

Chapter 6

Models and results

During visually guided reaching for objects, our brain transforms a visual stimulus on the retina into a finely coordinated motor act. This process employs the activity of neural populations distributed along several areas in parietal and frontal cortex (Section 3.2). An ubiquitous feature of neurons in the parieto–frontal region is their broad selectivity to movement direction (Sections 2.1.2). This directional tuning of the neuronal responses has two consequences. First, the generation of motor output is based on a population coding of movement direction (Sections 2.1.3). Second, co–activation of neurons with similar tuning properties can establish a mechanism for the transfer of visual information into the motor command required for reaching.

In this chapter, we address two basic questions on the cortical control of the direction of movement. First, we investigate by means of computational modeling the mechanisms whereby motor neurons develop directional selectivity. Section 1 describes the formation of a directional motor map based on a self–organizing process involving spiking neurons and using input patterns analogous to proprioceptive feedback. The neural responses of individual units are characterized and the resulting population code is analyzed and compared to experimental findings of motor coding of movement.

Second, the self–organized motor map is re–used for the simulation of a developmental process of eye–hand coordination acquisition. The model proposed in Section 2 explores the means by which visual directional activity can be correctly used to guide the generation of the desired motor action. Visual and motor network alignment is discussed in the light of the connectivity pattern organization and of the individual unit’s contribution to the map–

ping process.

6.1 Self-organization of neurons in motor cortex for coding the direction of movement

Modeling of the cortical coding of motion direction is relevant to understanding how control of movement is implemented in the central nervous system. It also represents a crucial obstacle to be overcome if the goal is the *understanding-for-reproducing* of the mechanisms involved in visual guidance of reaching movements. Given the importance of the subject, it is surprising that only a very small number of studies have been aimed so far at exploring the organization of directional motor maps (see Sections 2.2.3, 2.3). Previous work on coding of movement trajectories by neural population vectors (Lukashin and Georgopoulos, 1994) and on visuomotor mapping of direction (Burnod et al., 1992) represents only a starting point in modeling the organization of motor networks. What is still needed is a self-organizing neural network model to investigate how directional selectivity and population coding emerge in the motor cortex.

In this section such a neural network model is proposed. The simulation work is grounded in the experimental data showing that directional tuning is a prominent feature of motor neurons (Georgopoulos et al., 1984). The main findings of the experimental studies are: (1) neurons in the motor cortex are broadly tuned to preferred directions of movement (Section 2.1.2); (2) motion in a certain direction is determined by the activity of a large population of neurons; (3) the coordinated action of neurons can be characterized by using a neural population vector, which proves to be a good predictor of motion direction (Section 2.1.3); (4) particular directions are multiply represented in the motor cortex and cover a directional continuum; (5) the strength of the connection between two neurons tends to be negatively correlated with the angle between their preferred directions (Section 2.2.3).

A common approach in building neural networks that code motion directionality is to assign neurons with preferred directions from a directional continuum (see Section 2.3.2). This is done without considering the development of the neural selectivity or of the way neurons relate with each other as a consequence of having certain preferences. Instead of taking this view, we are interested in investigating the processes whereby directional selectivity emerges in a neural network. We want to explore how and why individual neurons learn to respond maximally for a particular movement direction, and not for another. Furthermore,




Figure 6.1: The motor system consists of a two-dimensional self-organizing feature map of spiking neurons. The input layer is fully connected to the output map. This has a heterogeneous composition of $N_e = 70\%$ excitatory neurons and $N_i = 30\%$ inhibitory neurons. The connectivity has a Mexican-Hat profile, represented by short-range excitatory synapses and long-range inhibition. The neural and connectivity parameters are shown in Table 6.1.

we want to discover the factors that influence the development of preferred attributes and the locus of acquisition of directional selectivity.

With respect to the representation of movement at the network level, the goal is to explore the possibility of obtaining in a stable self-organizing map, a distributed coding of a set of movement directions. The computational demand for the organizational process is to ensure the generation of distinct directional commands.

The model and the learning scenario proposed here are inspired by the simulation work on the self-organization of visual feature maps (Section 2.2.2, 2.3.1). This procedural transfer is grounded in the general idea that developmental principles described for sensory areas may reflect general mechanisms of cortical computation (Sections 2.2.3, 2.3). If this model proves successful in simulating the formation of directional motor maps, than it provides evidence for the generality of the mechanisms employed, while also helping to develop new hypothesis on the functional principles of the motor cortex.

6.1.1 Structure of the model

The motor system architecture is presented in Figure 6.1. It consists of a two dimensional heterogeneous self-organizing map with N_e excitatory and N_i inhibitory spiking neurons. Each neuron in the output map is fully connected to the input layer and is linked by probabilistic connections with other competitive neurons. The input unit activity is characterized

Model parameters
$max \Theta = 3, \Delta\Theta = 0.2$ - threshold parameters
$\tau_m = 5$ - membrane time constant
$u = 0.8, v = 3.5$ - refractory period parameters
$F_j = 10$ - number of spikes stored / neuron
$d_i = 2$ - input delay value
$\Delta d = 0.3$ - delay spread parameter
$\Delta f = 0.2$ - firing time noise
$N_e = 0.70$ - percent of excitation
$c_e = 10\%, c_i = 12\%$ - rates of connectivities
$\Delta c = 15\%$ - randomness on connections number
$\sigma_e = 4, \sigma_i = 3.5$ - connectivity decay parameters
$max_w_a = 0.5, max_w_l = 0.4$ - afferent and lateral weights parameters
$\Delta w = 0.1$ - weight spread parameter

Table 6.1: Neural and connectivity parameters for the self-organization of the motor directional map. The neural model is described by Equations 4.2, 4.3, 4.5 in Section 4.3. Detailed explanations of the connectivity parameters have been given in Section 5.3.2.

by a simple identity transfer function. The activation function of the output neurons is given by the membrane potential Equations 4.2 in SRM₀ model (Section 4.3). The neural response of an output unit evolves over time by combining the afferent signals with lateral excitatory and inhibitory feedback. The parameters used for setting the neural model are shown in Table 6.1 and their meanings have been described in Sections 4.3 and 5.3.2.

The self-organizing feature map (SOFM) has proven highly effective in modeling the formation of orientation and directional maps in visual cortex (Sections 2.2, 2.3). The basic principle of an SOFM is represented by the topology-preserving mapping of the input space to clusters of nodes in the output map. Learning of the afferent weights is assisted by a lateral feedback system, which leads to a cooperative organization of afferent and lateral connections (for the role of lateral connections in the development of cortical maps see also Section 2.2.4). In a network with a center-surround connectivity pattern (i.e., Mexican-Hat profile) the repeated exchange of localized excitation and long-range inhibition has an essential role in focusing the network activity, by enhancing the center and suppressing the activity at a remove from it.

It is also important to point at the sparseness of the connections between the competitive neurons. This is an important property of the network architecture, which is in agreement

with biological data showing that in the cerebral cortex each neuron is coupled to a reduced number of other neurons (Braitenberg and Schuz, 1998). In our network, the probability of a connection existing between two excitatory neurons i and j decays with the increase in the Euclidean distance $d(i, j)$. Conversely, for two inhibitory neurons the probability of connection increases with the distance between them. How fast the connectivity decays is given by the parameters σ_e and σ_i . Details on the procedure used to create the probabilistic connectivity pattern have been presented in Section 5.3.2. For this simulation, parameters are tuned such as each excitatory neuron connects probabilistically with neighbors within distance $d = 4$, whilst an inhibitory neuron acts upon a fraction of cells placed from distance $d = 4$ to the margins of the network.

6.1.2 Learning procedure

Normally, learning in a self-organizing map consists of three steps. First, the activation of the hidden layer is computed with respect to the input vector. Following this, a best-matching unit is selected and designated as *winner* of the learning step. Finally, the weight vectors are adapted for all units in the winner neighborhood. One way to modify the weights is as a function of the cells discharge rates, after the network has reached a stable firing state (Section 4.4). When computing with spiking neurons, a more appropriate alternative is to apply learning as a function of timing of single firing events (Section 2.2.5).

In the algorithm proposed here, the winner is randomly selected from the subpopulation of units that fire the quickest in one simulation step. After choosing a winner, learning is applied as follows. The afferent weights of a competitive neuron i are adapted in such a way as to maximize their similarity with the current input pattern j . A measure of the similarity is the difference between the postsynaptic potential ϵ_{ij} that encodes the input stimulus and the connection weight w_{ij} . Furthermore, a spatial and a temporal neighborhood of the winner are created, such that only the neurons inside the N_c area and which have fired up until a reference time T_{out} are subject to learning. The learning rule is adapted from Ruf and Schmitt (1998) and is given by the formula:

$$\Delta w_{ij} = \eta (\epsilon_{ij} - w_{ij}) \frac{T_{out} - T_i}{T_{out}}, \quad \text{any } i \in N_c \quad (6.1)$$




Figure 6.2: Modification of the synaptic weight w_{ij} is applied as a function of the parameters: T_j - time of the presynaptic spike, T_i - time of the postsynaptic spike, d - delay of transmission, and τ_m - time membrane constant.

with

$$\epsilon_{ij} = \begin{cases} \exp(-\frac{s}{\tau_m}), & \text{if } s = T_i - T_j - d \geq 0, \\ 0, & \text{if } s < 0. \end{cases} \quad (6.2)$$

In equations above, η is the learning rate, τ_m the membrane time constant, ϵ_{ij} is the postsynaptic potential from presynaptic unit j , and T_i and T_j are the times of the first spikes of neurons i , respectively j (see Figure 6.2). Note that the temporal neighborhood $T_{out} - T_i$ is larger for the winning neurons and decreases for units that fire later. Hence, the largest change in the afferent weights occurs in the case of the winning neuron.

