

Modeling the development of organization for orientation preference in primary visual cortex

Judith S. Law



Doctor of Philosophy
Institute for Adaptive and Neural Computation
School of Informatics
University of Edinburgh

2009

Abstract

The cerebral cortex of mammals comprises a series of topographic maps, forming sensory and motor areas such as those in the visual, auditory, and somatosensory systems. Understanding the rules that govern the development of these maps and how this topographic organization relates to information processing is critical for the understanding of cortical processing and whole brain function.

Previous computational models have shown that topographic maps can develop through a process of self-organization, if spatially localized patches of cortical neurons are activated by particular stimuli. This thesis presents a series of computational models, based on this principle of self-organization, that focus on the development of the map of orientation preference in primary visual cortex (V1). This map is the prototypical example of topographic map development in the brain, and is the most widely studied, however the same self-organizing principles can also apply to maps of many other visual features and maps in many other sensory areas.

Experimental evidence indicates that orientation preference maps in V1 develop in a stable way, with the initial layout determined before eye opening. This constraint is at odds with previous self-organizing models, which have used biologically unfounded ad-hoc methods to obtain robust and reliable development. Such mechanisms inherently lead to instability, by causing massive reorganization over time. The first model presented in this thesis (ALISSOM) shows how ad-hoc methods can be replaced with biologically realistic homeostatic mechanisms that lead to development that is both robust and stable. This model shows for the first time how orientation maps can remain stable despite the massive circuit reconstruction and change in visual inputs occurring during development. This model also highlights the requirements for homeostasis in the developing visual circuit.

A second model shows how this development can occur using circuitry that is consistent with the known wiring in V1, unlike previous models. This new model, LESI, contains Long-range Excitatory and Short-range Inhibitory connections between model neurons. Instead of direct long-range inhibition, it uses di-synaptic inhibition to ensure that when visual stimuli are at high contrast, long-range excitatory connections have an overall inhibitory influence. The results match previous models in the special case of the high-contrast inputs that drive development most strongly, but show how the behavior relates to the underlying circuitry, and also make it possible to explore

effects at a wide range of contrasts.

The final part of this thesis explores the differences between rodents and higher mammals that lead to the lack of topographic organization in rodent species. A lack of organization for orientation also implies local disorder in retinotopy, and analysis of retinotopy data from two-photon calcium imaging in mouse (provided by Tom Mrsic-Flogel, University College London) confirms this hypothesis. A self-organizing model is used to investigate how this disorder can arise via variation in either feed-forward connections to V1 or lateral connections within V1, and how the effects of disorder may vary between species. These results suggest that species with and without topographic maps implement similar visual algorithms differing only in the values of some key parameters, rather than having fundamental differences in architecture.

Together, these results help us understand how and why neurons develop preferences for visual features such as orientation, and how maps of these neurons are formed. The resulting models represent a synthesis of a large body of experimental evidence about V1 anatomy and function, and offer a platform for developing a more complete explanation of cortical function in future work.

Acknowledgements

First and foremost I would like to acknowledge my supervisor Jim Bednar for his invaluable support and encouragement throughout the past four years. Thanks also to Tom Mrsic-Flogel for providing me with interesting data and discussing ideas about the modelling work. I would also like to thank the other members of the Institute for Adaptive and Neural Computation for their advice and discussion, particularly members of the Bednar lab, Chris Palmer, Jan Antolik, Chris Ball and Stuart Wilson. Within the Neuroinformatics DTC I would like to thank Pat Ferguson (who is probably the best administrator in the world!), Mark Van Rossum, David Willshaw and Peggy Seriès.

I would not have been able to get through the past four years without my fellow PhD students, Finlay, Anisha, Andrea, John, Ruth and many others. Also thanks to my friends outside of academia, Lynn, Evelyn, Serena, Lorna, Nicola and Julie; without the regular wine and food nights I would surely have gone insane.

Most importantly I would like to thank my family, especially my mum, dad and nanny who have given me an enormous amount of both financial and emotional support throughout my seemingly endless academic pursuits.

Finally, I would not have been able to complete this work without the love and support from my husband Alan. Marrying you has been the reward for all this hard work.

Judith Kirkwood-Law

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Judith S. Law)

To my nieces and nephew, Jennifer, Matthew and Beth

Table of Contents

1	Background	1
1.1	Introduction	1
1.2	The early visual system	5
1.3	Primary visual cortex	7
1.4	Development of orientation maps	14
1.5	Relationship between topographic organization and adult function . .	17
1.6	Previous models of orientation map development	19
1.7	Summary of aims	24
2	Homeostatic requirements for stable development	27
2.1	Introduction	27
2.2	Methods	30
2.2.1	ALISSOM model architecture	31
2.2.2	Sub-cortical model architecture	33
2.2.3	Afferent connections to V1	34
2.2.4	V1 lateral connections	36
2.2.5	Learning	39
2.2.6	Model comparisons	42
2.2.7	Analysis.	42
2.2.8	Parameter values	44
2.3	Results	47
2.3.1	Orientation map stability in ALISSOM in comparison with previous SOM models	48
2.3.2	Homeostatic regulation of balance.	53
2.3.3	Targets for homeostatic regulation depend on input frequency.	56

2.3.4	Stable orientation map development before and after eye opening.	59
2.4	Discussion	61
3	A model of orientation map development consistent with the known wiring of V1	63
3.1	Introduction	63
3.1.1	Motivation	65
3.2	Methods	70
3.2.1	Simple di-synaptic inhibition	70
3.2.2	LESI model architecture	75
3.2.3	Model parameters	77
3.2.4	Learning	82
3.3	Results	82
3.3.1	Excitatory and inhibitory orientation maps in LESI model . .	83
3.3.2	Receptive fields and connectivity in LESI model	84
3.4	Discussion	85
4	Organization of orientation preference in mouse V1	91
4.1	Introduction	91
4.1.1	Imaging of organization in rodent visual cortex	93
4.1.2	Modeling species differences in organization	97
4.2	Analysis of two-photon calcium imaging of retinotopy in mouse visual cortex	99
4.3	Methods	106
4.3.1	Mouse model architecture	106
4.3.2	Cat model architecture	112
4.3.3	Input patterns and model parameters	116
4.3.4	Retinotopic scatter	116
4.3.5	Analysis	118
4.4	Results	119
4.4.1	Development of ‘salt and pepper’ maps in mouse driven by sparse activation of V1	119
4.4.2	Orientation maps persist despite retinotopic scatter in cat model	130
4.5	Development using natural images	133
4.6	Discussion	137

5	Mechanisms underlying development of disorder in mouse V1	139
5.1	Introduction	139
5.2	Methods	140
5.2.1	Mechanisms of retinotopic scatter	140
5.3	Results	142
5.3.1	Development of retinotopic scatter driven by sparse clouds of activity	143
5.3.2	Subcortical scatter	144
5.3.3	Salt and pepper maps arising from sparse connections to V1	149
5.3.4	Sparse afferent connections also lead to ‘salt and pepper’ maps in cat model	156
5.4	Discussion	158
6	Overview and future work	161
6.1	Introduction	161
6.2	Homeostatic requirements for stable development	161
6.2.1	Immediate future aims	162
6.2.2	Future outlook	166
6.3	A model of orientation map development consistent with the known wiring of V1	168
6.3.1	Immediate future aims	169
6.3.2	Future outlook	171
6.4	Mechanisms underlying topographic organization of mouse V1	172
6.4.1	Immediate future aims	173
6.4.2	Future outlook	176
6.5	General points	177
6.6	Conclusion	178
7	Conclusions	179
	Bibliography	181

List of Figures

1.1	The mammalian early visual pathway	4
1.2	Receptive field types in the retina, LGN and V1	6
1.3	Maps in the mammalian visual pathway	8
1.4	Measurement of orientation maps	10
1.5	Orientation map features	11
1.6	Lack of orientation map in rodent visual cortex	12
1.7	Contrast invariant orientation tuning in V1	12
1.8	Extra-classical receptive fields	13
1.9	Development of orientation maps revealed by chronic optical imaging	15
1.10	The effect of dark rearing and lid suture on orientation map development	16
1.11	Lateral connections in the tree shrew orientation map	18
1.12	Basic LISSOM model of the primary visual cortex	21
1.13	Mexican-hat connectivity in LISSOM	22
2.1	The problem of stability in feedforward networks	28
2.2	Cellular mechanisms of homeostatic plasticity	29
2.3	Architecture of the ALISSOM model for orientation map development	31
2.4	Coordinate systems in the Topographica simulator	32
2.5	ALISSOM and LISSOM activation functions	37
2.6	Comparison of orientation map development in three self-organizing models	49
2.7	Stability of map development	50
2.8	Automatic adjustment of intrinsic excitability in the ALISSOM model	52
2.9	Homeostatic regulation of balance rescues map selectivity and organi- zation	54
2.10	Afferent scaling in ALISSOM	55
2.11	Homeostatic regulation of balance rescues stability of map development	56

2.12	Map development in ALISSOM for different types of input	57
2.13	Adjustment of the frequency parameter (ν) for different types of input results in stable map development	58
2.14	A realistic model of stable map development before and after eye opening	60
2.15	Comparison of selectivity and stability in ALISSOM and ferret	60
3.1	Local connections between V1 neurons	65
3.2	Di-synaptic Inhibition	66
3.3	Contrast dependent receptive field measurement	67
3.4	Perceptual level surround modulation	68
3.5	Example of unsuccessful model architecture	72
3.6	Combined interaction of connections in unsuccessful model architecture	73
3.7	Effects of learning in first example of model architecture	74
3.8	LESI model architecture	75
3.9	Combined interaction of connections in LESI model	77
3.10	Activity patterns in the LESI model	83
3.11	Matching maps in the LESI model	84
3.12	LESI receptive fields	85
3.13	Long range lateral excitatory connections in LESI	86
4.1	Receptive fields of mouse V1 neurons	92
4.2	Distinct visual areas in mouse cortex	93
4.3	In vivo two-photon calcium imaging	94
4.4	Organization for orientation in mouse and cat, measured using two- photon imaging	95
4.5	Optical imaging of the retinotopic organisation of mouse visual cortex	96
4.6	Two photon calcium imaging of retinotopy in mouse	100
4.7	Stimulus configuration and two-photon imaging results in mouse . . .	101
4.8	Illustration of retinotopic scatter in azimuth and elevation in mouse . .	104
4.9	Scatter in receptive field position	105
4.10	Analysis of retinotopic scatter in mouse	105
4.11	Mouse early visual system	107
4.12	Mouse model architecture	108
4.13	Analysis of retinotopic scatter in the mouse model	110
4.14	Cat model architecture	113
4.15	Mouse model without retinotopic scatter	121

