

Chapter 7

Discussion and future directions

The final chapter of this thesis includes a discussion of the simulation results obtained and of the theoretical and practical significance of our work. In Section 1 the neurophysiological implications of the models on motor cortex self-organization and visuomotor mapping are discussed. In Section 2 we discuss the psychological relevance of the visuomotor model to the imitation issue and we present future extensions and possible applications of our work.

7.1 Neurophysiological implications

7.1.1 Emergent vs. innate directional selectivity of motor cortical neurons

Despite the vast body of knowledge currently relating to the motor control of the arm, there are conflicting explanations of how the neocortex participates in motor control (see the debate on muscles vs. movements encoded in primary motor cortex in Section 2.1.2). One impediment to a complete explanation of the function of the M1 is that the fundamental organizational principles of the cortical motor areas are yet not clear (Sanes and Donoghue, 1992). For instance, there is a large body of evidence indicating that directional tuning is an essential feature of motor cortical neurons (Section 2.1.3). However, it is not yet known whether this neural characteristic is acquired by experience or genetically encoded.

Our modeling work on the self-organization of motor cortex represents a first attempt to provide a learning scenario for how motor cortical neurons develop directional selectivity.

Put generally, it demonstrates that a self-organizing map can learn to distinctively represent and command 12 directions of movement, by extracting the similarity relationships from the input space. The success of the self-organization process is dependent on two factors: the input patterns and the feedback connectivity system.

A self-organizing feature map is a means of visualizing in a reduced dimensional space (usually two) the spatial relations existing in a multidimensional space. Hence, if we aim to approach the formation of the motor directional maps from the self-organization perspective, the crucial aspect consists of the characterization of the input signals that are available to the training process. From this view, the population coding operating in the motor cortex and the functional connectivity that has formed, are the result of a self-organization process and reflect up to a point, the peculiarities of the input space. However, as it was pointed out elsewhere in this thesis, in the case of the motor cortex as opposed to the visual cortex, it is less clear what precisely might be the input (i.e., training) data to the self-organization process. Our hypothesis is that the formation of the directional motor map is driven by proprioceptive feedback from those muscles involved in movement. This idea will be the starting point of our future work on modeling motor cortex organization (see details in Section 7.2.2).

In our model, a directional feature map emerges through unsupervised learning from a random initialization of the afferent and lateral connection weights. There is one built-in constraint in the shape of the network connectivity: the short-range distribution of excitation and the long-range inhibition. The connectivity with a Mexican-Hat profile is a general feature of a self-organizing feature map. That is, because short-range excitation is needed to focus the activity in the immediate neighborhood of the winning neuron, while the long-range inhibition helps to suppress the network response in the contralateral direction of the movement.

There is experimental evidence for the existence in the motor cortex of adult animals of excitatory connections which link nearby neurons with similar neural responses and of distal inhibitory connections between neurons with different tuning curves (Georgopoulos et al., 1993; Hatsopoulos et al., 1998). However, it is not clear if this is a result of a developmental process taking place in the motor areas or represents a built-in feature. We believe that these kinds of questions can be easily explored through modeling work, with promising and valid results. Part of our future work is aimed at exploring the influence on the map organization of the variation of connectivity pattern parameters (i.e., the rate of the connectivity and the

spatial distribution of the lateral feedback).

Can be the motor cortex modeled like the visual cortex?

Another observation with respect to the theoretical assumptions of our model, concerns the fact that it was mainly inspired by models of the visual cortex (see also discussion in Section 7.1.2 below). The power of the visual model was twofold. Firstly, at its heart, the modeling of visual cortex organization challenges the idea of innate cortical features (Hubel and Wiesel, 1962; Gilbert and Wiesel, 1992). We believe that if it is possible that the visual directional maps are shaped by experience, then it may also be feasible to model the developmental processes of motor cortex. This hypothesis is supported by experimental studies that indicate that there is a self-organization capacity in the cortex of adult animals, which is perhaps part of the original developmental organizing processes (Merzenich et al., 1983; Hess and Donoghue, 1994; Rioult-Pedotti et al., 1998).

Secondly, the modeling studies of the visual cortex put forward in the last decade a core of hypothesis on the computational and design principles of the brain. First, the self-organizing feature map (SOM, Kohonen, 1994) has been very successful in modeling the development of sensorial maps. It has become almost a *de facto* standard in the biological modeling of brain self-organizing processes. Secondly, computational studies pointed out the essential role played by the horizontal connectivity in the formation of orientation, binocular, or directional maps (Section 2.2.4). Placed in this context, our simulation work has the advantage of a bi-directional knowledge transfer. On one hand, our study has been largely inspired by existing models and data from the sensory cortices. On the other hand, if our model proves successful in simulating the formation of motor directional maps, then it provides computational evidence of the learning mechanisms and the functional principles of the motor cortex.

Our findings (Section 6.1) indicate that the self-organizing feature map represents an appropriate modeling framework for the developmental processes taking place in the motor cortex. Furthermore, we found that the lateral feedback system plays an essential role in the organization process, in a similar way to the role it plays in visual cortex development. Plasticity of both excitatory and inhibitory connections is essential for self-organization to occur, by finely adjusting cells tuning level to the input space features. Exploring the effects of learning in terms of single spike timing represents our original contribution in the area of

self-organization with spiking neurons. Our model indicates that directional information may also be read out from the timing of the first spike of fast responding neurons. Surprisingly, this observation comes out as a possible common feature of information processing in the visual and motor brain.