A similar rule is applied for learning in the lateral weights. The synaptic efficacy of a lateral connection is modified depending on the activity of the connected neurons and upon the arrival time of the presynaptic spike. The weight of an excitatory synapse is updated only if both neurons i and j are in N_e , if they both have fired and if the presynaptic spike has been emitted before the postsynaptic spike. For example, if the postsynaptic neuron i fired at T_i , then for all excitatory presynaptic neurons j that fired before T_i and fall inside a certain spatial neighborhood, the connection strength w_{ij} is potentiated. Note that strengthening of excitatory connections predominates over synaptic weakening. The later occurs only when the weight exceeds the postsynaptic potential value. Weights are also kept from excessive increase through normalization.

The learning algorithm for the inhibitory synapses was adapted after the rules proposed by Levy and Desmond (1985). If a presynaptic activity occurs before the firing of the postsynaptic neuron than we have a loss of the strength of inhibitory synapse with the value given

by Equations 6.1 and 6.2, where $T_i - T_j - d > 0$ and $T_j, T_i \notin 0$. If we have a presynaptic activity paired with a postsynaptic inactivity than we have a potentiation of the inhibitory synapse with

$$\Delta w_{ij} = \eta [-(\epsilon_{ij} - u_i \cdot w_{ij})] \frac{T_{out} - T_j}{T_{out}}, \quad \text{any } j \in N_c \text{ and } T_j > 0 \quad (6.3)$$

where u_i is the activation of the postsynaptic silent node, averaged over the T_{out} period. With the advancement of the simulation time, the neural response becomes more selective, hence more neurons are silent for other directions than the preferred direction. In this case, the rules above tend to favor an excessive increase of the inhibitory synapses, which is avoided through normalization by a fixed N factor.

6.1.3 Training patterns

In the theoretical introduction of this work it was suggested that a significant advance in the modeling of motor cortex development will be made by identifying the input signals that drive the organization process (Section 2.2.3). Hence, the milestone of our simulation work was the formation of a training set, which encodes the directional information needed for learning and moreover, is biologically plausible.

A source of inspiration in creating the input patterns was represented by the study of Lin et al. (1997) on neural trajectory computation (Section 2.3.2). Their work investigates how directional information can be extracted from the firing rates of motor cortex neurons by using a self-organizing feature map (SOFM). The SOFM is used to establish a correspondence between the discharge firing patterns of a number of $n = 81$ neurons computed at different moments in time, and the resulting direction of movement. Besides reading the directional information encoded, the self-organizing map allows a visualization of the similarity relationships between the input vectors. Thus, it indicates that neural discharge vectors for similar directions share a higher degree of similarity than vectors coding for opposite directions. This is a consequence of the fact that for similar directions, neurons fire at approximately the same rates.

Our hypothesis was that input vectors, which code directional information in the way described above, may be used for the self-organization of a directional motor map. The main assumption we made was that input vectors, which represent opposite directions, are highly dissimilar. Conversely, topologically close to each other directions are encoded by vectors

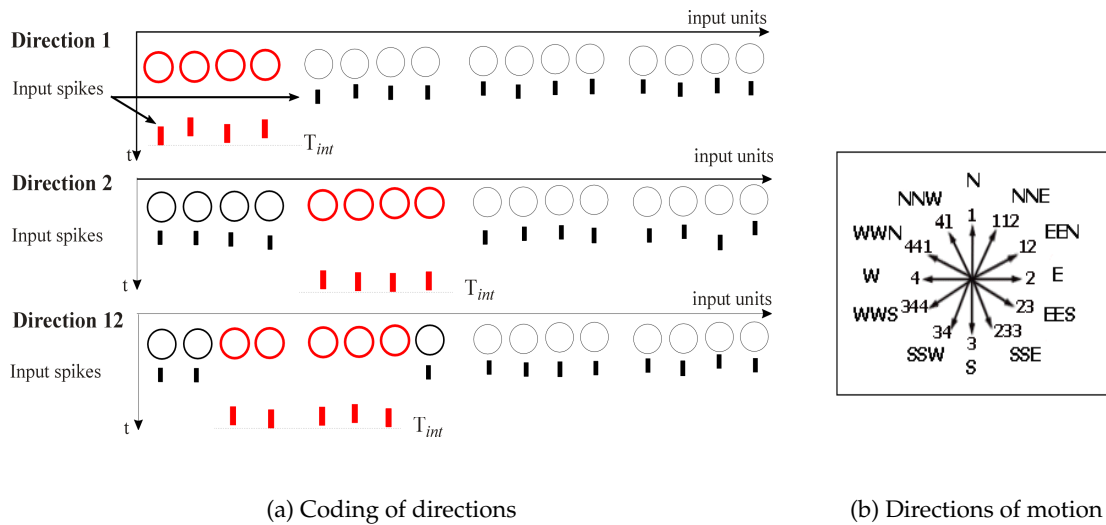
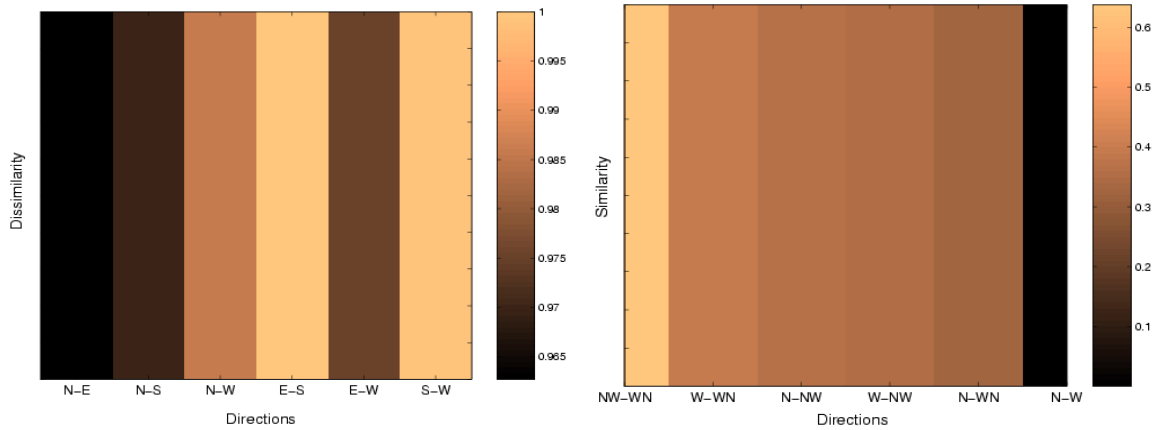


Figure 6.3: (a) Coding of three movement directions by the input vectors. The red units are referred to as principal units, since they code relevant information within a pattern. The rest of the units transmit noisy values. Orthogonal directions such as direction 1 and direction 2, are encoded by different subsets of input units. Direction 12 is topologically located between the two directions. Hence, it is encoded by a subset of units shared with the other two directions' representations. (b) Directions of movement which the self-organizing map learns to stably represent.

with similar values. A recent review of the literature on muscular control of movement has provided compelling support for our hypothesis. It indicated that the training inputs of the type we have used can be provided by proprioceptive signals arriving from the muscles involved in a movement (Theeuwens et al., 1996; Bolhuis et al., 1998). This is, because movements in opposite directions entrain activation of different sets of muscles (i.e., agonist/antagonist) that in turn provide feedback signals with distinct representations. An important future direction of our work is to use experimental data collected from arm muscles during directional movements to train the motor network (see also Section 7.2.2).

For simplicity we have defined only 12 directions of movement as represented in Figure 6.3b. Each direction is encoded by an input vector of dimension 16. The input values represent the firing times of the units (see Section 4.3.4). The values are generated in such a way that directions which form an angle larger than 90° are coded by vectors that are almost orthogonal (i.e., have a low degree of similarity). Examples are: North and East, or NNEast and SSWest. Directions that are situated inside an arc circle of 90° are encoded by input vectors with higher degrees of similarities.



(a) Dissimilarity measure

(b) Similarity between patterns

Figure 6.4: (a). Dissimilarity values computed between input patterns representing four orthogonal directions: North, East, South, West. (b). Similarity values between directions situated within a arc circle of 90^0 : North, North-West, West-North, West. The similarity (or dissimilarity) measure is given by the Euclidian distance between the vectorial representations of the patterns.

