Local Dynamic Interactions in the Collicular Motor Map: A Neural Network Model

Lina L. E. Massone and **Tony Khoshaba**

Technical Report 4/94

Neural Information Processing Laboratory
Northwestern University

Lina L. E. Massone (*massone@eecs.nwu.edu*) is with the Department of Electrical Engineering & Computer Science and the Biomedical Engineering Department of Northwestern University, 2145 Sheridan Road, Evanston, IL 60208-3118.

Tony Khoshaba (*khoshaba@eecs.nwu.edu*) is with the Biomedical Engineering Department of Northwestern University, 2145 Sheridan Road, Evanston, IL 60208-3118.

Accepted for publication in *Network*

Abstract

In this paper we explore the possibility that some of the dynamic properties of the neural activity in the gaze-related motor map (located in the intermediate layers of the superior colliculus) might be mediated by local interactions between movement-related neurons and fixation neurons. More specifically, the goal of this research is to demonstrate, from a computational standpoint, which classes of dynamic behaviors of the collicular neurons can be obtained without the intervention of feedback signals, and to hence begin exploring to what extent the gaze system needs feedback in order to operate. We modeled: (a) The collicular motor map as a dynamical system realized with a recurrent neural network. (b) The dynamics of the neural activity in the map as the transients of that system towards an equilibrium configuration that the network learned with a recurrent learning algorithm (recurrent backpropagation.) The results of our simulations demonstrate:

(1) That the transients of the trained network are hill-flattening patterns as observed by some experimenters in the burst-neuron layer of the superior colliculus of rhesus monkeys. This result was obtained despite the fact that the learning algorithm did not specify what the network's transients should be.

(2) That the connections in the trained network are excitatory within the fixation zone of the motor map and inhibitory elsewhere.

(3) That the results of the learning are robust in the face of changes in the connectivity pattern and in the initialization of the weights, but that a local connectivity pattern favors the network's stability.

(4) That nonlinearity is required in order to obtain meaningful dynamic behaviors.

(5) That the trained network is robust to abnormal stimulation patterns such as noisy and multiple stimuli and that when multiple stimuli are utilized the response of the network remains a stereotyped flattening one.

The results of the learning point out the possibility that the dynamics of the burst-neuron layer of the superior colliculus might be locally regulated rather than feedback-driven, and that the action of the feedback be confined to the layer of the buildup neurons. The results of the multiple-stimulation experiment support the hypothesis, already put forward by one of the authors in a previous work (Massone in press), that the averaging of the direction of movement following double stimulation of the motor map (Robinson 1972) does not occur at the level of the motor map. This paper constitutes also a study of the properties and responses of recurrent backpropagation under various choices for the network's and algorithm's parameters.

1 Introduction

It is known that the superior colliculus (SC) plays a major role in gaze shifts, i.e. combined movements of the head and the eyes aimed at rapidly displacing the visual axis. A number of electrophysiological studies (e.g. Cynader and Berman 1972, Robinson 1972) demonstrated that the intermediate layers of SC contain a topographically-organized motor map, in which the location of active neurons represents the motor error, i.e. the angular distance yet to be covered by the eyes in order to reach the target. These studies also showed that the topography of the motor map is such that the horizontal motor errors are encoded mainly along the rostro-caudal axis in a contralateral fashion, while the vertical motor errors are encoded mainly along the medio-lateral axis (see Figure 1a.) The more caudal (medio-lateral) the activity is located, the larger the horizontal (vertical) motor error. Recording studies revealed that, while the location of

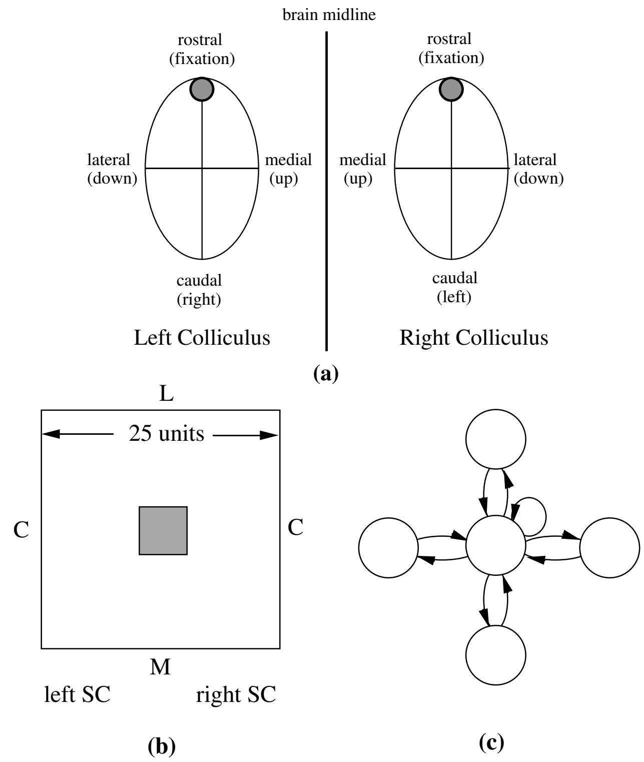

Figure 1: (a) Topographical organization of the motor map in the superior colliculus. The dark circles represent the right and left fixation zones. (b) The motor maps for the right and left superior colliculi were represented with an array of 25x25 units. The shaded square in the middle represents the right and left fixation zones located in the rostral pole. C: caudal, M: medial, L: lateral. (c) The structure of the network's connectivity.

active collicular neurons encodes the motor error, their firing rate is instead correlated with the movements velocity (Munoz and Guitton 1986, Berthoz et al. 1987, Roher et al. 1987.) Other recording studies showed that the rostral pole of the map contains the so-called fixation neurons i.e. neurons that are active when the eyes are on target and inactive during movement (Munoz and Wurtz 1992, Munoz and Guitton 1991.) Furthermore it was shown that when a motor error occurs a population of collicular neurons outside the fixation zone becomes active and the population's activity has the shape of a gaussian hill (Ottes et al. 1986, Munoz and Guitton 1991.)

The mechanisms by which the collicular motor map is involved in gaze control are a controversial matter. For a long time the colliculus was thought to be located upstream from the programming of the details of the movements. However recent experimental observations suggested that SC might be part of the postulated feedback loop (Robinson 1975) that the saccadic system uses to progressively reduce the error between current and desired orientation of the eyes. Some experimenters proposed that in primates the collicular motor error is generated by the difference between a cortical representation of the target and the distance already covered by the eyes (Waitzman et al. 1991.) Moreover they postulated that the hill of activity in the collicular motor map might progressively flatten during movement, until it disappears. The activity was then observed to reappear in the fixation zone (Munoz and Guitton 1991). Other researchers observed, mainly in cats, that the collicular activity travels through the motor map from its initial location towards the rostral pole where fixation neurons are located (Munoz et al. 1991b.) In a recent study, aimed at reconciling the apparent species-dependence of the collicular dynamics, Munoz and Wurtz (1993) identified two different movement-related types of cells in the SC of rhesus monkeys: burst neurons (BNs), located in the dorsal intermediate layers, and buildup neurons (BUNs), located in the ventral intermediate layers. Their study also showed that the neural activity in the BN layer declines during movement as the motor error decreases (a hill-flattening effect), while activity in the BUN layer spreads during movement towards the fixation neurons located in the rostral pole of the motor map (a hill-shift effect.) These results indicate that in primates both effects (hill-flattening and hill-shift) are simultaneously present and are likely to contribute in concert to the resulting motor acts.