4.16	Mouse model with ‘built-in’ retinotopic scatter	122
4.17	Mouse model with larger spread in ‘built-in’ retinotopic scatter	123
4.18	Mouse model with increased inhibition and decreased excitatory radius	125
4.19	‘Salt and pepper’ orientation maps with sparse activation	126
4.20	Receptive fields in ‘salt and pepper’ mouse model	127
4.21	Lateral connections in ‘salt and pepper’ mouse model	129
4.22	Summary of mouse model results	131
4.23	Cat model results	132
4.24	Comparison of cat model and experimentally measured orientation map	133
4.25	Mouse model results using natural images	134
4.26	Cat model results using natural images	135
5.1	Different types of scatter	141
5.2	Development of retinotopic scatter driven by sparse lateral connections	144
5.3	Orientation preference and receptive field development driven by sparse lateral connections	145
5.4	Measurement of retinotopic scatter in mouse LGN	146
5.5	Effect of subcortical scatter mechanisms on mouse model retinotopy .	148
5.6	Relationship between subcortical scatter and retinotopic scatter	149
5.7	Effect of subcortical scatter mechanisms on mouse model receptive fields	150
5.8	Effect of subcortical scatter on mouse model orientation maps	151
5.9	Effect of afferent connection sparsity on retinotopy in the mouse model	152
5.10	Relationship between sparsity and retinotopic scatter	153
5.11	Effect of afferent connection sparsity on orientation preference maps in the mouse model	154
5.12	Effect of afferent connection sparsity on mouse model receptive fields	155
5.13	Reverse correlation of receptive fields with different stimuli	156
5.14	Combination of sparse afferent connections and LGN receptive field scatter in the mouse model	157
5.15	Cat model results	158
6.1	Contrast responses in real neurons in comparison with ALISSOM model	164
6.2	Coupled excitatory and inhibitory homeostatic regulation	170

List of Tables

2.1	ALISSOM model parameters	46
3.1	LESI model parameters	80
4.1	Parameters for mouse ALISSOM model	109
4.2	Scatter properties for mouse model with ‘built-in’ scatter	111
4.3	Parameters for cat ALISSOM model	114
4.4	ALISSOM model parameters for mouse and cat models	117

Chapter 1

Background

1.1 Introduction

Our brains process and integrate sensory information and use this information to interact successfully with the environment. This processing occurs via trillions of connections between individual neurons, within a larger system of interconnected brain areas. This highly complex system is initially formed, and subsequently adapted, through a product of genetic and environmental factors. Intriguingly, the end product of this developmental process in the cerebral cortex is a set of topographic maps, forming sensory and motor areas such as those in the visual, auditory, and somatosensory systems. Understanding the rules that govern the development of these maps and how this topographic organization relates to information processing is a central question for neuroscience research.

Computational modeling is invaluable in this field, where experimental data is particularly difficult both to obtain and to interpret. A computational approach can bridge the gap between theory and experiment, providing a test-bed for theories, predictions for experiments, and in some cases highlighting the questions that experimenters ought to be asking. Computational models of visual information processing and topographic map development have been particularly successful, and have played an important role in developing our understanding of the fundamental mechanisms involved in neural map formation (Goodhill, 2007).

One reason for the success of computational models of vision is that in comparison with other sensory systems, the visual system is the most widely studied. Manipula-

tion of the external visual environment is relatively easy compared with other sensory systems, electrophysiological techniques can be used at several processing stages in the visual pathway, and imaging techniques such as optical imaging (Zepeda et al., 2004) and two-photon microscopy (Gbel and Helmchen, 2007; Stosiek et al., 2003) are providing a window on the anatomical and functional organization of visual processing areas in the brain. This makes the visual system an excellent starting point and test case for the investigation of the highly complex process of cortical development and information processing.

Previous computational models have shown that topographic maps can develop through a process of unsupervised self-organization, if spatially localized patches of cortical neurons are activated by particular stimuli. This thesis presents a series of computational models, based in part on this principle of self-organization, that focus on the development of the map of orientation preference in primary visual cortex (V1). This map is the prototypical example of topographic map development in the brain, and is the most widely studied, but the same self-organizing principles can also apply to maps of many other visual features and maps in many other sensory areas.

Experimental evidence from cat and ferret indicates that orientation preference maps in V1 develop in a stable way, with the initial layout determined before the animal has any normal visual experience (i.e. before eye opening). This constraint is at odds with previous self-organizing models, which have used biologically ungrounded ad-hoc methods to obtain robust and reliable development. Such mechanisms inherently lead to instability, by causing massive reorganization over time. The first model presented in this thesis (ALISSOM; chapter 2) builds upon a previous successful model of map development (LISSOM) by replacing ad-hoc methods with biologically realistic homeostatic mechanisms that lead to development that is both robust and stable. The ALISSOM model shows for the first time how orientation maps can remain stable despite the reconstruction of neuronal circuits and changes in visual input occurring during development. Implementing realistic homeostatic mechanisms in ALISSOM also highlights the requirements for homeostasis in the developing visual circuit. Eliminating ad-hoc mechanisms also makes ALISSOM a much simpler starting point for subsequent models of map development.

In previous self-organizing models, spatial restriction of activity has been achieved by a combination of short-range excitatory connections and long-range inhibitory connections between neurons. However, experimental studies suggest that this type of

connectivity is not observed in the layers of visual cortex where topographic maps can be measured (layer 2/3). A second model (LESI; chapter 3) shows how development can occur using circuitry that is consistent with the known wiring in layer 2/3 of V1. This new model, LESI, contains Long-range Excitatory and Short-range Inhibitory neural connections, which are known to be present in V1. Instead of direct long-range inhibition, it uses di-synaptic inhibition to ensure that when visual stimuli are at high contrast, long-range excitatory connections have an overall inhibitory influence. The results match previous models in the special case of the high-contrast inputs that drive development most strongly, but show how the behavior relates to the underlying circuitry, and also make it possible to explore effects at a wide range of contrasts.

The final part of this thesis explores the differences between rodents and higher mammals that could lead to the lack of topographic organization for orientation in rodent species. Analysis of retinotopy data from two-photon calcium imaging in mouse (provided by Tom Mrsic Flogel, University College London) shows that at single cell resolution there is local disorder in both orientation and retinotopy. In an initial set of simulations (chapter 4), this measured retinotopic disorder is built into a model of the early mouse visual pathway. These simulations indicate that disorder in orientation preference may arise via activity driven development, provided activity patterns in response to oriented stimuli in mouse V1 are sparse. Comparison with a model of the early cat visual pathway indicates that due to a difference in physical constraints (larger magnification factor and greater receptive field overlap) that activity patterns in the cat may be less likely to be sparse during development, thus ensuring that orientation maps form in this species. A second set of simulations (chapter 5) also investigate how the measured retinotopic scatter in the adult mouse could arise earlier in development. These simulations result in an alternative mechanism for the formation of disordered orientation maps resulting from initial wiring of sparse connections between the LGN and V1.

The rest of this chapter provides a review of the background necessary for understanding the detailed description of this work in subsequent chapters.




Figure 1.1: The mammalian early visual pathway. The left panel shows a schematic drawing of the pathway from the retina to the primary visual cortex (V1) through the lateral geniculate nucleus (LGN) of the thalamus. The right panels highlight the major anatomical structures. Light entering the eye passes through the ganglion cells and is focused on the photoreceptor layer. Signals from photoreceptors pass through bipolar cells to ganglion cells, the axons of which form the optic nerve, which projects principally to the LGN. Ganglion cells from the temporal retina (furthest from the nose) project to the ipsilateral (same side) LGN (red lines) and those from the nasal retina (nearest the nose) project to the contralateral (opposite side) LGN (green lines). Within the LGN, the projections from the two eyes are aligned, so the same topographic map (of the contralateral half of the visual field) is found in all layers. The axons of LGN neurons project almost exclusively to V1, where they terminate primarily in layer 4. The termination site within layer 4 depends on the layer in which the LGN neuron is found: parvocellular (P) cells project mainly to layer 4C β , magnocellular (M) to layer 4C α , and koniocellular (K) cells to layer 4A and lower layer 3. Reprinted from Solomon and Lennie (2007).

1.2 The early visual system

The early visual system (figure 1.1) begins with the light-sensitive photoreceptors in the retina. These cells make connections via retinal ganglion cells (RGCs) to neurons in the lateral geniculate nucleus (LGN). Visual information from the left half of the visual field (from both eyes) is processed on the right side of the brain and vice versa. In the optic chiasm, connections from the left and right retina that correspond to one half of the visual field cross to ensure connections reach the opposite side of the brain. Connections from the retina to the LGN are arranged retinotopically, so that nearby LGN cells respond to nearby portions of the visual field. Visual signals from the two eyes also remain segregated in different layers in the LGN.

The classical *receptive field* of a neuron in the visual system can be defined as the region of visual space to which that neuron responds. The receptive field can also describe the type of stimulus on the retina (in that preferred region) to which the cell responds maximally. RGCs and LGN cells have On-center or Off-center receptive fields (as shown in figure 1.2a,b), and therefore respond to increments (On-cells) or decrements (Off-cells) in light presented to the center of their receptive field. On-center LGN receptive fields also have an Off receptive field surround that increases the response of the cell to dark regions outside of the receptive field center and vice versa for Off-center receptive fields. In this way, RGC and LGN receptive fields can act (as a first approximation) like an edge detection system for the retinal image. Different types of LGN cells with, for example, different sizes of receptive fields or different temporal response properties are often separated into different layers in the LGN (although this can vary between species). For example, in the macaque monkey, magnocellular (M) cells (which have large receptive fields), parvocellular (P) cells (which have smaller receptive fields), and koniocellular (K) cells are separated into different, eye-specific layers in the LGN (figure 1.1).