The degree of similarity between two vectors is given by the number of units shared in the representation (Figure 6.3a). For example, if direction North (direction 1 in Figure 6.3a) is encoded by the firing times of four input units u_n , than a direction similar with N, such as NNE (direction 12 in figure) will be encoded by a set of units u_k , so that $u_k \wedge u_n = 2 \vee 3$. Conversely, an opposite direction such as South (direction 2 in figure) is encoded by a set of units u_s , with $u_s \wedge u_n = nil$. Note that information on each pattern is distributed over all 16 input units. Thus, a number of four up to five principal units encode the training information provided by the input pattern and the rest of units code noisy values. The similarity and dissimilarity values between four directions of movement are shown in Figure 6.4. The similarity between two patterns is computed by calculating the Euclidian distance between the input values on each of the 16 pairs of neurons in the vectors.

After creating a prototype pattern for each of the 12 directions of movement we generate a training set of 12x20 patterns by adding noise to each value in the original vector by the formula: $u = u_i + R(m)$, where $R(m)$ is a random number generated uniformly in the interval $[0, 0.5]$ for a principal unit, and $[0, 2]$ for the other units. A small training set was sufficient for our simulation and it favored multiple representations in the network of each directional pattern.

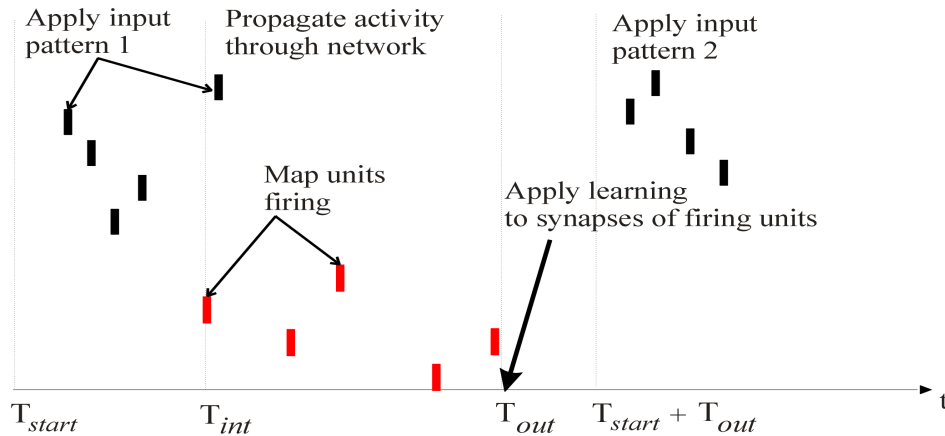


Figure 6.5: Phases of learning in a pulsed neural network: apply a time-coded input pattern starting from T_{start} ; integrate competitive units activity beginning with the time T_{int} ; propagate activity through the output layer until T_{out} ; and update connections weights according with the learning rules. See also Section 5.2 for more details on the implementation of this scenario.

6.1.4 Simulation

The main phases of a simulation with spiking neural networks and time-coded input patterns have been described in detail in Section 5.2. Figure 6.5 summarizes these stages as input pattern application, propagation of the network activity and adaptation of learning synapses. One learning step lasts for $T_{out} = 30$ ms and is divided into the following stages (Figure 6.5). Input patterns are applied starting with the moment $T_{start} = 0$. First integration of network activity is done at $T_{int} = 9$ ms. Network activity is propagated into the output layer and neurons are allowed to fire for approximately 20 ms, until T_{out} is reached. Afterwards, modifiable synapses are adapted according with the learning rule, network activity is reset and the process continues with another input pattern.

For this simulation, the weights of afferent and lateral connections are initialized around the midpoint of the input pattern values (e.g., 0.5). The learning rate starts with a value of 0.5 and is decreased by a fixed percentage $\eta = 0.999\eta$. An initial radius of N_c is set to 5 and is reduced over time to the minimum value of 2. Training was performed for 1000 cycles, equivalent to applying 240,000 directional sequences. After a coarse organization of the network weights occurred, its characteristics were analyzed, as presented below. The network training was discontinued for two reasons. First, learning with spiking neurons is faster than with the traditional, rate-coding model, hence self-organization may occur earlier (see also

Farkaš and Miikkulainen, 1999). Second, activity in a self-organizing map decreases rather fast with the advance of the simulation time. This is due to the increase of selectivity in the neural response. Consequently, the advancement of network training would favor the formation of narrow tuning responses in a small set of neurons, associated with a complete silence in the others.

6.1.5 Results

Self-organization of afferent weights

Training of a pulsed neural network start with neurons being equally responsive to all input patterns. This leads initially, to the activation of the entire network for the presentation of any input pattern. With the advance of simulation time, the number of neurons that respond to a particular input vector slowly decreases, as selectivity of the neural response increases. In the end, in the trained self-organized map, only a small number of patches of activity occur for each directional command. These are usually organized around the winners of the corresponding direction. Hence, as a result of the self-organization process aimed at increasing the similarity between input patterns and weight vectors (Eq. 6.1), the afferent weights of a neuron that wins mostly for a particular direction resemble closely the corresponding input vector.

The labeling of the map is performed by marking each output unit with the inputs for which it is the most representative. Figure 6.6 shows the neurons' preferred directions labeled according with their winning behaviors while a test is run. The trained network was tested with the same pattern set used for learning. This was presented for a number of 10 times, while the winning behavior of each neuron was recorded. If for instance, neuron i wins mostly for a particular input pattern, than it is labeled with the corresponding direction. For the correspondence between directions and numbering labels see Figure 6.3.

The trained network bears a set of features, which have been induced by the self-organizing network property of a topology-preserving mapping. That is, the similarity relationships from the input space are translated into neighborhood relations in the output map. We may say that the network found a way to map directions according to their similarity relations (see similarity measures in Figure 6.4b). It resulted in most similar patterns being mapped to neighboring units. Note the succession of neurons labeled with 1, 12, 2, 23, 3, 34, 4, 41 in




Figure 6.7: The afferent weights of seven neurons in the output map. (a) Initial random weights with an average value of 0.45. (b) As self-organization progresses, the weights organize into smooth hill-shaped profiles. Each weight vector has a peak (up to 0.9 value) for the input units which code the pattern best learned.

in first 4 input units), node 3 is broadly tuned to direction NEast (input units 2 to 6) and last 4 nodes are more sharply tuned to direction East (inputs 4 to 8).

Rate coding of directional information

Figure 6.6 shows the calibration of the output map with respect to the winning behavior of neurons. That is, it reflects the organization of afferent connection weights. Apart from this type of labeling, normally performed for a SOM, we want to characterize the selectivity of each neuron as a function of its discharge rate during a time interval. Thus, for each of the 12 directions we run a validation procedure consisting of the following: while 30 input patterns are applied with a frequency of 60 Hz, the activity of the network is allowed to oscillate and the firing rates of neurons are recorded. The preferred direction of the unit is given by the input, which evokes the highest discharge rate.

The network labeled according with the firing rate behavior is presented in Figure 6.8. Only neurons with tuning levels above 30% are represented and the tuning value (normalized) is indicated by the length of the thick line. Note that compared to Figure 6.6 the number of unselective neurons has decreased to approximately 25% of the total number of neurons. The existence of approximately 15% neurons which are not directionally selective has also been described experimentally in the motor cortex (Georgopoulos et al., 1984).




Figure 6.8: Self-organizing map labeled with the neurons' preferred directions, computed as the movement direction for which a neuron's discharge rate is highest. The level of neural tuning (normalized) is represented by the thick line. Two types of neural behavior are shown: *winning* neurons (in black) and *lateral* neurons (in blue). See the text for details on the functions of these neural behaviors. In the left side, an assembly of neurons is delimited, composed of winning and lateral neurons whose collective firing encodes the movement direction towards North. The gray arrows indicate the excitatory connections from the neurons which fire first (winners) to the neurons which fire later (lateral units).

The increase in the number of selective neurons compared to the previous calibration suggests that neurons other than winners have also developed directional preference. Thus, the analysis of the discharge rate behaviors revealed a set of neurons (drawn in blue in Figure 6.8) that had remained invisible during the first calibration of the map. These neurons have small values on the afferent weights and never won during the testing procedure for any of the directions involved. However, when recording their firing behavior in time it occurs that they spike later than the winners and their discharge rates are tuned to the direction of movement.

We refer to the latter units as *lateral* neurons, due to the fact that their spiking activity is mainly caused by the integration of the lateral excitation in addition to the afferent input. The lateral synapses from the winners (which fire quickest) to the later firing neurons are illustrated for a small assembly of neurons in the left upper corner of Figure 6.8. Because

neither the lateral nor the afferent stimulation solely can cause firing, the lateral neurons need to integrate both sources of excitation before they emit a spike. This constraint leads to a more selective directional response. Their tuning curves are less broadly than those of the winner neurons, with a mean width of 30° . Later in this chapter, a third category of neurons is described that need to integrate three types of inputs in order to spike.