Based upon some of these experimental findings, a number of researchers developed computational models of the saccadic/gaze system. A distinctive feature of practically all the recently-developed ones (e.g. Droulez and Berthoz 1991, Lefèvre and Galiana 1992, Dominey and Arbib 1992, Van Opstal and Kappen 1993, Massone in press) is the assumption that the driving force behind the collicular dynamics is a feedback that the motor map receives, for example, from the brainstem circuitry that processes the collicular output – a velocity or a position feedback. The models differ in the hypothesized mechanisms by which the feedback signals update the collicular activity. For example, Lefèvre and Galiana (1992) proposed a one-dimensional model of the collicular motor map in which the activity shift is achieved by an inhibitory projection of the velocity feedback onto the motor map and by a predefined connectivity pattern among the collicular neurons. Dominey and Arbib (1992) implemented the activity shift based on change-in-position signals calculated by differences of delayed versions of the neural integrators in the brainstem. The shift is then realized with a convolution operation. The model proposed by Van Opstal and Kappen (1993) realizes a flattening of the motor-map activity by means of a graded topographic projection of the velocity feedback onto the motor map. Droulez and Berthoz (1991) modeled the motor map as a retinotopic memory map that integrates the target displacement. The architecture and the structure of the connectivity pattern of their networks was derived from an analytical formulation of the activity shift. Massone (in press) proposed a scheme for learning in a supervised fashion the dynamics of a two-dimensional motor map based only upon velocity signals produced by the brainstem circuitry that processes the output of the motor map. In that model the activity of the motor map is updated by a multi-layer network placed in the feedback path. Such network was trained in a supervised fashion to realize a shift of the collicular activity.

Although the projection from downstream brainstem structures to the superior colliculus was anatomically investigated (e.g. Stechison et al. 1985), no definite experimental evidence exists as to which of the above-mentioned (or other) mechanisms the gaze system does, in fact, use in order to update the activity of the collicular motor map. In this paper we explore the possibility that some of the dynamic properties of the neural activity in the collicular motor map might be mediated by local interactions between movement-related neurons (BNs, BUNs) and fixation neurons, via the so-called intracollicular network, i.e. the network that connects the collicular neurons to each other. A recent electrophysiological study (Munoz and Wurtz 1993) demonstrated in fact that local interactions between movement neurons and fixation neurons do indeed exist and might play some role in coordinating the neurons' firing. More specifically, the goal of this research is to demonstrate, from a computational standpoint, which classes of dynamic behaviors of the collicular neurons can be obtained without the intervention of feedback signals so as to begin exploring to what extent the gaze system needs feedback in order to operate. To achieve this goal we undertook an approach based on recurrent learning that we will describe in detail in the next section. Consequently this research constitutes also a study of the properties and responses of a learning algorithm, namely recurrent backpropagation, under various choices for the network's and algorithm's parameters.

2 The Network and the Learning Algorithm

We represented one layer of the collicular motor map as a two-dimensional recurrent neural network composed of two interconnected pools of units, a fixation pool and a movement pool. Each unit in the network belonged to one of the two pools. The main idea behind our model is that a recurrent neural network is a dynamical system, and as such it exhibits all the properties that dynamical systems possess. Specifically, a recurrent network can be trained to settle into specific stable equilibrium points that correspond to specific configurations of activity of the network's units. Once the network is trained to settle into desired stable equilibria, whenever the network's units are in a non-equilibrium configuration, they undergo a transient between their initial state and one of the stable equilibria. Which of the equilibria the network will reach depends on the network's initial state. In this framework, the dynamics of the collicular neurons during a gaze shift can be viewed as the transient of a dynamical system from an out-of-equilibrium configuration to an equilibrium configuration (the idea of movement as the transition between the equilibrium states of a dynamical system was proposed by Schöner and Kelso (1988).) The one and only equilibrium configuration of such a collicular network is that in which all the neurons in the fixation zone are active and all the neurons outside the fixation zone are inactive – the configuration that was experimentally observed as the one corresponding to the end of a gaze shift (Munoz and Guitton 1991). An out-of-equilibrium configuration occurs whenever other areas of the motor map become active due to the occurrence of a motor error. In other words the occurrence of a motor error can be seen as an event that perturbs the system's equilibrium: it causes the intracollicular network to exit its equilibrium state and follow a transient in order to return to its equilibrium. Consequently, in such a system the duration of the transients is directly related to the duration of the gaze shifts.

We hence modeled the intracollicular networks for the right and left superior colliculi as a square array of 25x25 units (see Figure 1b) with the central 5x5 area representing a common right/left fixation zone [1]. For the connectivity pattern we initially chose the structure shown in Figure 1c: each unit projected to its four neighbours with bidirectional connections and received also a connection from itself (self-connection) [2]. This choice seemed a reasonable one to begin our study because local projective fields are a commonly encountered feature in many areas of the brain (Ramon y Cajal 1911, Gilbert 1983.) We did however experiment with different connectivity patterns, specifically with larger projective fields up to a fully connected network, and the results of these experiments will also be discussed in this paper. All the units in our network had a sigmoidal activation function, i.e. they computed their output as follows:

$$z_i = \sum_{j=1}^{N} w_{ij} y_j \tag{1}$$

$$y_i = f(z_i) = \frac{1}{1 + e^{-z_i}} \tag{2}$$

where w_{ij} is the weight from unit j to unit i, y_i and y_j are the outputs of units i and j respectively, N is the total number of units in the network. The network was trained with recurrent backpropagation (RBP) (Pineda 1989), which is a learning algorithm that makes it possible to set a recurrent network's equilibrium points. RBP is a well-known algorithm used to train recurrent networks. We will summarize here its basic principles for the sake of clarity and in order to illustrate our choices for the algorithm's parameters. The network's weights are initialized with random positive and negative values and the output of all the units is initialized to values that represent an initial state. Then three operations are performed sequentially: (a) a forward pass, (b) a backward pass, (c) a weight-update phase. The forward pass consists of computing the output of all the units based on the initial values of the outputs and on the initial values of the weights. This computation is repeated until the network relaxes into an equilibrium configuration [3], i.e. until the output of the units stops changing. The

[1]Our choice of representing the two colliculi as a single array of units was dictated by purely computational considerations (namely easier handling of the network) and does not imply that the two colliculi are anatomically joined as shown in Figure 1b.

[2]The total number of connections in this network (taking into account the degenerate cases at the edges) is 3025.

[3]RBP is based on the assumption that the recurrent network will reach a stable equilibrium during

forward pass for one unit can be described by the following differential equation:

$$\tau_{f_i} \frac{dy_i}{dt} = f(\sum_{j=1}^{N} w_{ij} y_j) - y_i \tag{3}$$

where τ_{f_i} is the time constant of unit i. When the forward pass is terminated, one can compute the error between the desired equilibrium configuration $\underline{y}^*_\infty$ and the equilibrium configuration reached by the network – $\underline{y}_\infty$. This error measure is used in the backward pass, which consists of relaxing to the equilibrium an ancillary network defined, for one unit, by the following differential equation:

$$\tau_{b_i} \frac{d\delta_i}{dt} = \sum_{j=1}^{N} f'(z_j)_\infty w_{ji} \delta_j + e_i - \delta_i \tag{4}$$

where τ_{b_i} is the time constant of unit i in the ancillary network, $f'(z_j)_\infty$ is the first derivative of $f(z_j)$ with respect to z_j computed at the equilibrium, $e_i = (y^*_{i,\infty} - y_{i,\infty})$. The values $\delta_{i,\infty}$ to which the ancillary network relaxes are then used in the weight-update phase as follows:

$$\tau_{ij} \frac{dw_{ij}}{dt} = f'(z_i)_\infty y_{j,\infty} \delta_{i,\infty} \tag{5}$$

where τ_{ij} is the time constant of weight w_{ij} and its inverse represents a learning rate. Various different forms exist for equations (4) and (5). The forms for equations (4) and (5) presented here are those given in Pineda (1987.) The sequence (3), (4), (5) is repeated until the error between the desired equilibrium state and the actual equilibrium state is lower than a predefined threshold. In practice one can run the three phases in parallel rather than in a sequential fashion by choosing the time constants such that $\tau_{f_i} \ll \tau_{b_i} \ll \tau_{ij}$. In our initial simulations we used values of 0.01, 0.1 and 10 for $\tau_{f_i}, \tau_{b_i}, \tau_{ij}$ respectively, $\forall i, \forall j$. RBP performs a gradient descent along a cost function

$$J = \frac{1}{2}(\underline{y}^*_\infty - \underline{y}_\infty)^T(\underline{y}^*_\infty - \underline{y}_\infty) \tag{6}$$

The algorithm just described can be extended to learning multiple equilibrium points, although in this case a number of problems might arise (Pineda 1988). It can also be extended to take into account the presence of external inputs. In our case, however, we limited our learning to the case of no external inputs and a single equilibrium point: a hill of activity in the fixation zone and no activity outside of it [4]. We therefore trained

---

the forward pass. However no formal proof of stability exists for generic connectivity patterns. The algorithm bases its validity on the heuristic that if the network stabilizes during the initial forward pass and if the weight-change process occurs slowly over time, then the network is likely to remain stable (Pineda 1989.) This was always the case in all our training sessions that used sigmoidal activation functions. The stability of the backward pass is guaranteed by the stability of the forward pass (Almeida 1987.)

