A Novel Computing Architecture for Cognitive Systems based on the Laminar Microcircuitry of the Neocortex – the COLAMN project

Michael Denham
Centre for Theoretical and Computational Neuroscience
University of Plymouth, UK
mdenham@plym.ac.uk

Abstract

Understanding the neocortical neural architecture and circuitry in the brain that subserves our perceptual and cognitive abilities will be an important component of a "Grand Challenge" which aims at an understanding of the architecture of mind and brain. We have recently embarked on a new five-year collaborative research programme, the primary aim of which is to build a computational model of minimal complexity that captures the fundamental information processing properties of the laminar microcircuitry of the primary visual area of neocortex. Specifically the properties we aim to capture are those of self-organisation, adaptation, and plasticity, which would enable the model to: (i). develop feature selective neuronal properties and cortical preference maps in response to a combination of intrinsic, spontaneously-generated activity and complex naturalistic external stimuli; and (ii) display experience-dependent and adaptation-induced plasticity, which optimally modifies the feature selectivity properties and preference maps in response to naturalistic stimuli. The second aim of the research programme is to investigate the feasibility of designing VLSI circuitry which would be capable of realising the computational model, and thus demonstrate that the model can form the basis for a novel computational architecture with the same properties of self-organisation, adaptation, and plasticity as those displayed by the biological system. A basic premise of the research programme is that the neocortex is organised in a fairly stereotyped and modular form, and that in this form it subserves a wide range of perceptual and cognitive tasks. In principle, this will allow the novel computational architecture also to have wide application in the area of cognitive systems.

.

1 Introduction

The neocortex of the brain subserves sensory perception, attention, memory and a spectrum of other perceptual and cognitive functions, which combine to provide the biological system with its outstanding powers. It is clear that the brain carries out information processing in a fundamentally different way to today's conventional computers. The computational architecture of the brain involves the use of highly parallel, asynchronous, nonlinear and adaptive dynamical systems, namely the laminar microcircuits of the neocortex. The neurons which make up a neocortical microcircuit (Silberberg et al, 2002; Mountcastle, 1997) are precisely connected to each other and to their afferent inputs through synapses in specific layers of the laminar cortical architecture, and on specific locations on their dendritic trees Thomson and Bannister 2003; Callaway, 1998). Each synapse acts as a unique adaptive filter for the transmission of data into the circuit and between pairs of cells. Thus whilst a single neuron may connect to many hundreds of other neurons, a signal sent by one neuron will be interpreted by each target neuron in a unique way. Furthermore, these connections are not static but change their transmission characteristics dynamically and asynchronously, on a millisecond timescale, partly determined by their highly precise spatial location in the dendritic tree (Häusser et al, 2003) but also in relation to the function of the different neuronal types that they connect. In addition, both the synaptic connections and the transmission properties of the dendritic tree have the remarkable ability to continuously adapt and optimise themselves to meet the requirements of novel tasks and environments. This takes place both through unsupervised, self-organising modification of their dynamic parameters, and through optimisation of the synaptic and dendritic dynamics by specific adaptation-induced and experience-dependent plasticity mechanisms.

Capturing the fundamental information processing properties of the laminar microcircuitry of the neocortex in the form of a computer model could provide the foundation for a radical new generation of machines that have human-like performance in perceptual and cognitive tasks. Such machines would be capable of using self-organisation, adaptation, and plasticity mechanisms which are inherent in the neocortex, in order to deal with complex, uncertain and dynamically changing information. They would potentially be much more powerful, require minimal programming intervention, and be resilient to failures and errors. Creating the necessary understanding of these properties of the neocortex, expressing them as a computational model of minimal complexity, and translating this model into the design of a computer architecture capable of realisation in VLSI, will require the collaborative efforts of neuroscientists, computer scientists, mathematicians, and engineers.

2 The aims of the research programme

The aim of this research programme is to create a new “brain-inspired” computational architecture which possesses the basic properties of self-organisation, adaptation and plasticity manifest in the laminar neural microcircuitry of the neocortex. The principal objective is a functional model of a “stereotypical” cortical microcircuit which captures these basic properties of the neocortex, and provides the basis for the design of a novel, modular computational architecture capable of realisation in a combination of analogue and digital VLSI circuits. The ultimate goal of this avenue of research is a “brain-inspired” architecture which will deliver human-like levels of performance for a wide range of perceptual and cognitive tasks, and deal with all sensory modalities.