The joint activation of the two types of neurons give rise to a sort of collaborative cell assembly (Reilly, 2001), which enhances the strength of excitation between neurons tuned to similar directions of movement and suppresses the response in opposite directions. On the left side of Figure 6.8 the approximate boundaries of the cell assembly coding for movement direction North are indicated. The most important effect of collaborative cells assembly formation is the emergence of a population coding (see Section 6.1.5). Both the emergence of collaboration between neurons and the formation of a population code are supported by the horizontal feedback system. That is why, the organization of the lateral weights deserves a particular analysis.

Self-organization of lateral weights

The formation of the directional map shown in Figure 6.8 was driven by two principles. First, the SOM property to preserve the topological relations between input patterns led to the formation of an ordered representation of the directional space. Second, the lateral connectivity pattern with a Mexican-hat profile played an essential role in shaping the network activity. Thus, it favored the formation of distributed, stable representations for each direction of movement.

The lateral connection strengths are not static, but they evolve together with the afferents. For each firing neuron the lateral weights are adapted by a Hebbian learning rule (Equation 6.1) according to how well the neuron's activity correlates with the activities of other neurons. As the afferent weights organize into a smooth profile (see Figure 6.7), lateral correlations between neurons fall off with distance and become stronger only between neighbors with similar directional selectivity. Thus, the excitatory weights between neurons which develop preferences to opposite directions decrease through normalization and end up being pruned. Conversely, inhibition (decorrelation) increases between neurons that become selective to orthogonal input patterns.

To extract the common features of the lateral synaptic strengths we calculated the mean




Figure 6.9: The dependence of the mean value of the lateral connection strength on the angle between the preferred directions of the neurons in a pair. The mean value of the weight was calculated by averaging over all weights, which connect neurons, whose preferred directions form the same angle.

value of the connection weight as a function of the angle between the preferred directions of two neurons in a pair. Figure 6.9 shows the mean values calculated considering all horizontal synapses in the network. We found that the mean connection strength is negatively correlated with the difference between preferred directions of the neurons in a pair. Our findings are in good agreement with the experimental estimation of lateral synaptic weights between motor neurons, as being a function of the difference between preferred directions (see Section 2.2.3).

Different profiles of afferent and horizontal strength vectors develop in winners as opposed to lateral neurons. One has to remember that during learning the change in all synaptic weights is constantly scaled by the difference between the time of the neuron's first spike and a fixed arbitrary time out (see Equation 6.1). Hence, the later the lateral neurons fire, the smaller the adaptation of their afferent weights. Instead, their firing is triggered by the receipt of a lateral spike, therefore, the more the lateral connections will be strengthened. Conversely, because in winner units the first spike is caused by the afferent stimulation, they develop an opposite profile to the weight vector, characterized by strong afferent connections and small values of the lateral weights.

The strength of inhibitory connections also evolves over time. Inhibition increases whenever the firing of a presynaptic inhibitory neuron is correlated with postsynaptic inactivity (see

Section 6.1.2). Because inhibitory synapses run over large distances, the firing of an assembly of neurons for one direction of movement will result in lateral inhibition spread all over the network. Consequently, neurons belonging to assemblies that are not involved in coding the current direction receive a high proportion of inhibition. Because lateral neurons rely mainly on lateral excitation to fire, they end up by remaining silent. Instead, winner neuron activity is due to afferent stimulation, hence they fire (even accidentally) in more cases and for more directions than the lateral neurons. Therefore, the responses of lateral neurons are more suppressed (inhibited) than those of the winners. Consequently, it enhances their narrow tuning to the input stimuli (see Section 6.1.5).

It is important to explore with future models the influence on the lateral connection organization of variants of learning procedure and of different rates of connectivity. For instance, a learning algorithm based on the timings of a spike series would allow the winner neurons to integrate the effects of lateral feedback. This would cause a significant strengthening of the entire lateral connectivity system. The increase in the number of lateral neurons can be obtained by implementing higher connectivity rates, for instance excitatory neurons fully connected with neighbors within a certain area. A strong recurrent feedback is essential for the generation of neural oscillations in the network that have a role in maintaining the movement representation in time (see also Section 7.1.3).

Population coding

To this point, results have been presented in terms of selectivity of individual cells to the direction of movement. With few exceptions, our findings indicate that neurons (both winners and laterals) are broadly tuned to several directions of movement. Moreover, the command in each movement direction elicits in the output map the activity of an entire population of neurons (i.e., the collaborative cell assembly). These observations suggest two ideas. Firstly, directional information cannot be read out from the individual response of the majority of neurons (e.g., 80%). Secondly, it is possible that directional information is represented in the collective activity of subsets of neurons broadly tuned to similar directions. The later hypothesis is illustrated in Figure 6.10(a) where the joint activity of neurons with various preferences, coding one direction of movement, is shown. That is, more than 40 neurons are firing at different rates while representing direction North.

In order to test the population coding hypothesis we computed for each of the 12 direc-

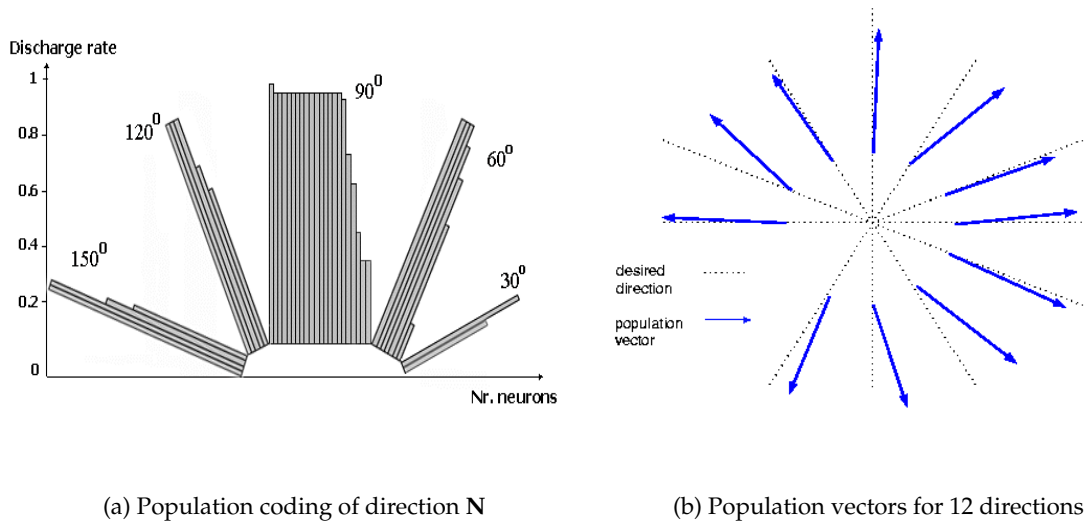


Figure 6.10: (a). Discharge rates (normalized) of the population of neurons coding direction North. Each bar represents the contribution of single neurons and its orientation corresponds to the preferred direction of the neuron. (b) Population vectors for 12 directions of movement. Each vector is a resultant of individual neuron contributions (only the neural population for direction N is shown).

tions of movement the neural vectors, by using the formula (adapted after Lukashin and Georgopoulos, 1994):

$$P_x(t) = \sum V_i(t) \cdot \cos PD_i, \quad P_y(t) = \sum V_i(t) \cdot \sin PD_i. \quad (6.4)$$

where V_i is the discharge rate of the neuron i during a testing interval of 600 ms. In Equation 6.4 we consider that the preferred direction of the neuron PD_i is given by its discharge behavior, as described in Section 6.1.5 (see also Figure 6.8). The populations vectors yielded by the vectorial summation $P_x + P_y$ are shown in Figure 6.10(b). Note that the neural vectors resulted point very closely to the desired directions of movement.

The population code is an effect of the distributed representation of direction, caused by the broad selectivity of output map cells. Each cell's activity is highest for a movement in a particular direction and decreases with movements further away from that direction. Figure 6.11 shows the discharge rates (normalized) of four neurons for different directions of movement. In the absence of a directional continuum, we cannot present an exact estimation of the tuning curves. However, we can approximate the shape and the width of the neural response. Neurons have unimodal tuning curves that differ from cells with the width of the curve of 120° (neuron 3 responds to maximum 5 directions of movement) to cells with a

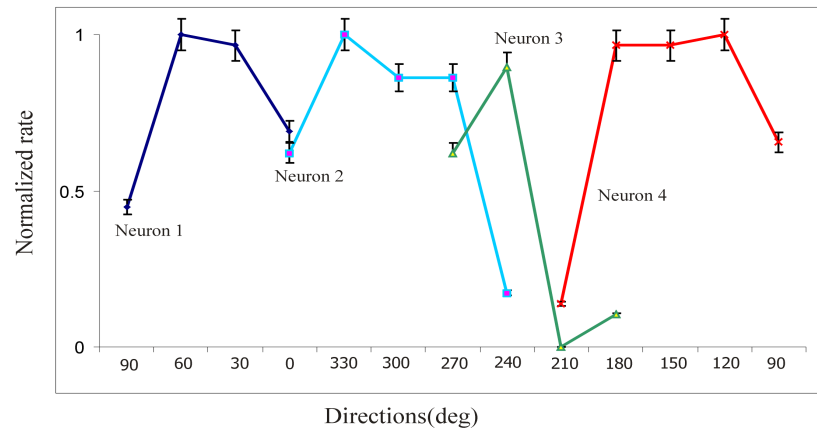


Figure 6.11: Discharge rates (normalized) of four directional selective neurons plotted as a function of the movement directions. The preferred directions of neurons are located at approximately 90° from each other. Note that neurons show different shapes and widths of the tuning curves.

curve width of 30° (not shown in figure). The median of the curve width is at 50° . A given neuron participates in coding of movements in directions which form an angle no larger than 90° with its preferred direction. Since directions forming an angle larger than 90° are encoded by orthogonal input vectors, the participation of the neuron in coding movement in such directions is due mainly to the successive propagation of excitation in the lateral feedback system.