[4] We chose a gaussian shape for the activity of the fixation zone because this is how this activity is often depicted in the literature (Munoz and Wurtz 1992, Munoz and Guitton 1991). Our training procedure could however be applied to different shapes of the fixation-zone activity (e.g. constant value for all fixation units including the baseline value) as long as the equilibrium configuration is defined accordingly.

the network shown in Figures 1b and 1c to reach such equilibrium. As the initial state we chose a configuration of the network that corresponded to a hill of activity in a randomly selected location of the map. The hill covered a 5x5 area of the map and we modeled the firing rate of the neurons in the hill with real numbers between 0.2 and 0.9. We chose 0.2 as the value of the neurons baseline rather than a value close to zero for two reasons: (i) designing a network with unrestricted dynamic behaviors including oscillations around the baseline value, (ii) avoiding the tail of the sigmoid, which would make the learning harder and perhaps impossible. The network was trained to reach the equilibrium from only one initial state, i.e. from only one position of the hill on the map.

It is important to point out that the only effect of RBP is to train the network to reach the desired equilibrium configuration, as stated by equation (6). There is no specification whatsoever in the algorithm as to what the transient from the initial state to the final state should be. Because the network is computationally redundant (the output of each unit is determined by five weights) and gradient-descent methods are only locally optimal, there is theoretically an infinite number of solutions that the learning algorithm can produce in order to reach the equilibrium. In other words the resulting transients cannot be predicted.

3 Results

3.1 Basic Results

This subsection presents the results of training our recurrent network with RBP. In our first training attempt, we initialized the weights with random values between -0.5 and 0.5 [5]. The network learned to relax to the desired equilibrium configuration in approximately 28,000 learning iterations, one iteration being defined by a forward pass, a backward pass and a weight-update phase. After 28,000 iterations the value of J was 1.0. The error could be lowered to 0.01 by significantly increasing the number of iterations, but with $J = 1.0$ the basic behavior of the network was already well defined. Consequently we chose a value for J near 1.0 as the threshold for all our learning sessions. After the learning phase was over we studied the behavior of the network (recall phase). The recall phase consisted of placing a hill of activity in the same location used during learning and running the network's forward pass defined by equation (3). Figures 2 and 3 show the results. Figure 2a is a plot of J versus the number of forward pass iterations during recall and shows that the network reached a configuration close to the equilibrium after approximately thirty iterations. Twenty additional iterations were required in order for the network to stabilize completely. In all our simulations one iteration corresponded to 1 msec. Figure 2b shows the output of all the units in the network during recall, i.e. the network's transient. The trajectories that decrease over time are the trajectories of the units that were initially active, i.e. the units on which the hill was initially placed. The trajectories that increase over

[5] We provide this information here because in later simulations (see Section 3.4) we studied the relationship between the weight initialization and the outcome of the learning algorithm.

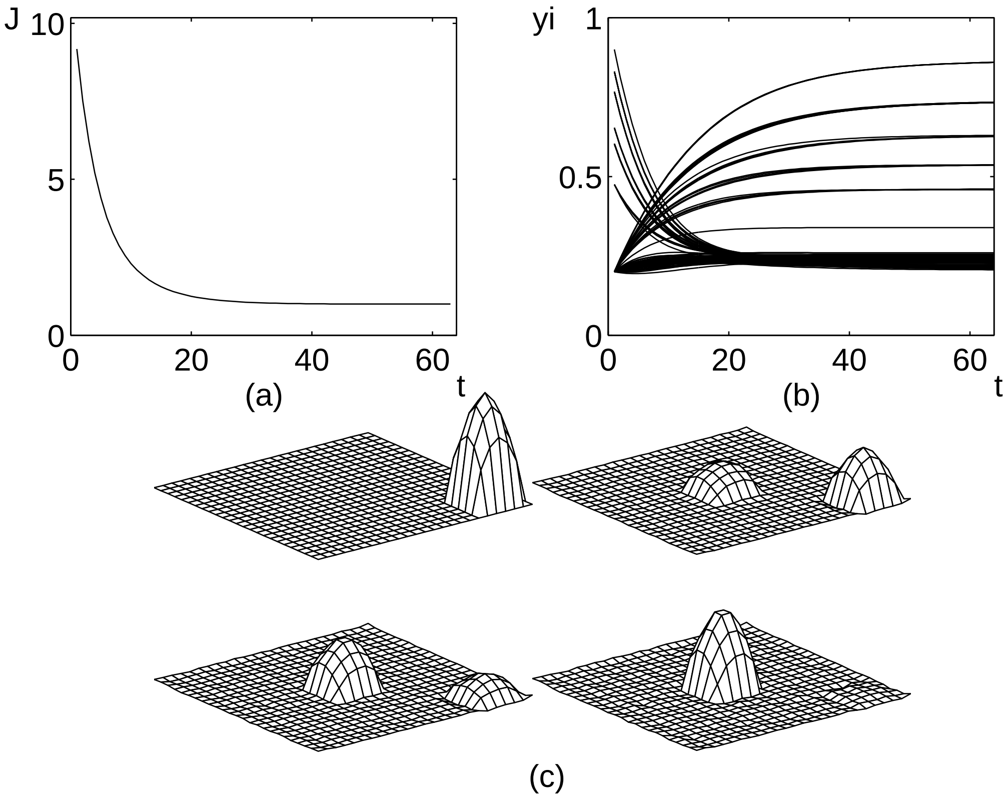

Figure 2: (a) The cost function J versus the number of forward pass iterations during recall. (b) The output y_i of all the network's units versus the number of forward pass iterations during recall. (c) The motion of the hill during recall at iterations number 0 (top-left), 5 (top-right), 10 (bottom-left), 20 (bottom-right). The starting location of the hill is the same used during training.

time are the trajectories of the units in the fixation zone. The trajectories that remain near the baseline are the trajectories of the units outside the fixation zone that were initially inactive. Figure 2c shows the corresponding motion of the hill. These results demonstrate that the transient of the trained network is a hill-flattening pattern as observed in the BN layer of the superior colliculus of rhesus monkeys (Munoz and Wurtz 1993.) This result was obtained despite the fact that the learning algorithm did not specify what the network's transients should be. This result points out the possibility that the dynamics of the burst-neuron layer of the superior colliculus might be locally regulated rather than feedback-driven, and that the action of the feedback might be confined to the layer of the buildup neurons. As far as fixation neurons are concerned, Figure 2b shows that the activity in the fixation zone increases from its baseline value to the final values throughout the transient. This is not, however, what was observed by the experimenters, who showed that fixation neurons become active only towards the end of the movements. In Section 3.6 we will demonstrate how the rise of activity in the fixation area can be delayed in order to obtain a better match with the experimental data.

