Dynamics of receptive fields in the visual system: plasticity of intra-cortical connections

G. Mato and N. Parga
Departamento de Física Teórica C-XI
Ciudad Universitaria de Cantoblanco
Universidad Autónoma de Madrid
28049 Madrid
Spain

Abstract

In this work we study the dynamical changes of receptive fields in a system in which the input has a lesion. A two-layer architecture representing the retina and V1 is introduced and the values of the horizontal connections in the second layer are updated in such a way that the (two-point and higher) correlations between the activities of different cortical neurons are minimized. We find that this algorithm, with a simultaneous dynamics for the activities of cortical neurons and for their horizontal connections, predicts several experimental facts: the expansion of the receptive fields, the bias in feature localization experiments and the *filling-in* phenomenon (in which the lesion is "filled" with the surrounding pattern after a time of adaptation). We find that non-Hebbian terms in the updating rule for the horizontal connections are essential to account quantitatively for the experimental results.

Submitted to *NETWORK*

1 Introduction

The classical concept of receptive field (RF) as the part of the visual space that activates a cell in a static and context independent way has to be changed to take into account experimental evidence showing that the neuron response does depend on the context in which the stimuli are presented. This is a manifestation of the dynamic nature of the RF. A reorganization of RFs was observed in cases where some time (of the order of minutes) after the occurrence of a retinal lesion the area of the RF increased by a factor order 5 [16], and cortical cells with their classical RF inside the damaged region recovered their activity.

These changes occur not only in the visual system, modification of cortical representation of sensory inputs induced by lesions have been also observed in the somatosensory system: digit amputation [29] shows that the area of somatosensory cortex where the amputated piece was originally represented becomes, after some time, sensitive to stimulation of nearby regions.

It was soon realized that similar effects could be obtained without the existence of real lesions. Stimuli can emulate the lesion if they are "localized"; that is, if there is some small part of input space that receives stimulation strongly different from their surrounding. Lack of stimulation in a small region of the visual space produces an effect similar to a scotoma. But now there is also the possibility to stimulate the small region more strongly than the rest of the visual field. Experiments with localized stimuli have been done in both the visual [36, 32, 25] and the somatosensory systems [23].

In the case of the visual system this was done by presenting stimuli that have an "artificial scotoma". If a stimulus consisting of a pattern of dynamic lines with a central square of uniform luminosity is presented, cells in the cortical scotoma initially decrease their activity. Using classical RF concepts, these cells should remain inactive regardless of what is going on in other places of the visual field. However their RFs expand, allowing a recovery of activity. Measurements show that after presenting the stimulus the RF area is found to increase by a factor of order 5. Apart from this expansion the response itself of the neuron to a point stimulus at the center of the RF increases by approximately a factor 2 [32]. The effect can be reversed by stimulating the central region. No attempt to study the increase of the RF size as a function of stimulus presentation time was done.

It has been speculated that this expansion can account for the perceptual *filling-in* effect [36] and other visual illusions [25, 14]. In perceptual *filling-in* an stimulus with an "artificial scotoma" is presented and it is observed that after a few seconds the central region is filled by the surrounding pattern [36]. A related phenomenon, more similar to a permanent lesion, is the fact that under normal conditions the blind spot on the retina is not perceived but it also appears filled by the surrounding part of the visual scene.

The changes observed in the RFs of cortical neurons when these localized stimuli are applied can be quantitatively studied with psychophysical experiments. For instance, changes in RF sizes are reflected in a systematic bias in feature localization tasks. It has been found [25] that the ability to determine the relative position of a short line segment in the middle of another two, presented close to the border of the artificial scotoma, was strongly biased in a way that is consistent with the expansion of RFs of neurons in the cortical scotoma. The magnitude of this visual illusion can be studied as a function of the presentation time of the stimulus. Surprisingly enough it was found that the shift in the perception appears already after about one second of presentation of the artificial scotoma and it roughly keeps its value for longer times. Kapadia et al. have also observed that the *filling-in* illusion appears later, about 10 seconds after the presentation of the stimulus.

In the somatosensory system it is found that behaviorally controlled tasks, that stimulate a limited sector of the skin in one or more fingers, generate an expansion of the area in the cortex that represents that region and a *contraction* of the RFs inside the stimulated area [23]. This result is analogous to the one mentioned above for the visual system, where a lack of stimulation in some region leads to an *expansion* of the RFs. Although these phenomena in visual and somatosensory cortex seem to be of a similar nature, in this work we will be mainly concerned with the visual system.

To explain these phenomena one first needs to know the anatomical substrate for the reorganization of RFs. The first question in this regard is whether these changes are generated at the cortex or at some earlier level in the sensory pathway. It has been observed that the scotoma in the lateral geniculate nucleus (LGN) does not disappear [16], suggesting that the changes take place after the LGN. Also supporting this claim, it has been found that injecting different retrograde tracers inside and outside of the original cortical scotoma generates two non-overlapping clusters in LGN. This suggests that the *filling-in* of the cortical scotoma is not mediated by the spread of geniculate afferents on the cortex [16].

More direct evidence of the role of intracortical connections in the remapping of the RFs was obtained by comparing the optical and spike point spreads (PS) before and after the retinal lesions. The optical PS are much larger than the spike ones, and are thought to be generated by long range horizontal connections radiating from the spike PS. These connections are found to be effectively excitatory at short range and effectively inhibitory at long range. After the lesion is presented for some time it is found that the spike PS has expanded to a size approximately equal to the size of the optical one before the lesion [11], suggesting that some long-range inhibitory intracortical connections have become excitatory. The conclusion that the locus of the changes of the RF sizes is the cortex has been also reached for the somatosensory system [23].

We model the visual cortex as a set of neurons described by their firing rate. They receive inputs from both feedforward connections coming from previous layers in the visual path and from horizontal connections coming from other neurons in the same area. Cortical activity and horizontal synaptic plasticity will be assumed to follow a simultaneous dynamics, although with different time scales. For the activity dynamics we propose a rather standard evolution equation and a fast integration time. Horizontal connections are treated in more detail because we want to identify the relevant properties of their plasticity changes that might account for the experiments that we described above. We then update these connections according to a rule that contains not only a Hebbian term but also higher order non-linear contributions; these terms are local and they depend on the activity of both neurons. Let us notice however that not all the contributions can be instantaneous. If so, immediately after presenting the artificial scotoma one cortical neuron looking at the lesion will be silent and remain in this way until this stimulus is taken away. As a consequence no plasticity changes will occur. A possible (an natural) way out of this problem is to use a generalized covariant rule where changes in horizontal connections depend on the difference of the firing rate with respect to a (time dependent) average activity.

As we said the emphasis in the modelization is put in the way horizontal connections are treated. Other aspects of the model will be simplified: there is no dynamics for the feedforward connections, although they spread according to the classical RF size. Neurons in V1 are directionally selective, but we also disregard this fact. Finally, considering the evidence that cells in the LGN scotoma do not recover their activity [16] (at least to the same extend that cortical cells do it) we do not include an intermediate layer between the retina and V1.

Although one can consider the model just as an empirical proposal to analyze the kind of synaptic plasticity necessary to reproduce the phenomenology of artificial lesions (and the present work can be read at this level), one can also look at it from a more theoretical perspective. The basic hypothesis is that the cortical representation is an efficient code for the information content of the stimulus [4, 5]. In other words, it is constructed in such a way that it can discriminate well between the most frequent stimuli. Information theory [7] is the correct framework to deal with this approach and it has been used to describe RFs developed under normal conditions [2, 18, 19]. The claim here is that this efficiency criterion can also be taken as a guide to define the form of the synaptic plasticity changes responsible for the RF dynamics.

In the absence of damages or persistent localized stimuli, that we take as the normal situation in, e.g., natural visual scenes, the RFs become the classical ones. In the presence of retinal lesions or under a persistent stimulus with an artificial scotoma, the application of the same ideas suggests that connections to neurons inside the cortical scotoma must change (provided that the anatomical constraints allow it). In fact, since these neurons are not stimulated they do not participate in the construction of an efficient code of the current statistical properties of the environment. However the existence of horizontal connections makes it possible for these neurons to collaborate in the construction of a better cortical representation of the surrounding, non trivial visual space. From this perspective the simultaneous dynamics of neuron activity and horizontal connections that we mentioned before represents an adaptive algorithm for the suppression of redundancy in the cortical code [24, 20]. In particular the dynamics for synaptic plasticity has a fixed point where the temporal average of every term contributing to its modification (both the Hebbian and the higher order ones) becomes zero. This guarantees that the two neuron activity correlation and a higher order correlation (e.g. a third order one) become zero: redundancy is reduced. Once the algorithm has converged to this fixed point, one can check *a posteriori* if other higher order correlations have also disappeared.

There are several adaptive algorithms that implement a code with minimal redundancy [24, 20, 9, 31]. In this work we have chosen the one by Herault and Jutten both for its simplicity and its biological plausibility. It has been applied before to the olfactory system as an strategy to discriminate odors in a mixture using high order fluctuations [22, 38].

Our purpose here is to explore within the model the possibility to reproduce the experimental facts that we summarize bellow:

- Changes of RFs, including their correct expansion under localized stimuli, activity changes of neurons in the cortical scotoma and the recovery of RFs to normal size when the artificial scotoma is taken away,

- The relevance of the different model parameters. In particular the necessity of including non-linear terms in the rule for synaptic changes. This is an important element of the model when it is seen from the point of view of redundancy minimization. If they are excluded by the data then the algorithm performs only decorrelation (i.e., it only eliminates correlations between pairs of neurons)

- The perceptual shift in the Kapadia et al. psychophysical experiment as a function of the relative position of the stimulus and the damage,
- The time scale at which the previous effect becomes observable, and its relation to the time scale where the *filling-in* effect appears.

Making compatible all these points is not trivial. If the time constant for the evaluation of the mean values is too small the expansion of the RF sizes will be small. On the other hand choosing this time constant very large, we will obtain that the psychophysical shift will not grow fast enough. However, by increasing the speed of the updating of the horizontal connections this problem can be overcome, and we can find a region in the space of parameters where the experimental constraints are satisfied.

The way that long-range connections behave is complex, showing both excitatory and inhibitory properties according to the context. Contextual dependence of the RF of neurons in V1 has been observed in many physiological [21, 27, 17] and psychophysical [34, 35, 26] works, and in some cases the compatibility between the results of both approaches has been verified as is discussed in [15]. Complex, strong stimuli tend to suppress the response of a neuron to its preferred stimulus while a weak stimulus tends to facilitate it [21, 27, 17]. In this work we will be mainly concerned with complex stimuli to describe the effect of normal visual conditions and with zero stimulation of the classical RF to describe the artificial scotoma. We then expect a suppressive effect of the surround in the first case and the appearance of facilitation as the network adapts to the scotoma in the second situation.

Summarizing, in this work we intend to apply the criterion of minimization of redundancy to a situation in which the input patterns contain an artificially damaged region. We will numerically evaluate the horizontal connections that minimize the correlations and then we will investigate the effect of this modification on the RFs. In the next Section we will introduce the model, given by the network architecture, the neuronal and synaptic dynamics, and the statistics of the inputs. In Section 3 we will show the results concerning the changes of RF sizes, the reversibility of these changes, the psychophysical shifts and the *filling-in* effect. An analysis of the effect of changes in the activation function is also presented in this section. In Section 4 we present the conclusions and discuss some possible extensions of this work.

2 The Model

2.1 Architecture

The architecture was chosen making several simplifying hypothesis,

- The model used in the simulations is a two-layered neuronal network. The first layer represents the retina, where the different stimuli are applied. The second layer represents the primary visual cortex. We do not include an intermediate layer to represent the LGN. It is known that cells in the LGN scotoma do not change appreciably their RF [16]. As we said in the Introduction we want to simplify all the elements that do not have a strong effect on the problem and concentrate instead on the horizontal connections.

- The orientation selectivity of cortical neurons is not explicitly included in our model. We assume that we are modeling a small portion of cortex with a fixed preferred orientation. Let us notice that this is compatible with the results by Das and Gilbert [11] that suggest that the response of a neuron with its classical RF inside a permanent lesion, recovers a preferred orientation identical to the one it had before the lesion.

- Both layers consist of an square array of N^2 units each. We are also assuming that each of our units does not represent one neuron but a small population of them. This is mainly because of computing limitations that do not allow us to simulate very dense networks. Consequently, synaptic connections have to be interpreted as effective ones. As we will see in the presentation of the simulations they can change their sign, from inhibitory to excitatory or viceversa. We will however discuss the effect of increasing both the size and the density of the layers.

- The cortical scotoma will be placed at the center of the second layer. We assume that the layer boundaries do not affect significantly the center of the layers and we take periodic boundary conditions.

The inputs are denoted by v_j and the state of the units in the second layer by u_i. The feedforward connections from the first to the second layer are excitatory and will be denoted by J_{ij}. These connections are fixed and perform a topographic mapping between the two layers (see Fig. 1). The divergent connections from a unit in the input layer have a maximum value when they reach a unit in the equivalent position in the second layer, and their strength decays according to a Gaussian function. If $d(i, j)$ is the distance between input unit i and output unit j:

$$J_{ij} = \frac{1}{2\pi\sigma_f^2} e^{-d(i,j)^2/2\sigma_f^2}. \tag{2.1}$$

The dispersion σ_f describes the spread of the feedforward connections. $d(i,j)$ is measured in terms of the number of lattice spacings between cells. The first factor was chosen in such a way that the total strength of the divergent connections is one (for each input cell) in the limit of a dense input layer.

