Cogn Neurodyn (2008) 2:171–177
DOI 10.1007/s11571-008-9046-0

RESEARCH ARTICLE

Tracking population densities using dynamic neural fields with moderately strong inhibition

Thomas Trappenberg

Received: 4 December 2007 / Accepted: 30 March 2008 / Published online: 17 April 2008

Abstract We discuss the ability of dynamic neural fields to track noisy population codes in an online fashion when signals are constantly applied to the recurrent network. To report on the quantitative performance of such networks we perform population decoding of the 'orientation' embedded in the noisy signal and determine which inhibition strength in the network provides the best decoding performance. We also study the performance of decoding on time-varying signals. Simulations of the system show good performance even in the very noisy case and also show that noise is beneficial to decoding time-varying signals.

Keywords Dynamic neural field model · Population decoding · Noise reduction

Introduction

It is increasingly obvious that information processing in the brain is very different from information processing in a digital computer. In particular, information processing in the brain is based on highly distributed representations of information. For example, it is well known from the work of Hubel and Wiesel (1962) that some neurons in the primary visual cortex are selective to edges with specific orientations. While those neurons fire most vigorously to edges with their so-called preferred orientation, the same neurons also fire to edges with other orientations, resulting in broad tuning curves (see for example, Henry et al. 1974). Accurate decoding can thus only be achieved by combining the information of many neurons.

Several questions arise when thinking about information processing with such population representations. For example, it seems at first puzzling that broad tuning curves are common in the brain as it is more difficult to decode information from such broad representations, in particular when considering the large noise present in spike trains. While it turns out that decoding errors can be very small even with brought tuning curves, when combining information from several neurons (see for example, Trappenberg 2002, Chapter 5.4.4), the presence of large fluctuations in neuronal firing are a serious challenge to accurate decoding. Indeed, depictions of tuning curves in common publications are somewhat deceiving as they represent averages over many, sometimes hundreds, of trials.

Population decoding has recently sparked new interest in the computational neuroscience community as it was realized that population decoding in very noisy circumstances can be achieved through very efficient implementations based on recurrent radial basis function networks that resemble possible cortical processing (Pouget and Zhang 1997; Pouget et al. 1998, 2000; Deneve et al. 1999; Wu et al. 2002, 2004; Li and Wu 2007). This research has shown that this implementation is statistically efficient and equivalent to maximum likelihood decoding (Wu et al. 2002), and that these methods also work with correlations in the noise (Wu et al. 2004). In this paper we will introduce this method in the formulation of dynamic field models and show that these models work quite well in practice without the commonly applied global operation of divisive inhibition. We also demonstrate that this method works well online, that is, when noisy signals are constantly applied. An important aspect of the model is that of

---

T. Trappenberg (✉)
Faculty of Computer Science, Dalhousie University,
6050 University Avenue, Halifax, NS, CanadaB3H 1W5,
e-mail: tt@cs.dal.ca

global inhibition that provides the necessary competition between noisy population components. We investigate in this paper which strength of activity-dependent inhibition is optimal for the online decoding. Interestingly, it turns out that the best performance is reached in a transition region of the model, where the model turns from a forgetful mode into a mode with sustained activity. Some preliminary results of these findings have been reported in (Trappenberg 2008). However, in this paper we will not only present these results in more detail, but we also show results of additional studies that have explored the ability of online tracking of time varying signals with the simplified model proposed in this paper. We found that the method works well even with fairly fast changing stimuli. Surprisingly, we also found that some slow changing stimuli provided a more challenging case for decoding. However, this was solvable by including noise in the system.

While we demonstrate fast population decoding in this paper, it is important to keep in mind that the decoding of a parameter such as the orientation of a stimulus in a visual scene is not necessarily the goal in cortical information processing. That is, we subscribe to the interpretation advanced by Pouget et al. (2000) that population responses in the brain should be viewed as probabilistic density information in the brain, that is, the brain cares not only about specific events but also about their likelihood. It is thus not desirable to extract only single parameters from the representation. The online method discussed here does provide ongoing information of complex signals in the system, which we believe is essential for further processing of the information in cortex. Thus, the extraction of a parameter such as the most likely orientation embedded in a signal should be seen as only a means of demonstrating the performance of the system.