This goal is well beyond the scope of the currently envisaged research programme; however, as a first step towards this goal, the programme will aim at capturing the fundamental properties of self-organisation, adaptation and plasticity of the neuronal circuitry in the primary visual area of the mammalian neocortex. This will allow us to build on the wealth of current neurobiological knowledge concerning the properties and interconnectivity of neurons and the behaviour of local and long-range neuronal circuitry in this area of neocortex in response to visual stimuli. It must be stressed that our aim is not to build a detailed, biologically-precise model of neocortex, but rather it is to identify and capture in a minimally complex model these key fundamental properties that underlie its remarkable information processing capabilities.

The specific aims of the proposed research programmme can therefore be summarised as follows:

1. To build a computational model of minimal complexity that captures the fundamental information processing properties of the laminar microcircuitry of the primary visual area of neocortex. Specifically the properties we aim to capture are those of self-organisation, adaptation, and plasticity, which would enable the model to:
   i. develop feature selective neuronal properties and cortical preference maps in response to a combination of intrinsic, spontaneously-generated activity and complex naturalistic external stimuli, and
   ii. display experience-dependent and adaptation-induced plasticity, which optimally modifies the feature selectivity properties and preference maps in response to naturalistic stimuli.
2. To investigate the feasibility of designing VLSI circuitry which would be capable of realising the computational model, and thus demonstrate that the model can form the basis for a novel computational architecture with the same properties of self-organisation, adaptation, and plasticity.

3 The research programme

The research programme involves a high level of integration of activities in neurobiological modelling, experimental neurobiology and the VLSI circuit design. It is organised into a set of such activities, each of which addresses a well-defined aim of the research programme, as described below.

3.1 Novel neocortical neuron and circuit connectivity models

A basic premise of the research programme is that the neocortex is organised in a fairly stereotyped and highly modular fashion. Although much is already known about the structure and functional connectivity of microcircuits in the neocortex, the current state of knowledge is only sufficient to inform the initial design and construction of the proposed computational model. Recent work eg Thomson and Bannister (2003), has contributed important and detailed insights into the synaptic connectivity and the dynamic and plastic aspects of information transmission along these synaptic connections within a cortical column. The research will draw on this and on further, on-going work in order to more fully elucidate the neuronal and synaptic connectivity which it is necessary to capture within the computational model in order to endow it with the self-organisation, adaptation and plasticity properties of cortical microcircuits.

In particular, the behaviour of circuit models of spiking neurons strongly depends on the properties of their constituents, the individual neurons, as well as on the synaptic connectivity between them. For example, the phase diagrams describing the dynamics of sparsely connected networks of excitatory and inhibitory neurons, which can exhibit different synchronous and asynchronous states (Brunel, 2000) change fundamentally when current-based integrate-and-fire neurons are replaced by conductance-based integrate-and-fire neurons as the constituents of the network. In order to provide the cortical modelling and the VLSI designs with the best possible description of single neurons (in terms of both accuracy and computational efficiency), the plan is to construct new types of integrate-and-fire neuron models that represent the biophysical mechanisms operating in biological neurons in a more realistic way, including the role of neuronal dendrites in the transformation of synaptic input into spike output.

Models will be validated by direct comparison with experimental data from experiments in brain slices and in the intact animal in vivo which describe the input-output relation of different types of neurons both at a functional, eg Chadderton et al., 2004, and a biophysical level, eg Häusser et al., 2001. We will focus on those characteristics of real neurons that are currently not, or only with insufficient accuracy, captured by the integrate-and-fire or spike response models currently available. We expect that models including the subthreshold dynamics of voltage-dependent conductances, including oscillatory behaviour, as well as the shunt conductances associated with action potential firing, which provide only a partial reset of the membrane potential in the neuron, will lead to more realistic yet compact descriptions of the input-output relations of different types of cortical neurons.

Single-neuron models will be complemented by three-dimensional geometric models of synaptic connectivity based on anatomical and physiological data from a large dataset of anatomically and physiologically identified, synaptically connected neurons which are being generated in a number of laboratories. Together these will provide an intra- and interlaminar wiring diagram of the cortical microcircuit. The functional properties of the synaptic connections between different types of neurons will be described by statistical distributions of the amplitudes and time courses of the synaptic conductances, including a representation of short- and long-term synaptic plasticity.

