Chapter 1

Introduction

Major advances in science often consist in discovering how macro–scale phenomena reduce to their miscroscale constituents. The 'astonishing hypothesis' is that our minds can be explained by understanding the detailed behavior of neurons in the brain and their interactions with each other (Crick, 1994). Brains are collections of billions of interconnected cells, each of them being an individual machinery which receives, process and transmits information (Kandel et al., 2000). The human brain generates complex patterns of behavior at different scales of organization based on the large amount of neural components that interact simultaneously in a rich number of parallel ways. In order to understand the inherent complexity of such a system many complementary research strategies are employed. Today, in this attempt, experimental and theoretical neuroscience studies are accompanied by mathematical theories of complex systems organization, evolutionary and developmental studies of the brain organization, all assisted by computational modeling means.

The work presented in this thesis lies at the intersection of three domains: neuroscience, artificial intelligence and developmental psychology. This work makes use of computer simulations with networks of neuron–like elements in order to understand how cognitive phenomena can be grounded at the neural level. In this introduction, we motivate our approach and we present the problem statement. In Section 1, a short overview of several research strategies of brain function is provided. We believe that our methods are best described as computational neuroscience. An introduction to this field is given, accompanied by a discussion of the advantages of computer modeling, as a methodology. The computational approach is based on information processing with spiking neurons, which are motivated in Section 1.1.2. In Section 2 we propose a number of challenging problems in motor control, which are investigated in this thesis through biologically inspired modeling. The specific objectives of our work are introduced in Section 3. Finally, Section 4 outlines the structure of the thesis.

1.1 What is Computational Neuroscience?

Computational neuroscience is an evolving approach that draws on neurobiological data, but uses computational modeling and computer simulations to investigate the principles of operation governing neurons and networks of neurons (see Bower and Beeman, 1998; De Schutter, 2001). The domain is rapidly growing, involving researchers with backgrounds ranging from psychology and cellular biology to artificial intelligence. What is new about this discipline and how does it relate to other research techniques of the nervous system?

Firstly, brain function can be investigated with dedicated tools from experimental neuroscience. Using imaging techniques (PET, fMRI), intra- and extra-cellular recording, patch clamp and neuron staining techniques promises to allow researchers to visualize the 'brain in action' (Toga and Mazziotta, 1996; Frackowiak et al., 1997; Kandel et al., 2000). Today, this experimental technology has advanced to the point that biological information can be readily obtained, hence, this knowledge has accumulated and generated huge databases of neurophysiological information (Bower and Beeman, 1998). Supported by advances in computer software and hardware technology, experimental biologists oriented towards the test of mathematical models and the simulation of neurobiological details (see Koch and Segev, 1998). This research direction has impelled the emergence of the computational neuroscience field, where it advocates the use of structurally realistic neural models in the simulation of the brain phenomena (Bower and Beeman, 1998).

The main goal of neuroscience, since its early beginnings, has been to bridge the gap between the behavior of single neurons and complex behavior emerging from their cooperative function. The theory of complex systems is specifically aimed at understanding this issue. The dynamics of complex systems are characterized by local, spontaneous ordering tendencies, which can lead to global self–organized states. The brain is such a self-organizing system and the modeling of how it can learn by itself represents a successful paradigm in the brain research (Kelso, 1988; Kohonen, 1995).

Dynamical systems theory is also increasingly exploited as a means of understanding brain function, both at a neural and cognitive level. At the neural level, a number of researchers have argued that the dynamical properties of firing neurons may have a central role to play in explaining how the brain computes (i.e., synchronous oscillations may play a crucial role in cognitive binding, Singer, 1994). At the cognitive level, a recent view is that cognitive processes are behavioral patterns of non-linear dynamical systems and are best studied using the mathematics of dynamical modeling and dynamical systems theory (Kelso, 1988; Port and Van Gelder, 1995). This paradigm has recently challenged the computational perspective of the brain function (Van Gelder, 1995). Despite the current debate of whether the brain *computes* or *integrates* (Chauvet, 2002), the dynamical systems theory remains a valuable, generally accepted mathematical approach for analyzing patterns of behavior generated at different scales of organization.