The horizontal connections between cortical neurons i and k will be constructed as a product $D_{ik}T_{ik}$ of a truly dynamical part T_{ik} and a static factor D_{ik}. The first connects all possible pairs of cortical cells while the static piece D_{ik} effectively limits the range of the horizontal connections. These are chosen once and for all according to a gaussian dependence on the distance between neurons

$$D_{ik} = \frac{1}{2\pi\sigma_l^2} e^{-d(i,k)^2/2\sigma_l^2}. \tag{2.2}$$

Again distances between cells on this layer are measured in terms of the number of lattice spacings between cells. Let us notice however that in general lattice spacings on the first and the second layers are not necessarily equal when expressed in physical units. The parameter σ_l denotes the spread of horizontal connections.

We will discuss the spike and the optical point spreads and the evaluation of RFs at the end of this section, once the dynamics of the horizontal connections is presented. Let us only mention for the moment how the parameters σ_f and σ_l affect these quantities. Both parameters contribute to the optical point spread which, under normal visual stimuli, should have a linear extension of roughly $\sigma_l + \sigma_f$: all cells within this radius would produce either EPSP's or IPSP's detected by optical imaging techniques [11]. On the other hand the spike PS, again under normal visual conditions, can not be much larger than σ_f; this is because at least beyond a excitatory core horizontal connections are inhibitory. According to Das and Gilbert [11] the area of the spike PS is about 5% of the optical PS, this means that σ_l is larger than σ_f by roughly a factor 2.

When the visual stimuli contain a scotoma, the dynamic factor in the horizontal connections, T_{ik}, will change. This will modify the spike PS but will leave the optical one almost unchanged. If it happens that the dynamics of the T_{ik}'s is such that in the presence of the scotoma some of them become less inhibitory or excitatory then the spike PS will increase in size.

The classical RF is not determined only by σ_f, although of course it does depend on this parameter. The range σ_l of the horizontal connections will also define the size of the classical RF *together with* the dynamics of the horizontal connections resulting from the network adaptation to the statistics of an ensemble of visual stimuli. Then both feedforward and horizontal connections contribute to define the classical RF. The difference between this RF and the dynamic one is that the latter results from adaptation to some persistent features in the stimuli. The problem here is how to choose the ensemble of visual stimuli that defines "normal" visual conditions and determines the classical RF. This point will be made more precise in the next section when discussing the statistics of inputs.

2.2 Dynamics and synaptic plasticity

The dynamics for the activity of units on the second layer is given by:

$$\tau \frac{du_i}{dt} = -u_i + H\left(\sum_{j=1}^{N^2} J_{ij} v_j - \sum_{k=1}^{N^2} D_{ik} T_{ik} u_k\right) \tag{2.3}$$

where v_j denotes the strength of the stimulus on cell j in the input layer and u_k is the firing rate of cortical neuron k. Let us observe that as D_{ik} is always positive, a positive value of T_{ik} means an inhibitory interaction and a negative value an excitatory one. The parameter τ represents the membrane time constant and H is an activation function that has the form

$$H(x) = \begin{cases} 0 & \text{if } x < 0 \\ x & \text{if } \theta > x \geq 0 \\ \theta & \text{if } x \geq \theta \end{cases} \tag{2.4}$$

Saturation is introduced to avoid divergence of the network activity in the case where excitatory horizontal connections are generated by the dynamics. This can happen for certain values of the parameters of the updating rule for the horizontal connections.

The membrane time constant is known to be of order 10 msec. This can be assumed to be faster than the other time constants present in the problem (modification of synapses, presentation of visual stimuli). Therefore we can assume that the system is always at the fixed point of eq.(2.3):

$$u_i = H\left(\sum_{j=1}^{N^2} J_{ij} v_j - \sum_{k=1}^{N^2} D_{ik} T_{ik} u_k\right) . \tag{2.5}$$

Before giving the equation for changes of the synaptic plasticity of horizontal connections we need to define the mean value of a generic function f of the neuron firing rates in a time window of size T

$$< f >_T (t) = \frac{1}{T}\int_{-\infty}^{t} f(t')e^{(t'-t)/T}dt', \tag{2.6}$$

and its instantaneous separation from this mean

$$\delta_T f(t) = f(t) - < f >_T (t). \tag{2.7}$$

The updating of the horizontal connections will be done according to

$$\tau_s \frac{dT_{ik}}{dt} = \alpha\delta_T u_i \delta_T u_k + \beta\delta_T u_i^2 \delta_T u_k + \gamma\delta_T u_i \delta_T u_k^2. \tag{2.8}$$

This rule will be applied for $i \neq k$. The self-interaction terms $D_{ii}T_{ii}$ will be kept constant at a value proportional to $J_{ii} - 1$ (see Appendix). Let us remark that this updating rule has two time constants. One is given by the averaging time window size T and the other corresponds to the typical value of the constants α, β and γ.

Even if the stimuli v_j were uncorrelated, the activity of the units on the second layer will be correlated because of the spread of the feed-forward connections. Eq.(2.8) is designed to eliminate these correlations. It implements an algorithm proposed by Jutten et al. to perform independent component analysis [24, 20]; this means that it has as a fixed point a solution for the T_{ik}'s such that the cortical neurons retrieve the statistically independent components of a signal (the stimulus v_j)[1].

This fixed point gives an efficient code of the typical stimuli because all redundancy has disappeared [30]. Strictly speaking this is not the ideal solution because in practice the network is noisy and a good cortical code needs some redundancy to compensate for the noise. For the sake of simplicity we are neglecting this effect and we think of our cortical units u_k as a population of true cortical cells with similar properties. A complete algorithm that builts an efficient code in a noisy environment (e.g. an adapted version to our problem of the one proposed in ref [28]) would find these populations automatically, but it could be computationally costly. Let us also remark that eq.(2.8) could also converge to a situation where its right hand side is zero globally instead of being zero the average of each term. Because of this we will explicitly check the kind of fixed point where it arrives.

[1] Actually ours is a slight generalization of the Herault and Jutten algorithm in that the full horizontal connection is not T_{ik} but its product with D_{ik}. By multiplying both sides of eq.(2.8) by D_{ik} one notices that this coupling obeys an updating rule where the constant coefficients α, β and γ are replaced by constant matrices α_{ik}, β_{ik} and γ_{ik}.

2.3 Statistics of inputs

An important point to completely determine the RFs is the description of the visual stimuli. The equations above give the evolution of cortical activity and synaptic weights in terms of the inputs v_j and they will undergo strong changes each time their statistics changes. Here we will describe the different kind of visual stimuli that we need and the way the statistics associated to each of them will be represented.

As we mentioned in the Introduction we want to understand the predictions of this model for experiments where visual stimuli present an artificial scotoma. This defines one of the ensembles of images that we will use. According to the experimental work [32, 25] these images consist of a dynamic background of a given mean luminosity and a central region of uniform intensity, but with the same mean value as the surrounding. Different variants for the dynamic part of the image are random dots, bars, etc.

However the experiment is done starting from an initial condition given by the values that activities and horizontal connections take under "normal" visual stimuli. We want to make here more precise what we understand by normal conditions. As horizontal connections change they induce modifications in the RF, however if these changes occur in a time scale large compared to the time scale in which the stimuli change, it makes sense to talk about a static RF and this is determined by the statistics of those stimuli. If now a portion of the visual space is kept fixed during a time comparable to the time scale of synaptic plasticity this change in the statistics will affect the RF. What is important in the experiment is the change from a stimuli statistics with a scale shorter than the one of the synaptic dynamics (normal conditions) to another with a persistent feature such as the artificial scotoma. One then compares the RF obtained in the first case (the classical RF) with the one obtained with the localized stimuli. However we will be more precise and define normal conditions in terms of a statistics characterized by a power spectrum $S(f)$ behaving as

$$S(f) \propto |f|^{\eta-2} \tag{2.9}$$

where η is a small number. This power spectrum was found in recent studies of the statistics of natural images [12, 39]. The ensemble power spectrum, that is the Fourier transform of the correlations between pairs of image pixels (averaged over orientations), is found to behave as in eq.(2.9) with $\eta \approx 0.2$ [39].

Summarizing, two kinds of image ensembles will be presented to the system. The first will represent normal visual stimuli, selected with a power spectrum as in eq.(2.9), that the retina receives before the lesion. It will give rise to classical RFs. The second kind represents images with an artificial scotoma. These images consist of uncorrelated inputs (i.e. random dots) with a given mean luminosity and a uniform central region of intensity equal to that mean.

2.4 Analysis of the dynamics

The system of coupled equations (2.5) and (2.8) is not analytically solvable. But we can analyze them qualitatively in order to identify the influence of the different parameters on the dynamics of the RFs.

After the system has adapted to the statistics of normal visual conditions the horizontal connections are inhibitory, with a range σ_l. When stimuli with an artificial damage are introduced the activity of a neuron inside the cortical scotoma will fall to zero because the input is absent, the horizontal connections give an inhibitory contribution and the activation function gives zero output for any negative argument. A neuron with its RF only partially inside the scotoma will decrease its activity in proportion to the fraction of its RF inside the scotoma. Let us call i (j) the index of a neuron totally inside (partially outside) the scotoma. Initially the modification of the weight between neurons j and i will be, approximately:

$$\tau_s \frac{dT_{ij}}{dt} \approx -\alpha < u_i >_T \delta_T u_j + \beta < u_i >_T^2 \delta_T u_j - \gamma < u_i >_T \delta_T u_j^2 \tag{2.10}$$

Since the introduction of the scotoma decreases the activity of neuron j, $\delta_T u_j$ will be initially smaller than zero. This means that the first term will have the sign of α, while the second and the third will have signs opposite to β and γ respectively. Therefore a positive α will (at least at the beginning) generate even more *inhibitory* connections that instead of expanding the RF will tend to contract it. The opposite is found for β and γ. This analysis suggests that one should search in a region with small positive or even negative α and β and γ positive.

We can immediately see that the averaging time T is essential for the expansion of the RF. If it were zero or very small the average activity of neuron i, $< u_i >_T$, would go very fast to zero and there would be no modification of the connections impinging on this neuron. The necessity of having some kind of memory effect was also noticed by several authors. Clothiaux et al. [8] use a memory effect in the context of synaptic changes of feedforward connections based on the BCM model [6] and it is also a relevant ingredient in some models for invariant object recognition [13, 41].

We can expect that the expansion will increase as T increases, although a quantitative analysis requires numerical simulations. On the other hand a high value of T will have the effect of making that RFs (and the psychophysical shifts) to grow more slowly, because it will take a longer time to "forget" the state of the system previous to the presentation of the scotoma. This effect can be partially compensated by decreasing τ_s, that regulates the speed of growth of the horizontal connections.

3 Numerical Analysis

3.1 Implementation of the algorithm

The inputs were generated according to the statistics defined in the previous section. A new input was presented every 10 time steps, both for normal visual stimuli and for the artificial scotoma. This fixes the time scale unit as the number of presented patterns and allows the comparison with psychophysical experiments. For instance the dynamic background of random dots used with the artificial scotoma in ref.[25] was refreshed at a rate of 30 Hz, this means that one time unit expressed in physical units is about 33 msec. The time scale T (used in the evaluation of activity averages) will be also expressed in the same units. For instance, $T = 50$ means that averages will be evaluated in a time window corresponding to the presentation of 50 patterns or, again in the case of ref.[25], about 1.6 sec. The numerical results presented in the next section were obtained by integrating eq.(2.8) using a first order Euler integration scheme. The time step was set at $\delta t = 0.1$. We checked that using smaller time steps (keeping the total integration time fixed) or a fourth order Runge-Kutta integration scheme does not change substantially the results. The mean value of the inputs v_j in the stimuli with scotoma was chosen to be 1.2 and the dispersion 0.3 (except in the scotoma where the dispersion is 0). The stimuli with the statistics of the natural images have the same mean value. The saturation value of the transfer function H was chosen as $\theta = 2$. With this choice, the internal field $\sum_{j=1}^{N^2} J_{ij} v_j - \sum_{k=1}^{N^2} D_{ik} T_{ik} u_k$ is typically in the linear region of the activation function.

Let us remark that an arbitrary scaling factor in the values of u_i can always be introduced by rescaling the inputs and the constants α, β and γ in a suitable way. For this reason the interpretation of u_i in terms of firing rates is quite arbitrary.

Most of the simulations have been performed in systems with 16×16 units in each layer. The artificial scotoma was taken as a square of 7×7 units. We have already discussed how the size of the RFs and point spreads are limited by the range of the feedforward and horizontal connections. These are chosen as $\sigma_f = 2$ and $\sigma_l = 4$. The spread of the feedforward connections is taken to be smaller than that of the horizontal ones because thalamic connections have a shorter range than intracortical connections [16] and also to represent properly the relative size of the spike and optical PS. The absolute value of the spreads is chosen in such a way that the RF are smaller than the size of the artificial scotoma by a factor approximately 2. These relative sizes are quite compatible with those used in some of the experiments [32].

3.2 Simulation of the experiments

The simulations are composed of several operations that depend on the experiment one wants to simulate:

- The first thing we need is to put the system in a state corresponding to normal visual conditions, as defined before. To do this the horizontal connections are first initialized to zero (except the self-interaction terms) [2] and stimuli generated with the statistics corresponding to eq.(2.9), are presented at the input layer. Lateral connections are then updated according to eq.(2.8). We check convergence and the nature of the fixed point by measuring activity correlations up to fourth order as a function of time.

- Once the network is in its normal state we evaluate the (classical) RF of some of the output neurons by presenting a point stimulus on the input layer at several positions and simultaneously measuring the response of the selected unit on the second layer.

- In another operation, performed after the network converged under normal conditions, the inputs display the statistics of images with an artificial scotoma. The change of ensemble induces a sudden increase in the correlations, that should decrease again to zero as the algorithm converges. In this case we are interested in the dynamics itself of the network.