Until very recently, the common belief in computational neuroscience was that information in the brain was carried mainly by the neuron's discharge rate (see Section 6.1.5). Recent experiments on visual categorization revealed the existence of a very fast processing of information in the visual cortex, possibly based on the order or timing of a single spike per neuron (Thorpe et al., 1996; Thorpe and Gautrais, 1998). In the case of the motor system, the influential work of Georgopoulos and his co-workers (1984, 1986) proposed the population coding scheme as the main paradigm used to interpret and predict movement based on the motor cells' discharge rates. Based on our modeling results, we suggest that a fast response of the motor cortical areas, read out from the timing of the first spike of optimally tuned neurons is certainly advantageous and quite likely implemented by the motor system. The only restriction is that such an answer has a very limited precision and only further processing of the directional information by a large population of cells can give rise to an accurate reach movement.

It is beyond the scope of this thesis to offer an answer to the question of whether directional selectivity is a genetically encoded feature of the motor cortex or is acquired by experience. The model proposed here provides only a number of computational ideas on what it takes to develop neural selectivity and population coding in a biologically plausible system, by unsupervised means. We believe that by developing optimal responses in its elements, rather than having them pre-wired, a system can show a flexible and plastic architecture that adapts to the resources available and to the particularities of the input space.

7.1.2 More evidence for the importance of horizontal connections

It was pointed out that in our model, an essential role in the organization of the motor map was played by plastic lateral connections (see Section 6.1.5). Our findings show the formation of functional connectivity in the motor area with a similar profile to the patterns of connections described in other brain areas. Thus, experimental data from primary visual cortex shows that horizontal projections link columns with common ocular dominance and orientation selectivity's (Gilbert and Wiesel, 1992). In the auditory cortex dorsoventral

connections link regions with matched characteristic frequencies (Read et al., 2001). Recent modeling of eye-saccades planning in the lateral intraparietal area (Xing and Andersen, 2000) have shown that in order to hold memory activity for a saccades, the neural population develops excitatory connections between units with similar preferred saccade directions and inhibitory connections between units with dissimilar directions.

Previous modeling results similar with ours have been obtained by Lukashin and Georgopoulos (1994). They found that during a supervised learning process, the strength of connection between directionally tuned motor neurons becomes negatively correlated with the difference between their preferred attributes. This sort of experimental data and computational work, suggest the manifestation in the brain of a general principle for horizontal connections organization. It is generally believed that this is reflected in the correlation between the strength of interaction and similarity among units' preferences.

With respect to the computational function of the lateral feedback system, our model of motor cortex organization and visuomotor mapping, indicate a crucial role of the horizontal connections in shaping the activity of the network and in favoring the formation of stable attractors of motion directions (see Section 6.1.5 and 6.2.4). On the short scale, the lateral excitation increases the collaboration within a cell assembly, while the lateral inhibition suppresses the answer in the opposite direction. On the large scale, the plastic connections implemented in the visuomotor system, mediate the transfer of information and synchronization over a large distance (i.e., 50 ms delay). Our findings suggest that correlated activity in motor and visual networks is a result of both organization of long-range connections and collaboration mediated by the local lateral pattern.

Similar observations have been made by Usher et al. (1996), who studied the role of long-range connections for visual binding and line completion. They used a network of leaky integrate-and-fire neurons with long-range connections implemented only between cells with similar orientation preference. Their findings revealed a clear tendency for synchronization between cells with same orientation preference separated by large distances. In their model, if two distal cells placed in the range of clustered connections receive the same stimulus (even if they are not optimally tuned themselves), they indirectly synchronize via the intermediate synchronization with their optimally tuned neighboring cells.

Compared to Usher and co-workers work, our simulation has the advantage of developing the profile of the long-range connections. In our case, the network coupling consists of a full connectivity from the visual to the motor network, initialized with weight values

nearby 0. By doing this, we do not arbitrarily restrict which is the visual directional information perceived by the populations of motor neurons that encode all movements directions (i.e., in our case, whole motor network). The coupling of neurons according to their preferences, should be an emergent feature, rather than a built-in property that limits the network plasticity.

In our model, as a result of learning, clustering of connections occurs in a similar way to the pattern implemented by Usher and colleagues. The strength of long-range connections cluster in spatial neighborhoods that correspond to the directional cell assemblies in the motor area. Accordingly, correlated activity between visual and motor neurons is induced not only via the long-range synapses, but also through the mediation of the visually-related motor neurons optimally tuned to the direction of movement (see discussion in Section 7.1.3 below). Future work will take into account a more realistic scenario in which the cortico-cortical projections start with some initial non-zero (i.e., biased) values and weakening of synapses, besides strengthening, will be allowed.

7.1.3 Dynamics of single neuron activity in the motor cortex

Up to this point, the discussion has focused upon describing the main requirements for self-organization of the motor map and the alignment of visual and motor neural representations. However, an important co-lateral effect of modeling these developmental processes was to observe the emergence in the motor network of different patterns of neural behavior. These may reflect various functions in the preparation and execution of movement, which are discussed below.