Methods

We consider the problem of decoding information from very noisy signals represented by populations of neurons. For simplicity we consider representations where a single parameter is encoded by the activity of nodes, although the method can be generalized to higher dimensions (Deneve et al. 2001). Each node represents thereby a collection of neurons and the activity represents the population rate (Wilson and Cowan 1972; Gerstner 2000). Each node is considered to have a different response profile in dependence of the encoded parameter. We use Gaussian tuning curves similar to orientation tuning curves found in V1 that ride on top of a constant background field. We further added noise and only consider periodic parameters to avoid boundary effects. Specifically, a raw signal s^{raw} with orientation x_o is calculated from

$$x_s = MIN(|x - x_o|, 2\pi - |x - x_o|), \tag{1}$$

$$s^{\mathrm{raw}}(x) = b + \frac{1}{\sqrt{2\pi}\sigma}\exp\left(-\frac{x_s^2}{2\sigma^2}\right), \tag{2}$$

with $b = 0.1$ and $\sigma = 2\pi/10$. To simulate noise we applied independent Gaussian noise to each node,

$$s^{\mathrm{noisy}}(x) = s^{\mathrm{raw}}(x) + n_\eta \eta(x; \mu, \sigma^2), \tag{3}$$

where η is a normal distributed random variable with mean $\mu = 0$ and variance $\sigma^2 = 1$, and n_η is a noise parameter setting the strength of the noise. The final noisy signal was calculated by rectifying and normalizing s^{noisy},

$$s(x) = a\frac{s^{\mathrm{noisy}}\,\Theta(s^{\mathrm{noisy}})}{\int s^{\mathrm{noisy}}\,\Theta(s^{\mathrm{noisy}})\mathrm{d}x}, \tag{4}$$

where $\Theta(x)$ is the step function and the parameter a was set to $a = 0.2$. An example of the noiseless signal and a corresponding signal with $n_\eta = 0.5$ is illustrated in Fig. 1. Note that a simple 'max' estimate would result in poor estimates of the center of the Gaussian signal.

As mentioned in the introduction, a common method to decode population signals is to use the them as the starting state of a radial-basis recurrent network and to iterate the network until a stable state is reached (Pouget and Zhang 1997). The continuous time dynamics of a standard recurrent network is thereby given as

$$\tau\frac{\mathrm{du(x,t)}}{\mathrm{d}t} = -u(x,t) + \int w(x,y)r(y,t)\mathrm{d}y + I^{\mathrm{ext}}(x,t), \tag{5}$$

where τ is a time constant and I_i^{ext} is the external input applied to the network. We formulated the network here with spatially continuous neural fields that are commonly discretised in numerical simulations to a network with N nodes with

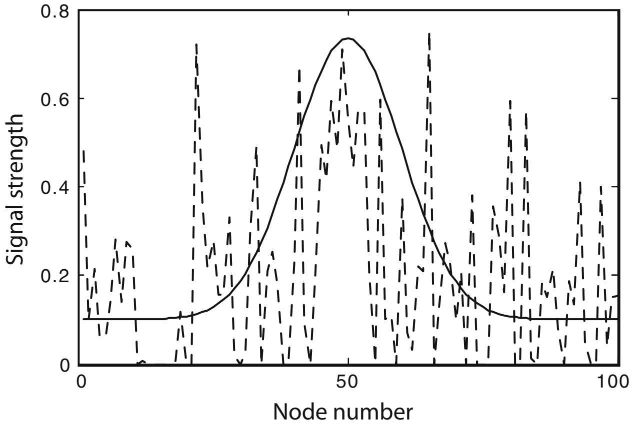

Fig. 1 Noisy population signals (dashed line) of a Gaussian of width 10 centered at node number 50

$$x \to i\Delta x, \tag{6}$$

where $\Delta x = 2\pi/N$ is a scale factor. The rate $r(x,t)$ in Eq. 5 is related to the internal state variable $u(x,t)$ by a nonlinear gain function $r = g(u)$. In the recent population decoding literature it is common to use divisive inhibition (Deneve et al. 1999),

$$r = \frac{u}{\int u \, \mathrm{d}x}, \tag{7}$$

which has the advantage of resulting in Gaussian activity packets. However, we propose to use the simpler local sigmoidal gain function,

$$r = \frac{1}{1 + \exp(-\beta u)}, \tag{8}$$

with slope parameter $\beta = 0.1$. The results in this paper demonstrate that decoding does not rely critically on the specific form of the gain function so that a more efficient local function can be used in practice.