Figure 3 shows in a qualitative fashion the distribution of the weights before training (3a and 3b) and after training (3c and 3d). Figures 3a and 3c show a 25x25 array of squares in which each square corresponds to a unit's self-connection. In Figures 3b and 3d each of the 25x25 squares was divided into 4 subsquares representing the four weights impinging on one unit. Lighter colors correspond to larger values of the weights. A quantitative analysis of the values of the weights after training showed that: (1) The range of the final weights was between -1.9 and 0.9. (2) All the weights outside the fixation zone were negative (inhibitory connections). (3) All the weights inside the fixation zone were positive (excitatory connections). This result suggests that the dynamics of the burst neurons might be regulated by an inhibitory network. The presence in the superior colliculus of an inhibitory network of interneurons was hypothesized before (see e.g. Van Opstal and Van Gisbergen (1989), Van Gisbergen and Van Opstal (1989)) although with a different function than regulating the neurons' dynamics. As far as fixation neurons are concerned, the presence of an excitatory network in the fixation zone seems consistent with the findings by Munoz and Wurtz (1993), who showed that microstimulation of the rostral SC (where the fixation zone is located) led to antidromic and/or ortodromic activation of some fixation neurons.

We tested the generalization ability of the network by initially placing the hill in locations of the motor map different from the location used during learning. Figure 4 shows the motion of the hill for two such cases. Specifically, Figure 4a refers to a hill placed in a different zone of the map but at the same distance from the fixation zone as the location used during training. This corresponds to a gaze shift of the same amplitude in a different direction. Figure 4b refers to a hill placed in a different location of the map and closer to the fixation zone than the training location. This corresponds to a gaze shift of a smaller amplitude in a different direction. The results of these simulations show that in both cases the network's transients were the same hill-flattening patterns observed for the training location.

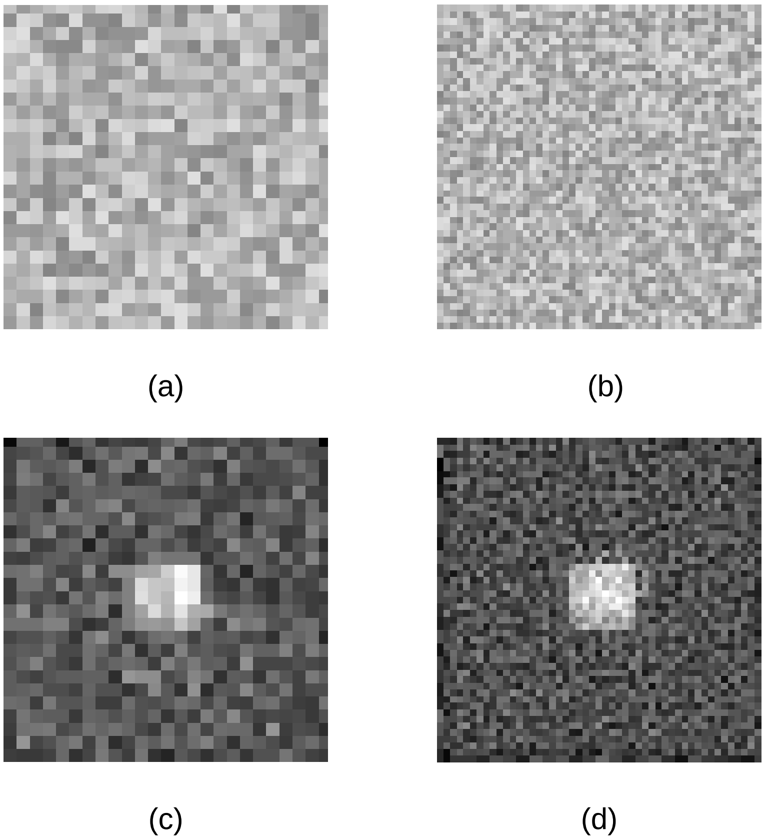

Figure 3: A qualitative representation of the network's weights before and after training. Lighter colors correspond to larger weights. (a) Self-connections before training. Each square represents one connection. (b) Inter-unit connections before training. Each group of four squares represents the four connections impinging on one unit. (c) Self-connections after training. (d) Inter-unit connections after training.

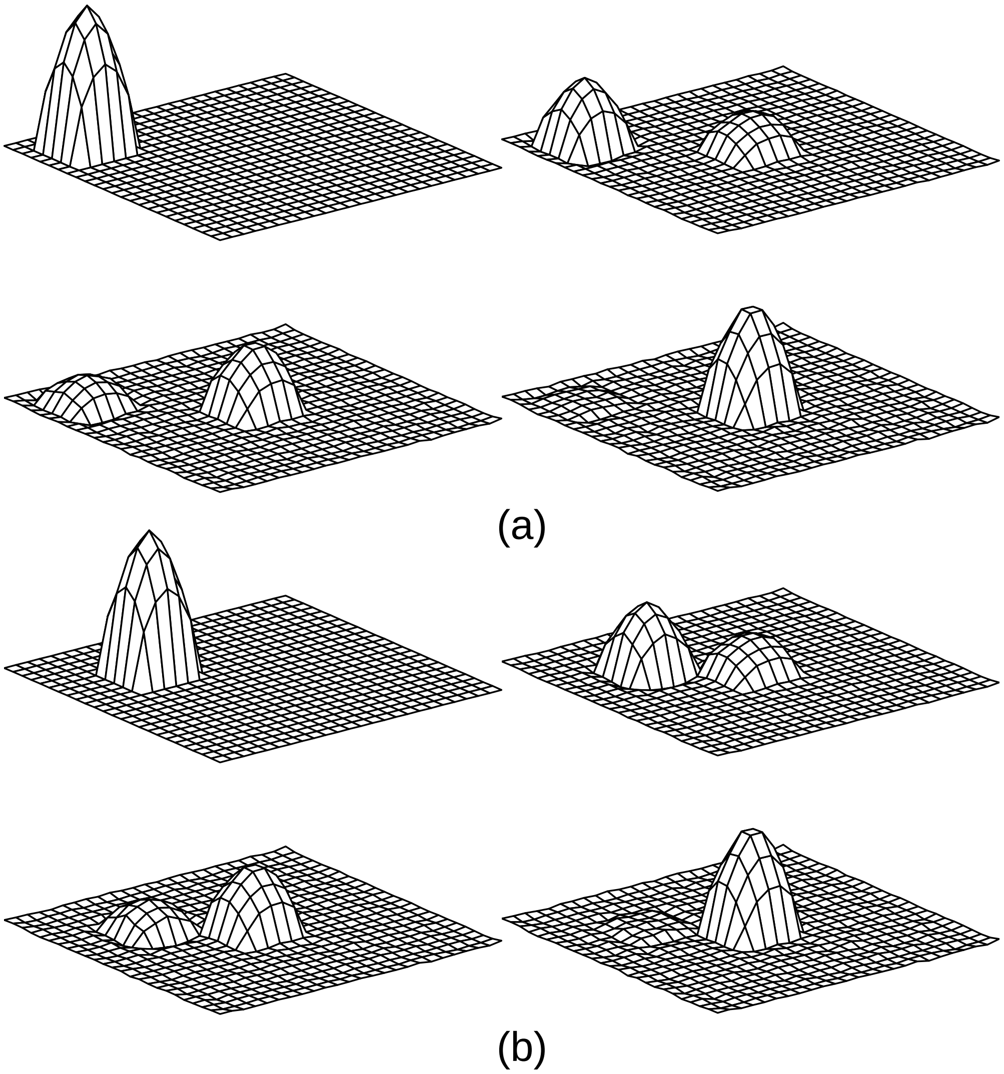

Figure 4: (a) The motion of the hill during recall at iterations number 0 (top-left), 5 (top-right), 10 (bottom-left), 20 (bottom-right). The starting location is different from the training location but at the same distance from the fixation zone. (b) The motion of the hill during recall at iterations number 0 (top-left), 5 (top-right), 10 (bottom-left), 20 (bottom-right). The starting location is different from the training location in terms of both position and distance.