During initialization, the neural spiking model is set up in such a way that all motor cells begin the simulation equally selective to all motion directions. However, learning in a self-organization map (SOM) is a competitive process and takes place by amplifying any small differences in the neural response. If one neuron wins for a certain direction, its synaptic strengths are modified to increase its chances of winning again for that pattern. In this tuning process, the variability of the neural response is an essential factor and is given by the level of noise (i.e., in the threshold value, the firing time, the spike transmission delay) and the pattern of connectivity. Due to the fact that input signals can arrive in a synchronous or asynchronous way, this neural variability causes the possible operation of neurons in two modes. Thus, a neuron is capable of switching between computational modes, from the

integration of firing rate input received from a large number of neurons, to the detection of coincident spike arrivals (see also experimental evidence for neurons acting in different computational modes in Destexhe and Pare, 1999). Hence, we can describe the existence of two main classes of neurons.

The first class, that of coincidence detectors, is mainly formed by the *winner neurons*, which represent about 40% of the neurons in the motor area (i.e., 110/264) (see Section 6.1.5). These neurons respond very rapidly to the input signals that are emitted synchronously and affected by similar values of noisy delays. Hence, their afferent weights become highly tuned to one input pattern. Note an important difference between learning with a SOM of continuous, rate-coding neurons and a SOM consisting of spiking neurons. While in the former an input pattern is mapped onto a single neuron that has the maximum activation, in the later, any pattern similar to the best-matching pattern will determine the firing of the winner neuron. Hence, if there is no increase of the neural threshold, what we have obtained in our spiking SOM, were neurons highly responsive to several (i.e., maximum three) directions of movement.

It is noteworthy that the preferred directions of each neuron, when represented on a circle, occupy an arc of maximum 60° . Similar results have been described experimentally by Battaglia-Mayer and colleagues (2000) on a study of early coding of reaching in parieto-occipital cortex. The authors have found that reach-related cells in this area have about three preferred directions of movement. Consequently, they characterized the neural response through a 'field of global tuning', defined as the sector of the directional continuum within which all its preferred directions lie (e.g., approx. a quarter of a circle).

A second class of behavior is represented by the integrators, which are neurons that are commonly needed to integrate a large number of inputs in order to fire. If a neuron did not spike at the coincident arrival of the input signals, then due to the exponential decay of postsynaptic potentials, the accumulation of several excitatory potentials will be required before the postsynaptic spike will occur. Hence, these neurons' activity strongly depends on the strength and the number of their lateral excitatory connections. In the motor condition, about 20% (i.e., 50/264) of all neurons show a significant directional tuning while operating in the integration domain (see Section 6.1.5). These cells, referred to as *lateral neurons*, need to integrate both motor input and local lateral excitation in order to become active.

Another subclass of integrators is the neurons, which need to sum excitation from three sources: afferent, local, and long-range connections. Directionally tuned activity occurs

in these motor neurons only during the visual condition. Accordingly, they were named *visually-related* neurons and they constitute about 10% of all motor units (i.e., 25/264). A remark here is important with respect to the degree of directional tuning of these neurons. It was shown in Sections 6.1.5 and 6.1.5 that most of the winners are broadly tuned to several directions of movement, while lateral neurons responses are less broader. In the light of the above discussion, we can say that the selectiveness or tuning of the neural response increases with the number of inputs integrated. Thus, the lateral neurons are significantly more directionally tuned than the winner neurons, and the visually-related neurons are optimally tuned (i.e., most selective) compared to both previous categories.

Even if our simulations represent a drastic simplification of the mechanisms involved in neural control of reaching, several hypotheses are presented here, with respect to the functional roles of neurons. Studies of visuomotor processing in the parieto-frontal network involved in reaching demonstrated the existence of various types of neural activity. During an instructed delay task followed by a pursuit tracking task, Johnson and colleagues (1999) have analyzed the directional discharge of neurons in monkey's premotor and primary motor cortex. From 240 neurons, in 132 cases, significant directional tuning was found for both the cue and track periods. In 26 neurons, directional tuning was found only during the cue period, and in another 54 the directional tuning was significant only in the track period. These neural behaviors can be classified as: (1) *visuomotor neurons*, whose activity show the co-existence of visual and movement control signals; (2) *signal neurons*, defined as motor neurons with visual properties, which respond transiently to the onset of the visual cue; (3) *movement-related neurons* that fire only for movement control.

We have obtained similar dynamics for cell activities in the motor network as a consequence of learning the visuomotor mapping task. From 264 motor neurons, about 60% developed visual properties, from both the winner and lateral neurons. In the absence in our model, of a behavioral task comparable to the instructed-delayed task, the *visuomotor neurons* are represented by those neurons which show directional activity during movement execution and under visual stimulation. The *signal neurons* correspond to our visually-related motor units (10%), which fire only during the visual stimulation period and are almost silent during movement execution. Finally, we have also found about 8% of motor neurons that are involved in the control of directional movement, but remain silent when stimulated by visual signals. These correspond to the *movement-related neurons* described experimentally.