To understand brain function it is important to consider that it is the end product of an evolutionary and developmental process. With regard to brain evolution from a simple to a more complex system, a viable approach is to consider that it has developed incrementally over evolutionary time (Crick, 1994). An incremental approach will tend to favor the addition of small specialists modules rather than a re–engineering of the entire system (Crick, 1994). For instance, taking this view, what evolution has done in constructing our sensory systems was to provide additional sources of constraint on the possible identity of an object (Reilly, 2001) or to facilitate the apparition of a new response to the environment (Goodale, 2000).

Compared to the evolutionary perspective, the developmental study of the brain is more within our grasp. It represents a primary methodology for the analysis and the understanding of the evolution of human cognition and behavior during a lifetime (see the developmental psychology and the epigenetic approach, Piaget, 1969). It also represents a fruitful paradigm in which artificial systems can be constructed to develop by their own means, complex behaviors based on simpler components (Zlatev and Balkenius, 2001; Weng et al., 2001). A developmental approach provides a structured decomposition of complex tasks. It divides high–level processes into computationally simple, and developmentally earlier, behaviors (Scassellati, 1998). From this perspective, language and cognition may be understood through the incremental development of more sophisticated structures on the foundations of preexisting low–level, sensory–motor programs (Reilly, in press).

To conclude, in order to understand the complexity of brain processes, several research disciplines are employed in a complementary manner. The new field of computational neuroscience is rooted in neuroscience, by drawing on biological data about information processing in the cells of the brain. In this attempt, the mathematical means of dynamical systems theory is valuable and can be employed for the analysis of behavioral patterns emerged. Regarding the developmental approach, in this thesis it is believed that only by proceeding in an incremental manner can we overcome the complexity issues facing any attempt to model the brain. The core methodology of computational neuroscience is represented by means of computer modeling. The advantages of cognitive modeling in general and those of simulation with functionally realistic neurons, in particular, are outlined below.

1.1.1 Why cognitive modeling?

Experimental investigations in neuroscience can provide a detailed characterization of the chemical processes that underlie cognitive processes, but what matters for most scholars in cognitive science are not the details themselves, but the principles that are embodied in these details (McClelland, 2000). What is essential about the computational models is that they enable researchers to explore the nature of these principles, by implementing on a computer the underlying mechanisms (Levine, 2000; Rolls and Treves, 1998; O'Reilly and Munakata, 2000).

Furthermore, it is important to view the computational modeling of the brain processes from the *reconstructionist paradigm*. This concentrates on the process of constructing human cognition from the action of a large number of interacting components (McClelland, 2000). The emphasis is on investigating the emergent phenomena which arise from these interactions, and which is not obviously present in the behavior of individual elements (Cleeremans and French, 1996; O'Reilly and Munakata, 2000).

Computational modeling allows manipulation and control of variables more precisely than can be done with a real system. This enables the researcher to explore the causal roles of different components (see the modeling of dendrites role in auditory coincidence detection in Agmon-Snir et al., 1998). Computer modeling allows testing of hypothesis, but it is also a powerful means of generating new and original hypothesis. A computational model can provide novel sources of insight into behavior by providing alternative explanations of a phenomena (see the new hypothesis on pattern recognition in a Purkinje cell in Steuber and De Schutter (in press)). Finally, when creating a computational model, one has to be explicit about assumptions and about exactly how the processes work (O'Reilly and Munakata, 2000).

Computational neuroscience differs from previous approaches to neural modeling (see the connectionist paradigm, Rumelhart and McClelland, 1986) in that, researchers in the area believe that understanding the way the brain computes is very closely dependent on the anatomical and physiological details of the neural components. Hence, it focuses upon the simulation of the cortical functions by using neural models that are in agreement with the structure and physiology of real neurons.

1.1.2 Why spiking neurons?

Cognitive models developed in the previous decade, dominated by the connectionist psychology paradigm (Rumelhart and McClelland, 1986) have neglected the spiking nature of the neurons. These models have been usually constructed based on a neural model with sigmoid activation function that gives a continuous, real-valued output. Networks of such neurons have proven very powerful computationally (O'Reilly, 2001). Moreover, they have been supported biologically by the long-standing belief in neuroscience that information in the brain is carried mainly in the neurons discharge rates (for a review see Recce, 1999).