- We will simulate the experiment of ref.[25] at several intermediate steps. This means that taking the current state of the network a stimulus made of three small horizontal bars is presented and the response of cortical neurons is measured. This operation is repeated placing the stimulus at different relative positions with respect to the scotoma.

- We can also evaluate the RF as a function of time by presenting again a point stimulus at different places on the input layer, but in particular we are interested in the RF when a

[2]This is somewhat arbitrary, we have checked that initialization to randomly chosen values does not affect the stationary state of the network.

stationary state is reached. This should give an expansion of the RF that can be compared with ref.[32].

4 Results

4.1 Changes of receptive field size

We first show the results obtained when the parameters of the algorithm are $\alpha = 0.2$, $\beta = 0.8$, $\gamma = 1$ and $T = 200$. In Fig. 2a we can see the time evolution of the activity correlations of two and three cortical neurons given by the average of each of the terms that appear in the updating of the horizontal connections. The artificial scotoma was introduced at $t = 200$, after adaptation to normal conditions had occurred; we can see that the change in the statistics makes the correlations to increase but after some time the network adapts to the new stimuli and they again converge to zero. In Fig. 2b we show that also third and fourth order correlations cancel, suggesting that the algorithm has indeed succeeded in reducing the redundancy in the output units.

The RFs obtained after adaptation to the statistics of normal visual conditions and after adaptation to the inputs with the scotoma are shown in Fig. 3a (for a unit in the cortical scotoma) and in Fig 3b. (for a unit with the RF well outside the scotoma). We can see that a unit inside the cortical scotoma has increased its maximum response by about 100 %, while the other is unchanged.

The size of the RF expansion depends on the parameters of the model. It is easy to understand the effect of T on this expansion. When the artificial damage is presented the activity of units in the cortical scotoma goes immediately to zero. If T is relatively small, their average activity $< u_i >_T$ would also vanish and there would be no substantial modifications of the horizontal connections affecting those units and no expansion of their RFs. Fig. 4a shows the dependence of the ratio between the maximum responses after and before the adaptation to the scotoma on the time constant T. We can see that as T decreases the ratio tends to 1, meaning that no expansion has occurred. In Fig. 4b we show the dependence of the expansion on τ_s. Increasing τ_s is equivalent to taking smaller values for α, β and γ, and leads to a smaller expansion of the RFs.

The effect of the other parameters is more complex. In Fig. 5 we show the effect of varying α fixing $\gamma = 1$ and $\beta = 0$; and also for $\beta = 1$ and $\gamma = 0$ and for $\beta = 0.5$ and $\gamma = 0.5$. For all these cases we observe that the expansion is a monotonously decreasing function of α, for positive values of this parameter. Let us also remark that these results are consistent with the first-order analysis performed in Section (2.4), where it was predicted that a positive value of α would lead to inhibitory connections and contraction of the RF. Negative values of α generate a contribution that tends to *correlate* the activity of the output units. This means that excitatory connections will be generated. If $-\alpha$ is large enough these excitatory connections become so strong as to make the network unstable, in the sense that even a very small input will lead the whole system to its saturation value. In this situation the output does not contain any information about the input images. This gives a lower bound on the admissible values of α.

The higher order terms in the updating rule, i.e. the ones proportional to β and γ, *are indeed necessary* in order to obtain an expansion of the RFs as large as the one observed in the experiments [10]. Fig. 6 shows the results for the relative expansion as a function of α when $\beta = 0$ and $\gamma = 0$ for different values of T (for values of α smaller than those shown in the figure the algorithm fails to decorrelate). We can see that even a large T does not give a sufficiently large expansion. By running simulations with networks of size 32×32, we have checked that this result is not affected by finite size effects.

4.2 Reversibility and contraction of receptive fields

We have also considered the effect of re-stimulating the region of the artificial damage after the network has adapted to it and the RFs of cells in the cortical scotoma have expanded. In this case we find that after normal visual conditions are restored the RFs collapse to their original size (see Fig. 7a). Stimulating the central region more strongly than the background there is a further contraction of the RFs, (Fig. 7b). In experiments of sensory stimulation or training on a tactile discrimination task, contractions of RF size have been reported [23, 37]. According to our simulations similar results should be also observed in the visual system.

4.3 Shifts

In this section we analyze some psychophysical consequences of the RF expansion. In particular we will be concerned with the alterations that this effect induces in spatial localization task. In the experiment described by Kapadia et al. [25] the subject first adapts to stimuli with an artificial scotoma, after some time these are replaced by another stimulus stimulus composed of three horizontal bars disposed along the vertical axis. The central bar is positioned in one of seven positions, either centered between the top and bottom bars or shifted by a fixed distance one, two or three units above or below the center. The task consists in saying whether the middle bar appears closer to the top or to the bottom bars. The experiment shows that if the test bars are presented in a position close to the place where the border of the artificial scotoma was located there is a systematic bias in the perceptual responses. The central bar appears to be shifted upwards, if it is presented near the lower border of the scotoma.

To relate the results of the simulations with these psychophysical results it is necessary to propose how the system estimates the position of the bars from the activity of cortical cells. One estimator of the position along one of the axis (say, y) is the *center-of-gravity* of the neuronal activation profile [3]. It can be obtained in the following way: let us de denote by w_y the average of the activities of all the units whose RF center is at coordinate y. A estimator E of the position of any stimulus will be

$$E = \frac{\sum_y y w_y}{\sum_y w_y}. \tag{4.11}$$

If we present as a stimulus three bars with the one in the middle at the center of the other two, denoting the estimation of the three positions by E_1, E_2 and E_3, the shift in the estimation will be

$$S = E_2 - \frac{E_1 + E_3}{2}. \tag{4.12}$$

In Fig. 8 we show the result of the time evolution of this quantity. It grows fast as soon as the stimuli with lesion are presented, reaching saturation after a time that depends on T and on the values of the parameters in the updating rule (α, β and γ). For the same parameters as in Fig. 2, we find that the shift saturates about 30-40 time units after the stimuli with scotoma are presented. If the patterns are presented with a rate of 30 Hz, as in [25], this means a convergence time of about 1-1.2 secs., which is roughly the same as in the experimental results [25]. Increasing α, β and γ (or equivalently decreasing τ_s) or decreasing T makes the convergence faster. On the other hand a smaller T implies that the time window over which the changes take place is smaller and that the total maximum value of the shift will be therefore smaller. The relation between this time scale and the one at which the *filling-in* effect occurs will be discussed in the next subsection.

We also checked the dependence of the shift with the position of the stimulus. In Fig. 9 we show that the shift is maximum when the stimulus is centered just at the border of the scotoma, agreeing with the results of [25]. If the test stimulus is presented at a suitable distance from the scotoma we find that the sign of the shift is reversed. We checked that this result is not a finite size effect by running simulations with a network of 32×32 neurons in each layer. We also checked that this phenomenon is robust to changes in the parameters of the updating rule. The psychophysical results do not clearly show this effect, possibly because the test stimulus was not placed far enough from the scotoma.

4.4 Filling-in effect

We now examine the question of the recovery of the activity of a neuron in the cortical scotoma. We interpret that the *filling-in* effect occurs when the adaptation of the horizontal connections to the artificial damage leads the activity of a neuron in the cortical scotoma to a value comparable to the one of the neurons outside it.

In Fig. 10 we show the level of activity of one neuron in the cortical scotoma compared with the average activity of the whole system for $\alpha = 0.2$, $\beta = 0.8$, $\gamma = 1$ and $T = 200$. We can see that they become similar about 70 time units after the stimuli with the scotoma are presented. In Fig. 8 we show the time evolution of the shift for the same parameters. We can see that it reaches its saturation value about 30-40 time units after the presentation of the scotoma. The comparison of the time scales shows that it takes longer to achieve the *filling-in* than to saturate the shift, although the difference is smaller than the factor 5-10 found in [25]. The value of the saturation activity of the neuron in the scotoma depends on the parameters α, β, γ and T. For fixed values of α, β and γ it increases monotonically with T. If T is too small the activity will saturate to a value smaller than the average. If it is too large it will saturate to a value larger than the average. For that reason the *filling-in* phenomenon imposes a constraint on the model parameters.

Let us remark that for the set of parameters used above all the experimental constraints (expansion of RFs, psychophysical shifts, *filling-in*) are satisfied. However a precise tuning to these values is not necessary to account for the experiments. There is a whole region in parameter space that does it. For instance, all those experiments are still explained by the model if α is taken in the range $[-0.25 : 0.5]$. On the other hand, changes of T affect more strongly the results, specially the time when the *filling in* effect occurs. Depending on the precision required for the similarity between the firing rate of the neuron in the cortical scotoma and the its average value on the layer, a change in T will have to be compensated by changes in the rest of the parameters.

4.5 Discussion of the results

As we have seen in Sec 4.1 , the lowest order term in the equation for the synaptic plasticity, eq.(2.8) can not account by itself for the phenomena we want to explain. For the values of the time constants required for the appearance of the perceptual shifts or the *filling-in* effect one can not find a value of the parameter α giving the correct expansion of the RF if $\beta = \gamma = 0$. However the maximum expansion that can be achieved for $\beta = \gamma = 0$ is not very far from the wished value. For this reason we have also explored the effect of changing the activation function used in eq.(2.3).

Keeping this function as in eq.(2.4) we first varied the slope g of its linear piece. Taking $\alpha = -0.1$ as the only non-zero parameter in the learning rule (this value of α favors a large expansion of the RF) we observed a monotonous increase in the RF size as a function of the gain but again saturation occurs before the correct expansion is reached.

We have also analyzed the effect of taking the activation function as a sigmoid function:

$$H(x) = 1/(1 + \exp(-g(x - x_0))). \tag{4.13}$$

The different non-linearity of this function does not have the desired effect in the value of the expansion of the RF. Similarly to the previous case the algorithm fails to converge before a convenient expansion can be obtained.

We conclude that, within the framework of this model, the usual covariant rule can not explain all the phenomena. To obtain the correct expansion of the RF within a time scale of the order of a few seconds, as observed in psychophysical experiments, requires the higher order terms in the learning rule. In the next section we discuss some possible modifications of the learning rule that could be done in order to check the robustness of this conclusion.

5 Conclusions

In this work we have analyzed the dynamics of RFs on the basis of some experimental results on damaged (artificial or real) sensory inputs [16, 32, 25, 36, 23, 29, 37]. The dynamics we used can be put into a wider theoretical framework based on information theory criteria such as minimal redundancy [4]. Redundancy is minimized by updating the horizontal connections with a rule that has a fixed point where second- and third-order correlations are zero.

We have found that synaptic plasticity can predict the expansion of RF sizes observed in the experiments, provided that higher order terms are taken into account in the updating rule, eq.(2.8). The relevant time scales for the expansion are given by T and τ_s. If T is too small the average activities go to zero too fast after the stimulus with the scotoma is presented and there is no time to generate the additional horizontal connection that will expand the RF. We can compare this time scale with the one obtained in the experiments. In ref.[25] the images with the artificial scotoma are presented at a rate of 30 Hz. On the other hand an expansion of the RF (as in Fig. 3), shifts in the psychophysical task and the *filling-in* effect (as in Fig. 12), that are qualitatively compatible with the experimental results, are obtained for $T = 200$, which means about 6 secs.

Besides the expansion of the RF size we have also analyzed the generation and time evolution of bias in feature localization tasks and the *filling-in* of stimuli with an artificial scotoma. We have found that for a suitable region in the space of parameters it is possible to reproduce most of the experimental results, such as the time evolution of the perceptual shift and the dependence of the shift on the distance between the test stimulus and the scotoma. We also found that the *filling-in* appears at a later time than the saturation of the shift although the difference is not a factor 5-10 as in the experiments.

The expansion of the RF has been also studied by other authors, although no attempt has been done to consider the constraints imposed on the modeling by the perceptual shift and the *filling-in* effects. Changes in the RF of neurons in the cortical scotoma have been modeled by Xing and Gerstein [42]. In that work several mechanisms are considered to explain the expansion of the RFs. The ones that involve plastic changes of the cortical synapses are a pure Hebb-like modification of the synaptic strength and a non-specific activity-related learning rule. Both fail to reproduce the experimental results. For the Hebbian rule the reason is that a neuron in the cortical scotoma will be inactive after the lesion. As a synapse is strengthened when both pre- and post-synaptic neurons fire within some time window, all the afferent synapses to a neuron will be unaffected, and a recovery of activity will be impossible. In the second case high frequency firing of one neuron causes the strength of all its synapses to increase. This rule will cause an expansion of the RF of all the neurons (both inside and outside the scotoma) and it will fail to generate the reversal of the expansion after the presentation of a center stimulus.

RF dynamics based on information theory ideas have been also proposed by Sirosh et al. [40]. Their rule however only decorrelates the output activities instead of lifting the redundancy. This is because the higher order terms are not included in their learning rule. An expansion of the RF is found in this work but the problem of the time scale where the perceptual shifts occur is not addressed in this work.

According to a different strategy suggested in [33] the output code is constructed by maximizing the information shared by the RF and the context (the *coherent infomax* criterium); in this scheme two hypercolumns with non-overlapping RFs are thought of as representing the classical RF and the modulating context.

The model proposed in this paper can be improved in many directions. After adaptation to the statistics of normal visual stimuli (that we defined as the ones with the power spectrum of natural images) the horizontal connections are all inhibitory except the self-interaction term that is excitatory. This is because the algorithm that generates the weights tries to make the activities of neurons on the second layer independent. In general maximization of information in noisy systems require some degree of redundancy. For this more general case Linsker [28] has proposed a mixed learning rule: in one phase the network is presented with the input signal plus noise and the weights are updated in order to decorrelate the outputs. In a second phase pure noise is presented and a modification rule that *increases* correlations is applied. It is possible to prove that such mixed rule actually implements the "infomax" principle. The application of such a rule when the system is adapting to the statistics of normal visual conditions leads to the generation of additional excitatory connections with a range that depends on the range of the feed-forward connections. This is because when noise is presented in the input layer, the length of the correlations in the second layer is determined by those connections. At longer distances there will be correlations only when the images are presented, and in this phase only inhibitory connections are generated. The final result will be excitation at short distance and inhibition at long distance. Let us observe this is compatible with the results of, e.g., ref. [11].