The next (and first cortical) stage in visual processing is the primary visual cortex (V1), also called striate cortex. Like the rest of the cerebral cortex, V1 is divided into six layers, labelled 1 through 6, although these labels have been increasingly subdivided and regrouped as additional aspects of the properties of V1 neurons have been discovered. The detailed connectivity in V1 is reviewed in depth in Binzegger et al. (2004); Chalupa and Williams (2008); Lund (1988); Van Hooser (2007) and Douglas and Martin (2007). In summary, the anatomical architecture of V1 is a complex mix of




Figure 1.2: Receptive field types in the retina, LGN and V1. Each diagram shows a receptive field on the retina for one neuron. Areas of the retina where light spots excite this neuron are plotted in white (On areas), areas where light inhibits it are plotted in black (Off areas), and areas with little effect are plotted in medium gray. The size of the RFs varies, but in the retina and LGN they typically have the same basic shape and are all spatially localized, i.e. their On and Off areas cover a small, specific portion of the retina. (a) On cells are found in the retina and LGN, and prefer light areas surrounded by dark. (b) Off cells have the opposite preferences, responding most strongly to a dark area surrounded by light. RFs for both On and Off cells are isotropic, i.e. have no preferred orientation. Starting in V1, most cells have orientation-selective RFs instead. V1 simple cell RFs can be classified into a few basic spatial types, of which the two most common are shown above: (c) A two-lobe arrangement, here favoring a 45° edge with dark in the upper left and light in the lower right, and (d) a three-lobe pattern, favoring a 135° white line against a dark background. Both types of RF are well represented by Gabor functions (Daugman, 1980; Jones and Palmer, 1987) and are variable in both orientation and phase. Complex cells in V1 respond regardless of the phase of the oriented stimulus (i.e. black and white regions are inter-changable, without changing the response of the cell). Many neurons are also sensitive for the direction of movement of these patterns, i.e. their RFs are spatiotemporal. For such a neuron, successive snapshots of the spatial RF at different times are shown in (e); together they form a spatiotemporal RF selective for a vertical light bar moving to the right. Reprinted with permission from Miikkulainen et al. (2005).

excitatory and inhibitory, inter-laminar, intra-laminar, feedback, and long-range connections. There may also be variations in V1 connectivity between species, although in general the laminar architecture appears to be similar. In macaque monkey, input from the LGN is mainly received in layer 4, where the input from M, P and K cells remains separated in layers. Cortical neurons in layer 4 then send projections upwards to layer 2/3. Layer 2/3 contains many excitatory projection neurons that send axons laterally within layer 2/3, but also intracortically to extrastriate visual cortical regions (e.g. V2, V3, V4, MT, etc.) and are also a target for feedback from extrastriate areas.

Neurons in layers 5 and 6 provide feedback connections to the LGN. A more detailed description of the functional properties of V1 cells is given in the next section.

1.3 Primary visual cortex

The responses of individual neurons in primary visual cortex (V1) depend with varying degree upon specific properties of parts of images presented to the retina, for example: which eye the image is presented to, retinal position, motion direction, line orientation, spatial frequency, temporal frequency or color. Across the cortical surface in primate and carnivore species, these response properties are smoothly topographically organized. In a vertical penetration of the cortex, cells also appear to have similar receptive field properties, leading to the notion that the cortex is comprised of anatomical and functional units termed ‘cortical columns’. Retinotopy of V1 ensures that neighbouring columns on the cortex respond to stimuli presented to neighbouring points on the retina (figure 1.3a), and ocular dominance columns separate cells that prefer input from either eye (figure 1.3b). The anatomical basis of retinotopic and ocular dominance maps appear to be reasonably well explained in terms of the distribution of inputs from the LGN to cortex because LGN neurons are already selective for retinal location and eye of origin. On the other hand, since cells in the LGN are typically not selective for orientation or motion direction, properties such as orientation (figure 1.3c) and direction (figure 1.3d) selectivity are ‘emergent’ in the cortex. As such, orientation and direction maps have become the subject of much greater interest. The most widely studied topographic feature map, and the example on which I focus in this thesis, is the map of orientation preference in V1 (figure 1.3c).

Orientation maps are usually measured as described in Zepeda et al. (2004) using optical imaging of intrinsic signals from cortical layer 2/3, acquired during the presentation of drifting sinusoidal gratings (figure 1.4a,b). This technique has produced robust and reproducible orientation maps for various different species including primates (Blasdel, 1992), ferret (Rao et al., 1997), cat (Bonhoeffer and Grinvald, 1993; Ohki et al., 2005) and tree shrew (Bosking et al., 1997). Orientation maps have a characteristic layout with features of the orientation topography (illustrated in figure 1.5a) such as *linear zones*, *saddle points*, *pinwheels* and *fractures* that can be observed in most cases. The Fourier transform of experimentally measured orientation maps (figure 1.5b) is ring shaped, indicating that orientations in the maps repeat in all directions with a spatial




Figure 1.3: Maps in the mammalian visual pathway. Functional properties are smoothly topographically organized in most mammalian species. (a) Preference for position in the visual field measured using functional magnetic resonance imaging (fMRI). The average response on an unfolded human V1 is displayed. Responses are topographically organized, with nearby neurons in the cortex preferring nearby eccentricities in the visual field (as color coded in the inset). Reprinted from Dougherty et al. (2003). (b) Ocular dominance bands in the visual cortex of macaque monkey, measured by injecting a tracer dye into one eye. White areas indicate responses to the right eye, and black areas indicate left eye responses. Reprinted from LeVay et al. (1985). (c) Smooth organization for orientation preference (as shown in the color key) in the tree shrew visual cortex. Reprinted from Bosking et al. (1997). More details on the measurement and production of these false-color maps are shown in figure 1.4a,b. (d) Carnivores such as ferrets have an additional map of direction selectivity layered on top of the orientation map; arrows indicate direction preference. Reprinted from Weliky et al. (1996). Maps for other features (such as spatial frequency and color) have also been found in various species.

frequency that corresponds inversely to the radius of the ring. Chapter 2 of this thesis examines in more detail how these map patterns develop.

Although smoothly varying orientation maps are found in primates, carnivores and tree shrews, studies of rodents and lagomorphs such as mouse (Schuett et al., 2002), rat (Ohki et al., 2005), hamster (Tiao and Blakemore, 1976), rabbit (Murphy and Berman, 1979), and grey squirrel (Van Hooser et al., 2005) have identified no orderly orientation map in V1. The lack of map is not reflective of a lack of orientation selective cells; in fact, there are only slightly fewer orientation selective neurons in squirrel (65%) in comparison with cat (87%), and these cells have similar tuning characteristics. In recent studies using two-photon microscopy and calcium sensitive dye imaging in the rat (Ohki et al., 2005) and mouse (Tom Mrsic Flogel, personal communication), the possibility that maps are present at a much smaller scale than can be observed using traditional optical imaging has also been ruled out. These studies showed that there truly is a ‘salt and pepper mix’ of orientation selective cells even at single-cell resolution (see figure 1.6).

This raises the question of whether there really is a general organizing principle for the cortex. Is it possible that ‘salt and pepper’ organization can arise via a different parameterization of the same general connectivity pattern? Or is there something inherently different about rodents that results in this different organization? Is there any functional significance to the smooth topographic organization in primates and carnivores? These questions are addressed in detail in chapters 4 and 5 of this thesis.

Typical V1 receptive fields are shown in figure 1.2c,d,e. The receptive fields of cells in V1 have elongated On and Off sub-regions and are therefore selective for the orientation of a contrast border. Orientation selective cells in V1 can be classed as ‘simple’ or ‘complex’ depending on whether they are sensitive to the phase reversal of a light or dark bar. Simple cells are sensitive to phase, whereas the response of a complex cell does not strongly depend on phase. The phase responses of V1 neurons are not considered in any detail in this thesis, but may be important to consider in future extensions of this work.

The original feed-forward model for how orientation selectivity arises in V1 was proposed by Hubel and Wiesel (1962) over 40 years ago, and suggests that the feed-forward afferent inputs from the LGN On and Off-center cells are specifically arranged in order to generate orientation selectivity of cells in V1. There is a large amount of




Figure 1.4: Measurement of orientation maps. (a) Optical imaging of topographic maps in primary visual cortex. Experiments involve removing part of the skull of a laboratory animal, exposing the surface of the visual cortex. Visual patterns are then presented to the eyes, and a video camera records either light absorbed by the cortex or light given off by voltage-sensitive fluorescent chemicals that have been applied to it. Reprinted with permission from Miikkulainen et al. (2005). (b) Orientation map measurement in cat visual cortex. Single condition maps (outer images) show the average activity recorded during the presentation of a drifting sine grating stimulus with a particular orientation. Depending on the neural activity, there will be small differences in the emitted or reflected light, and these differences can be amplified by repeated presentations and averaging. The resulting images are an indirect measure of the average two-dimensional pattern of neural activity resulting from each stimulus. Measurements are combined and color coded (as shown in the color key) in a false color orientation map (center image). Scale bar, $500\mu\text{m}$. Reprinted from Crair et al. (1997).




Figure 1.5: **Orientation map features.** (a) An orientation map measured by optical imaging in macaque monkey, reproduced from Blasdel (1992). Nearby neurons in the map generally prefer similar orientations, forming groups of the same color called iso-orientation patches. Other qualitative features are also found. Linear zones are straight lines along which the orientations change continuously, like a rainbow; one linear zone is marked with a long white rectangle. Pinwheels are points around which orientations change continuously. They often occur in matched pairs: such a pair is circled in white. At saddle points a long patch of one orientation is nearly bisected by another; one saddle point is marked with a white bowtie. Fractures are sharp transitions from one orientation to a very different one; a fracture between red and blue (i.e., without purple in between) is marked with a white square. Reprinted with permission from Miikkulainen et al. (2005). (b) The two-dimensional Fourier spectrum of the map in a, calculated using methods described by Erwin et al. (1995) on orientation map data from Blasdel (1992); reproduced with permission from Miikkulainen et al. (2005). In this and subsequent Fourier spectrum figures, the center represents the DC component and the midpoint of each edge represents half of the highest possible spatial frequency of the image horizontally and vertically (i.e. the Nyquist frequency); the amplitude is represented in gray scale from white to black (low to high). The Fourier spectrum is ring shaped, indicating that the orientations repeat in all directions with a spatial frequency that corresponds inversely to the radius of the ring, with larger rings representing higher frequency of repeating patches of orientation preferences in the map.

supporting evidence for this model based on the connectivity and response correlations between LGN afferents and cortical simple cells (Alonso and Swadlow, 2005; Alonso et al., 2001; Jones and Palmer, 1987; Reid and Alonso, 1995; Tanaka, 1983).