3.2 Transient Duration

We observed that the duration of the flattening was always constant, irrespectively of the distance between the hill and the fixation zone, i.e. of the amplitude of the gaze shifts. Experimental data show, however, that the duration of the movements increases fairly linearly with the movements amplitude. This result was often reported in the experimental literature. An example of a duration-amplitude curve is given in Carpenter (1988). While many of the computational models introduced in Section 1 can account for the linear duration-amplitude relationship by means of various different assumptions and mechanisms, it is still unclear how such a relationship is implemented in the brain. We demonstrate here two computational mechanisms by which the duration-amplitude relationship can be controlled in our network. One possibility is to assume that the time constants of the units, τ_{f_i}, are not constant throughout the map, but they increase with the distance of a unit from the fixation zone. We then retrained the network with values of τ_{f_i} distributed between 0.005 and 0.05 according to the units' distance from the fixation zone. The resulting transients had the same hill-flattening behaviors but their duration was longer for hills placed farther away from the fixation zone, as shown in Figure 5a. A second possibility for controlling the duration-amplitude relationship

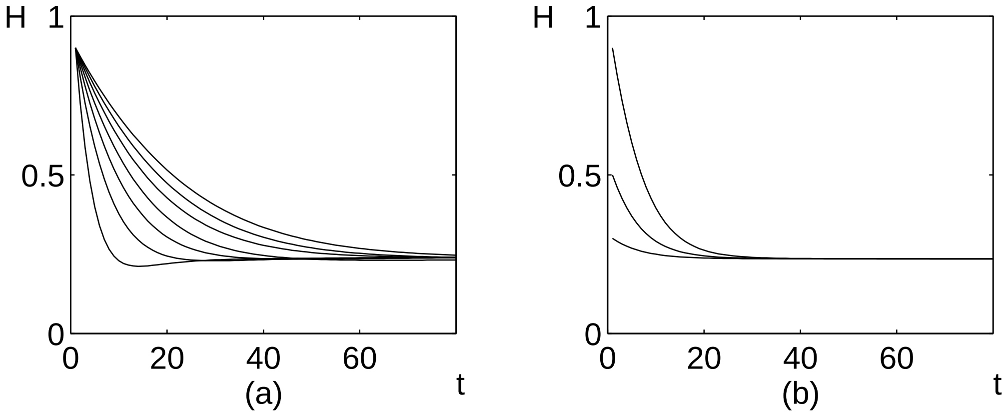

Figure 5: The height of the hill (H) at the initial location versus the number of forward pass iterations during recall. (a) The different curves refer to hills placed at different distances from the fixation zone. The units in this network had different time constants τ_{f_i} depending on their distances from fixation. The rightmost curve corresponds to $\tau_{f_i} = 0.032$, the leftmost curve to $\tau_{f_i} = 0.005$, the other curves to values in between. (b) The three curves refer to hills with initial height of 0.9, 0.5 and 0.3 respectively.

is the following. With the trained network we set the initial height of the hill to 0.5 and then to 0.3 (the initial height was 0.9 during learning.) The purpose of this experiment was to study the velocity of flattening of the hill as a function of its height. The results are shown in Figure 5b and they demonstrate that the taller the hill, the larger the number of iterations required for the hill to flatten and hence the larger the duration of the corresponding gaze shift. The results shown in Figure 5 demonstrate that in our network space-variant time constants and space-variant hill heights are effective computational mechanisms to achieve a desired relationship between duration and amplitude of the gaze shifts. As far as their biological plausibility is concerned, the latter mechanism is not likely to represent a biologically-plausible solution because, as pointed out in the introductory section, experimental data indicate that the firing rate of the collicular neurons is not correlated with movement amplitude. Wether the former mechanism - space-variant time constants - is a biologically-plausible one is left here as an open question.

3.3 Perturbation Experiments

We tested the robustness of the trained network to abnormal stimulations. We utilized a number of abnormal patterns including a double hill of activity and a hill of activity plus noise. Figure 6 shows the results for these two cases. It can be observed that in both cases the basic characteristic of the transient (hill-flattening) remains unchanged, and that the network reaches the same equilibrium configuration. This result is not surprising and is due to the fact that in a dynamical system with only one equilibrium configuration that configuration is reached independently of the initial state of the network. In other words, the process by which the network reaches its equilibrium is analogous to a pattern-completion process of a content-addressable memory (Hopfield 1982.) This property emphasizes the fact that a network that regulates its dynamic behavior through its own connectivity pattern is a very robust and reliable device. In the case of the double-stimulation (Figure 6a) we can notice that the two hills are treated by the network as two separate entities, and flattened independently. This independent processing of multiple hills is the same result reported by Massone (in press) for a motor map that was trained in a supervised (and totally different) fashion to shift the collicular activity towards the fixation zone. We will discuss the implications of this result in Section 4.

3.4 Connectivity Pattern and Weight Initialization

We studied the role of the weights initialization and the connectivity pattern in the outcome of the learning. The results presented in Section 3.1 to 3.3 were obtained with a local connectivity pattern (Figure 1c) and with the weights initialized randomly in the interval [-0.5, 0.5]. For this network, we analyzed the transient before learning, partially shown in Figure 7a. We can notice that with random values of the weights in the interval [-0.5, 0.5] the units' output decays or rises from an initial value to a final equilibrium value, and that the equilibrium values fall within a small portion of the output range. We then initialized the same network with larger values of the weights, namely in the interval [-5.0, 5.0]. During the transient caused by these large random weights (see Figure 7b) most of the units exhibited an oscillatory behavior and the equilibrium values were spread over the entire output range. We then trained this network on the same task that we used to train the network with small initial weights. The purpose of this experiment was to determine wether the characteristics of the initial transient (decays/rises vs. oscillations) would influence the characteristics of the transient of the trained network. Figure 8 shows the final transient for the network intialized with large random weights. We can notice that the solution found by the learning algorithm is still