Our results suggest that the formation of the motor network's response under visual guid-

ance take place in a few steps. First, the *visuomotor neurons* are activated via the long-range inter-cortical connections by the visual directional signals. Note that in our model, the strength of these connections reaches a peak for a difference between the distal visual and motor neuron preferred directions of 30° . Therefore, the motor network response evoked in this way is broadly tuned to the desired direction of movement. Instead, the *signal* (i.e., visually-related) neurons that occur in the motor network are optimally tuned to the visual direction of movement. Hence, they can play an essential role in finely adjusting the motor population response (see Section 6.2.4). The next stages in the formation of the network response involve a successive propagation of activation, started by the visuomotor and signal neurons and mediated by the motor lateral neurons. The joint activity of all these neurons leads to the formation of the desired direction attractor.

Up to this point, we focused upon discussing the immediate implications of our modeling results. In the remainder of this section and along the next section we will outline the theoretical relevance of our models and their possible application to real systems.

7.1.4 Theoretical significance of the visuomotor mapping model

In the theoretical background of this thesis we reviewed a number of recent neurobiological theories of visuomotor control of movement (Sections 3.1.2, 3.2). At that point, we introduced four main hypotheses:

- The sensorimotor cycle has a unitary nature;
- The visuomotor transformation is achieved gradually, supported by the combinatorial properties of the neurons;
- The existence of common motor programs for eye and hand movements can reflect the operation of cortical computational principle of 'program re-use'
- The alignment of motor and visual networks for the correct transfer of information can be learned through a simple Hebbian learning principle.

While the last principle has been directly addressed in our simulations, the implications of our work for the other points might not be immediately clear. Therefore, we propose below an integrative view, which presents the theoretical relevance of our computational results.

The progressive match framework

Despite the simplification of the visuomotor transformation process in our model, we believe it illustrates a number of basic computational principles of this process. In particular we consider that our modeling work is relevant to the theoretical framework proposed by Burnod and colleagues (1999), even more so because, at the time we implemented our simulations (2000) we were not aware of their work.

The basic computational demand for reaching is met by the alignment of the visual and motor neural representations. This was achieved in our model by implementing a Hebbian-like learning mechanism, that correlates activity in spiking neurons with some feature selectivity, that is, in our case, directionality. Burnod's et al. (1999) model is based on the operation of an equivalent mechanism. The visuomotor transformation is described in terms of a progressive match of different sets of sensory information by neurons with tuning properties. Matching takes place gradually, in several combinatorial domains. In each domain, an identical computational mechanism operates, through the co-activation of matching neurons tuned to the same preferred attribute (position or direction).

The contribution of our model resides mainly in the fact that it is based on computations with spiking neurons and implements a realistic population coding of motor directionality. The operation of the computational mechanism implemented is not restricted to any particular area. It can align neural representation coding for any type of stimulus features (orientation, direction, pitch, etc.). Moreover, it implements learning on two of the four combinatorial domains described by the authors (Section 3.2.3). If we consider the training input to the self-organizing motor cortex as proprioceptive afferent feedback coming from activated muscles, than we have in the motor network organization, the learning of the first domain, which relates muscle dynamics and arm command. By relating the gaze direction and hand movement direction in the visuomotor mapping process, the system learns hand-tracking in the third combinatorial domain. Note that in our model, and the Burnod et al. (1999) framework, motor control, i.e., referred to as a motor babbling stage in the progressive match framework, is developed prior to visuomotor mapping learning.

The relevance of our model is even more significant, if we consider that in the Burnod and coworkers proposal most concepts were inspired by neurophysiological data. Instead, our model started out with a minimum set of architectural assumptions and a number of equivalent concepts emerged in the network, through development. For instance, in the pro-

gressive match model, a key element is the set of units, which integrate information on the sensorimotor axis: sensory units, motor units, and matching units. It is clear that the type and functionality of these units has been implemented according to the experimental data (Johnson et al., 1996; Caminiti et al., 1998; see data described in Section 7.1.3). Conversely, in our model these types of behavior have simply emerged during the self-organization process.

Moreover, in the absence of a working model, Burnod and colleagues focus upon the role of the matching units within the learning process of sensorimotor correlations. Instead, we have seen that correlated activity in two networks is a result of synchronization via long-range connections but also through intermediate synchronization with neurons in the same cell assembly. That is, learning takes place in a more distributed manner and involves matching (i.e., visuomotor) units as well as signal (i.e., sensory) units and movement-related units. We believe that an important further step into the realm of biologically inspired modeling of arm-reaching will be represented by the complete implementation of the progressive match architecture. Our efforts will be aimed at implementing more conceptual elements of this framework. A particular goal will be to obtain the formation of *condition* or set units, which are involved in maintaining the neural representation during delayed tasks (Johnson et al., 1999; Burnod et al., 1999).

Cortical Software Re-Use

Another theoretical framework within which we can discuss our results is the cortical software re-use theory (CSRU, Reilly, 1997; Reilly, 2001). Put simply, CSRU states that a general principle of creative cognition is the appropriation of computational programs from one domain and their application to another. For instance, CSRU proposes that perceptual binding can be seen as a collaborative process between cell assemblies that are equally well developed (Reilly, 2001). The style of computation is best viewed as a process of dynamical entrainment, involving the synchronization of firing patterns in reciprocally connected cortical areas. We believe that the neural mechanism for 'binding' visual input to the relevant motor output for visuomotor mapping implements a similar type of computation. In CSRU terms, the visual neural activity is *re-used* to control the movement of the limb.