The width of the neural tuning found in our model is in partial agreement with the experimental findings of Amirikian and Georgopoulos (2000), which describe real neurons' tuning curves with a half-width from 30 to 90° and the median 56° . Their results indicate that motor cortical cells are more sharply tuned than previously thought (i.e., do not fit the cosine function). Our simulation findings point in the same direction (note that in our case the average of the midpoint of the curve is at 25°), but more modeling work is needed for a better fit of the results to the experimental data.

To this end, we can try to put everything together into a hypothetical scenario of how the map response evolves in time from the moment it is presented with a move command. First, winner neurons are activated very rapidly. By doing this, the first spikes occurring in the network offer a quick motor output available to initiate the movement (e.g, if this is desired). But, since winners' tuning to the intended direction of movement is rather broad, this initial specification of the direction is very rough. In our case it restricts movement in a sector of

30° around the desired direction. One can also compute the instantaneous population code contributed by the activated winners, but this information is not more accurate than the readout of single spike timings from a few sharply tuned winning neurons.

Following the winners' activation, the lateral neurons get involved due to the horizontal excitation spread. During successive propagations of activation in the neural assembly, the lateral neurons contribute to creating a steady state for the network activity. This state represents the attractor of the desired direction and has a completely distinct position in the space for each of the learned directions. At these immediately successive stages of movement initiation (or movement preparation) the directional information can be read out correctly to drive the movement accurately to the target.

6.1.6 Discussion

In this section, a neural network model was proposed for the emergence of directional selectivity in motor cortex, based on acquired experience (i.e., proprioceptive feedback). The self-organized network exhibits properties that are consistent with the experimental findings on biological motor control. Furthermore, we believe that the modeling process can provide us with valuable knowledge about the organization and developmental principles of the motor cortex. The main conclusions are summarized as follows.

First, the model shows spontaneous emergence of a feature map during unsupervised learning and starting from random afferent and lateral connections. The self-organization of the network leads to a stable, ordered representation of 12 directions of movement. Most of the neurons in the self-organized network develop directional selectivity. All directionally selective neurons have a preferred direction, computed as the movement direction for which the neuron's discharge rate is highest. Furthermore, directionally selective units can be divided into two categories. Winners are units whose afferent weights are tuned to the input vectors and this is reflected in a rapid response to the movement command. Lateral neurons are units whose firing is determined by the integration of lateral excitation with afferent stimulation.

With respect to winner neurons' directional tuning, we found that the vast majority of them are broadly tuned to several directions of movement. However, in the case of a small percent of winner neurons (about 15% of all units) directional information is accurately encoded in the timing of the first spike. It is noteworthy, that the original experiments of Georgopoulos

et al. (1982) also described the existence in the motor area of fast responding neurons, which are significantly tuned to one direction of movement. It is possible that the central nervous system uses this directional information for the initiation of fast movements, in a similar way the visual brain uses single spikes to take rapid decisions (Thorpe et al., 2001; see further discussion in Section 7.1.1).

To summarize these findings, we believe that the main outcome of our model is the emergence of directional selectivity in a motor self-organizing map of spiking neurons. Even if our model is a very simplified image of cortical motor computations and needs further refinements, we believe it represents a progress in the simulation of motor cortex organization. It was motivated by the absence in the motor control field of computational studies of how motor directional selectivity emerges (see Section 2.3).

The results obtained emphasize the advantages of modeling the self-organization process, as opposed to assigning neurons with preferred attributes (Lukashin and Georgopoulos, 1994). Thus, in our model, neurons can effectively 'gain' a preferred direction, as a function of the neural parameters, noise and location within the network. It shows that preferred attributes are not developed solely based on the neural features but also as function of its interactions with the network. Conversely, directional knowledge is encoded at the network level, that is in the horizontal connectivity, as well as in single neurons responses. We believe that by modeling the developmental processes, we can help to reveal the unknown functional principles of motor cortex organization.

An important issue concerns the function of the lateral connectivity system in the formation of the directional map. The role of horizontal connections in the organization of sensory feature maps have been emphasized recently by many studies on the visual cortex (see Section 2.2.2). However, much less is known about the profile and the function of lateral connectivity in the motor cortex. Our results indicate that neurons that develop similar directional selectivity become functionally correlated. In the trained map it was found that the strength of connection between neurons in a pair becomes negatively correlated with the difference between their preferred directions. This is an important modeling finding, which is in complete agreement with experimental estimations and previous simulation work (Georgopoulos et al., 1993; Lukashin and Georgopoulos, 1994. See further discussion in Section 7.1.2).

As a result of the distributed representation in the network of the directional information, we can observe the emergence in the model of a population coding. This result is even

more interesting, as it was obtained with a self-organization process based on the timing of single spike events. Note that during training any information regarding the discharge rate behavior of neurons was discarded. Furthermore, by its nature the self-organization mapping focuses upon learning winner neurons. Hence, the emergence of a distributed representation of direction is determined by the organization of afferent and lateral connections. Firstly, as the afferent weights of a winner become more similar to an input vector, the unit increases its sensitivity not only to the best matching input, but also to all similar inputs (see the similarity measure in Figure 6.4). Secondly, the network response to one direction is amplified by the lateral feedback system, which activates the lateral neurons that are not directly responsive to the input stimulation. We may conclude that the population coding is a function of both the properties of the input space and the built-in network's constraints which allow the formation of collaborative cell assemblies (see also discussion in Section 7.1.1).

In our population codes, the individual contributions of single neurons can be summed up using the population vector scheme (Lukashin and Georgopoulos, 1994). Even if the neural vectors slightly deviate from the intended direction of movement, the main achievement of the model is that it is able to provide distinct commands for each direction (Figure 6.10). This means that the self-organizing map managed, within a distributed representation, to form stable attractors for each of the directions involved. Future work is aimed at tuning the neural parameters and the network design constraints in order to increase the accuracy of where the population vector points.

6.2 Development of visuomotor alignment of directional codes

A fundamental operation of animal brains and robot controllers is the integration of visual information with motor commands. Visually guided reaching in primates is considered to involve a series of neuronal events that transform retinal information about target location into the metrics of arm movement (Section 3.1). Most traditional modeling efforts using neural networks for robotic motor control have focused upon developing various formalisms capable of doing the coordinate system transformation (Section 3.1.1). Nevertheless, these implementations are neither relevant to the biological motor control, nor are they very adaptive or capable of learning and development (see Kalaska, 1995). Conversely, it was pointed out that inspiration could be drawn from the way motor control is realized in living systems.

Thus, recent behavioral and electrophysiological findings gave rise to a conceptual framework for the understanding, and possibly implementing, the computational mechanisms underpinning visuomotor coordination (Section 3.2). The work presented is in agreement with these latter theories.

In this section, an artificial neural network model is proposed to address a basic issue of visuomotor coordination: what are the computational mechanisms that allow visual information on the direction of movement to evoke an appropriate motor response in the same direction? We believe that this operation meets the basic computational demand for visuomotor mapping and represents a building block of the perception–action cycle (see Section 3.1.2).

Our approach to the visuomotor coordination problem was inspired by experimental data showing that neural selectivity to movement direction is a ubiquitous feature of the parieto–frontal networks involved in reaching (Section 2.1.3). During visual analysis of movement and at all stages of motor control the activity of a substantial percentage of movement–related neurons depends upon the direction of movement (i.e., involves direction selective cells) (Kandel et al., 2000; see also Section 2.1.2). The idea was to use the *motion selective cell* as the basic information-processing element from which neural networks capable of visuomotor control are built. Moreover, it is currently believed that visuomotor mapping of information (including coordinate system transformation) is realized progressively, by the gradual involvement of several populations of neurons (Section 3.2). We propose that the correlated activity of motion selective cells in the parieto–frontal areas meets the basic computational demand in the chain of operations required for visually guidance of movement.

The learning algorithm proposed here consists of a sequence of processing steps. First, the visual perception of a moving stimulus is translated into the firing pattern of directionally selective visual neurons. Next, this activity is transmitted to the motor areas that are involved in the control of movement direction. Simply, by the correlated firing of visual and motor neurons with similar neural response, a communication pathway for the directional information is established. In our simulation, the visual system corresponds to some early stages of visual information processing (e.g., retina and primary visual cortex). We believe that an equivalent computational mechanism can operate for the coupling of two or more networks of directionally selective cells located at any stages along the parieto–frontal networks involved in reaching.