Most important for the use of recurrent networks for decoding is the form of the weight kernel $w(x,y)$. This is chosen to be shift invariant and only depends on the distance between locations in the neural field, that is

$$w(x,y) = w(|x-y|). \tag{9}$$

It is thereby common to used Mexican hat functions such as the difference of two Gaussian bell curves or a shifted Gaussian in a periodic feature space. Again, the precise form is not critical as long as this interaction kernel provides some local excitation and long range inhibition. We use a shifted Gaussian of the form

$$d = MIN(|x-y|, 2\pi - |x-y|) \tag{10}$$

$$w(x,y) = A_\mathrm{w}(\mathrm{e}^{-d^2/2\sigma_w^2} - C), \tag{11}$$

with an weight amplitude of $A_\mathrm{w} = 0.2/(\sqrt{2\pi}\sigma)$ and width $\sigma_w = \sqrt{2}\sigma$. The Gaussian form can be learned with Hebbian training on Gaussian pattern (Stringer et al. 2002), but most interesting for this paper is the inhibition constant C which provides the necessary, activity-dependent long range inhibition.

The dynamics of such center-surround neural fields has been analyzed by Amari in a classic paper (Amari 1977) that shows that this model exhibits several regimes characterized by different possible asymptotic states. For sufficiently large weight amplitude A_w, the specific regime is determined by the inhibition constant C. If C is low compared to excitation in the network, then the excitation will spread through the network resulting in runaway activity. In contrast, if inhibition is dominating (large C), then any activity in the field will decay without external reinforcement. In an intermediate regime it is possible to have activity packets where localized activity is stable. We call this regime of the model the memory regime.

The time evolution of the neural field approximated with 100 nodes is depicted in Fig. 2. The noisy input of Fig. 1 was applied until $t = 20\tau$, after which point it was removed. The field develops into a clean bubble around the middle node in Fig. 2a which was the center of the Gaussian signal chosen in this experiment. This demonstrates perfect decoding of the 'orientation' parameter. The bubble is maintained after the input is removed, which demonstrates the working memory regime of the neural field model. The clear asymptotic signal is commonly used for decoding, but in this paper we are specifically interested in fast online decoding. It is therefore important to note that some good estimates can be obtained much earlier, even during the time the external signal is still present. This can be seen when showing the time evolution until $t = 10\tau$ on a finer color scale (Fig. 2b). In Fig. 2c we show the time evolution with stronger inhibition. In this case, the inhibition is too large to maintain an activity packet after external input is removed. However, even in this case, we can see that some good estimates can be achieved during the time the signal is present and even shortly after it is removed. The main goal of the remainder of this paper is to investigate which inhibition strength C is best for decoding and to show that the neural field can track time varying signals.

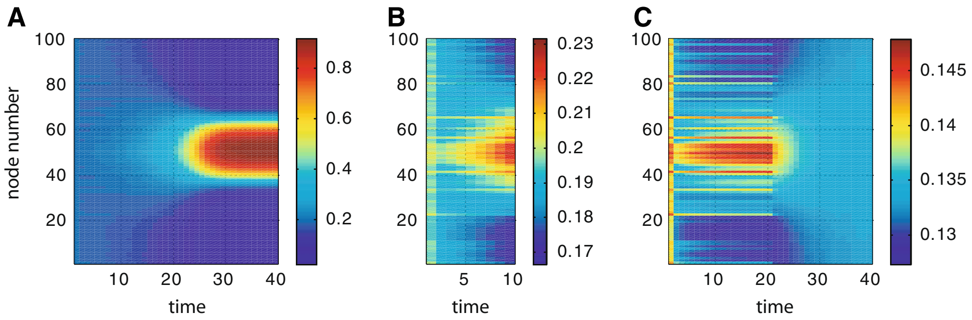

Fig. 2 Noisy population decoding with weak and strong inhibition in neural fields. The noisy input of Fig. 1 is applied for the first 20 time steps and then removed. (**a**) Time evolution of the neural field with *weak* inhibition ($c = 0.05$). (**b**) The first 10 time steps of panel a with a finer color scale. (**c**) Time evolution of the neural field with *strong* inhibition ($c = 0.07$)

Results

Online decoding and optimal inhibition regime

As was shown in Fig. 2, neural fields can be used to enhance noisy population decoding even during the time the noisy signal is applied to the system. However, the quality of decoding should be compared to the traditional case of decoding with asymptotic states. The question remains which inhibition constant is most useful for decoding. While moderately small inhibition in the memory regime of the neural field model seem to produce clear signals, stronger inhibition also increases the competition which could speed up the formation of clean signals.