Recent observations of how fast computations take place in the visual brain has questioned whether the firing rate interpretation alone can account for rapid neural information processing (Thorpe and Gautrais, 1997). Experimental evidence has been accumulated in the last years to indicate that biological neural systems use the timing of single action potentials to encode information (Abeles et al., 1993; Gerstner et al., 1996; Rieke et al., 1997). Consequently, learning that information can be encoded in the temporal pattern of neuron firing, the research on information processing in neural systems has focussed on investigating computations with *spiking neurons* (Gerstner and Van Hemmen, 1994; Maass, 1995; Rieke et al., 1997; Stevens and Zador, 1998).

This research stream has given rise to a new generation of neural networks, referred to as *pulsed neural networks* and has focused on providing a mathematical description of the computational properties of biological neurons (Gerstner, 1999; Maass, 1999). Several alternative theories to the rate–coding hypothesis have been proposed, which suggest different schemes for where the neural information may be contained, such as in the timing of the spikes, and in the correlated activity of neurons (Stevens and Zador, 1995; Recce, 1999).

The computational work of this thesis is based on a simplified spiking neural model that focuses upon several aspects of the biological neuron function. Compared to the classical rate–coding neuron, this model accounts for the spiking nature of real cells and allows different modes of computation and learning. On the other hand, in comparison with structurally realistic neural models, it accounts only for a limited number of computational aspects, with the advantage of being simple to analyze and to simulate in a large–scale network of neurons.

The development of the new generation of spiking neural networks has necessitated the design of dedicated simulation environments. Several simulation frameworks for the realistic modeling of biological neurons (Hines and Carnevale, 1995; Bower and Beeman, 1998) or for computation with simplified models (Delorme et al., 1999; Sougné, 1999) have been created in the last decade. In our case, the choice of a modeling environment was motivated primarily by the ability of the system to support the development of a family of models required by current and future work goals. What seemed to be the best solution was to use a general-purpose simulator, regardless of the neural model implemented, which was extended to allow computations and learning with spiking neurons. The simulator chosen is a classic neural networks simulator, SNNS: Stuttgart Neural Network Simulator (Zell et al., 1992), whose extension for modeling of spiking neurons is referred to, in this thesis, as SpikeNNS.

To conclude, this thesis advocates the use of the computer simulation methodology in order to investigate computational principles of the cortex. Our work is aimed at providing an illustration of how cognitive brain functions can be grounded at the neural level. In doing this, we employ simulations with biologically inspired neural models. The general aim is to explore within a developmental approach how emergent properties of the brain are high–level effects that depend on low–level computational properties of the basic constituents. This goal is investigated with a number of models of specific cognitive phenomena.

1.2 Identifying difficult cognitive problems

1.2.1 Neonatal imitation

Imitation is an essential behavior for cognitive development in infants, because it serves communication and rapid acquisition of adaptive behavior and is an alternative to expensive trial–and–error learning (Butterworth, 1999). Meltzoff and Moore (1977) have shown that infants as young as 12–days–old can imitate facial actions of caregivers, such as tongue protrusion, mouth opening and manual gestures. Their findings served as an essential testimony to the existence of imitation in newborns and suggested the possibility that immediate imitation is a fundamental mechanism of communication in humans (Nadel and Butterworth, 1999). Since the publication of their original results, a large number of studies have been issued to investigate the scale of the neonate imitation phenomena. A range of facial expressions and hand gestures, as well as deferred imitation capacity, have been described with respect to the age of onset and generalization capacity (Kugiumutzakis, 1999; Meltzoff and Moore, 1999).

The 'holy grail' for the theories of neonatal imitation has been to elucidate the mechanisms whereby infants are able to connect the felt but unseen movements of the self with the seen but unfelt movements of the others (Butterworth, 1999). This process is considered to require an inter–modal mapping, that is, a transfer of information between different perceptive modalities (i.e., visual and somatosensory) in order to control the imitative acts. A common belief is that proprioceptive feedback on self-produced movements can be compared to the visually specified target in a supra-modal (Metlzoff and Moore, 1999) or amodal framework (Trevarthen et al., 1999).

To us, neonatal imitation behavior represents the challenging problem. At a deeper analysis, it appeared to be the 'tip of the iceberg' of a more general problem of brain computation: sensory fusion or cross–modal matching of information. That is, the basic computational demand for imitation is met by the transfer of information between different modalities. In this way, it led us to the fundamental topics of perception–action coupling and sensorimotor coordination.