Subsequent models have also been proposed in order to overcome some of the main failures of this simple model. For instance, and unlike in a simple implementation of the Hubel and Wiesel model, the orientation tuning of V1 simple cells is invariant to contrast (Sclar and Freeman, 1982; Skottun et al., 1987), as illustrated in figure 1.7a. Although the response of a cell in V1 does increase with increasing contrast, the width of the orientation tuning curve stays the same. This property can also be illustrated by replotting the same data in figure 1.7a as contrast response tuning curves (figure 1.7b).




Figure 1.6: **Lack of orientation map in rodent visual cortex.** Orientation maps in approximately $300 \times 300 \mu\text{m}$ patches of visual cortex in cat (a) and rat (b). In this case, unlike the previous orientation maps shown, maps are measured for individual cells using two-photon imaging. Two-photon imaging is described in more detail in chapter 4. Visually responsive cells are coloured according to their preferred orientation as shown in the color key. (a) Two-photon cell based orientation map in a slice through cat visual cortex. Reprinted from Ohki et al. (2006). (b) Two-photon cell based orientation map in a slice through rat visual cortex. Reprinted from Ohki et al. (2005). Unlike in higher mammals, nearby neurons in rodent visual cortex do not appear to prefer similar orientations.




Figure 1.7: **Contrast invariant orientation tuning in V1**(a) The response of one V1 cell to stimuli of different orientations at 4 different contrasts (10, 20, 40, 80%). The cell is tuned to a preferred orientation of approximately 220° . (b) The same data in (a) is replotted as contrast tuning curves at different orientations. The response to the preferred orientation is always greater than non-preferred, even with increasing contrast. Reprinted from Sclar and Freeman (1982)




Figure 1.8: **Extra-classical receptive fields.** Example of surround suppression in one cell of cat area 17 (primary visual cortex). (a) Response of one cell to a drifting grating at its preferred orientation placed in the classical receptive field. (b) Influence of the most sensitive single modulatory region, here located at an oblique position of the surround, on the response to the central stimulus. (c) Influence of an annular surround. (d) Response to the surround stimulus alone. The surround alone does not elicit any response in the cell despite influencing the cells response when presented together with a central stimulus. Reprinted from Walker et al. (1999).

The response of a cell to stimuli at non-preferred orientations will always saturate with contrast at a lower value than stimuli at the preferred orientation, i.e. even if the contrast keeps increasing, the cell will never respond as much to a non-preferred orientation as it does to its preferred orientation. With the simple Hubel and Wiesel feed-forward model, orientation tuning curves would broaden with increasing contrast because the LGN response (and therefore V1 response) to non-optimal orientations would increase with increasing contrast.

Although a V1 cell may respond only to an appropriate stimulus within its ‘classical’ receptive field as defined above, it is often also possible to modulate the response of the cell by placing stimuli outside of the classical receptive field. These ‘center-surround modulations’ or ‘extra-classical receptive field effects’ have been shown to be present at multiple processing stages in the visual system and may also underlie many perceptual phenomena. Figure 1.8 illustrates an example of surround modulation in cat primary visual cortex. Despite not responding directly to the surround stimulus (figure 1.8d), the cell’s response to a stimulus in the center of its receptive field is reduced by placing the stimulus in the surround (figure 1.8a,c). Stimulation of different areas of the surround can also differentially affect the response of the cell (figure 1.8b). Extra-classical receptive field effects are discussed in more detail in Chapter 3, where I present a model V1 circuit that is compatible both with the known wiring of V1 and also with the main types of surround modulation.

1.4 Development of orientation maps

A small number of chronic optical imaging studies (where the same area of cortex is imaged over several days or weeks) have been used to investigate the development of orientation selectivity, primarily in cats and ferrets (Chapman and Bonhoeffer, 1998; Chapman et al., 1996; Godecke et al., 1997). Although the visual system of the ferret is similar to that of the cat, the ferret is born approximately 3 weeks earlier in development, making electrophysiological and optical recording and experimental manipulations of visual input much easier. Ferrets have therefore emerged as the preferred animal model for studies of development. Some degree of orientation tuning is present in ferret area 17 (V1) as early as visual responses can be recorded (approximately 2 weeks before the time of natural eye opening). Mature adult-like orientation-selective cells and orientation maps in ferret emerge just after eye opening, between postnatal days 32-36. The structure of the orientation map does not change significantly throughout this development, and the map appears to remain in a stable configuration as the optical signal becomes stronger over time (figure 1.9). This stability is remarkable, considering that during this time the cortex becomes exposed for the first time to external visual inputs, and the complex anatomical architecture described in previous sections is being constructed.

The extent to which visual development is driven by either a genetic program or the external environment has been a subject of much debate (reviewed in Huberman et al., 2008; Miikkulainen et al., 2005; Miller et al., 1999; White and Fitzpatrick, 2007). It is generally accepted that retinotopic maps in V1 are initially specified by chemical gradients that control axon guidance, and subsequently refined by an activity driven process (although the exact details still remain to be specified; see Huberman et al. 2008). As discussed briefly in the previous section, orientation selectivity is thought to arise from the specific arrangement of afferent connections from LGN On and Off cells. However, how this specific connectivity arises throughout development is still a matter for discussion.

Experiments comparing orientation maps measured in ferrets that have been reared in darkness show that orientation selectivity can still develop in the absence of normal visual input (although not to mature levels). In contrast, abnormal visual experience through lid-sutured eyes does disrupt development of orientation selectivity and the orientation map (White et al., 2001; figure 1.10). These results and others (Blakemore




Figure 1.9: **Development of orientation maps revealed by chronic optical imaging.** Orientation maps recorded at six different ages in one animal. Each row of the figure shows maps recorded in the left primary visual cortex of one ferret at the postnatal age indicated at the left of the row. Each column of single-condition maps shows orientation maps recorded in response to a particular orientation of a moving square-wave grating (0° = horizontal). In this example, the first clear orientation maps can be seen at P33. Individual iso-orientation domains remain stable over time and do not change their position, shape, or size. The four red arrows shown in each horizontal single-condition map highlight this stability by pointing to particular features in the map. In the combined angle maps, the color of each pixel indicates the preferred orientation of cells at that location in cortex. In polar maps, the color again indicates the preferred orientation, whereas now the brightness of the color also indicates the strength of orientation tuning. At early ages, the polar maps are almost completely dark, indicating that there were no regions of the cortex that showed tightly tuned orientation-specific responses. As the maps mature, iso-orientation domains become visible as colored regions in the map. These domains become more strongly responsive over time, and more domains appear as the maps mature, without changing the overall map pattern. Scale bar, $2mm$. Reprinted from Chapman et al. (1996).

and Cooper, 1970; Crair et al., 1998; Pettigrew, 1974; Sengpiel et al., 1999), suggest that normal visual experience is not required for the initial formation of the orientation map, but is required for its maintenance and maturation.

Exactly what drives the formation and early development of the orientation map before eye-opening is still a matter for debate. One possibility is that early development is driven by genetically determined internal activity patterns (e.g. coordinated




Figure 1.10: The effect of dark rearing and lid suture on orientation map development. The upper two rows show single-condition images evoked by horizontal (H) and vertical (V) grating stimuli; the third row shows difference images generated by subtracting horizontal and vertical single-condition images (scale bar = 1mm). The bottom row shows polar-magnitude maps generated by vector summation, where color represents orientation preference (legend below middle column) and brightness represents the magnitude of selectivity. Note the presence of well-formed but attenuated (less selective) orientation maps in dark-reared ferrets. In contrast, very little columnar structure was evident in lid-sutured ferrets, making maps of orientation preference nearly undetectable in this group. Thus, the impact of abnormal visual experience engendered by lid suture was far more devastating for the development of the map of orientation preference than was the absence of visual experience. Reprinted from White and Fitzpatrick (2007).

spontaneous activity that occurs in the developing retina; Wong 1999). The idea that spontaneous activity drives initial map development is supported by a studies in which spontaneous input patterns are disrupted. E.g., in Weliky and Katz (1997), spontaneous activity in the retina and LGN was disrupted with synchronous electrical burst stimulation in a wire cuff placed around the optic nerve. Although in this case the layout of the orientation maps appeared to be normal, orientation selectivity was reduced.

Other explanations for the initial development of orientation maps include scaffolding of lateral connections in V1 (Grabska-Barwinska and von der Malsburg, 2008), or determination by specific initial wiring of afferent connections from the retina and LGN that does not depend on spontaneous activity (Ringach, 2007; Wiesel and Hubel, 1974). Models that address these different theories are described in more detail in section 1.6.

The maturation of orientation selectivity, unlike the initial emergence of the map, is widely accepted to be driven by visual input, and is thought to be mediated by activity-dependent synaptic weight modification or ‘Hebbian learning’ (i.e. cells that fire together, wire together). However, the most straightforward implementations of Hebbian modification of synapses tend to destabilize the activity of neural circuits (due to positive feedback; Rochester et al. 1956). Recently, the concept of homeostatic plasticity has emerged to complement Hebbian learning (Turrigiano and Nelson, 2000), and several types of homeostatic regulatory mechanisms in the central nervous system have been observed. These mechanisms have been characterized in neuronal cultures and slices, and more recently *in vivo* in *Xenopus* and rodents. However, it is still not clear from these experimental studies how different homeostatic mechanisms work together, or the role of each in functional development of neuronal circuits in higher mammals. Homeostatic plasticity is addressed in more detail in chapter 2 of this thesis, where I develop a self-organizing model that uses realistic homeostatic mechanisms in order to reproduce stable map development.

1.5 Relationship between topographic organization and adult function

Functional properties, such as orientation selectivity, in V1 are a result of the information transmitted between cells, and therefore depend on their particular connectivity. The connectivity between neurons in V1 is highly dependent on the underlying topographic organization, and therefore it is important to consider how topographic organization may impact on function.