In order to report an estimate of the encoded information, in our case the location of the center of the Gaussian signal which we sometimes call ‘orientation’, we use a simple center of mass estimate based on the activations r of the neural fields at different times during their evolution. We thereby calculate a torque value,

$$m_i = \sum_{j=i-N/2}^{i+N/2} r_j (j - i) \tag{12}$$

for each node $i = 1, \ldots, N$ in the network and choose as an estimate of the ‘orientation’ the minimum,

$$\hat{x} = \operatorname{argmin} m_i. \tag{13}$$

The decoding error calculated as absolute minimum difference between the estimated orientation and the true orientation. In the following we show results of simulations in which we consider time varying noise in which a new noise term was chosen in each unit time interval. This represents a more plausible implementation of decoding conditions in the brain. We found similar results in the static noise case. Figure 3 shows results of simulations over 100 runs at different estimation times and for various inhibition constants C. Shown is the improvement ratio, the ratio of decoding errors, of the decoding when estimating the orientation from the neural field at various times relative to the center of mass estimation of the noisy signal at $t = 0$. Error bars in the figures represent the variance of the results over 100 runs. As can be seen, the estimation improves with ongoing evolution of the neural field, and the best improvements are achieved with an inhibition constant around $C = 0.06$. Interestingly, this is also the transition region from the memory regime to the regime with decaying activity.

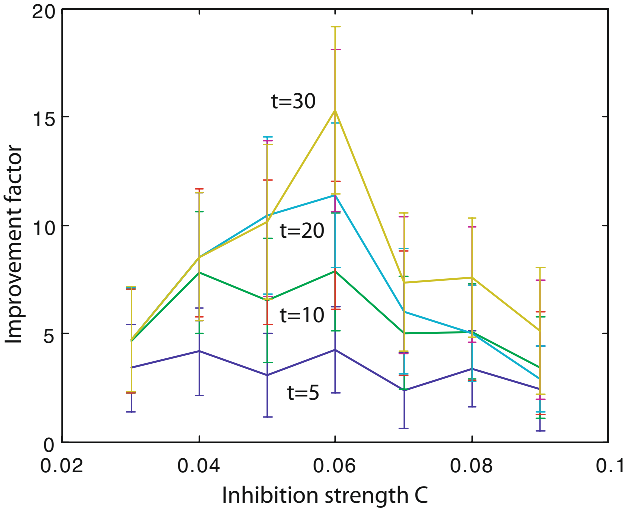

Fig. 3 Improvement factor relative to center of mass decoding of the original noisy signals with the neural fields at different times and different inhibition constants C

Further details of the estimation performance at different times and for different noise amplitudes from simulations with inhibition $C = 0.06$ are shown in Fig. 4. Figure 4a shows that the estimation error rapidly decrease with time and that there is only a small further improvement after the noisy signal is removed. It is also interesting to note that decoding with neural fields works well for very noisy signals. Figure 4b shows the decoding error with the center of mass estimation on the noisy signal (the curve labeled $t = 0$ in the figure) and the center of mass decoding at time $t = 20\tau$ based on the neural fields. The curves represent