Figure 6.12: Neural architecture for the alignment of visual and motor representations of direction. The left part represents the visual system, consisting of a retina and a cortical visual layer. On the right is shown the spiking self-organized motor map trained to encode 12 directions of movement. The visual cortical layer is connected to the motor map through a synaptic pathway subject to spike-timing dependent plasticity.

6.2.1 Architecture of the model

In order to simulate the computational mechanism described above, we have implemented a very simple architecture consisting of a visual system connected to a motor map. There are two input layers and two 'cortical' layers: the *motor area* and the *visual area*. The visual area is represented by a layer of directionally selective neurons which have receptive fields from an input, retina-like layer (see description in Section 6.2.2 below). The motor area is a self-organizing feature map trained to represent 12 directions of movement. The organization process has been presented in Section 6.1. All parameters describing the spiking neural model (i.e., postsynaptic potential, refractory period, transmission delays) are kept constant from previous simulation (Section 6.1, Table 6.1).

Each cortical neuron receives afferent connections from one input layer and lateral connections from other neurons in the same area. The visual cortical layer is fully connected to the motor map. The implementation of full cortico-cortical connectivity, rather than a coarse topographic ordering is motivated by the topology and the small scale of the networks. Due to the reduced dimension of the networks one can consider them as being tiny patches of cortical areas, large enough to comprise a complete representation of all features of input space (i.e., directions) at several positions. Full connectivity is required at this level to ensure that any population of motor neurons has access to all visual directions of movement.

In our view, the development of visuomotor coordination is achieved in two phases. During an initial organizational stage, the motor and visual systems develop independently, neural selectivity to the directional information. This stage ends up with the formation of two cortical maps, which stably encode direction of movement. Next developmental stage corresponds to the effective coupling of the two systems, which are allowed to interact and to learn, and eventually to coordinate. Experimental data concerning infants visual development indicate a prior emergence of cortical orientation mechanisms, followed at three month of age by the development of selectivity to motion (Atkinson, 2000). The later enables the cortical control of eye and hand movements. From this age until the age of 5 – 6 months the cortical mechanisms for reaching (i.e., for eye-hand coordination) are developed. These observations indicate a possible sequencing of the processes involved in the perception and control of motion and in the achievement of visuomotor coordination.

The separation of the developmental stages is usually implemented by models of visuomotor coordination, but most commonly the directional maps are obtained by assigning (rather than developing) preferred directions to the motor (and/or visual) neurons (see Bullock and Grossberg, 1993; Salinas and Abbott, 1995). An original feature of our visuomotor mapping model is that it is based on a motor network, which has self-organized to represent directionality. Hence, it allows natural interactions between neurons and a self-developed way to encode the direction of movement.

In the next subsections, we focus upon describing the visual coding of movement direction and we proceed to the modeling of visuomotor coordination development. The reader should refer to the description of the motor map organization from the Section 6.1.

6.2.2 Visual coding of movement direction

After several decades of research on visual processing of information, it is now well established that the cortical analysis of visual space relies on the functioning of a fundamental neural machinery referred to as the *hypercolumn* (Section 2.2.2). A hypercolumn represents a set of columns which are responsive to lines of all orientations from a particular region in space via both eyes, and to movements in directions orthogonal to the orientation axes (Kandel et al., 2000). Experimental results have described a precise organization with an orderly shift in axis of orientation (or direction) from one column to the next (Section 2.2.2).

The formation of orientation and directional selectivity in the visual cortex has been exten-

sively modeled using the self-organization paradigm (Kohonen, 1984). Most commonly, directional selectivity is developed in a map of cells with orientation preferences (Section 2.3.1). When the self-organized visual map is presented through the retina with a moving stimulus of a certain orientation, its response is represented by the firing of a certain set of neurons. These are the cells in the network which receive activation from retinal units located along the trajectory of the moving stimulus and whose preferred orientations (i.e., directions) are similar to the stimulus attributes. For an illustration of this process the reader can refer to a demonstration of how a moving oriented line is processed in a visual feature map (Bednar and Miikkulainen, in press; http://www.cs.utexas.edu/users/jbednar/sweeping_small.html).

Based on these findings, our visual area was implemented with a built-in, basic capacity of signaling motion direction. We have simulated a hypercolumnar organization based on directionally tuned neurons. One hypercolumn consists of 4x4 neurons, each firing for only one preferred direction and being silent for movements in different directions. Excitatory lateral synapses connect each neuron with the first order neighbors that have the same preferred direction.

The cortical map of area 18x18 is connected to a retinal layer formed from 6x6 neurons (left part of Figure 6.12). Each cortical neuron receives input from a fixed-size receptive field of dimension 3x3 units. The receptive field of one neuron (i, j) was centered at $(i/3, j/3)$ and the afferent weights were initialized within the range $[0.5, 1]$. The magnification factor, which is the ratio of cortical neurons per retinal neurons, is 9. In this way, it allows different receptive fields for each location in the retina (see description of implementation in Section 5.3.2). The lateral weights were initialized within the range $[0.3, 0.4]$. Parameters of the visual model were inspired by the simulations of visual cortex organization performed by Miikkulainen and colleagues (1996, 1998).

The scale of the hypercolumn dimension and the range of lateral connectivity are clearly a severe simplification. The built-in capacity for signaling motion was also implemented for simplicity reasons, as a substitute for the self-organization of the visual map. It was motivated by our intention to focus rather on the motor network organization and the visuomotor mapping process, than on the organization of a visual directional map. The main reason was that the formation of visual cortical maps has been extensively modeled compared with that of the motor cortex organization (see Sections 2.3). Nevertheless, this simple mechanism allows us to investigate the problem of directional information mapping, while retaining the principal characteristics of visual feature maps (i.e., narrowly tuned directional

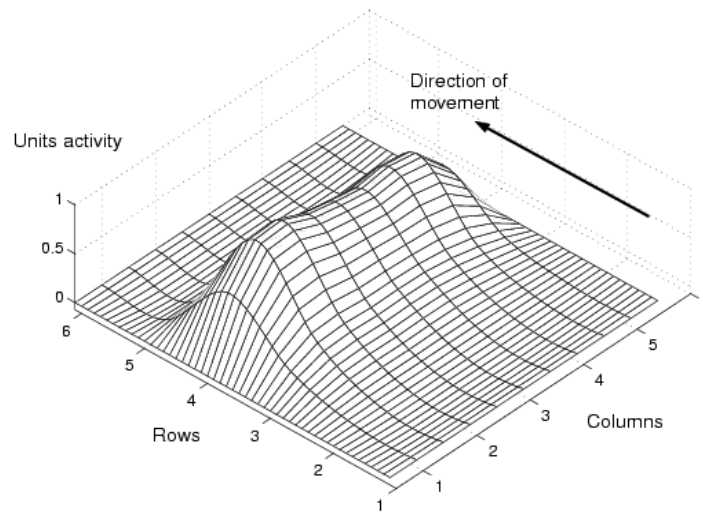


Figure 6.13: The visual input to the model consists of moving bars of a fixed length. There are 12 input patterns, corresponding to bars moving in the 12 directions of motion that can be controlled by the motor network. Movement of a bar through the retina elicits the activity of the cells located along the direction of motion.

neurons, lateral connections, receptive fields).

The retinal input to the system consists of directionally oriented moving bars of a fixed length (see Figure 6.13). A set of 12 input patterns was created corresponding to the 12 directions of movement controlled by the motor network. A bar moving in a certain direction in the retina determines the retinal neurons located along its pathway, to spike at a fixed continuous rate. The retinal activation reaches the visual layer after an average delay of 60 ms. This value was implemented in accordance with the delay existent in the brain between retina and primary visual cortex (Thorpe and Gautrais, 1997).

6.2.3 Learning procedure

The learning scenario was inspired by the sensorimotor stages that an infant progresses through to develop eye–hand coordination. When executing movements during the early motor–babbling period, infants perceive and learn contingencies between the motor activity and the visual image of their actions. By doing this, the simultaneous moving and tracking of the arm can provide the behavioral context for the development of visuomotor coordination. This becomes functional from about four months onwards, when infants start to make visually directed arm movements (Atkinson, 2000).

The behavioral framework described above is simulated by using the paradigm of direct inverse modeling (Kuperstein, 1988), also known as the *motor babbling* approach. This has been discussed in detail in Section 3.3. The general algorithm consists of three steps: (1) a movement in a random direction is generated, (2) the visual image of the arm moving is recorded, and (3) the system learns the correlation between the motor command and the visual effect of the movement. Most of the previous models of visually guided arm reaching have been based on different forms of error-correction mechanisms (Ritter et al., 1989; Fiala, 1995; Jordan, 1996). Only recently, it was pointed out that unsupervised learning means can also lead to an accurate mapping of sensory information (Salinas and Abbott, 1995; Burnod et al., 1999; see also Section 3.3).

Prior modeling work using correlation-activity associations was based on rate coding neurons, where learning is applied as a function of neural discharge rates (Salinas and Abbott, 1995; Baraduc et al., 1999). In this thesis, the focus is upon learning and computation with spiking neurons, based on the timing of single firing events (Section 4.2.3, see also Section 6.1). We believe this represents a simpler and more adequate framework for the study of visuomotor mapping through unsupervised means. That is, because detecting temporal correlations between populations of neurons is a natural computation with spiking neurons which has no equivalent in the rate coding networks.