1.2.2 Visuomotor coordination

Most human and animal movements are under continual sensory guidance. Even a simple task such as reaching and grasping an object, requires the analysis of visual information to trigger a reaching movement and the integration of proprioceptive information and tactile sensation to tune the grip to the weight, friction, and shape of the object (Kandel et al., 2000). That is, purposeful action is possible through the integration of sensory signals coming from various sources (vision, hearing, touch) and their translation into a set of motor commands to the muscles (Massone, 1995; Kandel et al., 2000).

Consequently, significant efforts have been dedicated in the last decades to the understanding of the computational mechanisms that support sensorimotor coordination in neural systems. Particular attention has been given to the development of coordination between the eye and the hand for reaching movements (see Caminiti et al., 1992). One reason is represented by the huge applicability, that a mechanism which allows the correct transformation of visual signals into motor output, would have in the implementation of artificial systems capable of online, adaptive motor control (Zeller et al., 1995; Weng et al., 2001). Another motivation is that by understanding the brain computations that underlie reaching movements, it will be possible for humans to control a robotic arm solely through the power of thought (Nicolelis and Chapin, 2002).

Despite the efforts of classical artificial intelligence solely, it seems that the behavioral capabilities of biological organisms can be simulated only by closely reproducing neural computational mechanisms. Consequently, several recent proposals have been made for a more biologically inspired modeling of visuomotor coordination (Burnod et al., 1992; Bullock et al., 1995; Zeller et al., 1995). These attempts are facilitated by the development of a new conceptual scheme of cortical control of reaching movements.

Research on the visual guidance of arm–reaching movement has made significant progress in the recent years, in explaining control theory formalisms such as the 'coordinate transformation' (i.e., sensorimotor transformation), in terms of computational properties of single neurons and networks of neurons (Andersen et al., 1997; Kalaska et al., 1997; Caminiti et al., 1998; Burnod et al., 1999). Details regarding the manner in which the brain solves the sensorimotor mapping have cumulated and have given rise to an integrative framework that links neurophysiological and computational aspects and allows the ready implementation of the latter in terms of the former (Bullock et al., 1995; Burnod et al., 1999).

What is still needed, are models whose architecture is supported by anatomical evidence, composed of elements that correspond as closely as possible to known neural cell types, and whose functionality meet psychophysical criteria (Bullock et al., 1995). It is this challenge that motivates part of the work presented within this thesis. The advantage of modeling the neurophysiological processes involved in the cortical control of reaching is bi-directional. Firstly, more adaptive and flexible models can be obtained. Secondly, models can be used to test computational principles and eventually to reveal unknown mechanisms of visuomotor control of movement.

1.2.3 Cortical control of motion direction

The human motor system is organized in a functional hierarchy, with each level concerned with different decisions (Kandel et al., 2000). Voluntary movements are organized at the highest level of the frontal cortical lobes, in the premotor and primary motor cortex. These areas are involved in the preparation, execution and adaptation of movement. Movements of the arm, such as reaching or grasping, involve multiple joints and require precise activation of the skeletal muscles. This raises the question of whether cells at the cortical level control muscle activation or do they encode more global features of the movement, such as direction, amplitude, or speed (Kakei et al., 1999). Today, there is substantial evidence that the direction of movement is represented at the cortical level, in the activity of large populations of cells that are broadly selective to the direction of motion (Georgopoulos et al., 1984).

Feature detection is not only a characteristic of motor cells, but represents an essential property of neurons in the brain. It means that neurons respond to a particular feature of the stimulus (i.e., orientation, direction, frequency of signal) and specialize in detecting a range of values of that feature in the input space (Hubel and Wiesel, 1962). The dedication of neurons is very widely distributed in the brain. For instance, directional selectivity has been described for neurons in the visual, motor, and touch cortex (Kandel et al., 2000). The specialization of the neurons generally occurs by the way each cell is connected with other cells, from the input layers or within the same layer, and it results in the emergence of cortical feature maps (see orientation visual maps in Blasdel, 1992). The development of these maps can be modeled through a process of self-organization and topographical mapping of the input space into the network nodes (Kohonen, 1984).