Lateral connections within V1 can extend up to 8mm (in macaque) and appear to connect locally to their neighbors and to other neurons with similar orientation preferences. Some examples of these connections mapped using tracer injections in tree shrew are shown in figure 1.11. These connections have been proposed to be capable of accounting for some of the extra-classical receptive field effects in V1 mentioned in the previous section (figure 1.8); (Angelucci and Bullier, 2003; Angelucci et al., 2002a,b). However, as shown in figure 1.11, individual neurons have a very specific connectivity depending on their position within the orientation map. In order to explain the full range of extra-classical receptive field effects, it is also necessary to consider




Figure 1.11: Lateral connections in the tree shrew orientation map. (a) The vertical and horizontal orientation preferences in a $8\text{mm} \times 5\text{mm}$ section of V1 in the adult tree shrew, measured using optical imaging. The areas responding to vertical stimuli are plotted in black and those responding to horizontal stimuli in white. Vertical in the visual field (90°) corresponds to a diagonal line at 135° in this plot. The small green dot in the middle marks the site where a patch of vertical-selective neurons were injected with a tracer chemical. The neurons to which that chemical propagated through lateral connections are colored red. Short-range lateral connections target all orientations equally, but long-range connections go to neurons that have similar orientation preferences and are extended along the orientation preference of the source neurons. (b) The same information plotted on a $2.5\text{mm} \times 2\text{mm}$ section of the full orientation map to the right and below the injection site. The injected neurons are colored greenish cyan (80°), and connect to other neurons with similar preferences. Reprinted from Bosking et al. (1997).

the underlying map structure. The dependence of extra-classical receptive field effects on topographic organization is discussed in more detail in chapter 3.

Another example of how map position may influence computation is the degree of orientation tuning of V1 neurons. Based on the presence of local isotropic connections (figure 1.11), neurons near pinwheel centers are expected to receive much more broadly tuned synaptic inputs from the local circuit than neurons within iso-orientation domains. Indeed, the membrane potentials of neurons in regions of homogeneous orientation preference do show sharper orientation tuning (Schummers et al., 2002, 2004). The tuning of spiking responses has previously been more controversial, with some previous studies suggesting that these are invariant across the map (Marino et al., 2005; Ohki et al., 2006) but a more recent study indicating that spiking responses at pinwheels are indeed more broadly tuned (Nauhaus et al., 2008).

Experimental studies have also reported differences in other response properties that are dependent on map structure. Neurons at, or near, pinwheel centers differ in their capacity for short term plasticity (Dragoi et al., 2001; Felsen et al., 2002; Sengpiel and Bonhoeffer, 2002) i.e. they are more likely to alter their response during adaptation to

a particular stimulus. This difference is also consistent with the more broadly tuned subthreshold response of neurons at pinwheel centers due to local isotropic connectivity.

In order to fully investigate the effect that the underlying map organization has on functional properties, a full description of the orientation map and its underlying connectivity would be required, together with the ability to test the system with many different experimental paradigms. A computational model of orientation map development can provide a mathematical description of this 2-dimensional orientation map structure together with the embedded realistic connectivity. Section 1.6 describes this type of model and chapter 3 of this thesis addresses this challenge in more detail.

1.6 Previous models of orientation map development

Models of orientation map development can be split into two groups. The first group aim to describe the initial development of the map before the onset of visual experience (as described in section 1.4), often without appealing to activity driven mechanisms. The second group describe the maturation and maintenance of these maps, which is widely accepted to be driven by visual experience. There is some overlap in these groups resulting from the theory that initial maps may also emerge by activity driven processes occurring before natural visual input, i.e. spontaneous activity in the retina, LGN or cortex. Many models that can explain the development of orientation maps via activity driven processes could therefore also explain their initial emergence by taking into account spontaneous activity.

One model that seeks to explain the initial emergence of the orientation map without appeal to activity driven mechanisms is the recent model of Ringach (2004, 2007). This model suggests that orientation preference is initially set up by the connections that V1 neurons receive from retinal ganglion and LGN On and Off cells. In the retina, retinal ganglion cells (RGCs) are arranged on a nearly hexagonal grid or mosaic. If On and Off RGCs and LGN cells do not initially overlap because of the quantization effects of the mosaic, and are offset from each other, then sampling from these cells can convey some orientation selectivity to V1. If nearby V1 neurons then sample from nearby LGN cells, orientation preference will be mapped smoothly in V1. Although this model reproduces many of the observed statistics of V1 neurons, there are several

outstanding issues. First, the maps produced by this algorithm do contain most features of animal maps, but they do not reproduce their characteristic periodicity (i.e. they do not have a ring shaped Fourier transform). Also, given that the map patterns remain stable (Chapman et al., 1996), it is unclear how realistic maps could emerge from such a starting point. Moreover, it is not yet clear how maps of orientation preference that are similar regardless of the eye to which the stimuli are presented (as found by Godecke and Bonhoeffer 1996) could be formed using this model, given that the retinal mosaic will differ in the two eyes.

Other models (Ernst et al., 2001; Shouval et al., 2000) have suggested that intra-cortical lateral connections could explain the initial orientation map structure. Shouval et al. (2000) suggests that lateral connections provide a framework for the developing orientation maps by inducing an initial broad orientation tuning bias in V1 neurons. Further activity driven development of afferent connections to V1 can then sharpen orientation tuning. A recent model has also addressed how this scaffold of long-range connections could develop in the cortex via spontaneous activity-driven processes (Grabska-Barwinska and von der Malsburg, 2008). A serious problem with these models is that developmental time-scales suggest that afferent connections from the LGN are already present in V1 before lateral connections begin to form crude clusters (Sengpiel and Bonhoeffer, 2002; White and Fitzpatrick, 2007). It is unclear why afferent connections to V1 would only begin to adapt after lateral connections have formed an initial scaffold. Also, such a mechanism would only lead to selectivity for long contours, a prediction which has yet to be tested in newborn animals.

Models of activity-dependent orientation-map maturation are based on a principle of self-organization. Competition and cooperation between neurons in a network will result in only a few units in the network responding to each input pattern. Different neurons, or clusters of neurons, in the network will learn to respond to different types of input, resulting in a map-like structure. These types of models range from more abstract models that do not include neurons or connections (Grossberg and Olson, 1994; Niebur and Wörgötter, 1993; Rojer and Schwartz, 1990), correlation based learning models that specify the inputs to the network as abstract functions (Erwin and Miller, 1998; Linsker, 1986; Miller, 1994; Tanaka, 1990; Yuille et al., 1989), models that substitute specific connections between neurons with interaction functions (Goodhill and Cimoneriu, 2000; Obermayer et al., 1990; von der Malsburg, 1973), to much more realistic models with specific modifiable connections that can learn through the pre-




Figure 1.12: **Basic LISSOM model of the primary visual cortex.** The core of the model consists of a two-dimensional array of computational units representing columns in V1. These units receive input from the retinal receptors through the On/Off channels of the LGN, and from other columns in V1 through lateral connections. The solid circles and lines delineate the receptive fields of two sample units in the LGN and one in V1, and the dashed circle in V1 outlines the lateral connections of the V1 unit. The LGN and V1 activation in response to a sample input on the retina is displayed in gray-scale coding from white to black (low to high). The V1 responses are patchy because each neuron is selective for a particular orientation, and only particular orientations exist at each location in this image. Reprinted with permission from Miikkulainen et al. (2005).

sensation of natural images (Burger and Lang, 1999; Miikkulainen et al., 2005). Many of these models reproduce orientation preference maps that are realistic in appearance, but are not based on a realistic V1 architecture. The many different types of models are critically reviewed extensively by Erwin et al. (1995); Miikkulainen et al. (2005) and Swindale (1996).

Given the aim of investigating how orientation maps can develop and give rise to realistic connectivity embedded within a topographic organization, it make sense to consider more realistic models of orientation map development, i.e., those with specific interactions between model neuron units which behave, to some approximation, like real neurons. One of the most realistic self organizing map models to date is the LISSOM (Laterally Interconnected Synergetically Self-Organizing Map) model (Miikkulainen




Figure 1.13: **Mexican-hat connectivity in LISSOM.** (a) A sheet of model neurons representing V1. Each single neuron unit (such as the central unit in the sheet, marked here with a cross) makes short range excitatory connections (with a radius denoted by the solid red line) and longer range inhibitory connections (dashed blue line) with other neuron units in the sheet. (b) Taking a slice through this two dimensional V1 sheet (along the dotted line in (a)) and looking at the profile of interaction strengths, reveals that the effective interaction between neurons has a profile that looks like a sombrero or ‘Mexican hat’ (solid green line). Nearby neurons will excite each other but at longer distances the overall interaction between neurons is inhibitory.

et al., 2005). This model (the architecture of which is shown in figure 1.12) is capable of reproducing many of the experimentally observed feature maps (orientation, spatial frequency, disparity, ocular dominance, color and direction), and variations of the model can learn either through the presentation of simplified retinal inputs or from natural images. One of the main improvements that LISSOM makes on other simpler self-organizing algorithms is that the lateral interactions are represented by specific modifiable connections rather than fixed connections such as isotropic interaction functions like Gaussians or Difference of Gaussian kernels.

The LISSOM model is based (as are the majority of self-organizing map models) on a ‘Mexican-hat’ type of intracortical interaction, with recurrent long range inhibition and short range excitation between neurons (figure 1.13). In the model, as images are presented to the photoreceptors, afferent connection weights learn correlations between the pre-synaptic LGN activity and the post-synaptic V1 response. In the case where inputs contain elongated contrast borders, oriented receptive fields will form. Recurrent short-range lateral excitatory interactions and long-range lateral inhibitory interactions force nearby neurons in V1 to have similar responses, leading to a smooth orientation map organization. The lateral connections between correlated regions in V1 are strengthened, whereas those between un-correlated regions grow weaker. This

leads to the development of patchy long range connections similar to those observed biologically (see figure 1.11).

During development in the LISSOM model, orientation selectivity is initially quite unstable before settling upon the final map structure. This instability is inconsistent with the experimentally observed development of orientation maps, which arise as a stable structure that is maintained throughout development (figure 1.9). However, the inclusion of automatic mechanisms based on the principles of homeostatic plasticity can reconcile the LISSOM model with the stable development of maps. The issue of stability and implementation of a new stable LISSOM like model (Adaptive LISSOM, ALISSOM) is considered in detail in chapter 2.