In our algorithm, one learning step consists of a sequence of actions as follows (see Figure 6.14):

1. The motor input units generate a command for movement in a random direction d . This input elicits, in the motor map, the activation of the neural assembly involved in the representation of the direction d . The activity in the motor cortical population is maintained for a time interval of 400 ms, by the activation of the input command with a constant firing rate of 30 Hz.
2. Following the movement onset (i.e., at $D1 = 100$ ms after the motor command is issued) the retina is presented with a bar moving in direction d . This motion persists until the end of the learning cycle. After an average delay of $D2 = 60$ ms, the retinal information reaches the visual network, where the motion selective cells fire and signal the direction.
3. The cortical visuomotor circuit transmits the neural activity from the visual to the motor area with a delay of $D3 = 50$ ms.
4. During a time window of approximately 200 ms while both networks are active,

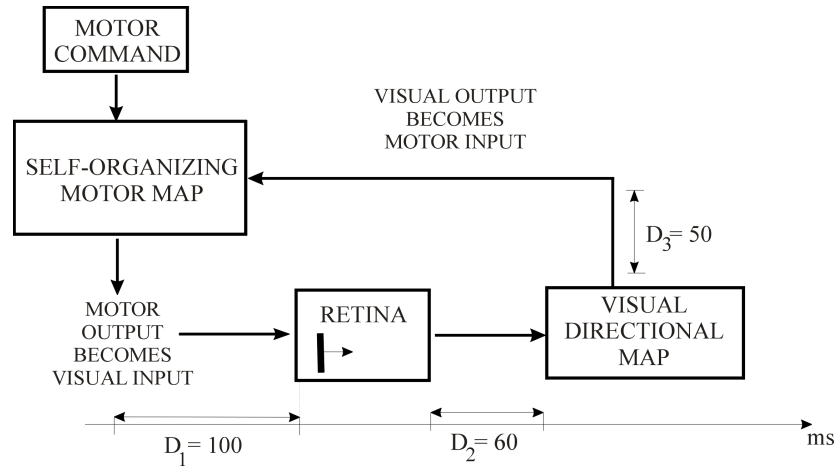


Figure 6.14: Schema of the training procedure used for the alignment of visual and motor directional representations. Note that the motor output becomes visual input and the visual activation is fed back into the motor network. The delay $D_1 = 100$ ms corresponds to the interval between the time moment when the motor command was issued and the eye starts seeing the arm movement. The delay $D_2 = 60$ ms represents the time required by the visual input to get from the retina to the visual network. The third value corresponds to the transfer delay from visual area to the motor area of $D_3 = 50$ ms.

spike-timing dependent learning is applied to the inter-cortical connections. The connection between a visual presynaptic neuron j and a motor postsynaptic neuron i is increased as a function of the time difference between the arrival of the postsynaptic potential t_j and the firing moment of postsynaptic neuron t_i :

$$\Delta w_{ij} = \begin{cases} \eta \left(\sum_{t_j} \epsilon_{ij} - w_{ij} \right), & \text{if } (\exists) \epsilon_j \neq 0 \\ 0, & \text{otherwise} \end{cases} \quad (6.5)$$

with

$$\epsilon_{ij} = \begin{cases} \exp\left(-\frac{t_i - t_j}{\tau_m}\right), & \text{if } t_i - t_j \geq 0. \\ 0, & \text{if } t_i - t_j < 0. \end{cases} \quad (6.6)$$

In Equation 6.5, η represents the learning rate. The sum is calculated over the decayed values of postsynaptic potentials ϵ_{ij} coming from the visual neuron j . Each potential is computed by using the Equation 6.6, at the moment t_i when neuron i fires. The summation represents one way to ensure that the last presynaptic spike that has reached the neuron i is considered for learning. Difficulty arises from the large delay ($D = 50$ ms) between

the visual presynaptic and the motor postsynaptic neurons, correlated with the exponential decay of the ϵ values.

Furthermore, Equation 6.5 states that no modification of the synaptic weight w_{ij} is performed for time steps prior to the arrival of at least one visual spike to the motor neuron. That is, because about half of the time in a learning step, the motor neurons do not receive any activity coming from the visual network, due to the cumulated delay value $D1 + D2 + D3$. Hence, by applying anti-Hebbian learning a rapid pruning of all connections would occur. Our simplification is meant to speed up the learning process. More realistic scenarios that lead to similar results can be implemented, with the condition that the overall time when the visual synapses are active is much larger than the initial time interval when they are silent. In our case, a reduction of the strength of connection occurs only if the weight value increases over the value of presynaptic spikes effects.

Learning with the above algorithm happens quickly, mainly because the inter-cortical connection weights are initialized with low values between $[0, 0.1]$ and strengthening of synapses dominates over synaptic weakening. The learning rate η was set to 0.1. After 500 cycles a pattern of strong connections develops from the visual to the motor area, causing the alignment of visual and motor neural representations in such a way that permits visual information to drive motor movement.

6.2.4 Results

To test the learning of visuomotor mapping of direction, the retina is presented with a bar moving in a constant direction for a time period of 400 ms, while the activity evoked in the motor map is recorded. For each of the 12 possible direction of motion, this motor activity is analyzed with respect to the shape of the firing patterns elicited (i.e., neurons involved) and to the population vector resulting. Figure 6.15a shows the discharge rates of the motor neurons activated by the visual information coding the movement of a bar in direction North. An initial qualitative analysis reveals that the firing patterns elicited in the motor network in the visual condition (i.e., when the motor map is exclusively driven by the visual stimuli) and in the motor condition (i.e., when a movement command is issued, see Figure 6.10a) are very similar.

Moreover, if we compute the population vectors (PV) in the visual condition, we find that the motor PVs driven by visual information point very closely to the desired direction of

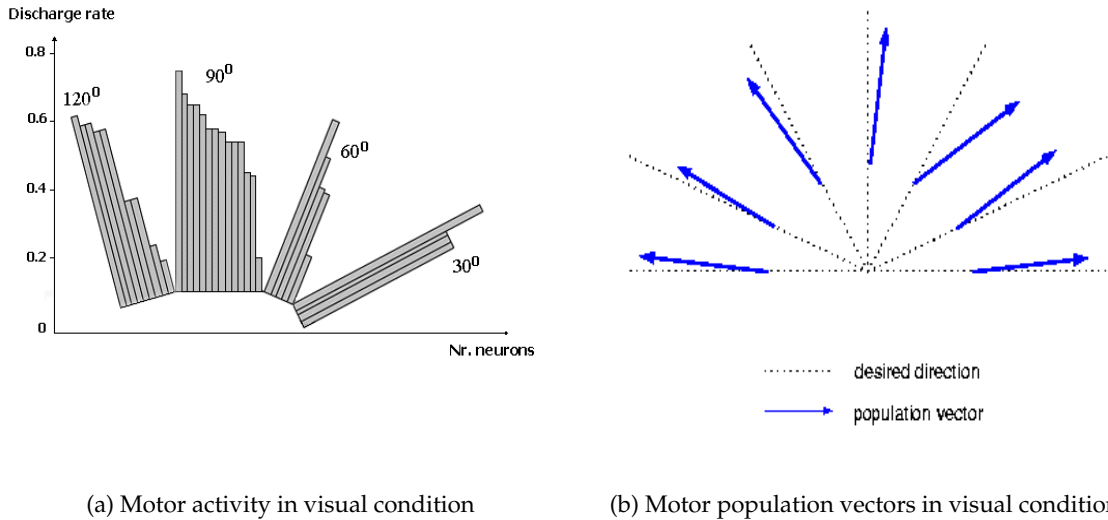


Figure 6.15: (a) Population activity occurring in the motor network in visual condition, while a movement in direction North is perceived. The network activity elicited by the visual input resembles very closely the population activity which controls movement in the same direction (see Figure 6.10a). (b) The motor population vectors computed during the visual condition for 7 directions of movement, by using Equation 6.4.

movement (Figure 6.15b). This means that the visuomotor system learned how to map correctly the visual information into directional movement. Alignment of the activity in the two networks is due to the selective strengthening of the cortico–cortical connections associating visual and motor neural populations with similar directional selectivity. In the following section, we try to explain how the alignment of the maps has occurred and what are the characteristics of the inter–cortical connectivity which allow it.

Organization of inter–cortical connectivity

The reader is reminded that one of the main results of the simulation on motor cortex self–organization was that the strength of lateral connections is negatively correlated with the difference between neurons’ preferred directions. With respect to the visuomotor model, we want to perform a similar analysis of the cortico–cortical connections. Hence, we compute the cumulated weight value for all synapses between neurons whose preferred directions form the same angle. Results are shown in Figure 6.16. To illustrate the result the cumulated (normalized) value was used, instead of the average measure, because of the large number of weights with very small values (< 0.1) resulting from initialization. However,




Figure 6.16: Dependence of the cumulated (normalized) weight value on the angle between preferred directions of neurons. Note that maximal connectivity strengths correspond to an angle between preferred directions of 30° .

both measures indicate the same finding: the negative correlation of inters-cortical connection weights with the difference between the preferred directions of the visual and motor neuron pairs.

