_+|$ by a quantity higher than $|\bar{\mathbf{W}}_\theta|$.

Another key fact at this point, is that *it does not depend on* δ (as soon as lower than 1). It means that we can use high δ and speed up the convergence. However, it does not mean that all δ lead to the same result. It does, in the linear case, since the closed-form solution writes:

$$\bar{\mathbf{u}} = [1 - \bar{\mathbf{W}}_\theta]^{-1}\,\bar{\mathbf{i}}$$

thus without δ dependence. But it does not in the non-linear case, since the fixed point is now a function of δ as easily readable on the master equation.

Anyway, it shows that it is not always a good idea to make all the analysis in the continuous case and then "discretize with a small ϵ".

Using bounded excitation in the non-linear case. The previous result is quite informative, but rather weak, since rectification is not taken into account.

A few algebra allows to improve this result and derive the following sufficient boundary condition, in the non-linear case:

$$|\bar{\mathbf{W}}_\theta \vee 0| < 1$$

i.e. in words: the solution remains bounded if the magnitude of the positive (excitatory) weights is lower than one. This captures the intuitive fact that the negative (inhibitory) weights action is always bounded by the rectification, thus only positive weights matter. An example of practical adjustment curve is given in Fig. 2.

Without this condition, as corroborated by numerical experiments, the system is in usual conditions divergent. However, it does not mean that, given a *restrained* set of inputs, higher parameter's values are not going to yield convergence.

4 Results about the bumps

Let us now discuss how to obtain, at the implementation level, bumps with controlled shapes. Here, only the linear case is considered.

Gaussian series approximate bumps. Radial symmetric profile means that $\mathbf{b}(\mathbf{p}) = \beta(|\mathbf{p}|^2)$, normalized iff $\int_{\mathbf{p}} b_i(\mathbf{p}) = 1$. They are positive iff $\beta_i(r) \geq 0$ and decreasing iff $\beta_i'(r) \leq 0$. This puts our specifications in equation.

For Gaussian series:

$$\beta(r) = \textstyle\sum_s \mathbf{b}^*(s)\, e^{-s\,r} = \mathcal{L}(\mathbf{b}^*)(r)$$

where $\mathcal{L}$ stands for the Laplace transform, while $\mathbf{b}^*$ has been extended to the continuum using a Dirac comb, while the Post-Bryan inversion formula[2] [12] allows to consider the Laplace transform of such positive functions.

This observation means that *bumps can be approximated using Gaussian series* for the simple reason that Gaussian series coefficients are nothing but the discrete approximation of the bump β profile's Laplace transform. This representation thus benefits from all the Laplace transform properties:

Bump are normalized iff $\int_{s>0} b_i^*(s)(\pi/s)^{n/2} = 1$. Several additional properties characterize such bumps: they are flat at 0 unless singular (i.e. $|\mathbf{b}(\mathbf{0})| < +\infty \Rightarrow |\nabla\mathbf{b}(\mathbf{0})| = 0$), vanish at infinity ($\lim_{r\to+\infty}\beta(r) = 0$) thus without constant component ($\lim_{s\to 0} b^*(s) = 0$), etc..

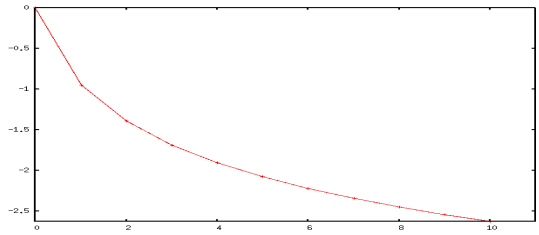

Figure 2. An example of maximal gain curve in the non-linear case. Given σ_+ in abscissa (in pixel unit for a 100×100 2D map), the normalized maximal value of $\log_{10}(|\mathbf{A}_+|)$ is drawn. Here, not a Mexican hat but a step-wise constant profile has been used, since the derivation can be performed with any W profile.