A stronger claim (i.e., hypothesis) is that the limb movement may make use of the eye motor activity, in a so called *motor-to-motor program re-use* (see Section 3.1.2, Metta et al., 1999;

Reilly and Marian, 2002). Note that this hypothesis states that visual (i.e., retinal) information is not necessarily required for the guidance of movement. The alternative to this process is that the motor program for eye movement is re-used to control limb movement in the same direction. This re-use has the advantage that eye-movement related signals can be read out at any processing stage, from various cortical and subcortical areas, and they are in head-coordinates, compared with the retinal information in Cartesian coordinates.

Some computational support for this hypothesis already exists. Metta and co-workers (1999) have implemented a model of visually guided reaching based on the alignment of the head map with the arm network (see Section 3.3.1). Similar results have been obtained by Marjanovic et al. (1996), who constructed a system that first learns to foveate a visual target and then re-use the saccade map to achieve ballistic reaching. Such modeling work provides a compelling example of how motor programs for eye movements or heading (see also Kolesnik and Streich, 2002) can support the development of visually guided reaching. From a developmental perspective, the program re-use makes much sense, as the 'software' for heading, eye movements and gaze focus develop *priori* to the control of reaching. Consequently, this hypothesis has a great potential in the robotics field.

The perception–action cycle

A final thought in this theoretical section, will be given to the unitary nature of the sensorimotor coupling. With respect to this rather abstract issue, much less can be inferred from our simple model of visuomotor mapping. One might say that our assumptions are rather a personal choice than scientifically proven facts. We have developed them while trying to find ways to implement the sensorimotor transformations.

The personal belief of the author of this thesis is that sensorimotor mapping is a fundamental, built-in property of any living nervous system. This means, that as a general principle of any nervous system, sensing–for–acting is implemented as one unitary computational operation. Hence, to characterize the task of transforming the sensory information from one modality to the motor output in another modality as the sensory–motor transformation 'problem', is perhaps to view it from the wrong angle. Moreover, because it is a general and ancient design principle of the nervous system, it is implemented at the lowest-level of the system and it is preserved in more complex variants of the system (i.e., primates or human brains). We consider that the apparent modularity of the human brain and the high

degree of sophistication of its circuitry conceal the functioning of some low-level, built-in mechanisms which implement fundamental computational operations.

With respect to the transformation concept, we believe that it owes much to Cartesian dualism (i.e., external vs. internal world). We are probably on the same line with the critique made by a roboticist of the general tendency to assume that description and implementation of a system must be equivalent:

We believe that classical and neo-classical AI make a fundamental error: both approaches make the mistake of assuming that because a description of reasoning/behavior/learning is possible at some level, then that *description must be made explicit and internal to any system* that carries out the reasoning/behavior/learning (Brooks et al., 1998, page 961, our emphasis).

Insights from our modeling work, and other more sophisticated models by Salinas and Abbott (1995), Burnod, Baraduc and colleagues (1999), showed us that a global complex operation, such as information transfer for the visual guidance of movement, may rely on the simple mechanism of correlated activity of single cells. The core of our model is based entirely on the ubiquitous feature of neurons to be directionally selective. A more general solution to the problem of sensorimotor transformation based on similar basic computational mechanisms was given by Salinas and Abbott (1995). Furthermore, it was discussed that combinatorial properties of directionally and positional selective neurons are the key element in the progressive match architecture for visually guided reaching (Burnod et al., 1999).

First, at the level of a single cell, several sources of information can be integrated along the sensorimotor axis (see Section 3.2.3). Second, the correlated activity of cells with the same preferred attributes (direction and position) can allow the correct transfer of information. Third, the coordinate transformation can be understood in terms of neural gain field, where the response of a neuron is a product of the receptive field and the linear gain field (see Section 3.2.2). The point we want to make here, is that neural information processing systems rely heavily, on the computational features of single units.

In the computational neuroscience field the ideas outlined above, are well known. However, in the field of artificial intelligence and robotic applications, almost no attention is given to the properties of the neural model. For instance, a very succinct comparison between types of neural models existing reveals the following. A continuous rate-coding neuron, that represents the computational unit of the classical neural networks, can compute a temporal

linear summation of inputs. A simplified model of the spiking neuron can in addition detect coincidence, can do multiplexing, and can compute in a temporal domain using delay codes (Maass, 1999). A compartmental model, which includes the dendritic tree, can perform spatial summation, nonlinear operations (division), can increase its discrimination and memory power up to thousand times that of the linear neuron, and can detect movement direction and binaural stimuli (Koch, 1999; Poirazi and Mel, 2000).

The above comparison represents a twofold argument. First, artificial systems may benefit enormously from paying more attention to the neurobiology of the living systems and to the way these implement perception and control of action. Second, the single neuron is indeed a very powerful computational device. Hence, we believe that by connecting these neurons in small size circuits, primitive operations such as the perception–action cycle can be implemented as an intrinsic feature of the system.