In psychophysical experiments have been also observed longer time scales, that are related to an improving in the performance of the task [25]. Such phenomena are not taken into account in our model. Longer time scales should be explicitly introduced to account for learning.

The dynamical changes have been studied in the context of the modifications of horizontal connections in V1. However similar changes can be observed when different kinds of inputs (that include more complex stimuli) are presented to the system, suggesting that connections in different regions of the cortex can also be modified. In [36] it is found that if the pattern surrounding the scotoma contains type-written text, the scotoma is filled-in by text (although subjects could not actually read the latters in the filled-in region). Moreover, if the surrounding pattern suddenly changes (for instance from twinkling dots to flickering lines) the process of filling-in has to begin again. This is consistent with a change of the locus of the synaptic modification.

It would be worthwhile to explore different models for synaptic dynamics to see how robust the need of higher order terms in this rule is in order to obtain the correct RF expansion at the correct time. One possibility would be to use a generalization of the model described in [1] including the contribution of higher order correlations similar to the rule in eq.(2.8). Another alternative would be to introduce the memory effect in a different way; this could be done for instance as in [13, 41].

Acknowledgements

We would like to thank Jean-Pierre Nadal and Antonio Turiel for discussions and to Edmund Rolls for interesting comments of the manuscript. G.M. thanks the Spanish Ministery of Education and Science for financial support. An E.U. grant CHRX-CT92-0063 and a Spanish grant AEN93-0673 are also gratefully acknowledge.

6 Appendix

If the inputs v_j are independent the solution to the problem of decorrelation is obtained when the outputs u_i differ from the inputs only by a permutation and a scaling factor. The permutation symmetry of the system is broken by the topographic mapping between the two layers. However the scaling factor is not fixed. As the input to the second layer is a linear combination of the v_j's and the feed-forward connections J_{ij} it is not possible to say how much of this input is coming from each factor. Therefore it is necessary to fix the scale of the outputs arbitrarily by fixing N^2 numbers. We choose to fix the N^2 self-interactions $T_{ii}D_{ii} = J_{ii} - 1$.

The reason for this choice is that if we are in the linear region of the transfer function H the fixed point of Eq.(2.3) can be written as:

$$\mathbf{u}(\mathbf{I} + \mathbf{DT}) = \mathbf{Jv} \tag{6.14}$$

where $\mathbf{I}$ is the identity matrix. In this case a scaling factor equal to 1 ($\mathbf{u} = \mathbf{v}$) is obtained when $\mathbf{I} + \mathbf{DT} = \mathbf{J}$. For the diagonal terms we obtain the expression of above. Of course, this arbitrary normalization can be also applied when the inputs v_j are not independent or if the system is not always operating in the linear region of the transfer function.

References

[1] Amit D. and Brunel N. 1995, Learning internal representations in an attractor neural network, *NETWORK* **6**, 359-388

[2] Atick J. J. 1992, Could information theory provide an ecological theory of sensory processing?, *Network* **3**, 213-251

[3] Baldi P. and Heilimberg W, 1988, How sensory maps could enhance resolution through ordered arrangements of broadly tuned receivers, *Biol. Cybern.* **59**, 313-318

[4] Barlow H. B. 1961 Possible principles underlying the transformation of sensory messages. In Rosenblith W., editor, *Sensory Communication*, page 217. M.I.T. Press, Cambridge MA

[5] Barlow H. B. and Foldiak P. 1989, Adaptation and decorrelation in the cortex, *The Computing Neuron* ed. R. Durbin, C. Miall and G. Mitchison (Cambridge, MA: Addison Wesley), 54-72

[6] Bienenstock E. L., Cooper L. N. and Munro P. W. 1982, Theory for development of neuron selectivity: orientation specificity and binocular interaction in visual cortex *J. Neurosci.* **2**, 32-48

[7] Blahut R. E. 1988, *Principles and Practice of Information Theory.* Addison-Wesley, Cambridge MA

[8] Clothiaux E. E., Bear M. F. and Cooper L. N. 1991, Synaptic plasticity in visual cortex: comparison of theory with experiment, *J. Neurophysiol.* **66**, 1785-1804

[9] Comon P. 1994, Independent component analysis, a new concept? *Signal Processing* **36**, 287-314

[10] Das A. and Gilbert C. D. 1992, Dynamic changes of receptive-field size in cat primary visual cortex *Proc. Natl. Acad. Sci. USA* **89**, 8366-8370

[11] Das A. and Gilbert C. D. 1995, Long-range horizontal connections and their role in cortical reorganization revealed by optical recording of cat primary visual cortex, *Nature* **375**, 780-784

[12] Field D. J. 1987, Relations between the statistics of natural images and the response properties of cortical cells, *J. Opt. Soc. Am.* **4**, 2379-2394

[13] Foldiak P. 1991, Learning invariance from transformation sequences, *Neural Computation* **3**, 194-200

[14] Gilbert C. D. 1992, Horizontal integration and cortical dynamics *Neuron* **9**, 1-13

[15] Gilbert C. D., Das A., Ito M., Kapadia M. and Westheimer G 1996, Spatial integration and cortical dynamics, *Proc. Natl. Acad. Sci. USA* **93**, 615-622

[16] Gilbert C. D. and Wiesel T. 1992, Receptive field dynamics in adult primary visual cortex *Nature* **356**, 150-152

[17] Grinvald A., Lieke E. E., Frostig R. D., and Hildesheim R. 1994, Cortical Point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex, *Journal of Neuroscience* **14**, 2545-2568

[18] van Hateren J.H. 1992, Theoretical predictions of spatiotemporal receptive fields of fly lmcs, and experimental validation, *J. Comp. Physiology A* **171**,157-170

[19] van Hateren J.H. 1993, Spatiotemporal contrast sensitivity of early vision, *Vision Res.* **33**, 257-267

[20] Hérault J. and Jutten C. 1994, *Réseaux neuronaux et traitment du signal*, Hermes, Paris

[21] Hirsch J. A. and Gilbert C. D. 1991, Synaptic Physiology of Horizontal Connections in the Cat's Visual Cortex, *Journal of Neuroscience* **11**, 1800-1809,

[22] Hopfield J. J. 1991, Olfactory computation and object perception, *Proc. Natl. Acad. Sci. USA* **88**, 6462-6466

[23] Jenkins W. M., Merzenich M. M., Ochs M. T., Allard T., and Guic-Robles E. 1990, Functional reorganization of primary somatosensory cortex in adult owl monkeys after behaviorally controlled tactile stimulation, *J. Neurophys.* **63**, 82-104

[24] Jutten C. and Herault J 1991, Blind separation of sources, part I: an adaptive algorithm based on neuromimetic architecture, *Signal Processing* **24**, 1-10

[25] Kapadia M. K., Gilbert C. D. and Westheimer G. 1994, A quantitative measure for short-term cortical plasticity in human vision, *J. Neurosci.* **14**, 451-457

[26] Kapadia M.K., Ito M., Gilbert C.D. and Westheimer G. 1995, Improvement in visual sensitivity by changes in local context: parallel studies in human observers and in V1 of alert monkeys, *Neuron* **15**, 843-856

[27] Knierim J. J. and van Essen D. C. 1992. Neuronal Responses to Static Texture Patterns in Area V1 of the Alert Macaque Monkey, *Journal of Neurophysiology* **67**, 961-980

[28] Linsker R. 1992, Local synaptic learning rules suffice to maximize mutual information in a linear network, *Neural Comp.* **4**, 691-702

[29] Merzenich M. M., Nelson R. J., Stryker P., Cynader S., Schoppmann A. and Zook J. M. 1984, Somatosensory cortical map changes following digit amputation in adult monkeys, *J. Comp. Neurol.* **224**, 591-605

[30] Nadal J.-P. and Parga N. 1994, Nonlinear neurons in the low noise limit: a factorial code maximizes information transfer, *Network* **4**, 565-582

[31] Nadal J.-P. and Parga N. 1996, Redundancy reduction and independent component analysis: algebraic and adaptive approaches, submitted to *Neural Computation*

[32] Pettet M. W. and Gilbert C. D. 1992, Dynamic changes in receptive-field size in cat primary visual cortex, *Proc. Natl. Acad. Sci. USA* **89**, 8366-8370

[33] Phillips W. A., Kay J. and Smyth D. 1995, The discovery of structure by multi-stream networks of local processors with contextual guidance, *Network* **6**, 225-246

[34] Polat U. and Sagi D. 1993, Lateral Interactions Between Spatial Channels: Suppression and Facilitation Revealed by Lateral Masking Experiments, *Vision Research* **33**, 993-999

[35] Polat U. and Sagi D. 1994, The Architecture of Perceptual Spatial Interactions, *Vision Research* **34**, 73-78

[36] Ramachandran V. S. and Gregory R. L. 1991, Perceptual filling in of artificially induced scotomas in human vision, *Nature* **350**, 699-702

[37] Recanzone G. H., Merzenich M. M., Jenkins, W. M., Kamil A. G. and Dinse H. R. 1992, Topographic reorganization of the hand representation in cortical area 3b of owl monkeys trained in a frequency discrimination task, *J. Neurophysiol.* **67**, 1031-1056

[38] Rospars J.-P. and Fort J.-C. 1994, Coding of odour quality: roles of convergence and inhibition, *Network* **5**, 121-145

[39] Ruderman D. L. 1994, The statistics of natural images, *Network* **5**, 517-548

[40] Sirosh J., Miikkulainen R. and Bednar J. 1996, Self-Organization of Orientation Maps, Lateral Connections, and Dynamic Receptive Fields in the Primary Visual Cortex. In *Lateral Interactions in the Cortex: Structure and Function.* Eds: Sirosh, J., Miikkulainen, R., and Choe, Y.. The UTCS Neural Networks Research Group, Austin, TX. Electronic book, http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96.

[41] Wallis G. and Rolls E. 1996, A model of invariant recognition in the visual system. Submitted for review to *Progress in Neurobiology.*

[42] Xing J. and Gerstein G. 1994, Simulation of dynamic receptive fields in primary visual cortex 1994, *Vision Res.*, **34**, 1901-1911

FIGURE CAPTIONS

Figure 1

Architecture of the system. Black circles represent neurons on the input layer and squares denote populations of neurons (columns) in V1. The architecture has been simplified neglecting intermediate layers. J_{ij} is the feedforward connection between input neuron j and a cortical unit i. T_{ik} represents the dynamical part of the horizontal connection between the presynaptic unit k and the postsynaptic unit i

Figure 2

a) Second and third order cumulants as a function of time, $\tau_s = 1$, $\alpha = 0.2$, $\beta = 0.8$, $\gamma = 1$ and $T = 200$.

Solid line: $\overline{< \delta_T u_i \delta_T u_j >}$

Dashed line: $\overline{< \delta_T u_i \delta_T u_j^2 >}$

The line denotes the average over the neurons. From time 0 to 200, the inputs have the statistics of the natural images. From 200 to 400 they have the scotoma.

b) Third and fourth order correlations. Same parameters as in the previous figure.

Solid line: $\overline{< \delta_T u_i \delta_T u_j \delta_T u_k >}$

Dashed line: $\overline{< \delta_T u_i \delta_T u_j \delta_T u_k \delta_T u_l >}$

Stimuli are presented regularly and one time unit corresponds to the presentation of a single pattern. To compare with the experiments in [25] this time unit has to be taken equal to 33 msec.

Figure 3

a) Change in receptive field for a neuron in the cortical scotoma. It is measured after averaging the internal field of the unit when the stimulus is presented at a distance r from the center of the system. Same parameters as in Fig. 2.

Solid line: receptive field after adaptation to normal visual conditions (statistics of natural images).

Dashed line: receptive field after adaptation to images with a scotoma.

b) Change in receptive field for a unit outside the cortical scotoma. Same parameters as in Fig. 2.

Solid line: receptive field after adaptation to normal visual conditions.

Dashed line: receptive field after adaptation to images with a scotoma.

Figure 4

a) Ratio between the final and the initial receptive fields at its maximum value as a function of T for $\tau_s = 1$, $\alpha = 0.2$, $\beta = 0.8$ and $\gamma = 1$

b) Ratio between the final and the initial receptive fields at its maximum value as a function of τ_s for $T = 200$, $\alpha = 0.2$, $\beta = 0.8$ and $\gamma = 1$

Figure 5

Ratio between the final and the initial receptive fields at its maximum value as a function of α ($T = 200$).

Solid line: $\beta = 0$ and $\gamma = 1$

Dashed line: $\beta = 1$ and $\gamma = 0$

Dot-dashed line: $\beta = 0.5$ and $\gamma = 0.5$.

In all these cases, if smaller values of α are used the algorithm does not converge to the fixed point with zero correlations.

Figure 6

Ratio between the final and the initial receptive fields at its maximum value as a function of α for $\beta = \gamma = 0$ and different values of T:

Solid line: $T = 50$

Dashed-line: $T = 100$

Dot-dashed line: $T = 200$

For each one of the curves if smaller values of α are used the algorithm does not converge to the solution with zero correlations.

Figure 7

a) Reversibility of receptive fields. Same parameters as in Fig. 2

Solid line: Receptive field after adaptation to normal stimuli.

Dashed line: Expansion after stimulus with artificial scotoma.