3.2 Functional analysis and modelling of the neocortical microcircuit

It will be essential to provide constraints for the proposed computational model. This is a non-trivial but essential task if we are to ensure that the modelling work does not result in "parameter explosion". In particular, it will be necessary to constrain the model on the basis of the functional properties of the neurons and their interconnectivity in the cortical microcircuit. *In vivo*, cortical cells receive input from several thousand synaptic connections simultaneously, and only a proportion of these are connections from other cells within the cortical microcircuit. Some aspects of the intra-columnar connectivity revealed by intracellular recordings will form an essential part of the function of the cortical column, while other aspects are unimportant details that are best ignored in the proposed computational model. Constraining the model therefore means deciding which aspects are important, and estimating the strength of their contribution relative to other, external inputs and influences. This will require combining new extracellular recording techniques and novel statistical analysis and modelling approaches. Silicon array electrode techniques make it possible to record spiking activity simultaneously from dozens of neurons throughout a cortical microcircuit, in the living brain while it is carrying out its natural information processing tasks. In the past, simple filter models have been used to predict responses of individual neurons in sensory cortex (Schnupp et al, 2001). The research programme will aim at a dramatic improvement in these simplistic models through the use of novel statistical modelling techniques.

3.3 Learning rules for the development of stable self-organised feature selectivity

The role of self-organisation in the stimulus-dependent development of orientation selectivity was first suggested by von der Malsburg and recently reviewed by Miller et al (1999) and Sur and Leamey, 2001. The latter suggest that spontaneous patterns of neural activity in the absence of visual stimuli may be sufficient in the early periods of development, after the initial cortical circuitry has been established, for the early development of orientation selectivity, but that the formation of orientation selectivity is strongly influenced by input activity to the developing cortex (Sur and Leamey, 2001; Sur et al, 1988). Experiments show that input activity has an influence on synaptic connections in the cortical circuitry which gives rise to orientation map development and long-range horizontal intracortical connections in layers 2/3. A cortical microcircuit model would thus need to embody development of the dynamical interactions provided by intracortical connections in an activity-instructed self-organising process of map development. As yet, it would appear that no biological models exist which implement this activity-dependent self-organising process of development.

It has been demonstrated experimentally that the self-organised modification of synapses depends on the precise timing of spikes, causing the neuron to evolve in such a way as to be driven by its fastest and most reliable inputs. Therefore it seems reasonable to hypothesise that the learning rules which govern the self-organised emergence of cortical orientation selectivity should yield populations of selective cells, large enough to perform fast and reliable computation, yet small enough to be efficient. The investigation of these issues will lead to an understanding of how learning rules can self-organise the synaptic interconnectivity in the cortical microcircuit to produce a stable, sparse coded orientation selective network. An important component of this work will be to investigate how the stability of feature selectivity might be helped by recurrent interactions between neurons. These connections could break the symmetry and stabilise the synaptic weights, improving the stability of feature selectivity. The precise details of the learning rules are expected to be of crucial importance for the final selectivity patterns learned. This holds for both rate based as for spike timing dependent learning rules (van Rossum et al, 2000).

3.4 Neural coding of feature selectivity properties of cortical circuits

Intimately related to the investigation of developmental self-organisation learning rules is the question of neural coding, i.e. of how neuronal populations represent sensory information. This is even more evident in the case of spike timing dependent learning rules. For instance, if stimulus features are coded across the cortical microcircuit by either precise spike times of individual neurons or by synchronous neuronal activity across neurons, it is of importance to know how such coding affects learning. Likewise, the developmental learning rules which result in specific patterns of synaptic connectivity have to support the selective neural coding of the stimulus feature set. A major objective of this part of the research programme will be to understand what advantages the laminar architecture of the neocortex offers in terms of efficiency of information representation.

By using mathematical analysis techniques based on the principles of information theory, the role of columnar organization in cortical information representation has recently been investigated (Panzeri et al, 2003), but it is clear that the laminar organisation can provide both advantages and constraints that are as important. It has been shown that real cortical neurons encode information by timing of individual spikes with millisecond precision (Panzeri et al, 2001a) and investigated what mechanisms are need to read out this information, eg dendritic processing must be important to decode information if most information is encoded by the “label” of which neuron fired each spike, and not very important if instead neurons can sum up all spikes at the soma and still conserve all information (Panzeri et al, 2003).