7.2 Applications and future directions

The central goal of this thesis was to offer a computational model that helps to bridge the gap between cognitive description and neural implementation of mental phenomena. That is, to understand the link between what a single computational element does and what many of them do when they function cooperatively. It was also pointed out, that understanding the way the brain organizes the control of movement can be largely beneficial to the design of artificial control systems. In general, the research dedicated to the understanding of computations in real nervous systems shares the same motivation: to apply what is learned from nature into the design of adaptive, intelligent, and eventually self-developing artificial systems. We will try to discuss, in this section, possible integration and future extensions of our models to systems of motor control. Up to the present these ideas are only at the stage of proposals. It remains future work to show to what degree their implementation can be beneficial.

7.2.1 Transforming plans in actions

A possible integration of the motor cortex organization model is within control systems based on a direction-mapping strategy. In this case the system implements a transformation from spatial trajectory to end-effector directions or rotations as opposed to end-effector

positions (Bullock et al., 1993; Fiala, 1995; Ritter et al., 1989). The idea of mapping spatial into motor directions is supported by experimental evidence on the directional selectivity of cells in premotor and primary motor cortex (Caminiti et al., 1991; Georgopoulos et al., 1986; see Section 2.1.3). There is also psychophysical evidence for a direction-based rather than a position-based transformation coming from studies on blind reaching. These experiments suggest that the magnitude of the error is dependent on movement amplitude, rather than on the end-point alone (Fiala, 1995).

One of the most efficient implementations of direction mapping for visually guided reaching is the DIRECT model proposed by Bullock et al. (1993) (see Section 3.3.2). To control arm movements, the system first performs an integration of position and visual directional information into a position-direction map. In our model this corresponds to the motor network which learns to align motor and visual directional information. Furthermore, the DIRECT model focuses on learning the mapping from motor directions in body-centered coordinates into joint-rotations in joint coordinates. The accuracy of a three joint arm movement in 2D and 3D space strongly relies on the way the visual directional and positional information are correlated in the motor map. The authors argue that only a sharp tuning of each cell in the map, to a visual direction in a particular joint position, can ensure the accuracy of reaching movements. Even if they acknowledge that this is a significant deviation from the neurophysiological evidence (see population coding of directionality in motor cortex Section 2.1.3) they justify it by the fact that in the case of broadly tuned cells, the model fails to generate correct reaching movements.

The accuracy of visually guided movements is not an easy task for our model either, nor for any model grounded on neurophysiological data (see also Baraduc et al., 1999). Our analysis of the network organization leads us to believe that mapping accuracy is strongly influenced by two factors: the parameters of the horizontal connectivity pattern and the quality of the motor population codes for directions. By the parameters of lateral connections we mean the rate of connectivity, the profile of excitation and inhibition, and the plasticity rules. By the quality of the motor coding of direction we mean the stability and the distributed nature of the neural representation of each direction. In other words, the accuracy of visually guided reaching not only depends on how well the visuomotor coupling is learned, but also on how precise the motor control of the movement is itself. That is why we consider that a separation of the visuomotor development process into two stages may be beneficial. First, a motor babbling or motor learning phase is required, to ensure the motor cortex organizes

for the precise control of movements. Learning the visuomotor correlations follows this.

To conclude, we believe that the advantages of integrating a developmental model similar to ours in a motor control system are many:

- It allows the emergence of population coding of directionality based on broadly tuned cells. This gives biological plausibility to the model, besides bringing all the benefits of a distributed representation (as opposed to a localist representation): flexibility, plasticity, reduced size.
- By exploring the way learning evolves in the lateral connections, our model allows the formation of stable attractors of movement directions, which in turn contributes to the accuracy of reaching.
- Only by modeling the developmental process, can one observe the emergence of different patterns of neural behaviors, with different functions in integrating and combining information, matching, conditioning or delaying the response.

Motor primitives and the equilibrium point hypothesis

Another direction to follow in order to translate our motor plans into actions, is to control the arm movement in conformity with the spring-like properties of muscles and reflex loops (Bizzi et al., 1992). This idea involves putting together the concept of motor primitives and the equilibrium point hypothesis, as an alternative to the inverse dynamics problem (i.e., the DIRECT model).

The motor primitives represent an appealing, rather theoretical concept, used by researchers on both artificial and biological motor control to reduce the complexity of movement generation to elementary units of action. They are defined as a set of *basis behaviors*, which are not further reducible to each other and which can be composed to produce the complete behavioral repertoire for the system (Brooks, 1986; Mataric, 1997; van Essen et al., 1996). On the other hand, the equilibrium point hypothesis is an experimentally-derived theory, according to which movement arises from shifts in the equilibrium positions of the joints. An equilibrium position is a consequence of the interaction of central neural commands, reflex mechanisms, muscle properties and external loads.

A recent extension of the equilibrium point theory, developed by Bizzi et al. (1991) and

Bizzi and Mussa-Ivaldi (1995) proposes that the muscles generate *convergent force fields* (i.e., equivalent to motor primitives), which direct the limb toward an equilibrium point in space. The vectorial superposition of these independent force fields can generate a vast repertoire of motor behaviors. The simulation studies of Mussa-Ivaldi (1999) have shown that by combining a small number of convergent force fields it is possible to reproduce the kinematics features of reaching arm movements.