Gaussian series yield "nice" bumps. The previous result is useful because it is very easy to shape the bump using suitable coefficients, for instance using rational profiles of the form $\beta_1(r) = 1/(1+|r|^{n/2})$, n even since:

$$\begin{aligned} b_1(s)^* &= \tfrac{1}{n}\textstyle\sum_{i=1}^{n} e^{-s\,\sin(\alpha_n^i)}\,\sin\left(s\,\sin\left(\beta_n^i\right) + \alpha_n^i\right) \\ \alpha_n^i &= \tfrac{\pi}{n}\left(2\,i + \tfrac{n}{2} - 1\right) \\ \beta_n^i &= \tfrac{\pi}{n}\left(1 - 2\,i\right) \end{aligned}$$

or of the form $\beta_2(r) = 1/r^n, n > 0$ since:

$$b_2(s)^* = s^{n-1}/n-1!$$

as illustrated in Fig. 3. Such profiles correspond to desired bumps in neural field applications, among many other suitable alternatives. They come with closed-form formula and show how malleable are such shapes.

Note that a simple change of scale allows to adjust the width and the amplitude of the bump, the challenge here being to adjust the *shape*.

Gaussian series are linear map fix points. Linear map fixed points $\mathbf{u}^\bullet$ are defined by $-\mathbf{u}^\bullet + \mathbf{W} * \mathbf{u}^\bullet + \mathbf{i} = 0$, where $*$ stands for the convolution operator. Since convolution and linear combination of Gaussian are Gaussian, *providing that the input* $\mathbf{i}$ *and the weights profile* $\mathbf{W}$ *are Gaussian series, the output* $\mathbf{u}^\bullet$ *is a Gaussian series.* This result is made explicit by the following formula for the Gaussian series coefficients:

$$-u^\bullet(\nu) + \pi \textstyle\sum_{r>\nu, s>\nu, \frac{r\,s}{r+s}=\nu} \frac{W(r)\,u^\bullet(s)}{r+s} + i(\nu) = 0,$$

[2] *Post-Bryan result:* A continuous function on $\mathcal{R}^+$ which is of exponential order for some c (i.e. $||b(r)|/e^{c\,r}|_\infty < \infty$) with Laplace transform $b^*(s)$ is non-negative iff $\forall k \geq 0, \forall s \geq c, (-1)^k\, b^{*(k)}(s) \geq 0$.

This is true for rather general weight profiles, not necessarily "Mexican-hat" profile only.

In this specific case, it is possible to derive a closed-form expansion, convergent as soon as $|A_\pm| < 1$, which is the case in practice. We obtain up to a given order r:

$$u^\bullet(\nu) = \sum_s i(s) \sum_{p,q,p+q<r} \eta_r(p,q) \frac{(-1)^q \pi^{p+q} A_+^p A_-^q \sigma_+^{2\,p} \sigma_+^{2\,q}}{1+s\,(p\,\sigma_+^2+q\,\sigma_-^2)} \delta\left(\nu - \frac{s}{1+s\,(p\,\sigma_+^2+q\,\sigma_-^2)}\right),$$

with $\eta_r(p,q) = \frac{(p+q)!}{p!\,q!}$ easily derived by symbolic calculation.

Note that the previous two formulas are written in the scalar case for the sake of simplicity, whereas it is straightforward to derive them in the vectorial case. It is also straightforward to derive (using symbolic calculation) for more general kernels.

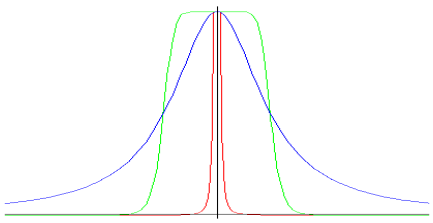

Figure 3. Three examples of rational profiles for n = 2 (Gaussian like, in blue) and n = 10 (square like, in green), and a singular profile (pointed like, in red). See text for details.

Practical implementation. From the previous results it becomes possible to design numerical routines in order to adjust at will the bump's shape. The previous result development means that we do not have to adjust the profile in the $2D$ space but simply in the $1D$ radial profile, thanks to the Gaussian series representation. This fact very likely generalizes to non-linear profiles.

For instance, given an input $i(\nu)$ and a desired output $b_1^\bullet(\nu)$ the following program:

$$\min_\theta \int_\nu \left| b_1^\bullet(\nu) - u_{\theta,i(\nu)}^\bullet(\nu) \right|_\Lambda$$

allows to find the optimal weight parametric value θ, using e.g. the explicit formulas given previously. The choice of the measure $\int_\nu$ and of the metric $|.|_\Lambda$ is application dependent.