Dot-dashed line: Receptive field after elimination of the scotoma.

b) Contraction of receptive fields in the presence of strong stimulation.

Solid line: Receptive field after adaptation to normal stimuli.

Dashed line: After stimulating the central region (size 7×7) with a stimulus 4 times stronger than the background.

Figure 8

Shift as a function of time. The test stimulus is centered on the border of the scotoma. Same parameters as in Fig. 2

Figure 9

Shift for different positions of the test stimulus relative to the scotoma. The distance between the center of the test stimulus and the center of the system is d. Same parameters as in Fig. 2

Figure 10

Same parameters as in Fig. 2.

Solid line: Average activity of the system.

Dashed line: Activity of a neuron in the cortical scotoma

Fig. 1

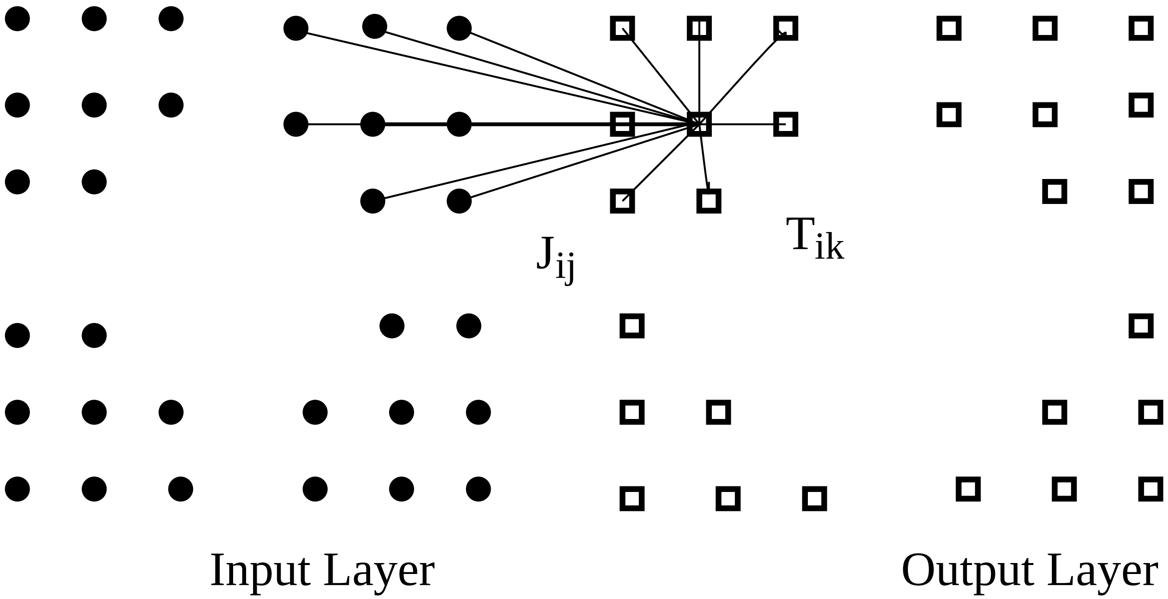

Figure 1:

Fig. 2a

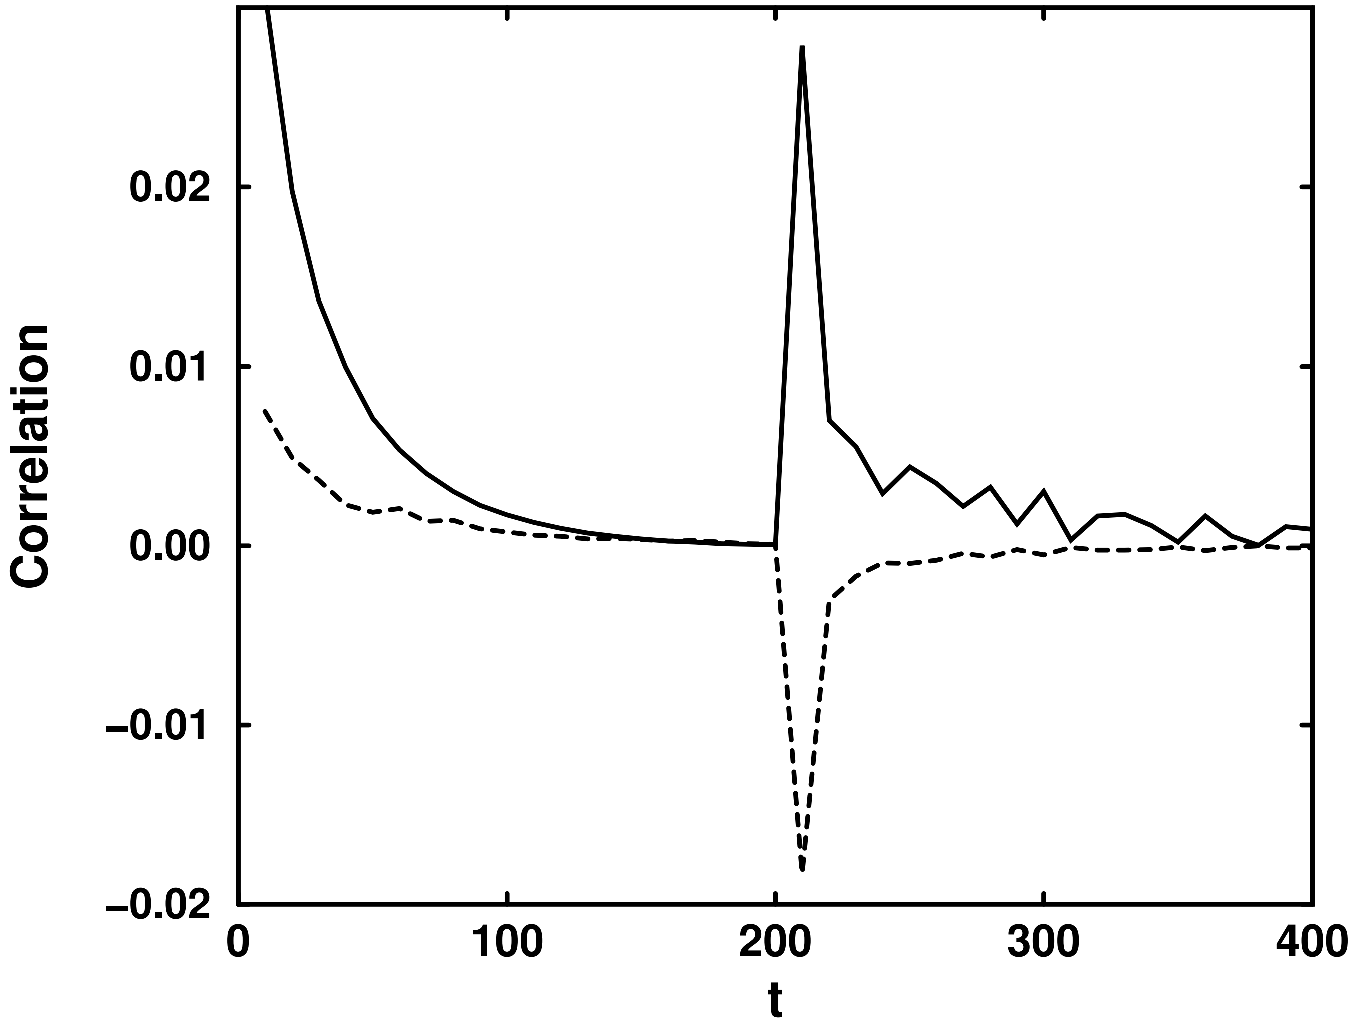

Fig. 2b

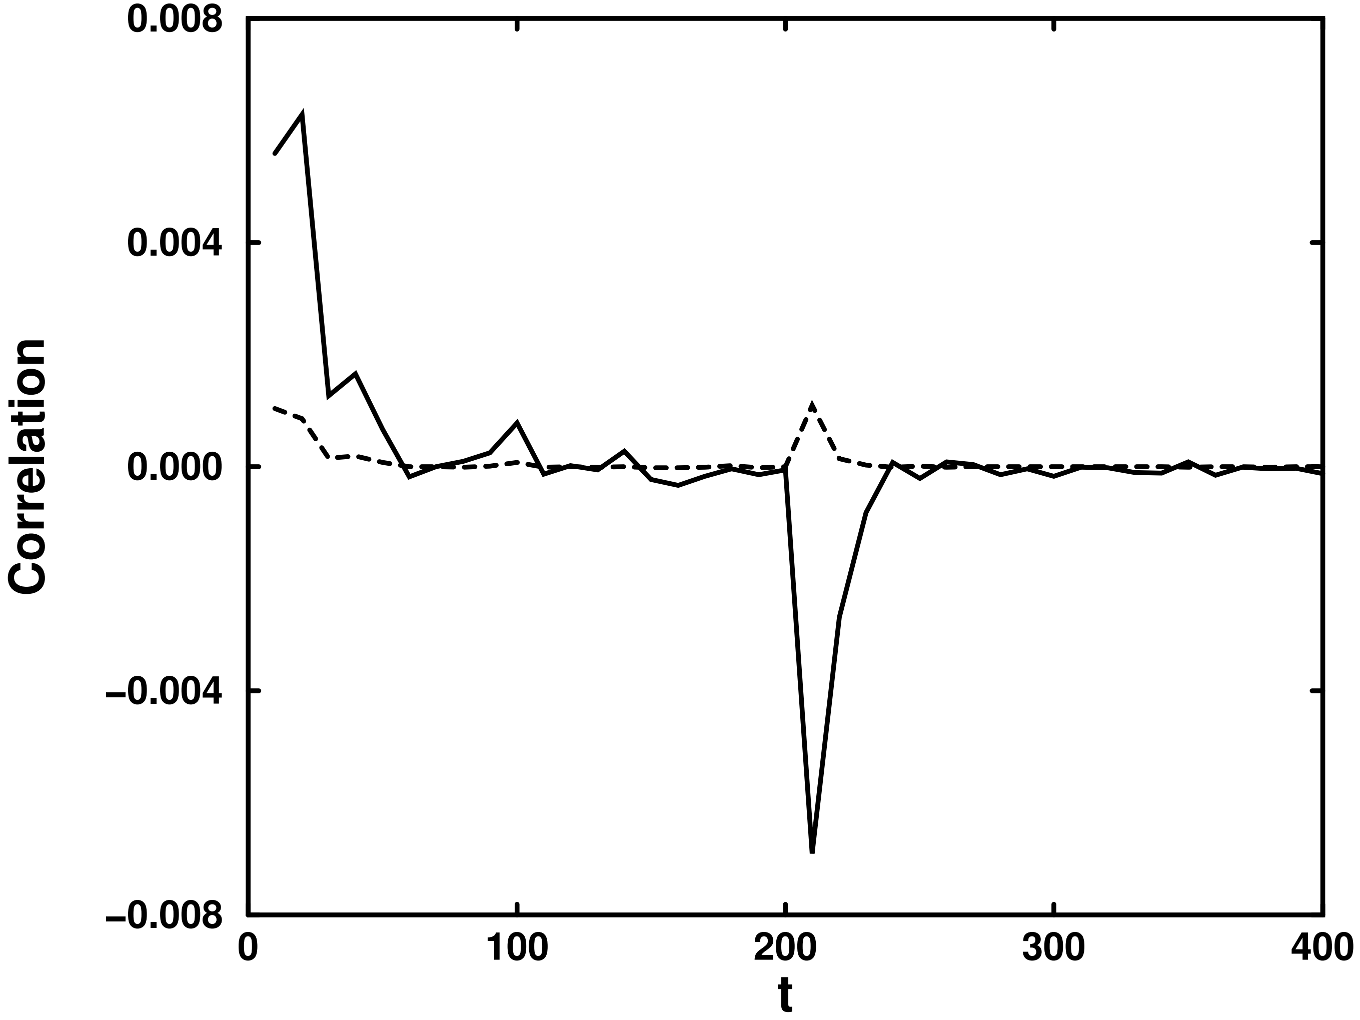

Fig. 3a

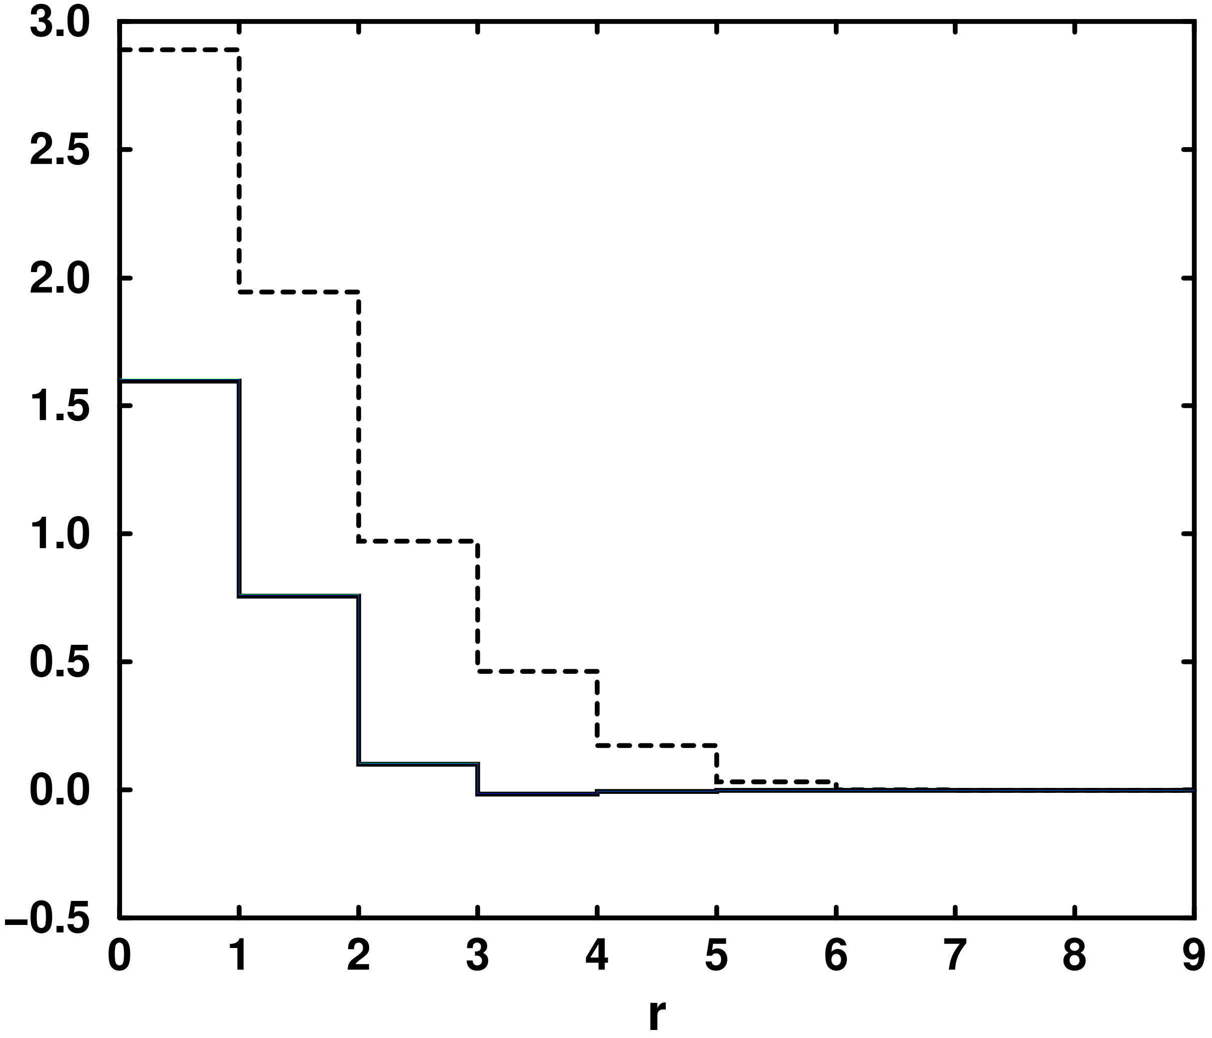

Fig. 3b

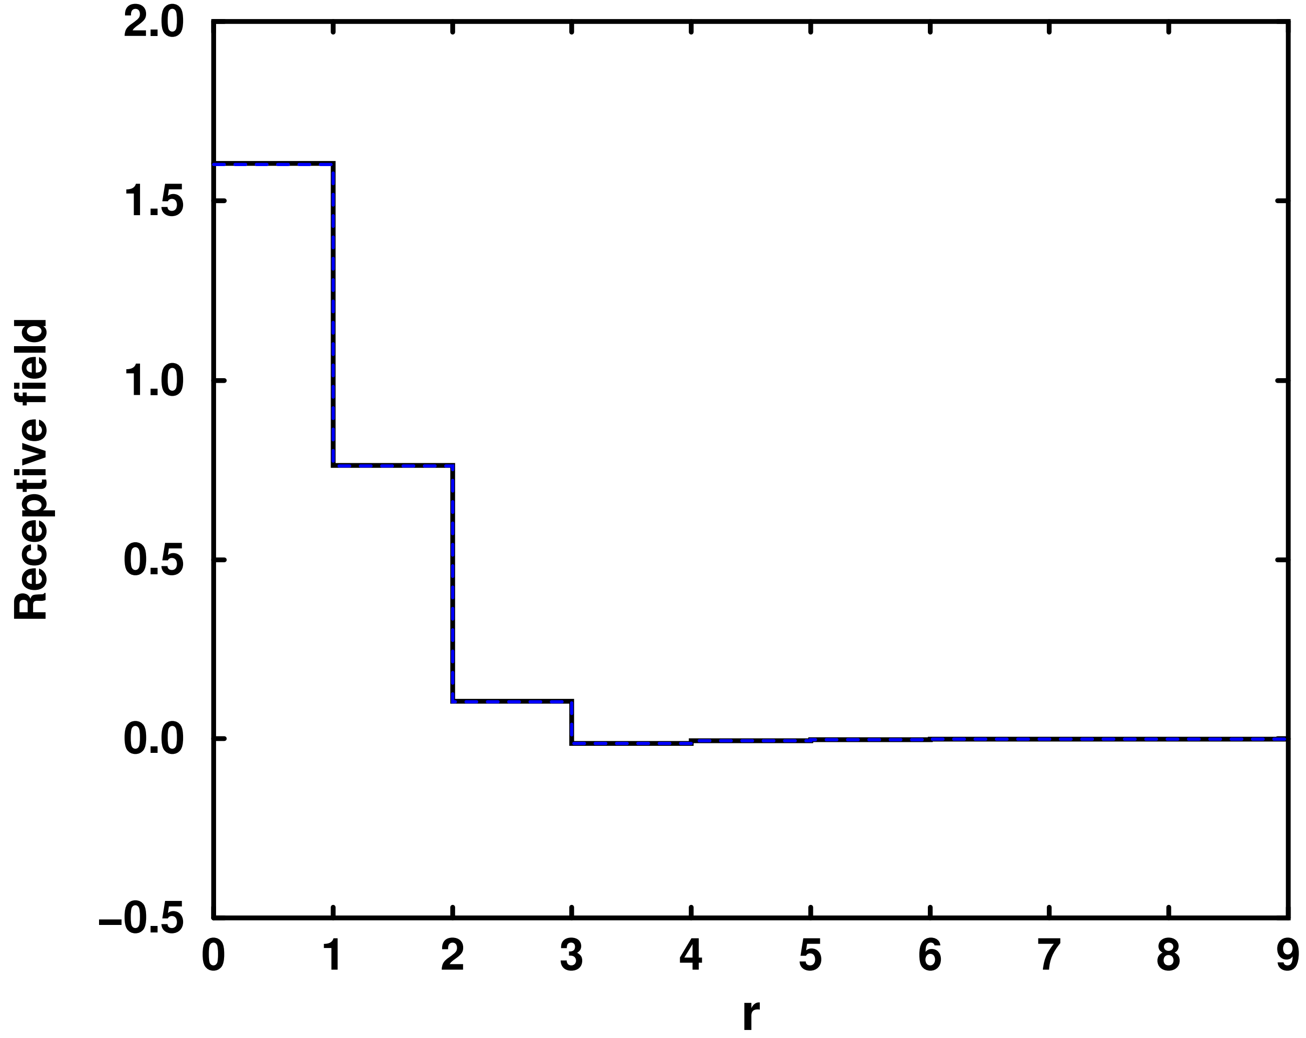