The motor primitives proposed by Bizzi and colleagues suggest that spinal mechanisms can serve as substrate for the operation of motor cortical activity, in order to produce a directed movement of the limb. Georgopoulos (1996) proposed an integrative account of how directionally tuned motor cortical commands can be translated in the activation of muscles. In his view, this mapping can be accomplished by connecting a population of central cortical neurons through a set of motor inter-neurons, with a number of spinal populations associated with different motor primitives. Then, the preferred direction of a cortical cell emerges as the vectorial, weighted sum of the force fields that act on the hand at a certain position in space.

For our simplified version of motor control, this idea can be more beneficial and easier to implement than dealing with the complexity of a multi-staged architecture, such as the one implemented in the DIRECT system (i.e., with nine layers and learning at four different stages). It also allows a bi-directional transfer of information in the system: an efferent path, from the cortical motor network to the muscles and a re-afferent path, which brings feedback on the muscles activation to the motor cortex. We believe that this can be the appropriate framework for our future modeling work of motor directional map organization based on training input coming from muscles (see Section 7.2.2 below).

7.2.2 Future model

Our future model of visuomotor learning will be developed with a specific goal. That is, it will represent the neural controller of an avatar, endowed with simple vision, action upon objects, and proprioceptive feedback on the effect of its movements. The first step in this extension of the actual version of the model is to provide the motor network with proprioceptive feedback. In our view, the information available for the formation of the motor directional map is represented by afferent signals from those muscles that are involved in movement.

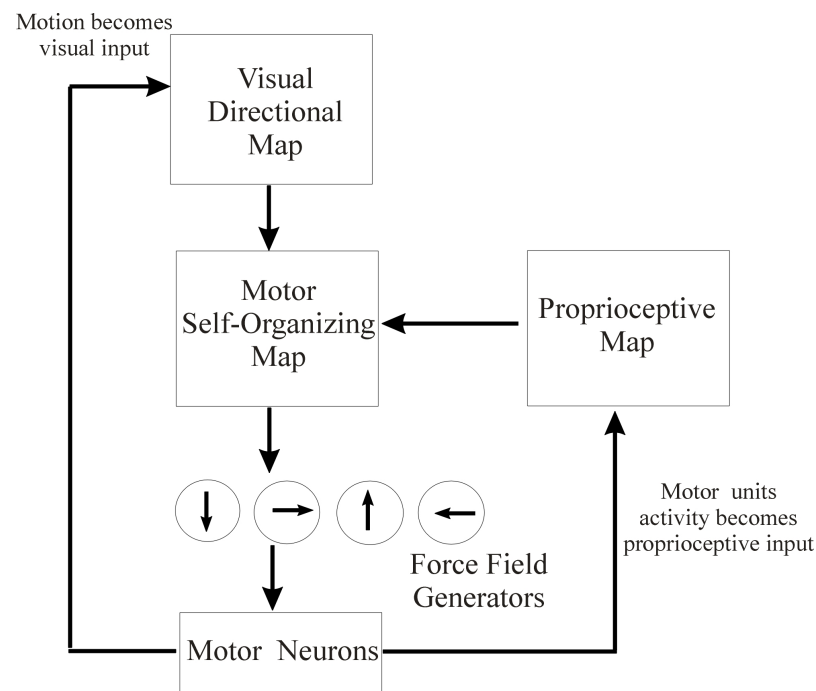


Figure 7.1: General architecture of the future work model. It includes several modules. A visual network of directionally tuned cells, with a retina-like input layer. A motor network which self-organizes for the control of movement direction. Its output is send to a set of force fields generators, which control the motor units in the muscles. Its input is provided by a proprioceptive network that receives afferent signals from the motor units involved in movement. These signals contain directional information, derived from the preferred directions of the muscles.

Recent research on the contribution of muscles to joint torque indicated that mono- and bi-articular muscles have different functional roles in the control of multi-joint movements (Bolhuis et al., 1998). Experimental data demonstrated that the activation of bi-articular muscles vary with the direction of force exerted, while mono-articular muscles show significant direction-dependent activation. Furthermore, it was shown that the mono-articular muscles have preferred movement directions, which cluster over subjects for both force direction and arm posture. To us, this data suggests that the motor unit activity may provide the directional information required for the organization of the cortical directional map.

Further, it is known that motor neurons in M1 make use of feedback information via afferent sensory pathways. At the contraction of muscles, information on which muscle contracts and how much tension it generates is fed-back to the motor cortex through the primary somatosensory cortex. Previous modeling work on this area was done by Chen and Reggia (1996) who studied the relation between the formation of motor and somatosensory feature

maps for arm-muscle control. They have shown that an alignment of the neural representations of muscle activation in the two areas occurs through correlated-activity means. We intend to explore this alignment in the case of the motor and somatosensory networks encoding directional information, derived from muscle tensions and contractions. The proprioceptive network will be mainly used in order to transmit the motor units' activations as input to the motor cortical network, in the hope that a directional motor map will form. The proposed architecture is outlined in Figure 7.2.2.