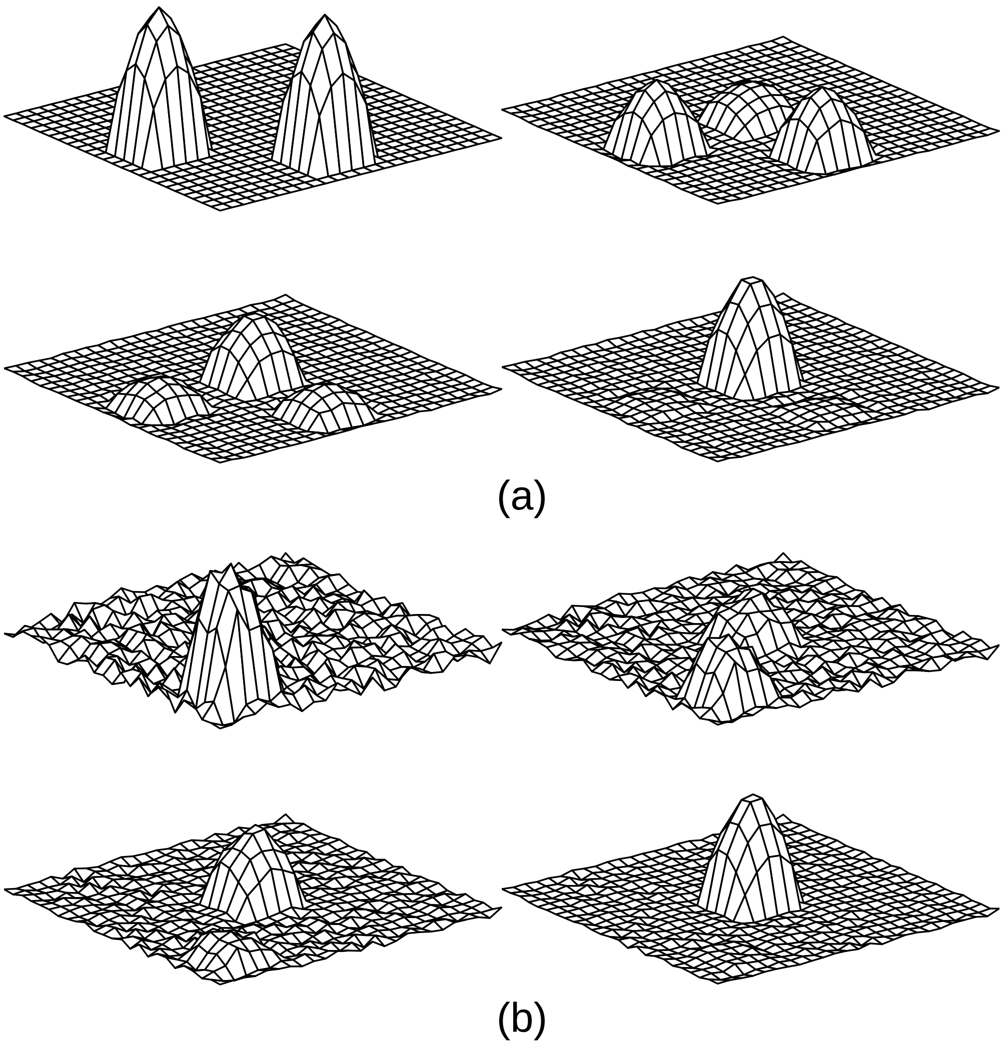

Figure 6: (a) The network's transient following stimulation of the motor map with two hills of activity. (b) The network's transient following stimulation of the motor map with a hill of activity plus random noise.

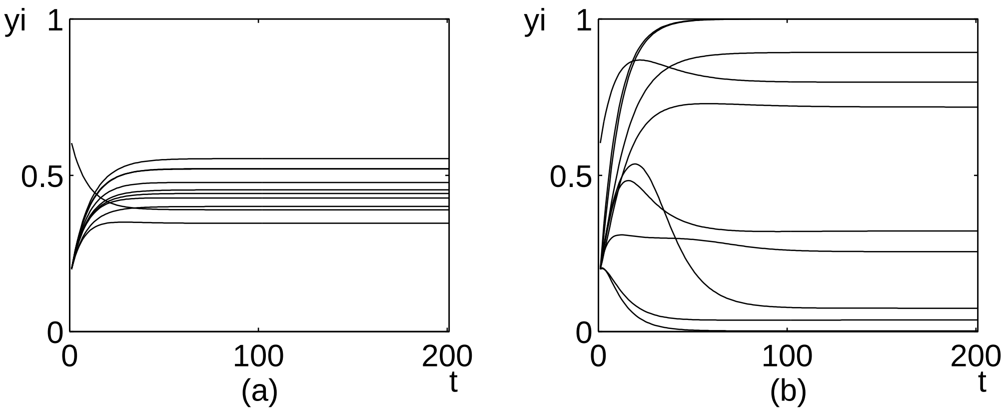

Figure 7: (a) A subset of the units' outputs during the transient produced by random weight values initialized in the interval [-0.5, +0.5]. (b) A subset of the units' output during the transient produced by random weight values initialized in the interval [-5, +5].

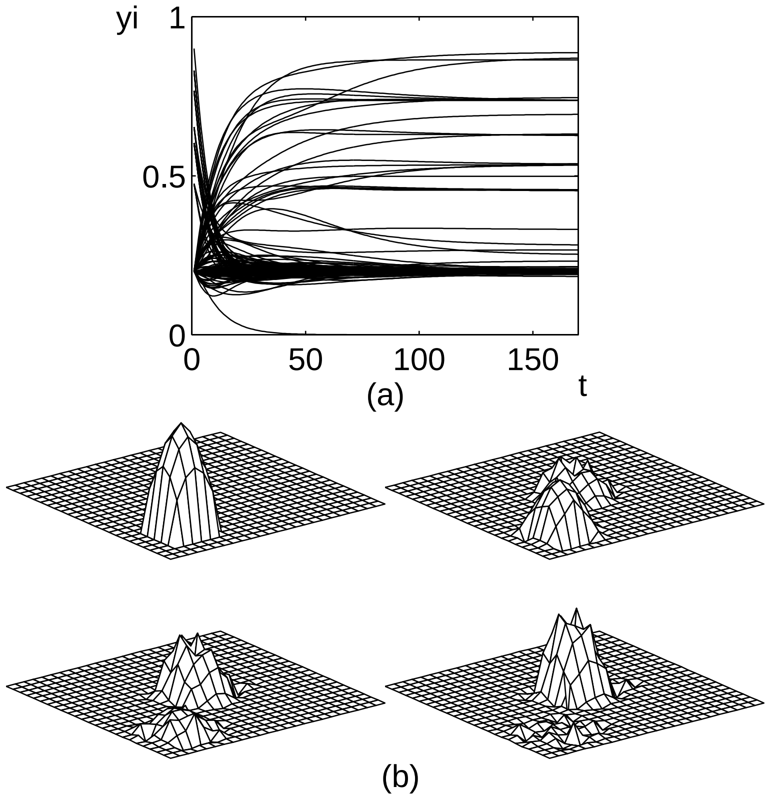

Figure 8: (a) The units' output for the network that was trained from initial random weight values in the interval [-5, +5]. (b) The corresponding motion of the hill.

a hill-flattening solution. We were able to eliminate the apparent noisy nature of the transient by further lowering the error during learning. We also found that the weights of the trained network were globally organized as in Figure 3 (excitatory in the fixation zone and inhibitory elsewhere.) This result indicates that the flattening transient that we observed for the first network as well as the resulting organization of the weights were not a consequence of the chosen small range for the weights initialization. We then investigated wether our results could be a consequence of the chosen connectivity pattern. First we eliminated all the self-connections and we retrained the network. The results were identical to the case in which the self-connections were present, indicating that self-connections are not crucial parameters to our network. We then carried out an additional learning session with a fully connected network, i.e. a network in which every unit was connected to every other unit. We began with random weights in the interval [-0.5, 0.5]. Figure 9a shows the transients of this network before learning, Figure 9b the transients of the same network after learning, and Figure 9c the motion of the hill. The transient before learning was an oscillatory one, even with small weight

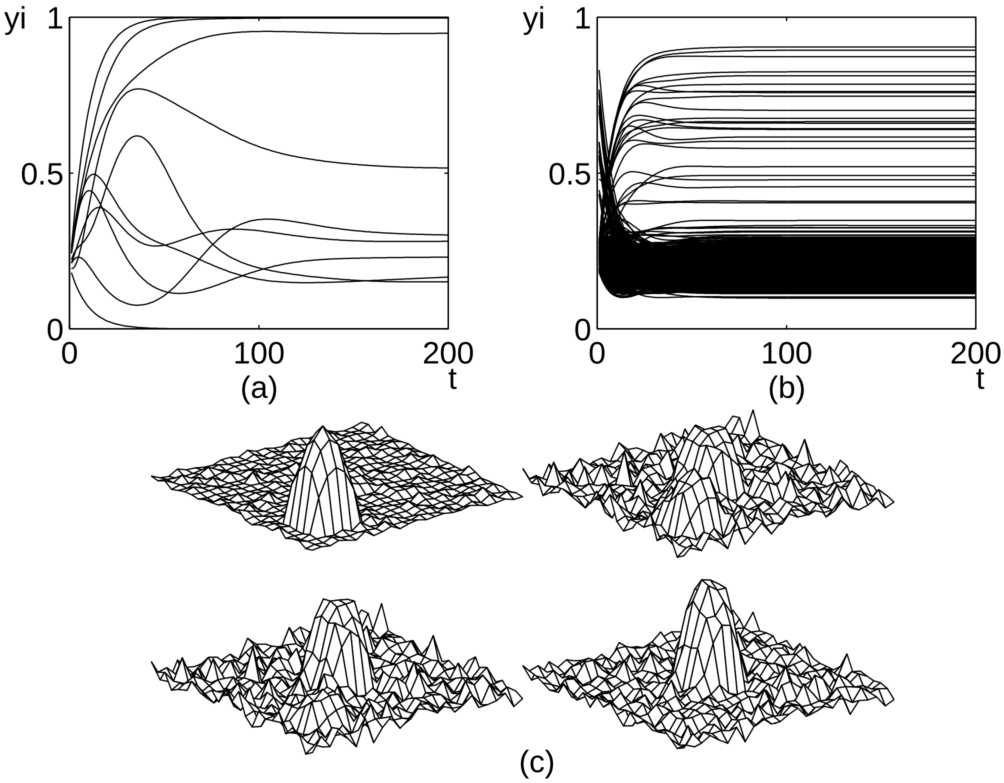

Figure 9: Fully-connected network. (a) A subset of the units' outputs during the transient produced by random weight values initialized in the interval [-0.5, +0.5]. (b) The units' output for the trained network. (c) The corresponding motion of the hill.