Despite the substantial evidence indicating that directional tuning is an essential feature of the motor cortical neurons, it is not clear yet how the cortical control of motion direction is developed. This contrasts with the detailed anatomical and computational knowledge existent on the development of sensory, and in particular, visual maps (Obermayer et al., 1990; Douglas et al., 1991; Blasdel, 1992; Sirosh, 1995). Nevertheless, it is believed that the development of visual preferences may be based on a few design principles that in turn rely on very general mechanisms utilizing the input structure of the system (Niebur and Wörgötter, 1992).

Part of the work presented in this thesis concerns the simulation of the process whereby neurons in the motor areas develop directional selectivity. Firstly, modeling this process can offer insights into whether directional tuning of motor cells is innate or acquired. Secondly, the investigation of the functional principles of motor cortex can show in what degree the organizational mechanisms of the sensory cortices also operate in the motor areas. Thirdly, understanding the way the motor cortex organizes itself to control the direction of motion is highly important for explaining the development of sensorimotor coordination as a whole. Furthermore, current models of motor control which implement the control of end–effector direction, lack a developmental model of how neurons at the cortical level acquire directional selectivity (Bullock et al., 1993). Hence, we believe that the modeling of the formation of directional motor map may have important implications for neurophysiology and robotics.

1.3 The problem statement

Most generally, the research presented in this thesis is aimed at investigating how cognitive functions in the brain can emerge from the properties of basic components, when these interact and function cooperatively. This objective is narrowed down to the study of two topics: cortical control of movement direction and visuomotor mapping of directional information. Our detailed research objectives are:

- Modeling of the self–organization of motor cortical neurons for coding the direction of movement.
- Modeling of the alignment of visual and motor neural representations for the guidance of directional movements.

Note that these objectives are related. That is, by achieving the self–organization of the motor directional map, the organized network can be used in the next stage for the development of visuomotor mapping.

The original contribution of our models is represented primarily, by the investigation of the two topics described above within the biologically inspired computational framework, of spiking neural networks. Secondly, the motor directional map model represents a first attempt to simulate the emergence of cortical directional selectivity within the self–organization paradigm (Kohonen, 1984). The resulting representation of movement will be compared to the neurophysiological data that describes the coding of direction of movement in the motor cortex (Georgopoulos et al., 1984; Georgopoulos et al., 1993). If the model succeeds in developing neural directional selectivity with a similar profile to that of real motor cells, than we have a computational hypothesis for how the population coding emerges in the motor cortex. Learning the visuomotor mapping of directional information in the conditions of a population coding in the motor area represents another innovative feature of our modeling work. This has computational and neurophysiological relevance for the learning of visuomotor mapping.

In order to achieve these main objectives, they are preceded by the design of a modeling environment for networks of spiking neurons. When modeling large–scale pulsed neural networks with plastic synapses, the time efficiency of the simulation becomes an essential issues in the design of the simulator (Jahnke et al., 1999). A number of strategies are implemented and compared and an innovative event–driven mechanism is proposed to reduce the time of simulation.

1.4 Thesis outline

The structure of this thesis is organized as follows.

Chapter 2 provides a biological framework of cortical control of motion direction, for use in our modeling work. It introduces the self-organization paradigm for biological modeling of cortical feature map formation. Finally, it reviews several recent neural network models that address issues of cortical coding of motion direction.

Chapter 3 introduces a number of new theories on the biological and computational mechanisms of perception–action coupling. Next, it focuses upon the description of an integrative framework of how the visual guidance of arm reaching is implemented in the brain. A review of several biologically inspired models of learning of visuomotor coordination is presented.

Chapter 4 focuses upon the description of the implemented spiking neural model. It presents how information can be communicated in the timing of single spikes, what types of computations are implemented and at what level of detail. Learning with spiking neurons is also discussed.

Chapter 5 presents the implementation of the SpikeNNS simulator. Design issues are discussed, with particular attention given to the strategies used to increase the time–efficiency of the simulation. The configurable features of the simulator are outlined.

Chapter 6 describes in turn, the model of the motor cortex self-organization and that of visuomotor mapping learning. The results of the simulations are analyzed and discussed in comparison with current neurophysiological data and with previous modeling work.

Chapter 7 is devoted to a final discussion of the neurophysiological and theoretical implications of our models. It also proposes the potential integration of the models within artificial motor control systems and it proposes possible avenues of future work. The relevance of our visuomotor mapping model to the imitation problem is discussed in the end.