7.2.3 The imitation challenge

At the end of this chapter, we want to turn our attention back to the initial motivation of this thesis, that is, the neonatal imitation phenomena. Here, we discuss this issue within the more general context of imitative behaviors, which represent one of our future modeling goals. This is because imitation plays a central role in human development and is currently being explored as a powerful, alternative mechanism for teaching robots (Schaal, 1999; Dautenhahn, 2000; Billard and Mataric, 2001). Hence, we would like to abstract some relevant ideas to the imitation modeling, from the work presented in this thesis.

The challenge posed by neonatal imitation is to understand the capacity of infants as young as few hours, to imitate facial expressions, such as tongue protrusion and mouth opening (Meltzoff and Moore, 1977). Meltzoff and Moore (1999) proposed that a key element in explaining the mechanisms of this behavior is that the imitative act can be differentiated into the *body part* and the *movement performed*. They consider that evidence suggests that neonates select what body part to use before they have determined what to do with it. Further, finding the correct action on the organ involves more effortful behavior preceded by a series of searching movements.

Certainly, being capable of organ identification is probably the most astonishing hypothesis regarding the newborns innate capabilities. Interestingly, this problem is less controversial when it comes to implementation in artificial systems. Generally, artificial systems avoid this problem, either by dealing with the imitation of a single body part (i.e., arm) or by implying a similar physical morphology between the demonstrator and imitator (Billard and Mataric, 2001). Such a built-in capacity is face detection, based on a direct mapping between the organ representation and the corresponding motor control area (Breazzal and Scassellati, 2002).

With respect to limb action identification, the recent discovery of the *mirror neural system* in the monkey's premotor cortex (Rizzolatti et al., 1999) has been proposed as the system responsible for the linkage of self-generated and observed actions (Arbib, 2002). The interpretation is that mirror neurons can allow the matching of the neural command for an action with the neural code for the recognition of the same action executed by another individual. Mirror neurons have been observed for reaching and grasping actions, and they are highly specialized for certain types of movements. It is possible that these neuron properties are innate, similar to face detection capacity, hence they may explain, up to a point, the neonatal imitation of gestures. Thus, they can facilitate recognition of hand manipulation and may be involved in mapping the hand sight into the hand self-motion.

However, even if mirror neurons functionality has been recently incorporated in several imitation modeling proposals (Billard and Mataric, 2001; Maistros and Hayes, 2002; Metta and Fitzpatrick, 2002), few attempts have been made, so far, to understand the way they develop such a highly specialized *matching* property (Arbib, 2001). We believe that exploring the way mirror neurons' functionality emerges can provide insights into their 'true' role in imitation and language development. We consider that an improved version of our visuomotor model, which is already capable of showing emergence of multi-modal neural behaviors and to give rise to different dynamics of the neural response, is in a good position to explore this topic.

With respect to the *movement component* of the imitative act, most researchers agree that it is not innately specified, but up to the present it is not known yet how this mapping is achieved with such specificity. Here, it is important to delimit the existence of two developmental stages in imitation. First is neonatal imitation, which mainly consists of facial gesture imitation and is probably supported by a subcortical system (Atkinson, 2000). This is followed by the emergence of a true imitative behavior, which occurs after a few months of post-natal life and is marked by the acquisition of eye-hand coordination (i.e., at about 3-4 months) (Butterworth, 1999; Atkinson, 2000).

Our view is that imitation in the former case can be best explained in the terms of dynamical systems theory. From this perspective, development is self-organizing around various attractors on which the configurations of the system tends to converge (Butterworth, 1999). From this view, the prenatal movements provide experience to link muscles activations to body configurations. In our terms (i.e., the terms of the progressive match architecture), this stage corresponds to learning in the first combinatorial domain, where co-activation of mo-

tor commands and organ configuration (i.e., muscles activation) establish the foundation of proprioceptive control of movement. What this stage does, it creates a set of attractors in the movement–somatosensory space.

Further, seeing the protrusion of the tongue between lips means, in terms of object identification, that both lips and tongues visual areas are activated. These are further mapped to the corresponding motor areas. Here, the oscillatory activity in the two areas can only evolve towards one of the existent attractors. One attractor, which comprises neural activity for both lips and tongue movement, is the one that places the tongue between lips. Note that, the other imitated behaviors, such as mouth opening or eye–movements are even simpler, as the activity here involves only one area and elicits a limited number of possible actions. The search in the space of possible behaviors (i.e., see above the effortful search of the correct action) is equivalent to the formation of the desired action attractor. We can compare this process, with the formation in the motor network of the directional response in the presence of visual stimulation. The task in the case of neonatal imitation is more difficult, because it does not appear to involve any learning of the visuomotor connection, but is actually learnt on–line, resulting in the convergent correction of the movement. This is possible, we believe, due to the combinatorial properties of the neurons, which allow proprioceptive, motor and visual information to be matched.

In the latter case (i.e., after 3 month of postnatal life), imitation represents the result of a self-organizing learning process. The beginning of true imitative behavior accompanies the emergence of eye-hand coordination. This suggests that both processes require the development of equivalent neurobiological mechanisms. Hand–eye coordination begins to develop between 2-4 months, inaugurating a period of trial–and–error practice at sighting objects and grabbing at them. When executing actions, infants perceive and learn contingencies between the motor activity and the visual image of the movements. Our work focused upon the modeling of this behavioral scenario in order to develop visuomotor coordination. We believe that the operation of the developed computational mechanism can establish the foundations for imitative behavior.