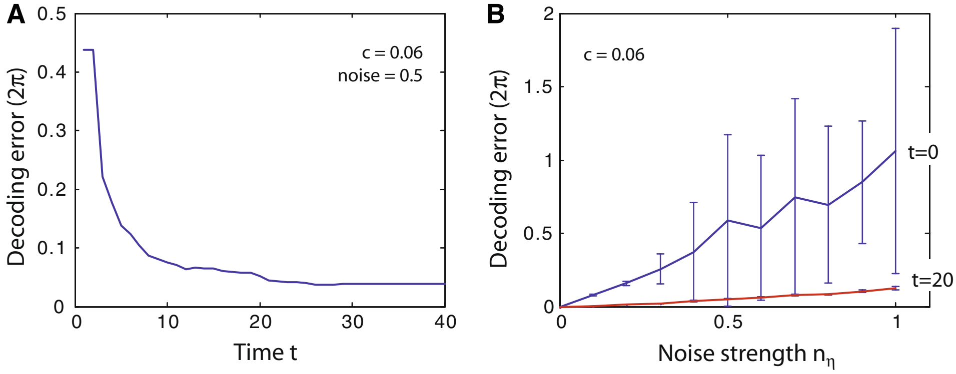

Fig. 4 (**a**) Temporal evolution of the decoding error with optimal inhibition $C = 0.06$ for a noisy signal with noise strength $n_e ta = 0.5$. (**b**) Decoding error with optimal inhibition $C = 0.06$ for different noise levels $n_e ta$ with the center of mass decoding of the original signal $t = 0$ and with the neural fields $t = 20$

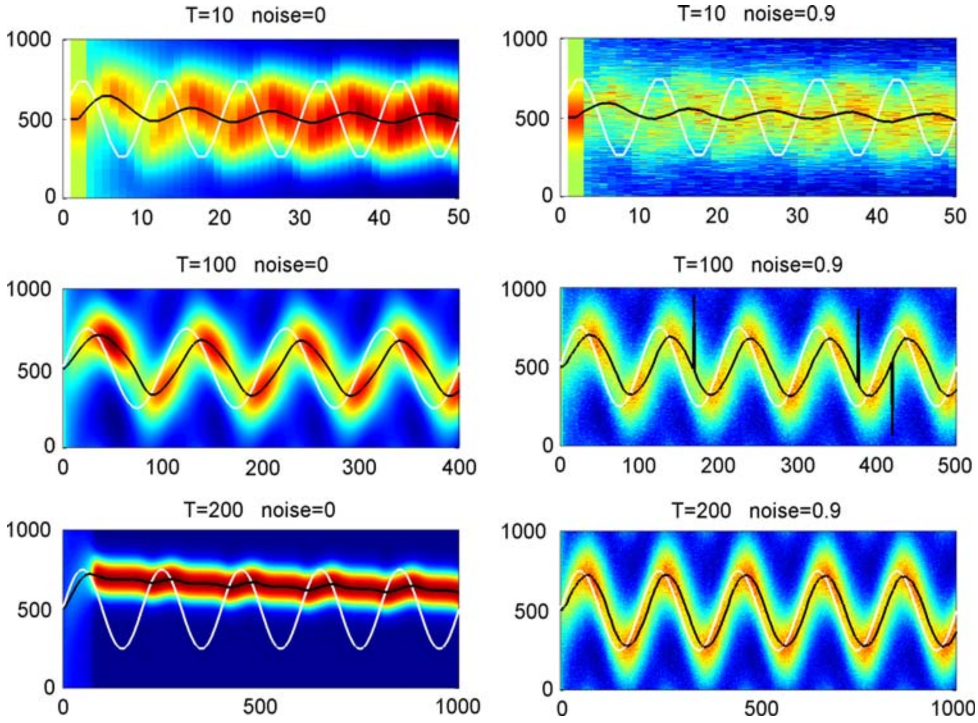

Fig. 5 Online population decoding with time varying signals with no noise ($n_\eta = 0$, left column) and with strong noise ($n_\eta = 0.9$, right column). The time window is chosen to always depict 2.5 times the period. The white lines depict the center of the Gaussian signal and the black lines are the instantaneous estimate based on the center of mass decoding with neural fields

averages over 100 runs and the error bars depict the standard deviations in these runs. While the decoding error does of course increase in both cases, it is interesting to note that the neural field decoding deteriorates much less with increasing noise.

Online tracking of time varying signals

The results reported so far have shown that fast and robust decoding is possible with neural fields even during the time when noisy signals are applied as external input. In this section we further explore the performance of this method with time varying signals. As an example, we consider Gaussian signals similar to the ones used above but where the center of the Gaussian is a piecewise constant function in time. Specifically, the center of the Gaussian is changing with a sine function of period T. The noisy input signal is then chosen at the beginning of each unit time interval from this function and applied as constant input during this small time interval, that is,

$$x_o(t) = \sin\left(\frac{[t]}{T}\right), \tag{14}$$

where $[t]$ denotes the integer component of the continuous time. During each unit time interval there is also a new noise component chosen for the input signal that is applied as external input to the neural field.