Although there is substantial evidence for the short range excitatory connectivity in LISSOM, the longer range patchy lateral connections in the cortex (up to 6-8mm in macaque, Angelucci et al. 2002a; shown in figure 1.11) are primarily excitatory rather than inhibitory. Local short range inhibitory connections (with a radius of approximately $500\mu\text{m}$) and/or patchy inhibitory connections that extend over a slightly wider spatial range than local connections (2mm at most) are observed (Kritzer et al., 1992), however, inhibitory connections do not appear to extend over a similar wide range as lateral excitatory connections, particularly in layer 2/3 where orientation maps are measured.

One way the LISSOM model could be justified is by assuming that the most long range excitatory connections do not play a part in orientation map development. Lateral connections in the LISSOM model would therefore represent medium range inhibitory neurons and short range excitatory neurons. However, this explanation cannot account for the long range patchy excitatory connections which are observed in layer 2/3 (figure 1.11); a major focus of the LISSOM model. Another possibility is that connections in LISSOM are a simple approximation to the more complex circuitry that is actually present in V1. Chapter 3 of this thesis shows how a model based on the successful principles of LISSOM can reproduce map development using an architecture which contains no long range inhibitory connections.

The LISSOM model has previously only been applied to species in which orientation maps are observed in V1. However, as mentioned in section 1.3, there are a number of species (particularly rodents) that do not have this map organization. There is an enormous amount to be gained in establishing the relationship between rodent visual cortex

and that of higher mammals. For example, the results of experiments that are much easier to perform in rodents (such as genetic manipulations or two-photon imaging) could be extrapolated to make more direct predictions about similar effects in higher mammals. The relationship between rodents and higher mammals is considered in more detail in chapters 4 and 5 where I have implemented models parameterized using specific information about the mouse and cat visual pathway.

1.7 Summary of aims

Topographic organization is a feature not only of primary visual cortex but also many higher visual areas, together with the auditory, somatosensory, motor and prefrontal cortex. Despite V1 being the preferred starting point for the investigation of the development and function of cortical topographic organization, a full understanding of these features in V1 is still distant.

As described above, experimental evidence indicates that orientation preference maps in the primary visual cortex (V1) develop in a stable way, with the initial layout determined before eye opening. In chapter 2 I ask how this stable development could occur despite the large amount of circuit reconstruction and changes in input that are happening at this time. With this model I aim to highlight the requirements for homeostatic plasticity in the developing V1.

Existing models of orientation map development have all been based on a ‘Mexican-hat’ connectivity of long-range inhibitory connections and short-range excitatory connections in V1. In reality, however, this type of connectivity is not consistent with the connectivity observed in the cortex. In chapter 3 I aim to show how topographic maps can develop in a model that is consistent with the known wiring in V1. This model is the first to be consistent with current models of extra-classical receptive field effects, and would therefore be a starting point for the future investigation of the relationship between topographic organization and function.

In chapters 4 and 5 I address the difference in functional organization between rodents (particularly the mouse) and higher mammals. I aim to understand the possible differences between these species that could lead to the lack of orientation maps in rodents, and generate testable predictions for future experiments.

Overall, this work should contribute to a better understanding of topographic organi-

zation in V1 and in the cortex in general. The resulting models represent a synthesis of a large body of experimental evidence about V1 anatomy and function, provide concrete predictions to be tested in future experiments, and aim to offer a platform for developing a more complete explanation of cortical function in future work.

Chapter 2

Homeostatic requirements for stable development

2.1 Introduction

As discussed in section 1.4, experimental evidence from ferrets and cats indicates that orientation preference maps in the primary visual cortex (V1) develop in a stable way, with the initial layout determined before eye opening. Stability of the orientation map is remarkable given the massive circuit reconstruction and change in visual inputs occurring at this time. This trade-off between stability and plasticity is a non-trivial problem faced in general by neurons in the developing nervous system. The ability of neurons to learn and adapt requires changes in synaptic connectivity, yet it is also essential that neurons are neither excessively active nor silenced. Figure 2.1 illustrates that assuming the cortex is (to a first approximation) a feed-forward network of neurons, then maintaining the transfer of information at each level of the network requires regulating the gain of the units in order to preserve the responses of neurons.

Homeostatic mechanisms such as adjustment of intrinsic excitability (Schulz, 2006), scaling of synaptic strengths (Turrigiano and Nelson, 2004), or sliding thresholds for plasticity (Abraham, 2008; Bienenstock et al., 1982) have been shown to allow individual neurons to maintain stability of overall activity levels, whether during activity-dependent synaptic weight modification (here called ‘learning’ for brevity), or, over longer timescales, during development (Davis and Bezprozvanny, 2001; Desai, 2003). Figure 2.2a illustrates the widely accepted idea that neurons use some measure of




Figure 2.1: The problem of stability in feedforward networks. (a) Schematic feedforward network consisting of five layers of neurons. (b,c) Expected firing rates of the red neuron in each layer in response to a sinusoidal input. (b) If each action potential in a preceding layer causes more than one action potential in neurons at the next layer (on average), the firing rates saturate and eventually all information about the stimulus is lost. (c) If each action potential causes less than one action potential in neurons at the next layer (on average), the signals eventually fail to propagate. Reprinted from Turrigiano and Nelson (2004).

their average activity over time in order to trigger homeostasis and regulate their responses. Figures 2.2b c and d show the effect of different types of homeostasis on the input/output function of an individual neuron. It is clear from this figure that the interaction of these diverse mechanisms and the events within a cell that trigger them is likely to be very complex, particularly within a developing network of neurons.

The developing orientation map network has been considered in several types of computational models, but these models have not always considered the mechanisms underlying map stability. For example, in ‘low-dimensional’ models (Goodhill, 2007) such as the Elastic net model (Carreira-Perpin et al., 2005; Durbin and Willshaw, 1987; Goodhill and Cimonieriu, 2000) the stability of map development has not been considered, however, it is likely to depend on the ‘receptive field size’ parameter (usually K) which is decreased (or annealed) during development. In more mechanistic models, such as the self organizing map (SOM; Obermayer et al. 1990; von der Malsburg




Figure 2.2: Cellular mechanisms of homeostatic plasticity. (a) Homeostatic plasticity uses some measure of activity (such as integrating average firing rate over some long time scale, indicated by the integral sign) to adjust excitatory and inhibitory synaptic strengths, as well as the voltage-dependent conductances (Na⁺ and K⁺) that control neuronal firing properties. (b) Regulation of intrinsic conductances (intrinsic plasticity) can modify the input/output curve of the neuron, by shifting the 'activation threshold', i.e., moving the curve to the left (so it will fire less for a given level of synaptic drive) or right (so that it will fire more for a given level of synaptic drive). (c) Intrinsic plasticity can also modify the slope of this curve, thus changing the level of the gain between input and output. Intrinsic plasticity will change the sensitivity of a neuron to both excitatory and inhibitory inputs, suggesting that this is a general mechanism for allowing neurons to maximize the detection of whatever input they do receive. (d) By scaling the strength of all the inputs to a neuron up or down (synaptic scaling), a neuron's response can be shifted up or down its input/output curve; this curve determines how much the neuron fires for a given amount of synaptic drive. Excitatory and inhibitory inputs can be regulated independently, which allows neurons or local circuits to adjust the balance of excitation and inhibition in an activity-dependent way. Reprinted from Turrigiano and Nelson (2000).

1973) lateral interactions are often still modelled using an interaction function which does not depend directly on the input. The SOM model also includes various ad-hoc methods that lead to unstable maps during development. Conversely, computational models that have considered stability of learning in a developing network, have not considered the underlying orientation map structure (e.g. Adaptive resonance theory; reviewed in Raizada and Grossberg 2003).

The hypothesis investigated in this chapter is that homeostatic mechanisms which un-

derlie the stability of activity levels in individual neurons could also underlie the stability of organization of the orientation map in V1. Some SOM models have already implemented biologically realistic homeostatic mechanisms in map development (Butko and Triesch, 2007; Sullivan and de Sa, 2006), but, as stated above, the stability of map development, the effect of specific modifiable lateral interaction strengths in V1, or changes in the strength or type of input during development are not taken into consideration in these models.

This chapter proposes a realistic model of stable orientation map development based on the LISSOM architecture described in section 1.6. The LISSOM model includes lateral connections that have specific interaction strengths that adapt during the course of development. The model shows that adjustment of intrinsic excitability and synaptic normalization can achieve a desired average activity in individual neurons, but this is not enough to develop a stable orientation map. Neurons must maintain their average responses to stimuli regardless of changes in the statistics of visual input. Neurons must also ensure a balance between afferent inputs and lateral interactions that is required for good map organization. The requirement for these mechanisms would not have been observed unless specifically modelling lateral interaction strengths and changes in inputs during development. This model is the first to reproduce stable orientation map development despite changes in the strength and pattern of inputs before and after eye opening.

2.2 Methods

The model architecture and analysis methods are described in detail in this section. The ALISSOM model was built using the freely available Topographica simulator (Bednar et al., 2004) and is based on the hierarchical LISSOM architecture (Miikkulainen et al., 2005). Automatic adaptive mechanisms have been introduced in this model, hence the name ‘Adaptive LISSOM’ or ‘ALISSOM’. The specific differences between ALISSOM and LISSOM are highlighted in each section. Where relevant, paths to specific code in the Topographica simulator (svn revision 10541) have been included in italics (for example *topo.patterns.basic.Gaussian*) to aid reproduction of the results.




Figure 2.3: **Architecture of the ALISSOM model for orientation map development.** The model consists of two-dimensional arrays of computational units representing local groups of neurons at each visual processing stage. Connections to one unit in each sheet are shown. The array of V1 units corresponds to approximately $4\text{mm} \times 4\text{mm}$ of ferret visual cortex. V1 units receive lateral excitatory (small circle) and lateral inhibitory (large circle) connections from nearby V1 units. Lateral interactions within V1 lead to 'bubbles' of activity in response to afferent input. V1 units receive afferent input from sheets representing the On and Off channels in the retinal ganglion cells (RGC) and lateral geniculate nucleus (LGN), which in turn receive input from the retinal photoreceptors. The difference-of-Gaussian afferent connections from the photoreceptors to the RGC/LGN units form a local receptive field on the retina, and cause On-center RGC/LGN units to respond to light areas surrounded by dark, and Off-center units to dark areas surrounded by light. Neighboring RGC/LGN neurons have different but overlapping receptive fields. Input to the retinal photoreceptors can be any type of patterned image or photograph, such as the small natural image patch shown here.