Another application is the input control of such calculation maps: Given parameters θ, find the input $i(\nu)$ yielding a given output $b(\nu)$. The solution is straightforward in the present framework, as illustrated in Fig. 4.

5 Discussion

The previous results do not "prove" the existence of such stable bump solutions, this being already known [1, 2]. Here the goal is to allow to calculate numerically the related parameters. This modest objective is attained since we are able to tell if a parameter set yields convergence and what is going to be the shape of the bump.

A slight modification of the previous derivation, allows to not only consider synchronous but asynchronous sampling (see [3] for details) since the map is still contracting when using asynchronous sampling. This has been numerically experimented.

In a recent work [10] a model has been designed that performs global competition, only using local connections, with diffusion of the inhibition throughout the network. This is far quicker to have a few local interactions when computing activity within the network and makes the model a real candidate for distributed computations. The next step is to adapt the present results to this paradigm.

acknowledgment Partially supported by the ANR grant MAPS.

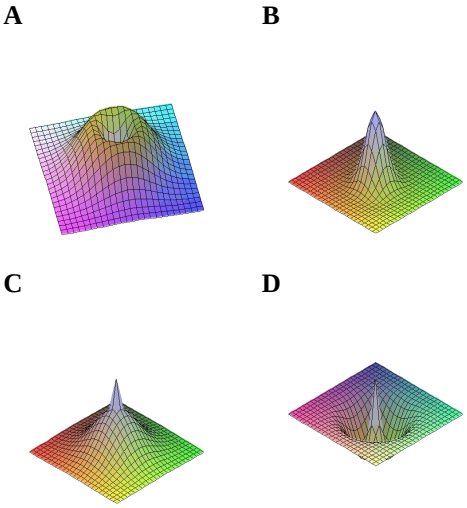

Figure 4. Four examples of input profiles yielding a Gaussian output $u(\mathbf{p}) = e^{-|\mathbf{p}|^2}$. **A**: $\theta = (1, 1, 0.2, 1)$ shown in a $[-5..5]^2$ window. **B**: $\theta = (0, 1, ., 1)$ shown in a $[-5..5]^2$ window. **C**: $\theta = (0, 0.1, ., 0.1)$ shown in a $[-20..20]^2$ window. **D**: $\theta = (1, 1, 0.02, 0.05)$ shown in a $[-20..20]^2$ window. It also shows the versatility of Gaussian series shapes.

References

[1] H.R. Wilson and J.D. Cowan. A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Biological Cybernetics*, 13(2):55–80, sep 1973.

[2] S.-I. Amari. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27(2):77–87, jun 1977.

[3] N. Rougier and J. Vitay. Emergence of attention within a neural population. *Neural Networks*, 19(5):573–581, 2006.

[4] P. Goldman-Rakic. Topography of cognition: parallel distributed networks in primate association cortex. *Ann. Rev. Neurosci.*, 11:137–156, 1988.

[5] F. Grimbert. *Mesoscopic models of cortical structures*. PhD thesis, University of Nice Sophia-Antipolis, feb 2008.

[6] Stephen Coombes. Waves, bumps, and patterns in neural fields theories. *Biological Cybernetics*, 93(2):91–108, 2005.

[7] Stefanos E. Folias and Paul C. Bressloff. Breathing pulses in an excitatory neural network. *SIAM Journal on Applied Dynamical Systems*, 3(3):378–407, 2004.

[8] F. Grimbert and O. Faugeras. Bifurcation analysis of Jansen's neural mass model. *Neural Computation*, 18(12):3052–3068, December 2006.

[9] T. Viéville, S. Chemla, and P. Kornprobst. How do high-level specifications of the brain relate to variational approaches? *J. Physiol. Paris*, 101, 2007.

[10] N. Rougier. Dynamic neural field with local inhibition. *Biological Cybernetics*, 94(3):169–179, 2006.

[11] T. Wennekers. Separation of spatio-temporal receptive fields into sums of gaussian components. *J. Computational Neuroscience*, 16(1), 2004.

[12] Kurt Bryan. Elementary inversion of the laplace transform, 1999. Web publication on `http://www.rose-hulman.edu/~bryan`.