Some example time evolutions of the neural field for different frequencies of the time varying signals for the cases of no noise ($n_\eta = 0$; left column) and very noisy signals ($n_\eta = 0.9$; right column) are shown in Fig. 5. In these figures, we indicate the true centers of the time varying signal as white lines, and the instantaneous center of mass decoding with black lines. Note that the illustrated time intervals are adjusted to always show a time span of 2.5 T after starting the simulations with a homogenous field.

It is probably not surprising that the neural field can follow the time varying signal fairly well for intermediate frequencies of the signal. There is, of course, some time lag (phase) between the true signal and decoded one since the neural field needs some time to respond to the stimulus. The signal can also follow fairly well the high frequency stimuli, although the amplitudes of the decoded signal are much smaller than the signal. Thus, the system predicts an underestimation of orientations with increasing frequencies. More surprising is the performance of the system with much lower frequencies. The lower left graph in Fig. 5 shows an example where the decoding breaks down. This problem is still present when increasing the size of the network, although it can not be ruled out that this effect disappears in the continuum limit. However, noise turned out to be very beneficial in this situation, as shown in the right column of Fig. 5. Indeed, the neural field is now able to track even the low frequency stimulus. Furthermore, it is important to note that the decoding performs very well even with these very noisy signals ($n_\eta = 0.9$). While there are some outliers for intermediate frequencies, typically at the locations of largest change in signal orientation, the overall quality is still comparable to the performance without noise.

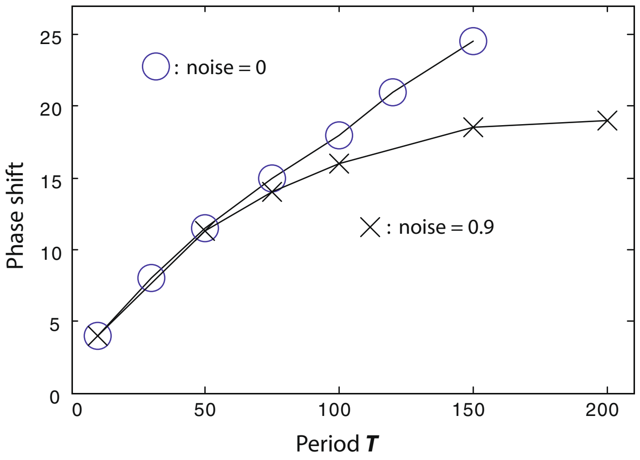

Fig. 6 The phase difference between the input signal and the response of the neural field. The noiseless case, $n_\eta = 0$, scales linear with the period *T*, while in the case with strong noise, $n_\eta = 0.9$, the phase shift approaches a constant

The improvements of decoding with noise for lower frequencies can also be seen in Fig. 6. This figure shows the phase delay between the original signal and the tracking with the neural fields. In the noiseless case this phase shift increases linearly with increasing period *T* of the signal until the tracking breaks down. In the noisy case the phase starts out at similar values to the noiseless case but levels off for lower frequencies. Such a constant delay could easily be absorbed in the processing of signals though backdating of the decoded signal.

Discussion

In this paper we discussed the tracking of noisy population codes with neural fields. This method has gained considerable interest as it not only performs well in the statistical sense (Pouget et al. 1998; Wu et al. 2002), but is also a biologically plausible implementation of how cortical networks with lateral interactions could process noisy information provided to them from lower or higher cortical areas. We showed that this method works well online, when external inputs are continuously applied. Also, we investigated which inhibition constant (which regulates the amount of long range competition relative to short range cooperation) is best suited for population decoding. Interestingly, we found that this inhibition constant should be chosen in a range where the neural field model goes from the memory regime to the regime with decaying activity packets. This seems to indicate that some balance between cooperation and competition is useful and illustrates why simple competitive networks might not provide the best implementation of winner-take-all mechanisms.