Fig. 4a

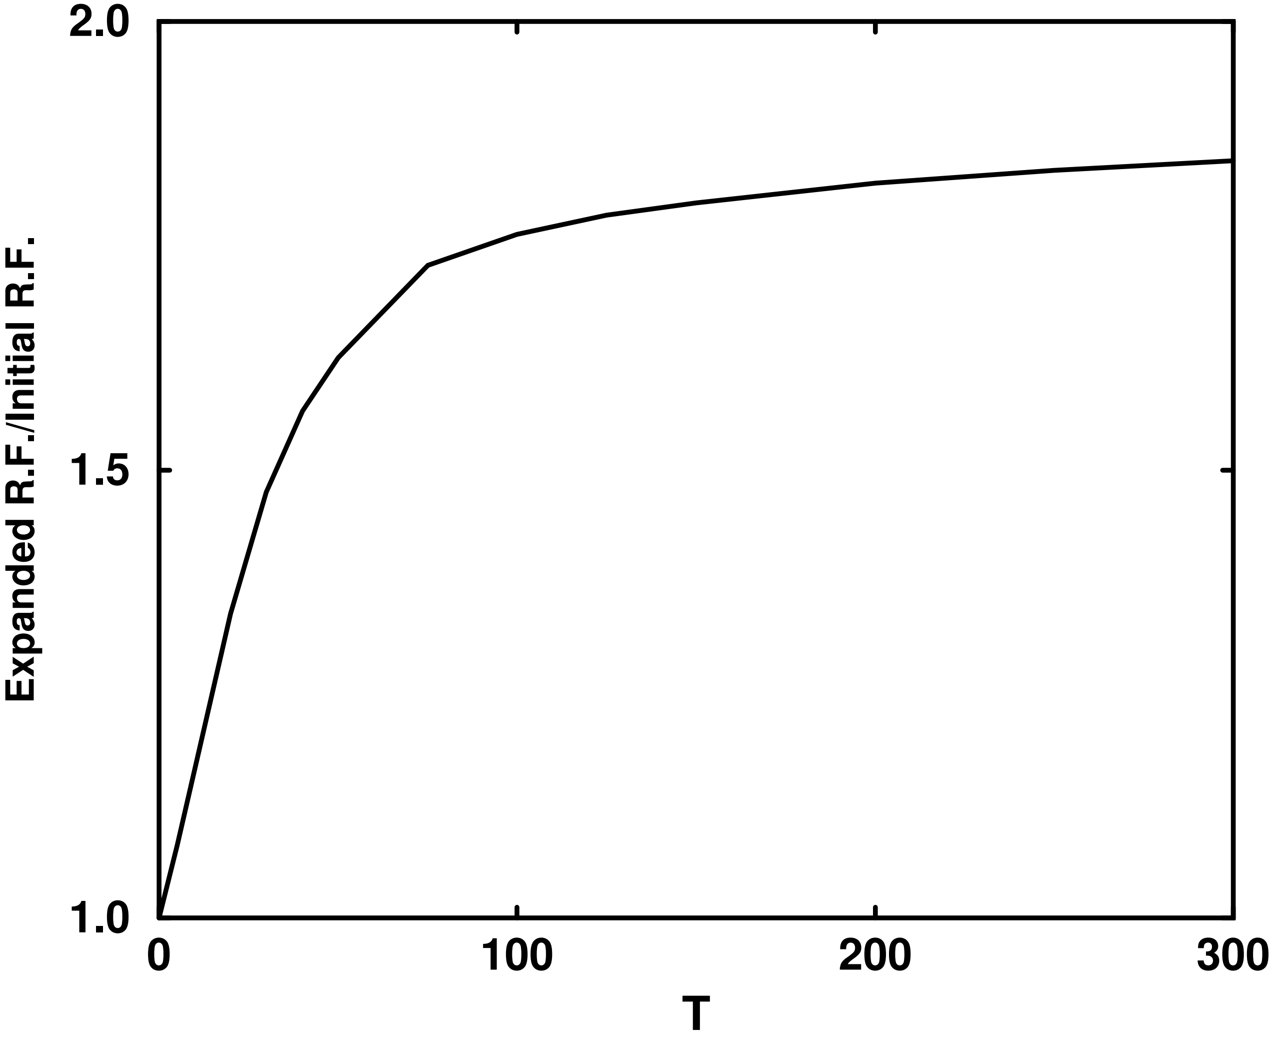

Fig. 4b

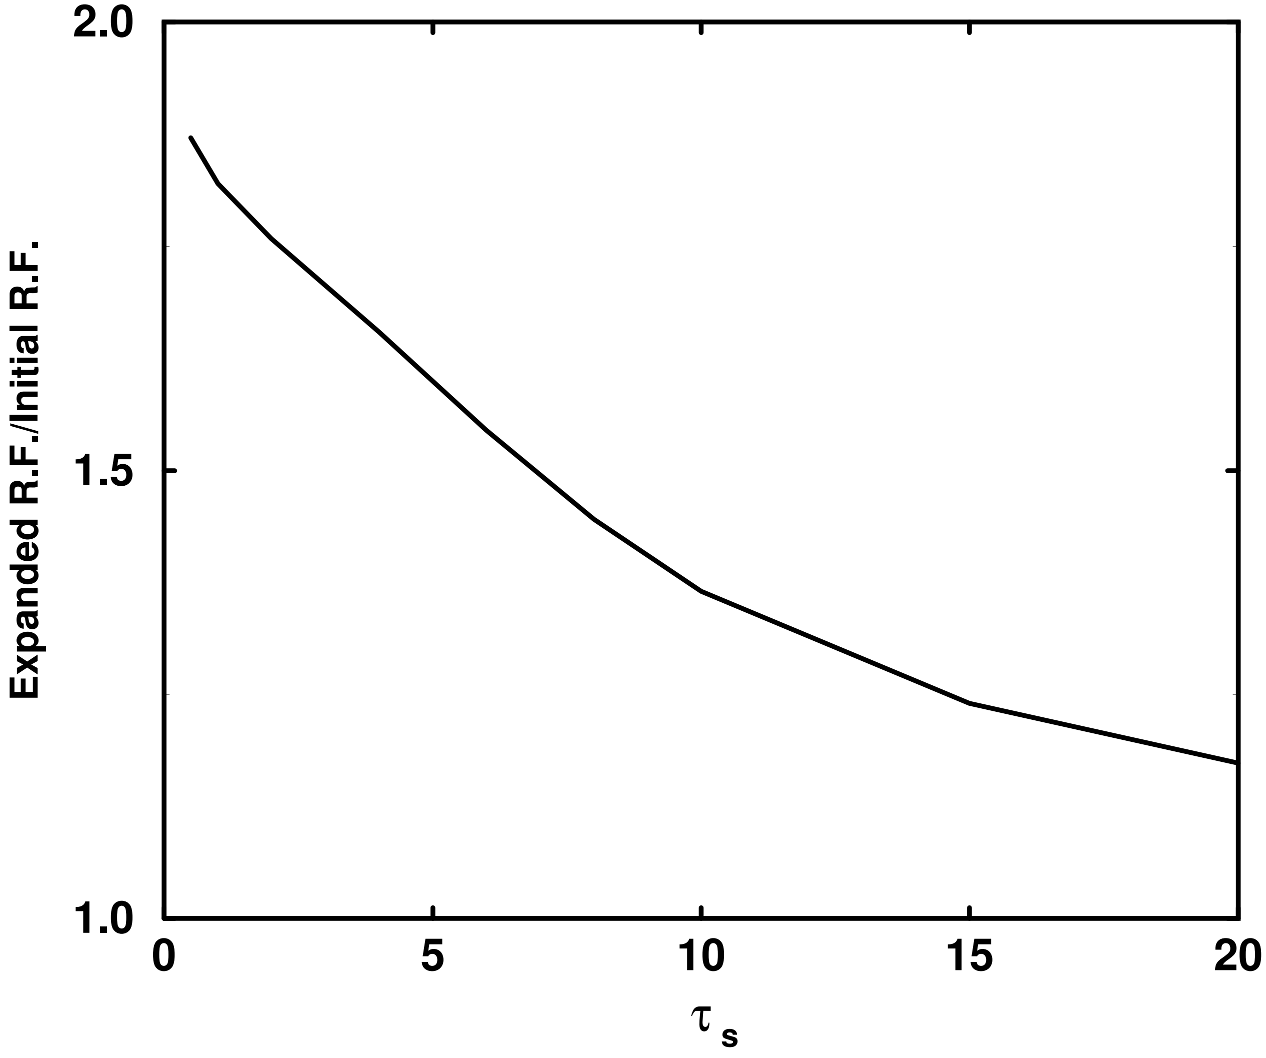

Fig. 5

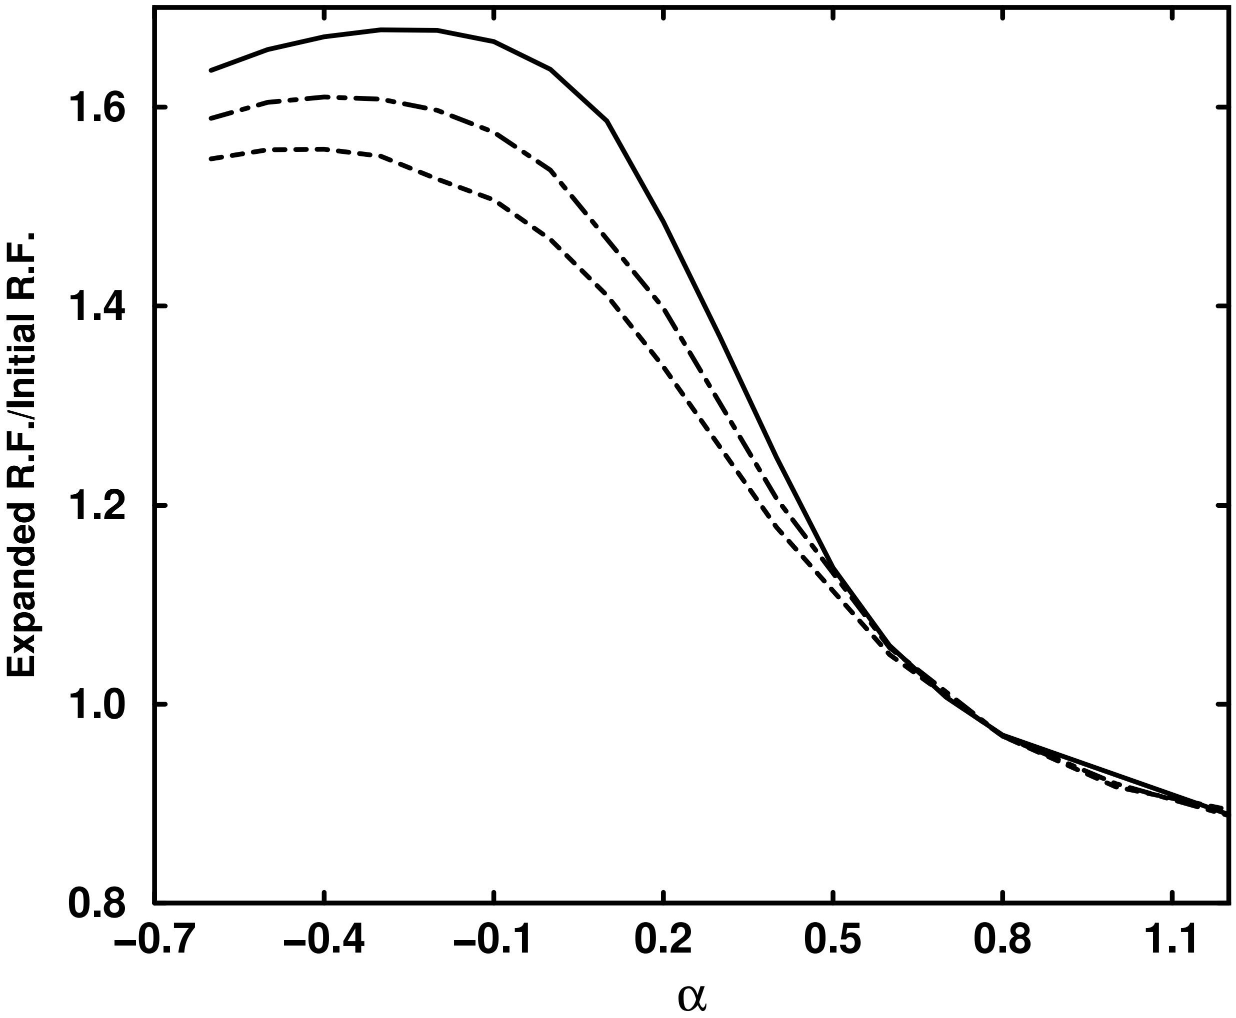

Fig. 6

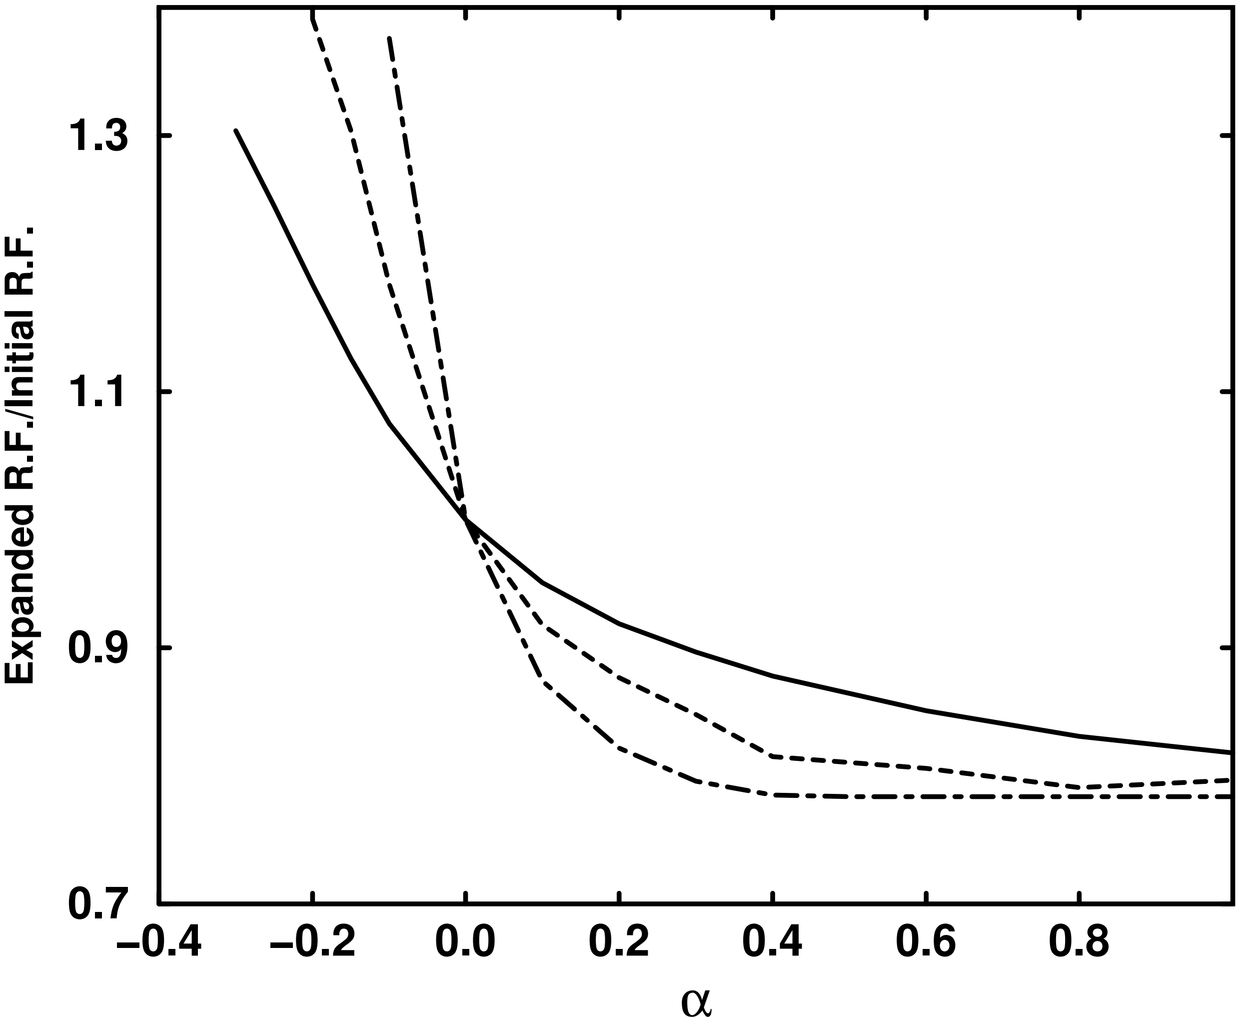