Research in this part of the programme will extend these ideas by investigating in detail the information processing capabilities of laminar cortical circuits. In particular, we will determine (i) the “neuronal code” used in different laminae, i.e. which of the features (e.g. spike count, precise spike times, synchronization) characterizing the responses of neuronal populations in different laminae convey the most sensory information (ii) whether the precise synaptic connectivity within the laminar neocortical architecture is to some extent “optimal” for fast information transmission from one neuron/layer to another neuron/layer. By “optimal” we mean that the observed wiring comes close enough in terms of transmitted information to the best possible one. By fast we mean that all of this information must be transmitted by the model synaptic system in time scales as fast as the cortical ones (Panzeri et al, 2001b).

3.5 Learning rules for experience-dependent and adaptation-induced plasticity in the developed cortical microcircuit

It is well known that the ability to detect small orientational differences can be significantly improved through training on a visual discrimination task over an extended period of time. This perceptual learning process is also seen to have a long-lasting effect, indicating that it must be the result of some form of long-term synaptic plasticity in the brain. Other characteristics of the learning, which can be psychophysically observed, such as the lack of transfer of the learning from one orientation to the orthogonal orientation or from one learned retinal location to a nearby nonoverlapping location, indicate that the plasticity must involve the primary visual cortex, where the neurons have localised orientation selectivity and small receptive fields. Orientation plasticity has also been demonstrated in response to continuous visual stimulation for a period of seconds to minutes, a process known as adaptation. The results suggest that adaptation-induced orientation plasticity involves changes in circuit connectivity which then define a new preferred orientation. As proposed in the review by Dragoi and Sur (2003), the changes in orientation selectivity following adaptation imply a circuit mechanism that reorganizes responses across a broad range of orientations, and suggest that adaptation-induced orientation plasticity in primary visual cortex is a self-organised emergent property of a local cortical circuitry acting within a non-uniform orientation map. Research in this part of the programme will investigate the learning rules necessary to support the proposed emergence of adaptation-induced modification of orientation selectivity, and whether such learning rules can also support long-term experience-dependent plasticity of orientation selectivity.

3.6 Novel neocortical neuron and circuit connectivity models

Spatiotemporal response properties of neurons in fully developed primary sensory areas are not static but can change on various timescales. Dynamic changes of response properties on long timescales have been assigned to adaptation and plasticity mechanisms. But responses also change on fast time-scales of a few to a few hundred milliseconds revealing rich dynamic features that result from the neural and synaptic activation dynamics and ongoing interactions between neurons within and across cortical microcircuits eg Bringuir et al (1999). Recent experiments eg Ringach et al (2002) indicate even more complex responses of cortical neurons and circuits to naturalistic stimuli. Spatiotemporal responses look similar to those for simple bar or grating stimuli, but there are also significant differences (Ringach et al, 2002). In part these differences seem to be related to influences from outside the classical receptive field: These experiments provide evidence that spatiotemporal response properties of cortical neurons are dynamically shaped in quite intricate ways by intrinsic neuronal and synaptic activation dynamics, interactions between neurons within the microcircuit, and longer ranging synaptic recurrent, feedforward and feedback circuits. These dynamical properties may underlie the surprisingly fast and adaptable information processing within the cortical microcircuit.

A general analytical approach has recently been described (Wennekers, 2002) that relates differently tuned enhanced and suppressed phases in a spatiotemporal response function to feedforward or recurrent pathways between participating cell classes. Although useful for some spatiotemporal phenomena, much of the complexity in real neural responses remains unexplained by such models. Models of complex spatiotemporal phenomena which incorporate the influence of ongoing and spreading activity, or responses to real-world stimuli, are still scarce.

3.7 Feasibility analysis for VLSI circuit design

A major aim of the research programme is to use the computational model of the neocortical laminar microcircuit to define an efficient and implementable VLSI “building block” for a novel computational architecture. The mapping of the model of the cortical microcircuit model into the VLSI circuit design for a novel computational architecture will require the investigation of detailed issues with respect to the numerical accuracy, performance, power consumption and area cost of novel analogue and digital circuit alternatives. These investigations will form the activity of this workpackage. We envisage a structure for the VLSI design based upon an analogue VLSI spiking neural substrate, interconnected via a digital VLSI address-event communication network, all controlled by a software configuration and control system. However, the integration of low-level neural models implemented by analogue VLSI circuits, with digital VLSI for signal routing and communication will need to go far beyond the simple protocols currently used by the neuromorphic engineers, and presents a major challenge. The research will also aim at understanding the implications, in both directions, for including or omitting certain components in the computational model, and assessing the relationship between the levels of description chosen for the computational model and the constraints of VLSI circuit design. In addition, the issues of optimisation (area, power) are present both in the neocortical microcircuit, and in silicon, so some direct analogies on a physical level will be investigated, eg the possible arrangement of the physical layout of devices in a way which is inspired by the 3-dimensional laminar architecture of the cortical microcircuit, in which connections between the cortical microcircuit “building blocks” are predominantly within or between certain layers.