The brain faces the challenge of combining a multitude of bottom-up sensory information with top-down expectations to achieve a coherent assessment of the world. This is consistent with the hypothesis that the cortex forms a predictive memory system with different levels of abstractions (Friston 2005; Hinton 2007; George and Hawkins 2005; Dayan et al. 1995). While in this paper we extracted (decoded) the ‘orientation’ from the noisy population code, it is important to note that this was mainly done to provide a quantitative illustration of the performance of the neural field’s ability to track noisy inputs. The extraction of a single parameter, which was done here with a center of mass scheme, is in this view not the main goal of the neural field processing in the brain. Indeed, it should be beneficial to represent some measure of likelihood of features with the population code, consistent with the interpretation of (Pouget et al. 2000). In this view it is also important to note that population coding is not restricted to low level sensory information with tuning curves where population decoding is commonly discussed. It is, of course, more difficult to establish this experimentally since it is much more challenging to measure such higher-level ‘tuning curves’ as the neurons that are involved in the representation of a stimulus might not be organized in the spatial way as the spatial organizations in hypercolumns found in early visual areas.

In this view we could also think about the noise in a different way. While in the experiments shown here we chose the noise term from stochastic processes, we could also imagine that these signals represent additional information from other features of the sensory scene encoded in the population. It is therefore desirable to not completely neglect these feature as would be the case when using the asymptotic states of the neural field without input for further processing. With the value of inhibition found in this study and the strength of input used in the simulations, it seems that we achieved some balance between the information provided by the external input and the activity generate from the neural field itself. A further step in establishing neural fields as component of cortical processing is to study their ability to integrate top-down information, which can itself change with time, with the bottom-up signals processing discussed in this paper.

Acknowledgments I thank Jason Satel for comments on the manuscript and Gregor Schöner for discussing the interpretation of decoding. This research was supported by NSERC (Canada).

References

Amari S (1977) Dynamics of pattern formation in lateral-inhibition type neural fields. Biol Cybern 27:77–87

Dayan P, Hinton GE, Neal R, Zemel R (1995) The helmholtz machine. Neural Comput 7:889–904

Deneve S, Pouget A, Latham P (1999) Divisive normalization, line attractor networks and ideal observers. Adv Neural Inform Process Syst 11:104–110

Deneve S, Latham PE, Pouget A (2001) Efficient computation and cue integration with noisy population codes. Nat Neurosci 4:826–831

Friston K (2005) A theory of cotrical responses. Philos Trans R Soc B 360:815–836

George D, Hawkins J (2005) A hierarchical bayesian model of invariant pattern recognition in the visual cortex. Int Joined Conf Neural Netw 3:1812–1817

Gerstner W (2000) Population dynamics of spiking neurons: fast transients, asynchronous states, and locking. Neural Comput 12(1):43–89

Henry GH, Dreher B, Bishop PO (1974) Orientation specificity of cells in cat striate cortex. J Neurophysiol 37:1394–1409

Hinton GE (2007) Learning multiple layers of representation. Trends Cogn Sci 11:428–434

Hubel D, Wiesel T (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J Physiol (Lond) 160:106–154

Li S, Wu S (2007) Robustness of neural coding and its implications on natural image processing. Cogn Neurodyn 1:261–272

Pouget A, Zhang K (1997) Statistically efficient estimations using cortical lateral connections. Adv Neural Inform Process Syst 9:97

Pouget A, Zhang K, Deneve S, Latham P (1998) Statistically efficient estimation using population coding. Neural Comput 10(2):373–401

Pouget A, Dayan P, Zemel R (2000) Information processing with population codes. Nat Rev Neurosci 1:125–132

Stringer S, Trappenberg T, Rolls E, Araujo I (2002) Self-organising continuous attractor networks and path integration: one-dimensional models of head direction cells. Netw: Comput Neural Syst 13:217–242

Trappenberg T (2002) Fundamentals of computational neuroscience. Oxford University Press, USA

Trappenberg T (2008) Dynamics of population decoding with strong inhibition. In: Rubin W, Fanji G, Enhua S (eds) Advances in cognitive neurodynamics. Proceedings of the International Conference on Cognitive Neurodynamics 2007, Springer

Wilson H, Cowan J (1972) Excitatory and inhibitory interactions in localized populations of model neurons. Biophys J 12:1–24

Wu S, Amari S, Nakahara H (2002) Population coding and decoding in a neural field: a computational study. Neural Comput 14:999–1026

Wu S, Amari S, Nakahara H (2004) Information processing in a neuron ensemble with the multiplicative correlation structure. Neural Netw 17:205–214