Fig. 7a

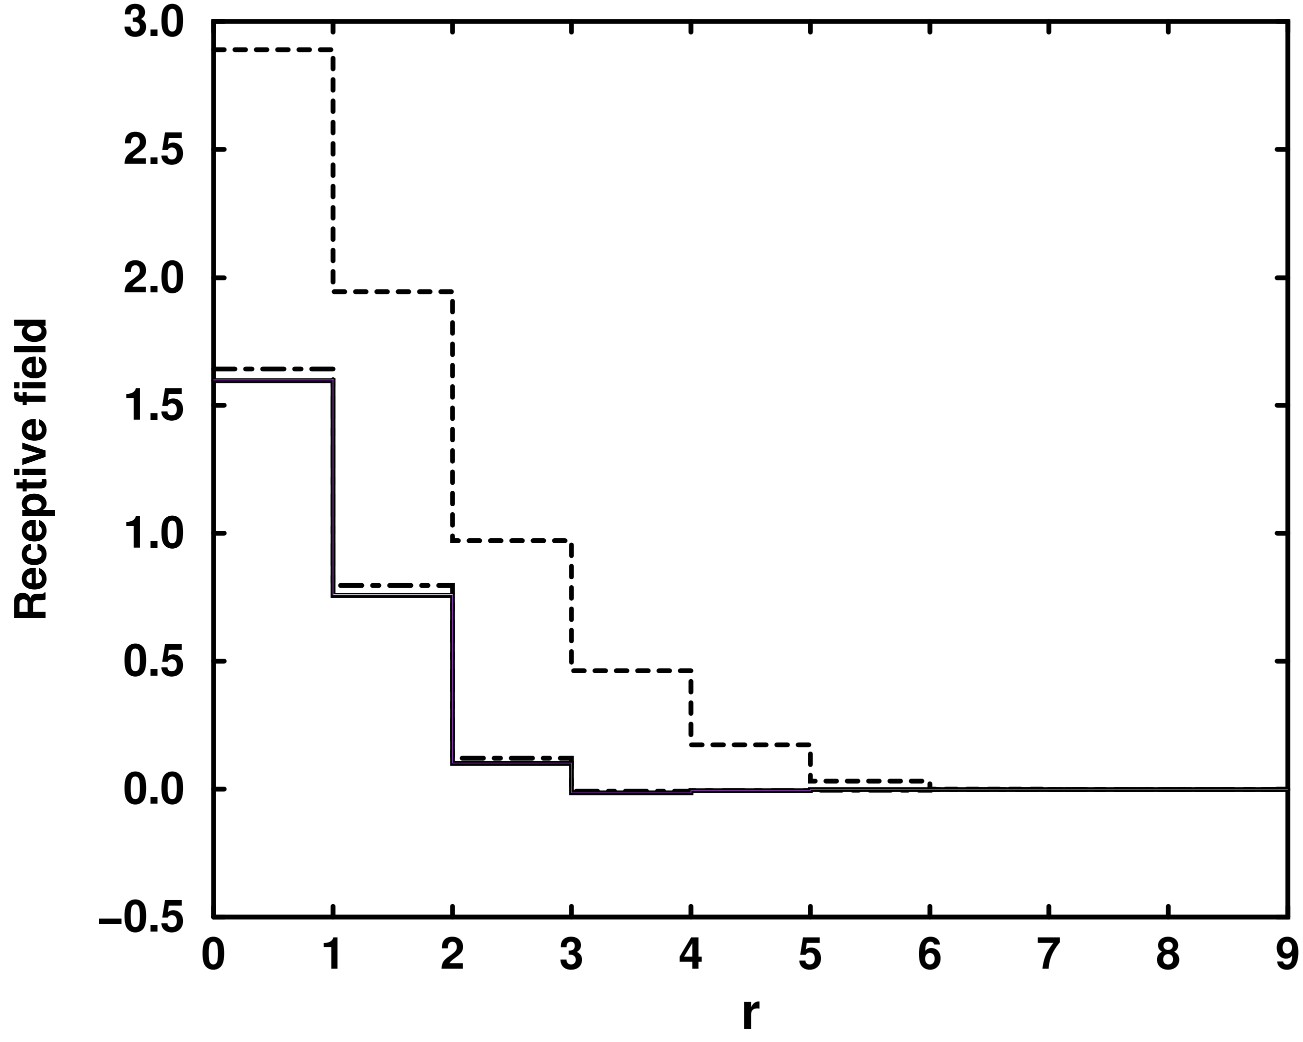

Fig. 7b

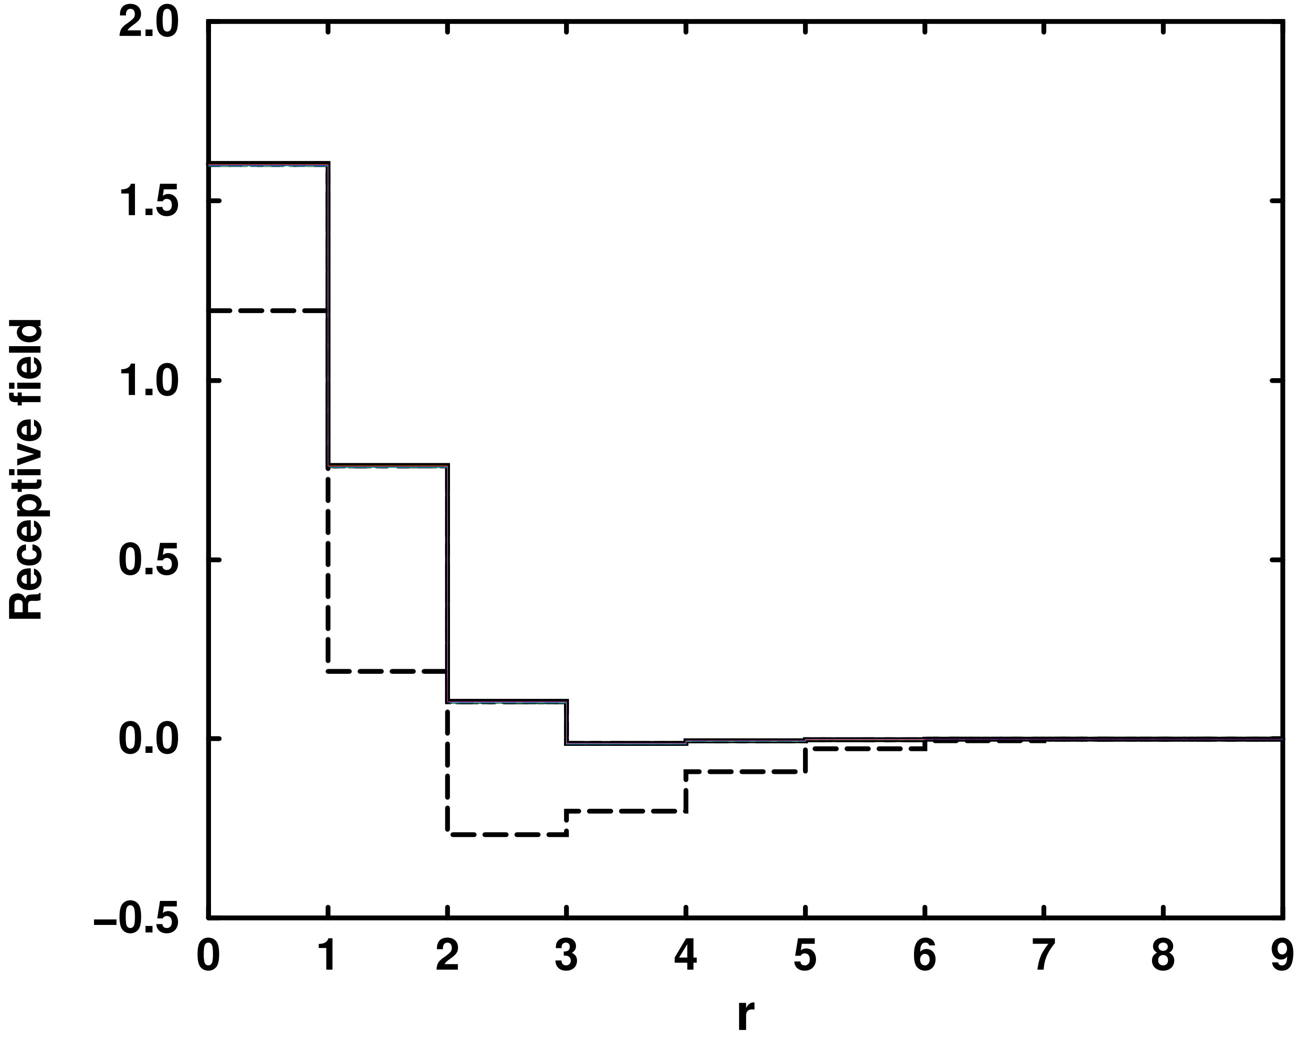

Fig. 8

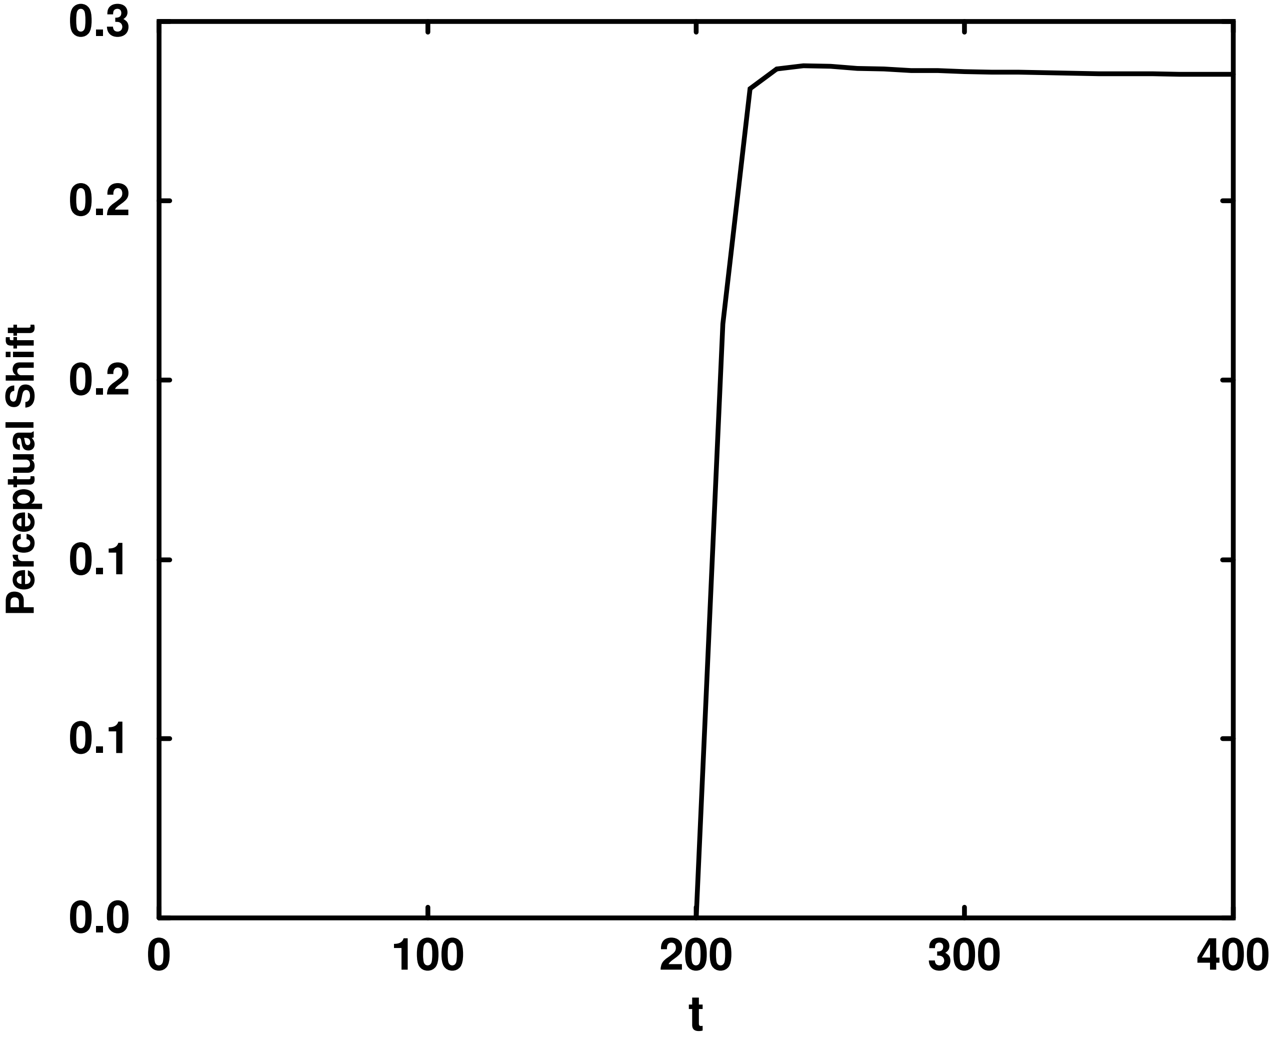

Fig. 9

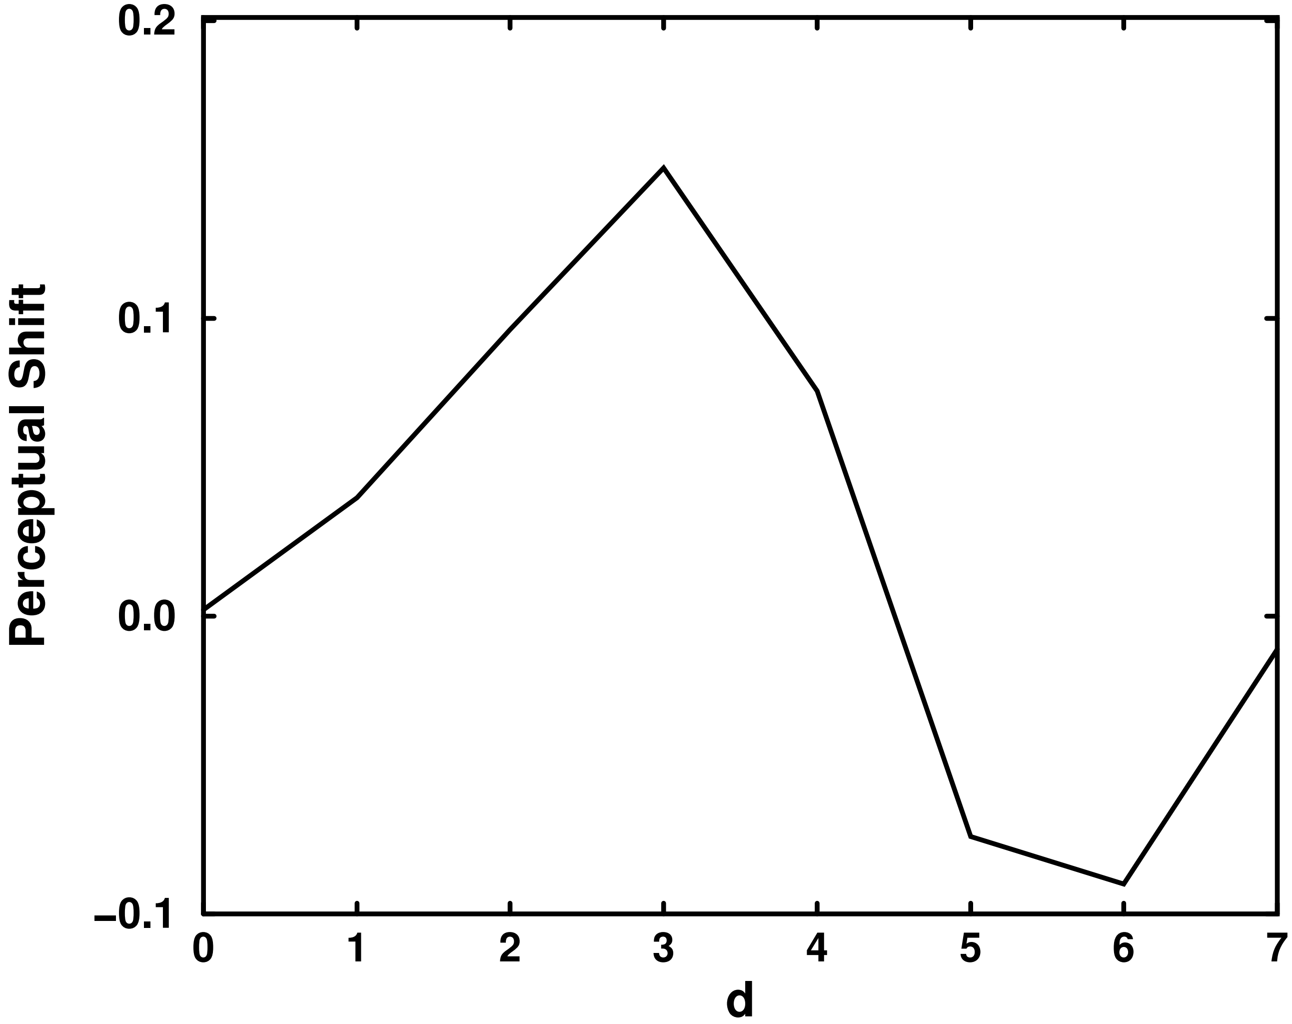

Fig. 10