2.2.1 ALISSOM model architecture

The model is a simulated network of single-compartment firing-rate units that represents the pathway from photoreceptors in the retina to neurons in V1 (figure 2.3). ALISSOM consists of four sheets that represent the input (retinal photoreceptors), the pathway from the retina to the retinal ganglion cells/lateral geniculate nucleus (RGC/LGN) On and Off cells, and V1. A sheet is a continuous abstraction of a finite, two-dimensional array of neural units and corresponds to a rectangular portion of a continuous two-dimensional plane.




Figure 2.4: **Coordinate systems in the Topographica simulator** (a) Locations in a sheet are specified using floating-point sheet coordinates (x,y) contained within the sheet's *BoundingBox*. The thick black line shows the *BoundingBox* of the default sheet, which extends from $(-0.5,-0.5)$ to $(0.5,0.5)$ in sheet coordinates. Any coordinate within the *BoundingBox* is a valid sheet coordinate. (b) The 1.0×1.0 area of the sheet in (a) can be represented by a 5×5 matrix, whose *BoundingBox* (represented by a thick black outline) corresponds exactly to the *BoundingBox* of the sheet to which it belongs. Each floating-point location (x,y) in sheet coordinates corresponds uniquely to a floating-point location (r,c) in floating-point matrix coordinates, and vice versa. Individual units or elements in this array are accessed using integer matrix index coordinates, which can be calculated from the matrix coordinate (r,c) as $(\text{floor}(\text{int}(r)), \text{floor}(\text{int}(c)))$. For the example shown, the center of the unit with matrix index $(0,1)$ is at location $(0.5,1.5)$ in matrix coordinates and $(-0.2,0.4)$ in sheet coordinates.

The Topographica simulator allows simulation parameters to be specified in units that are independent of the level of detail used in any particular run of the simulation. To achieve this, Topographica provides multiple spatial coordinate systems, called sheet and matrix coordinates. Quantities (such as sheet dimensions) are expressed in sheet coordinates. In practice however, sheets are typically implemented using some finite matrix of units. Each sheet has a parameter called its density, which specifies how many units (matrix elements) in the matrix correspond to a unit length in sheet coordinates. The transformation between sheet and matrix coordinates is illustrated in figure 2.4.

The reason for having multiple sets of coordinates is that the same sheet can at another time be implemented using a different matrix specified by a different density. Provided all dimensions for the simulation are given in sheet coordinates, these can be converted automatically to matrix coordinates, allowing different densities to be used at different times without changing any other parameters. For example, early parameter search-

ing can be performed using low density (and therefore less computationally intensive) simulations. Final simulations (for illustrative figures, for example) can then be run at higher densities without changing the simulation results.

In the ALISSOM model, the V1 sheet has dimensions (in sheet coordinates) 1.0×1.0 . If every sheet were to have the default 1.0×1.0 area, units near the border of higher-level sheets would have connections that extend past the border of lower-level sheets. This connection field cropping will often result in artifacts in the behavior of units near the border. To avoid such artifacts, lower-level sheets should usually have areas larger than 1.0×1.0 . The size of the RGC/LGN (1.5×1.5) and retina (2.25×2.25) sheets are therefore chosen to ensure that each unit in the receiving sheet has a complete set of connections, minimizing edge effects in the LGN and V1.

2.2.2 Sub-cortical model architecture

One training iteration (t) in the model is defined as the time at which the retinal input (i.e., the image presented to the retinal photoreceptor sheet) changes. Images are presented to the model at each iteration by activating the retinal photoreceptor units. The activation value of unit i in the retina sheet ($\Psi_{i,Retina}$) is given by the gray-scale value in the chosen image at that point.

The images used in this chapter are either oriented and elongated Gaussian patterns, noisy disk patterns or natural image patches (see figure 2.12). The center coordinates and orientation of each Gaussian pattern are chosen from a uniform random distribution and cover the full extent of the retina. Noisy disk patterns are activated in a circular region with a radius approximately equal to the size of the retina; the activation falls off at the edges according to a Gaussian. Disk centers are chosen from a uniform random distribution with bounds larger than the retina size to ensure that neurons at the edges of the map are equally likely to receive input from all parts of the disks. Uniform zero-mean random noise is then added to the disk pattern. Natural images are retina-sized patches from images of natural objects and landscapes from a dataset by Shouval et al. (1996). The pictures were taken at Lincoln Woods State Park and scanned into a 256×256 pixel image. No corrections were used for the optical distortions of the instruments and there is no pre-processing of the images.

At each iteration, the afferent connections from the retina sheet to the RGC/LGN sheets

are updated. The activation level η for a unit at position j in an LGN sheet at time t is defined as:

$$\eta_{j,LGN}(t) = f \left(\gamma_L \sum_{i \in F_j} \Psi_{i,Retina}(t) \omega_{ij} \right) \quad (2.1)$$

The activation function f is a half-wave rectifying function that ensures positive activation values, γ_L is a multiplier for the overall strength of connections from retina to LGN, $\Psi_{i,Retina}$ is the activation of unit i taken from the set of neurons on the retina from which LGN unit j receives input (its connection field F_j), and ω_{ij} is the connection weight from unit i in the retina to unit j in the LGN.

Weights of connections between the RGC/LGN On and Off sheets and V1 are set to fixed strengths with a standard difference-of-Gaussians kernel (equation 2.2), with On connection fields having a positive center and negative surround, and vice versa for Off. More precisely, the weight ω_{ij} between an On center cell at location (0,0) in the LGN On sheet and a unit in location (x,y) on the photoreceptor sheet is given by:

$$\omega_{ij} = \exp \left(-\frac{x^2 + y^2}{\sigma_c^2} \right) - \exp \left(-\frac{x^2 + y^2}{\sigma_s^2} \right) \quad (2.2)$$

where σ_c determines the width of the central Gaussian and σ_s the width of the surround Gaussian. The weights for an Off center cell are the negative of the On center weights (i.e. surround minus center). The center of each LGN connection field is mapped to the location in the retina corresponding to the location of that the LGN unit (i.e., the projection is retinotopic).

2.2.3 Afferent connections to V1

The sub-cortical architecture up to this point is identical to the LISSOM architecture. However, the connections from the RGC/LGN sheets to the V1 sheet in the ALISSOM model differ from those in the LISSOM model. Unlike in LISSOM, these connections are subject to a scaling mechanism which ensures a ‘gain control’ between the LGN and V1.

Each V1 neuron in the model receives connections from three different connection types or ‘projections’ (p), i.e., the afferent projection from the LGN (On and Off channels combined; $p = A$), the recurrent lateral excitatory projection ($p = E$), and the recurrent lateral inhibitory projection ($p = I$) from other V1 neurons. For individual

afferent projections from the LGN On and Off sheets ($p = A$) the contribution X_{jp} to the activation of unit j in the V1 sheet (from e.g. LGN On) is calculated as:

$$X_{j,LGN-On}(t) = \left(\frac{\kappa}{\overline{X_{jA}}(t-1)} \right) \sum_{i \in F_{j,LGN-On}} \Psi_{i,LGN-On}(t-1) \omega_{ij} \quad (2.3)$$

and similarly for LGN Off, where $\Psi_{i,LGN-On}$ (or $\Psi_{i,LGN-Off}$) is the activation of unit i taken from the set of neurons on the LGN sheet (On or Off) from which each V1 unit j receives afferent input (its connection field F_j), and ω_{ij} is the connection weight from unit i in the LGN to unit j in V1. κ is the target average contribution to the total V1 activation from the combined afferent projections. $\overline{X_{jA}}$ is a smoothed exponential average of the total afferent contribution to a V1 unit :

$$\overline{X_{jA}}(t) = (1.0 - \beta)\Lambda + \beta\overline{X_{jA}}(t-1) \quad (2.4)$$

where

$$\Lambda = \sum_{i \in F_{j-On}} \Psi_{i,LGN-On}(t-1) \omega_{ij-On} + \sum_{i \in F_{j-Off}} \Psi_{i,LGN-Off}(t-1) \omega_{ij-Off} \quad (2.5)$$

The smoothing parameter (β) determines the degree of smoothing in the calculation of the average. $\overline{X_{jA}}$ is initialized to the target $\overline{X_{jA}}(0) = \kappa$.

The effect of this scaling mechanism is to bring the average contribution of the afferent projection to the activation of each V1 unit closer to the specified target. If the activity in the afferent projection moves away from the target during training, it is scaled to bring it closer to the target. In ALISSOM I have assumed that this scaling is performed at the ‘post-synaptic’ level in each neuron, i.e., there is a local mechanism in each V1 neuron which is regulating its afferent activity from the LGN. There are other possible mechanisms which would have the same effect, such as gain control in LGN neurons or in retinal ganglion cells. The reason for using this particular mechanism is because I have chosen to implement all homeostatic mechanisms locally in V1. However, future implementations of this model may use a different mechanism.

The rate of averaging of afferent activity (determined by the smoothing parameter β) determines how quickly the system responds to changes. If β is too high the system is slow to adapt, leading to an inability to maintain the target average activity. If β is too low the system becomes very sensitive to random fluctuations in the afferent input, interfering with the balance and leading to poorly organized maps. If this is indeed the underlying mechanism which implements gain control for afferent input to

V1 neurons, it may be the case that during normal development, input strengths vary only in a small range that is compatible with the rate of homeostatic regulation. On the other hand, there is also some evidence (from chronic optical imaging in ferret; Chapman et al. 1996) that map organization is not always robust, which may indicate changes in strength that fall outside of the ability of the system to adapt.

2.2.4 V1 lateral connections

As well as afferent inputs, each V1 neuron receives input from recurrent lateral excitatory and inhibitory projections. Recurrent short-range lateral excitatory interactions and long-range inhibitory interactions force nearby neurons in V1 to have similar responses, while the lateral connections learn pairwise correlations between V1 neurons. Recurrent lateral connections are updated 10 times per iteration at every settling step t_s . The contribution X_{jp} to the activation of unit j in the V1 sheet at time t from each lateral projection type, i.e., recurrent lateral excitatory ($p = E$) or recurrent lateral inhibitory ($p = I$), is updated every settling step t_s :

$$X_{jp}(t_s) = \sum_{i \in F_{jp}} \eta_{i,V1}(t_s - 1) \omega_{ij} \quad (2.6)$$

Where $\eta_{i,V1}$ is the activation of unit i taken from the set of neurons in V1 to which unit j is connected (its connection field F_j) and ω_{ij} is the connection weight from unit i in V1 to unit j in V1 for the projection p .