4 Summary

The neocortex of the brain subserves sensory perception, attention, memory and a spectrum of other perceptual and cognitive functions, which combine to provide the biological system with its outstanding powers. It is clear that the brain carries out information processing in a fundamentally different way to today's conventional computers. The computational architecture of the brain clearly involves the use of highly parallel, asynchronous, nonlinear and adaptive dynamical systems, namely the laminar neural microcircuits of the neocortex. The fundamental aim of this research programme is to create a new brain-inspired" computing architecture which possesses the basic properties of self-organisation, adaptation and plasticity manifest in the neural circuitry of the neocortex. The objective is a modular architecture based on a representation of a "stereotypical" cortical microcircuit. The research will focus on the laminar microcircuits of the primary visual cortex in order to build on the wealth of neurobiological knowledge concerning the behaviour and interconnectivity of neurons in this area of neocortex. However the wider objective would be to use the laminar microcircuitry of primary visual cortex as an exemplar for a stereotypical neocortically-inspired architecture. This will allow the architecture to be deployed in a wide range of perceptual tasks, and potentially also in cognitive tasks such as decision making,, with minimal changes to the basic circuitry. The aim is not simply to build a detailed, biologically-precise model of primary visual cortex, but rather the challenge is to identify and capture the key fundamental principles and mechanisms that underlie the remarkable and ubiquitous information processing power of the neocortex.

Acknowledgements

The collaborators in this research programme are: Alex Thomson (School of Pharmacy, London); Michael Häusser and Arnd Roth (UCL); Jan Schnupp (Oxford); Mark van Rossum and David Willshaw (Edinburgh); Stefano Panzeri, Piotr Dudek and Steve Furber (Manchester); Thomas Wennekers and Susan Denham (Plymouth). I would like to fully and gratefully acknowledge their individual contributions to the description of the aims, objectives and research activities which I have set out on their behalf in this paper. Any mistakes herein are however my own. The COLAMN research programme is supported by a grant from the UK Engineering and Physical Research Council under the Novel Computation Initiative.

References

Silberberg G, Gupta A, Markram H (2002) *Trends Neurosci* 25, 227-230

Mountcastle V (1997) *Brain* 120:701-722

Thomson AM, Bannister AP (2003) *Cereb Cortex* 13, 5-14

Callaway EM.(1998) *Annu Rev Neurosci* 21:47-74

Häusser M, Spruston N, Stuart GJ (2000) *Science*, 290:739-744;

Brunel N (2000) *Network* 11:261-280

Chadderton P, Margrie TW, Häusser M (2004) *Nature* 428:856-860

Hausser M, Major G, Stuart GJ (2001) S*cience.* 291:138-41.

Schnupp JW, Mrsic-Flogel TD, King AJ (2001) *Nature* 414:200-204

Miller KD, Erwin E, Kayser A. (1999) *J Neurobiol.* 41:44-57

Sur M, Leamey CA (2001) *Nature Reviews Neurosci.* 2:251-262

Sur M, Garraghty PE, Roe AW (1988) *Science* 242:1437-41

van Rossum M, Bi GQ, Turrigiano GG (2000), *J. Neurosci.* 20, 8812-8821

Panzeri S, Petroni F, Petersen RS, Diamond ME (2003) *Cerebral Cortex* 13: 45-52

Panzeri S, Petersen R, Schultz SR, Lebedev M, Diamond ME (2001a) *Neuron* 29: 769-777

Panzeri S, Rolls ET, Battaglia F, Lavis R (2001b) *Network* 12: 423- 440

Dragoi V, Sur M (2003) In Eds. Chalupa LM and Werner JS, *The Visual Neurosciences*, MIT Press.

Bringuir V, Chavane F, Glaeser L, Fregnac Y (1999) *Science* 283, 695-699

Ringach DL, Hawken M, Shapley R (2002) *J Vision* 2, 12-14

Wennekers T. (2002) *Neural Computation* 14: 1801-25