In comparison to the organization of lateral weights in the motor cortex (Figure 6.9) the visuomotor connectivity is characterized by a larger distribution of weights values. In the latter case, non-zero weights exist between neurons with PDs forming an angle up to 150° and the peak of the synaptic strength corresponds to a difference between PDs of 30° (i.e., compared with 0 for the motor map). Both features suggest a broader coupling of neurons, which has an effect on the accuracy on the generation of the desired direction of movement.

The findings with respect to the strengths of the mapping connectivity are relevant to those inferred mathematically by Salinas and Abbott (1994). The authors proposed that for map alignment to occur, the strength of the synaptic connection between neurons has to depend only on the magnitude of the difference between their preferred locations (i.e., directions) (see Section 3.3.3). In our simulation we found that the function between the strength of connection and the angle is a negative correlation.

Analysis of neural properties

The organization of the inter-cortical connections as described above, indicates one of the causes of the networks alignment: the visual and motor neurons with similar directional preferences become selectively coupled. Strengthening of these synapses instead of others occurs due to the time-correlated activity of those visual and motor neurons involved in coding the same direction of movement.

The above analysis is quantitative. It shows that at the population level, neurons coding similar directions become selectively coupled, but does not tell us how many and which of the motor neurons respond to visual information. To answer these questions, we begin by examining the profiles of motor activity patterns in the visual and motor conditions. The analysis reveals differences with respect to which neurons are activated in each condition and in the tuning properties of those neurons involved.

First, in the visual condition a larger distribution of the preferred directions of neurons involved in the generation of movement is observed. Compared to the motor condition, where neurons participate in coding directions which form an angle no larger than 90° with their PDs, in the case of visually guided movement neurons with optimal tuning up to 150° from the current direction are activated. As we mentioned above, this broad activation is due to the development of inter-cortical connectivity that links more than only neurons with similar preferred directions.

The formation of this visuomotor coupling can be explained by considering the characteristics of the learning procedure and the specifics of the motor neural responses. In our learning algorithm, for a synapse between two neurons to be strengthened, the presynaptic visual spike which travels through the axon for about 60 ms (i.e., time interval whilst the postsynaptic neuron can fire up to 8 times) has to reach the postsynaptic motor neuron in a short time window before this one fires. Obviously, this large delay makes the synchronization of the presynaptic and postsynaptic activities more difficult. Accordingly, motor neurons, which fire at high discharge rates during the whole movement, rather than neurons which fire precisely timed, but fewer spikes, will be favored in establishing functional connections with the visual neurons.

The analysis of the dynamics of motor neurons reveals two things. First, the neurons that become responsive to the visual stimulation are indeed those neurons (lateral or winners) that exhibit high discharge rates in the motor condition. By firing during both the execution

and the perception of movement direction, these units give rise to a new type of *visuomotor* neural behavior (see discussion in Section 7.1.3). Second, for each direction of movement we found that about 20% of the motor neurons remain silent during the visual condition. That is, they participate in the action execution in the motor condition, but not when the stimulus is visual.

Third, an interesting finding is the activation during the visuomotor mapping of a particular subset of motor neurons, from those, which are not very active for movement execution. These are motor cells that hardly reach the threshold when a movement is generated in the absence of visual input. Hence, during the motor condition, they emit very few spikes due to random factors (i.e., noise) and usually they are in a sub-threshold state. A slight increase in the inter-cortical weights determines that the visual signal is integrated together with the local lateral excitation and the motor input and causes their constant firing. Note that in order to fire, these neurons have to receive intra-lateral excitation, hence they are situated inside the assembly currently involved in the execution of movement. Furthermore, because the visual spikes trigger their activity the inter-cortical synapses of these units are significantly increased. This leads to the fact that, eventually, the visual signals independently can activate them.

Results show indeed, that these signals related to neurons firing at high rates during the visually guided actions. By contrast, the motor neurons involved in the command of movement spend much of their time in the refractory period. Hence, they have less chances to synchronize with the arrival of the visual spikes and to increase significantly their inter-cortical weights. Furthermore, it is not a coincidence that the signal related neurons are placed in the center of a cell assembly, similarly to the way lateral neurons are situated. That is, because only in such positions characterized by a strong surrounding excitation, can these neurons accumulate enough stimulation to fire. In a reciprocal manner, the visual related neurons send excitation to the other cells situated in the same neighborhood.

Based on these observations, we propose a hypothetical scenario for how visuomotor mapping of direction takes place in our model. In the first place, a percent of the motor neurons are activated directly by the excitation coming from the visual neurons. Second, the visual signals trigger activity in the visually-related motor neurons, which in turn spread activation in their neighborhood and determine the population to fire. Note that the visuomotor neurons firing are placed inside the cell assembly, which we want to activate. This hypothesis is consistent with the neurobiological findings describing the existence in the primary

motor cortex of different types of units, including sensory-related, motor, and sensorimotor neurons (Zhang et al., 2000; see also discussion in Section 7.1.3).

Further work will be aimed at exploring the means by which we can improve the visuomotor mapping accuracy. These preliminary results suggest that a correct alignment of the two maps is favored by a large, distributed representation of directions, coded in the discharge rates of neurons, rather than in the precise timing of the spikes. We believe that a more accurate visuomotor mapping can be learned using a larger motor network, where each direction of movement is multiply represented in the discharge rates of an extensive population of neurons. We also aim to explore in more detail the role of visually-related motor neurons in the formation of the motor response.

6.2.5 Discussion

In this section a computational model for the visuomotor mapping of directional information was proposed. From our view, learning of the visual guidance of movement is a developmental process, which takes place by unsupervised means. Moreover, it occurs between two systems that have already developed specific capabilities for representing and generating directional movements. To elicit an appropriate motor response, the directional information must be transferred from the visually related areas to the arm-control areas. We believe that the basic operation in this process is carried out through the correlated activity of directionally selective cells placed along the reach-related areas of the parieto-frontal network.

Our model shows that the visuomotor transfer of coded information is supported by the development of inter-cortical connection weights negatively correlated with the difference between the preferred directions of the visual and motor neurons in a pair. This result confirms the estimation made by Salinas and Abbott (1994). Furthermore, we found out that not all motor neurons in the map become responsive to the visual input. Those which do, are represented mainly by neurons that are significantly tuned to directional information (i.e., fire with high discharge rates) during the movement command. Besides the visuomotor neural behavior, a type of signal-related neurons develop in the motor area, from those, which were not very active during the generation of movement. These neurons may play an important role in relaying information from the visual cortex to the motor areas (see also Section 7.1.3).

The idea of using the correlated activity of motion selective cells for visuomotor mapping has already been modeled by Burnod et al. (1982), Salinas and Abbott (1995), and more recently has been made the basis of a theoretical framework proposed by Burnod et al. (1999). Compared to this earlier work, the merit of the current study is to have obtained an alignment of the motor to visual information in conditions of realistically implemented motor coding of direction. In our scenario, by contrast with previous models, directional selectivity and population coding emerge in the motor network, as a result of a developmental process. Moreover, self-organization in our model occurs on networks of spiking neurons. In this respect, the simulation work described in this chapter revealed interesting findings on where directional information is coded. On the one hand, for the fast control of movements, directional information might be read out, if necessary, from the timing of the first spike of a subset of spontaneously activated neurons. On the other hand, the visuomotor mapping accuracy depends on the way the directional information is coded in the motor cortex. In this case, rate coding over a large population of neurons supports the correlation of activity in the two networks and allows learning. More work is needed in future to analysis in detail the synchronization of activity at the level of single spiking events.

Compared to the previous models of sensorimotor mapping based on the correlated firing of motion selective neurons our model has modest aims and attempts to address fewer issues. We do not account for coordinate system transformation (see Salinas and Abbott, 1995) and our simulations are not effectively implemented for motor control (Bullock and Grossberg, 1993). Rather, our model constitutes an illustration of how a *complex* problem, such as the visuomotor mapping can be implemented in a very simple manner by two networks of directionally selective spiking neurons. This model argues for the importance of the individual neurons low-level properties (i.e., as directional selectivity) in implementing large-scale phenomena. Our hypothesis that the neural selectivity for stimulus attributes lies at the foundation of visuomotor mapping is very similar to the theoretical concept of combinatorial domains proposed by Burnod et al. (1999). We consider that our simulation findings have computational relevance for the experimental framework developed by these authors (see discussion in Section 7.1.4).

Moreover, the model proposed supports the thesis that gaze (i.e., eye) movement neural activity can be *re-used* to control the movement of a limb. In our view, this means that the directional gaze information becomes correlated, through a mechanism similar to the one described here, with the arm movement activity, allowing the eye-hand coordination

through a motor-to-motor program re-use (see Section 7.1.4). In the final chapter, the relevance of our findings to the current theses of visuomotor development will be discussed (Section 7.1.4, along with the possible applications of our models (Section 7.2).