values. The transient after learning was again a flattening one. The resulting weights however, were not organized as in Figure 3. They actually did not exhibit any sort of organization. This is due to the fact that in order for such a large number of weights to organize themselves into structured patterns the training process should include a larger training set. In all the training sessions presented in this paper we limited the training to one initial location of the hill. We then tested the fully connected network with the weights initialized in the interval [-5.0, 5.0]. With the initial weights in this range the fully connected network became immediately unstable (see Figure 10a) and hence could not be trained. These results demonstrate that the connectivity pattern

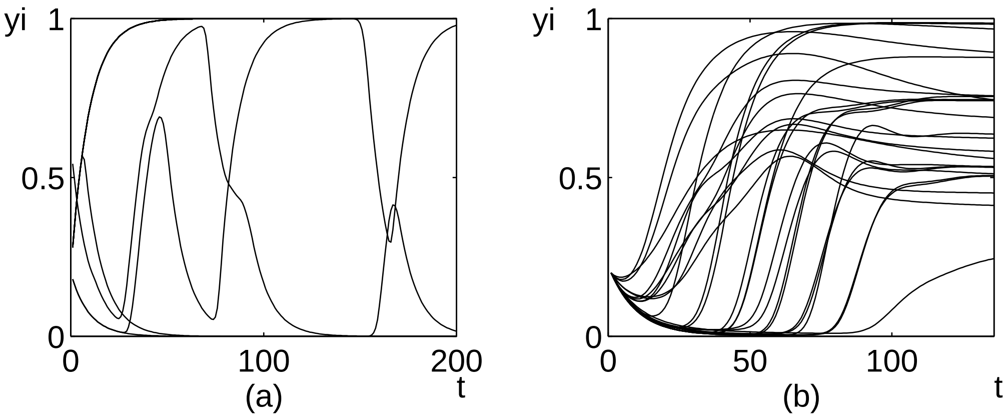

Figure 10: (a) Fully-connected network. A subset of the units' outputs during the transient produced by random weight values initialized in the interval [-0.5, +0.5]. (b) The output of the units that represent fixation neurons after the locally-connected network was trained in presence of fixed negative weights from each unit outside the fixation zone to each unit inside the fixation zone. The value of the fixed inhibition was arbitrarily set, in this case, to -0.05.

is not a crucial factor when the corresponding network is stable. It is however an important factor for the stability of the network: large initial values of the weights caused instability in the fully connected network but not in the locally connected one.

3.5 Activation Function

We carried out an additional learning session with the locally-connected network using a linear activation function rather than a sigmoidal one. We would like to point out that one of the differences between a linear dynamical system and a nonlinear one is that the nonlinear system can exhibit multiple equilibria while the linear system is bound to possess a single equilibrium point. Because our task implied learning a single equilibrium configuration, there was no theoretical obstacle to achieving our goal with a linear system. With a linear choice for the units' activation function, however, the network became immediately unstable, even with small initial values of the weights, and hence could not be trained. This result indicates that the saturation in the neurons firing rate plays a crucial role in the production of meaningful dynamic behaviors.

3.6 Fixation Neurons

In Section 3.1 we pointed out that in our trained network the activity of the fixation neurons increased gradually during the transient from the baseline value to the equilibrium value (see for example Figure 2b), while experimental data show that fixation neurons become active only towards the end of the movement. While one can assume that the rise of activity in the fixation zone is delayed and regulated by mechanisms other than interactions between burst neurons and fixation neurons (e.g. by mechanisms that occur in the buildup-neuron layer as postulated by Munoz and Wurtz (1993)), we wanted to investigate to what extent this task could be achieved with our network. In order to do that, we added to the locally-connected network a set of fixed inhibitory connections from each unit representing movement neurons to each unit representing fixation neurons. This organization is consistent with data by Munoz and Wurtz (1993) who reported that electrical stimulation of the caudal colliculus causes short-latency inhibition of the fixation zone. We then trained the new network. The values of the new inhibitory weights were kept fixed during learning, while the other weights were learned as previously described. The results are shown in Figure 10b and they demonstrate that in presence of the additional inhibitory connections the rise of activity in most fixation units is delayed. The results of the learning were not as clean as in the previous sessions (the equilibrium values for the units in the fixation zone could be learned with a larger approximation than in previous sessions, as Figure 10b shows, and with a much larger number of iterations) because with a subset of the weights kept fixed during training the learning algorithm no longer follows the exact direction of the gradient. We could obtain different delay patterns with different values of the fixed inhibition and avoid firing rates below the baseline value by moving the baseline closer to zero (e.g. 0.5).

4 Discussion

In this paper we explored the possibility that some of the dynamic properties of the neural activity in the collicular motor map might be mediated by local interactions between movement-related neurons and fixation neurons. In particular we studied, from a computational standpoint, which classes of dynamic behaviors of the collicular neurons can be obtained without the intervention of feedback signals, so as to begin exploring to what extent the gaze system needs feedback in order to operate. We modeled the motor map as a dynamical system realized with a recurrent neural network. The network contained two interconnected pools of units that represented fixation neurons and movement neurons respectively. We modeled the dynamics of the neural activity in the map as the transients of that system towards an equilibrium configuration that the network learned with recurrent backpropagation. Hence this paper constitutes also a study of the properties and responses of recurrent backpropagation under various choices for the network's and algorithm's parameters. The chosen equilibrium configuration was one in which the fixation zone contained a gaussian hill of activity and the rest of the units were inactive, i.e. active at their baseline value. As the initial state of the network, we chose a configuration in which all units but 25 fired at their baseline value. The remaining 25 units contained a hill of activity in a randomly-selected location. With this choice for the initial state of the network we implicitly assumed that fixation neurons were already in their "off" state and the gaze shift ready to begin. We

will elaborate on this assumption shortly. The results of our study can be summarized as follows:
(1) The transients of the trained network were hill-flattening patterns as observed by some experimenters in the burst-neuron layer of the superior colliculus of rhesus monkeys. This result was obtained despite the fact that the learning algorithm did not specify what the network's transients should be, the computational redundancy of the network, and the local character of teh gradient-descent method. We interpret this result as an indication that the dynamics of the burst neurons in the superior colliculus might be locally regulated rather than feedback-driven, and that the action of the feedback might be confined to the layer of the buildup neurons. This result also justifies our assumption on the initial state of the network, namely that the fixation neurons are initially in their "off" state. One can indeed attribute the function of turning off fixation neurons prior to a gaze shift to the other population of movement neurons – the buildup neurons. Experimental observations tell us, in fact, that the activity in the BUN layer precedes the onset of a gaze shift, while activity in the BN layer accompanies the gaze shift itself (Munoz and Wurtz 1993.) If the transients of the trained network had been shifts of the activity rather than decays, then our assumption on the initial state of the fixation neurons to be the "off" state would not be a plausible one.
(2) The connections in the trained network were excitatory within the fixation zone of the motor map and inhibitory elsewhere, suggesting that dynamic properties of the burst neurons might be regulated by an inhibitory network.
(3) The rise in the activity of the fixation neurons could be delayed by adding fixed inhibitory connections from the movement neurons to the fixation neurons. We are not, however, implying that this is necessarily the mechanism utilized by the real system: the activity of the fixation neurons might as well be regulated by mechanisms other than local interactions between burst neurons and fixation neurons, e.g. by mechanisms that occur in the layer of the buildup neurons (Munoz and Wurtz 1993). The purpose of this experiment was just to determine if delaying the rise of the activity in the fixation zone with a local inhibitory network was a feasible computational task.
(4) The duration of the flattening could be controlled with two computational mechanisms: space-variant time constants and space-variant hill height. The former might possibly constitute a biologically plausible mechanism, not the latter.
(5) The results of the learning were robust in the face of changes in the connectivity pattern and in the weight initialization, but a local connectivity pattern favored the network's stability. We found, in fact, that with a local connectivity pattern the network could tolerate a larger range of weight values without undergoing unstable oscillations. This result indicates that stability could be one of the reasons why local connectivity patterns are so widely utilized by the central nervous system.
(6) Nonlinearity was required in order to obtain meaningful dynamic behaviors. A linear version of the network could not be trained despite the fact that the training was theoretically possible.
(7) The trained network was robust to abnormal stimulation patterns such as larger, noisy and multiple stimuli. The robustness of the network was a consequence of the pattern-completion abilities of recurrent networks. In addition we found that, when the motor map was activated with a double stimulation, the two hills of activity were

treated as separate entities and flattened independently. When Robinson (1972) performed double-stimulation experiments in primates, he found that the direction of the resulting movement was quite consistently the average of the directions dictated by each stimulation. Robinson's results are often taken as an indication that the output of the collicular motor map is an average of its neural activity. Our simulation results indicate that, with a motor map represented as our recurrent network, the activity is not averaged locally and the dynamics of the motor map remains a stereotyped one even when the motor map is abnormally stimulated. Interestingly, we obtained the same identical result (separate handling of multiple hills) with another and quite different computational model of the gaze system (Massone in press) in which the collicular activity was trained in a supervised fashion to shift towards the fixation zone. That model also included the circuitry downstream from the superior colliculus as well as the eye plant. The results obtained with that model pointed out the possibility that the averaging of the direction of movement following a double stimulation of the motor map might be a consequence of the distribution of the weights that carry the output of the superior colliculus. In Massone (in press) we argued about the introduction of non linearities and anisotropies in that distribution. The results of the double-stimulation experiment presented in this paper – an independent flattening of the two hills – constitutes, in our opinion, additional support for arguing that the averaging of the collicular activity does not occur at the level of the motor map but is accounted for by an appropriate distribution of the outgoing weights. To the best of our knowledge, the collicular activity during a double stimulation has not been recorded. Only the resulting eye movements have. Our results point out that it is important that such an experiment be performed, in order to assess the validity or invalidity of what our two neural-network models suggest, namely that the averaging of the collicular activity does not occur in the motor map, but occurs, indeed, in downstream areas.

Acknowledgements
The authors would like to thank David Ferster, David Mogul, Mark Segraves, John Troy and two anonymous referees for their inputs and comments on early versions of this manuscript. This research was supported by the National Science Foundation, grant number BCS-9113455 to Lina L.E. Massone.

References

Almeida L B (1987) A learning rule for asynchronous perceptrons with feedback in a combinatorial environment *Proceedings International Conference on Neural Networks* (San Mateo, Ca: Morgan Kaufmann) pp 609-618

Berthoz A, Grantyn A and Droulez J (1987) Some collicular neurons code saccade velocity *Neuroscience Letters* **72** 289-294

Carpenter R H S (1988) *Movements of the eyes* (London: Pion Limited)

Cynader M and Berman N (1972) Receptive-field organization of monkey superior colliculus *Journal of Neurophysiology* **35** 187-201

Dominey P F and Arbib M A (1992) A cortico-subcortical model for generation of spatially accurate sequential saccades *Cerebral Cortex* **2** 153-175

Droulez J and Berthoz A (1991) A neural network model of sensoritopic maps with predictive short-term memory properties *Proceedings National Academy of Sciences U.S.A.* **88** 9653-9657

Gilbert C D (1983) Microcircuitry of the visual cortex *Annual Review of Neuroscience* **6** 217-247

Hopfield J J (1982) Neural networks and physical systems with emergent collective computational abilities *Proc. National Academy of Sciences U.S.A.* **79** 2554-2558

Lefèvre P and Galiana H L (1992) Dynamic feedback to the superior colliculus in a neural network model of the gaze control system *Neural Networks* **5** 871-890

Massone L (in press) A neural-network system for control of eye movements: basic mechanisms *Biological Cybernetics*

Munoz D and Guitton D 1986 Presaccadic burst discharge of tecto-reticulospinal neurons in the alert head-free cat *Brain Research* **398** 185-190.

Munoz D P and Guitton D (1991) Control of orienting gaze shifts by the tectoreticulospinal system in the head-free cat. II. Sustained discharges during motor preparation and fixation *Journal of Neurophysiology* **66** 1624-1641

Munoz D P, and Wurtz R H (1992) Role of rostral superior colliculus in active visual fixation and execution of express saccades *Journal of Neurophysiology* **67** 1000-1002

Munoz D P and Wurtz R H (1993) Interactions between fixation and saccade neurons in primate superior colliculus *Society for Neuroscience Abstracts* **19** 787

Munoz D P, Pelisson D and Guitton D (1991) Movement of neural activity on the superior colliculus motor map during gaze shifts *Science* **251** 1358-1350

Pineda F (1987) Generalization of backpropagation to recurrent neural networks *Physical Review Letters* **18** 2229-2232

Pineda F (1988) Dynamics and architecture for neural computation *Journal of Complexity* **4** 216-245

Pineda F (1989) Recurrent backpropagation and the dynamical approach to adaptive neural computation *Neural Computation* **1** 169-172

Ramon y Cajal S (1911) *Histologie du Système Nerveux de l'Homme et des Vertébrés* (Paris: A. Malsine)

Robinson D A (1972) Eye movements evoked by collicular stimulation in the alert monkey *Vision Research* **12** 1795-1808

Robinson D A (1975) Oculomotor control signals. *Basic Mechanisms of Ocular Motility and their Clinical Implications* eds P. Bach-y-Rita and G. Lennerstrand (Oxford:Pergamon Press) pp 337-374

Roher W H, White J M and Sparks D L (1987) Saccade-related burst cells in the superior colliculus: relationship of activity with saccade velocity *Society for Neuroscience*

Abstracts **13** 1092

Schöner G and Kelso J A S (1988) Dynamic pattern generation in behavioral and neural systems *Science* **239** 1513-1520

Stechison M T, Saint-Cyr J A and Spence S J (1985) Projections from the Nuclei Prepositus Hypoglossy and intercalatus to the superior colliculus in the cat: an anatomical study using WGA-HRP *Experimental Brain Research* **59** 139-150

Van Gisbergen J A M and Van Opstal A J (1989) Models. *The Neurobiology of Saccadic Eye Movements* eds. R. H. Wurtz and M. E. Goldberg (Amsterdam: Elsevier) pp 69-101

Van Opstal A J and Kappen H (1993) A two-dimensional ensemble coding model for spatio-temporal transformation of saccades in monkey superior colliculus *Network* **4** 19-38

Van Opstal A J and Van Gisbergen J A M (1989) A nonlinear model of collicular spatial interactions underlying the metrical properties of electrically elicited saccades *Biological Cybernetics* **60** 171-183

Waitzmann D M, Ma T P , Optican L M and Wurtz R H (1991) Superior colliculus neurons mediate the dynamic characteristics of the saccades *Journal of Neurophysiology* **66** 1716-1737