The contributions from all projections p described above are combined using equations 3.3 and 3.4 to calculate the activation of a neuron j in V1 at time t ($\eta_j(t)$):

$$\eta_j(t) = \sigma(t) = \frac{1.0}{1.0 + e^{(-a(t)x(t) - b(t))}} \quad (2.7)$$

where σ is a sigmoid transfer function, as illustrated in figure 2.5a. The input activity $x(t)$ is a sum of the contributions from all projection types (lateral excitatory, lateral inhibitory or afferent):

$$x(t) = \sum_p \gamma_p X_{jp}(t) \quad (2.8)$$

The projection strength scaling factors for each projection type p are $\gamma_A = 1.0$, $\gamma_E = 1.0$, $\gamma_I = -1.0$ for Gaussian inputs. In images where activity is more broadly distributed across the retinal photoreceptors (noisy disks and natural images) there are longer-range correlations in V1 activity which result in patchy long-range lateral connections in the model. Divisive normalization of weights (equation 3.2) causes weights




Figure 2.5: **ALISSOM and LISSOM activation functions.** (a) The transfer function for individual neurons in the ALISSOM model is a sigmoid (or logistic) function. The parameters of the sigmoid (equation 3.3) determine the slope (a) and threshold (b) of the curve. Changes in the slope parameter change the amount of response to similar input values. Changes in the threshold parameter increase or decrease the amount of input required before the neuron will respond. (b) The LISSOM model uses a piecewise linear transfer function which is an approximation of the sigmoid function. In this case, the neuron requires an input as large as the lower threshold before responding, and saturates at the higher threshold. In both cases, the output activation values are limited to $[0..1]$.

to be spread out more widely, reducing the ability of lateral inhibitory interactions to balance lateral excitatory interactions locally. Therefore $\gamma_I = -2.0$ to compensate in these (more realistic) cases.

The second homeostatic mechanism in the ALISSOM model is the maintenance of a target average response in V1 neurons. Because of the recurrent connections in V1, regulating the responses of V1 neurons also ensures that the overall strengths of lateral interactions are regulated.

The parameters of the sigmoidal transfer function $\sigma(x)$ for each unit are adapted during development using the homeostatic intrinsic plasticity rule proposed by Triesch (2005). This rule adjusts the parameters a (slope) and b (threshold) of the sigmoid (figure 2.5a; equation 3.3) at the end of each iteration using a gradient descent learning rule:

$$a(t) = a(t-1) + \epsilon \left(\frac{1}{a(t)} + x(t) - \left(2 + \frac{1}{\mu} \right) x(t)\eta(t) + \frac{1}{\mu} x(t)\eta(t)^2 \right) \quad (2.9)$$

$$b(t) = b(t-1) + \epsilon \left(1 - \left(2 + \frac{1}{\mu} \right) \eta(t) + \eta(t)^2 \frac{1}{\mu} \right) \quad (2.10)$$

This rule was derived by minimising the distance between the distribution of output activity for each unit (η_j) and an exponential distribution with desired mean μ . ϵ is the homeostatic learning rate.

This mechanism of automatic threshold adjustment is an implementation of homeostatic intrinsic plasticity, whereby a neuron controls the ability of synaptic inputs to lead to an action potential by changing the number and distribution of its ion channels (Daoudal and Debanne, 2003; Davis and Bezprozvanny, 2001; Schulz, 2006; Zhang and Linden, 2003); (see figure 2.2b,c). In this high-level implementation, the slope and threshold of the input/output function for each model neuron is adjusted directly in order to maintain a target average firing rate.

This threshold adjustment differs dramatically from LISSOM, which has a hand-tuned program of threshold changes during training. The initial activation function for all V1 neurons in LISSOM is a piecewise linear function with lower bound $x_l = 0.083$ and upper bound $x_u = 0.633$, as illustrated in figure 2.5b. The upper and lower bounds are both increased at set times throughout development in order to increase the threshold for activation of each neuron.

Individual neurons in LISSOM have identical activation function parameters at each stage of development and the threshold changes are made globally across the whole V1 sheet. Because this adjustment is a manual process in LISSOM, it is impractical to randomize the parameters of the initial transfer function. However, in ALISSOM, the initial slope and threshold parameters are chosen from a uniform random distribution with bounds a : [10,20] and b :[-8.0,-4.0]. The intrinsic plasticity adjustment in ALISSOM ensures that each neuron will adapt its transfer function so that the neuron will be activated during development. For example, if the neuron has an initially high threshold this will be reduced at the start of development to allow the neuron to become active. As the neuron learns and becomes more active during development, the threshold will be adjusted to ensure that the desired target average firing rate is obtained.

The target average firing rate μ and target average afferent activity κ for each V1 neuron (equation 2.3) are specified jointly in the ALISSOM model by a ‘frequency’ parameter v ($\mu = 0.0045v$, $\kappa = 4\mu$), which is based on how frequently a V1 neuron is expected to respond for a given input type. Adjusting this parameter ensures that the neuron’s

average response to a stimulus rather than just its average activity is maintained. For example, the frequency parameter ν is approximately doubled when the number of Gaussians presented per iteration is doubled from 2 ($\nu = 2.0$) to 4 ($\nu = 3.56$). The frequency value is initially chosen for 2 Gaussians, so the frequency value for 4 Gaussians does not equal exactly 4.0 because images are presented to the full area of the retina, therefore V1 will not always be activated by 4 complete Gaussian patterns at each iteration. The necessity for the ‘frequency’ parameter is discussed in detail in section 2.3.3.

The ratio between μ and κ is chosen to provide a balance between lateral interactions and afferent input which gives the best map organization and selectivity. Map organization is judged by observing how close the Fourier transform is to the ring shape that has been observed in animal maps (see figure 1.5); (Erwin et al., 1995). Specifically, lateral connections must be strong enough to form ‘bubbles’ of activation in V1, or else the model will not form a well organized orientation preference map. This balance is discussed in detail in section 2.3.2.

Balance in the model is maintained by scaling afferent inputs to V1 and adjusting lateral interactions within V1 to maintain a constant ratio between their strengths. This could also be achieved via several other mechanisms that would produce similar results. For example, afferent inputs could be regulated pre-synaptically, in the LGN or retina. Alternatively, V1 activity or lateral connection strengths could be regulated directly depending on the strength of afferent inputs. This is discussed in more detail in section 6.2. In any of these cases, the underlying goal would be to retain a balance between afferent and lateral input in V1 that ensures ‘bubbles’ of activity that drive organization are formed and that these do not change dramatically during development, despite changes in input.

2.2.5 Learning

In ALISSOM, connection field weights are fixed, isotropic 2D Gaussians for the lateral excitatory projection (equation ??, *topo.patterns.basic.Gaussian*), and initially random within a Gaussian envelope (equation 2.11, *topo.patterns.random.GaussianCloud*) for afferent and lateral inhibitory projections. Specifically, for a neuron located at (0,0):

$$\omega_{ij} = u \times \exp\left(-\frac{x^2 + y^2}{2\sigma_p^2}\right) \quad (2.11)$$

where (x, y) is the location of the postsynaptic neuron, $u = 1$ for the lateral excitatory projection ($p = E$) and u is a scalar value drawn from a uniform random distribution for the afferent and lateral inhibitory projections ($p = A, I$), σ determines the width of the Gaussian. Weights are initialized differently for the short-range lateral excitatory projection because early simulations indicated that the final weights achieved after learning in this projection will be Gaussian. Therefore, turning learning off and initializing the weights to their final configuration saves computational resources. Weights for each projection are only defined within a specific maximum radius r_p .

In the LISSOM model, unlike ALISSOM, connection field weights in the lateral excitatory projection are initially random within a Gaussian envelope (equation 2.11) with $\sigma_E = 0.104$ and a bounding radius of 0.104. Over the course of development the lateral excitatory bounds decrease to a final value of 0.00174. This decrease in the lateral excitatory bounds in LISSOM contributes to the instability of the orientation preference map (see section 2.3.1).

In the model, as images are presented to the photoreceptors, afferent and lateral connection weights are adjusted using a simple Hebbian learning rule. Afferent connection weights learn correlations between the pre-synaptic LGN activity and the post-synaptic V1 response, and lateral connections learn correlations in the patterns of V1 activity. Hebbian connection weight adjustment at each iteration is dependent on the pre-synaptic activity, the post-synaptic response, and the Hebbian learning rate:

$$\omega_{ij}(t) = \frac{\omega_{ij}(t-1) + \alpha \eta_j \eta_i}{\sum (\omega_{ij}(t-1) + \alpha \eta_j \eta_i)} \quad (2.12)$$

where for unit j , α is the Hebbian learning rate for the connection field F_{jp} in projection p . Unless it is constrained, Hebbian learning will lead to ever-increasing (and thus unstable) values of the weights (Rochester et al., 1956). In ALISSOM (as in LISSOM, see Miikkulainen et al., 2005) the weights are constrained using divisive post-synaptic weight normalization (equation 3.2), which is a simple and well understood mechanism. Afferent connection weights (from LGN On and Off) are normalized together in the model. A more biologically motivated homeostatic mechanism such as multiplicative synaptic scaling (Sullivan and de Sa, 2006; Turrigiano, 1999; Turrigiano and Nelson, 2004) or a sliding threshold for plasticity (Bienenstock et al., 1982) could be implemented to achieve similar results.

Learning rate parameters are specified as a fixed value α_p for each projection. The unit-specific values used in the equation above are then calculated as $\alpha = \frac{\alpha_p}{\tau_p}$, where τ_p

is the number of connections per connection field in the projection p .

Because afferent scaling in the model is performed at the post-synaptic level of V1, differences in pre-synaptic afferent activity levels (i.e., LGN activity) will affect the rate of learning adjustment of weights during development (equation 3.2). In order to compensate for this effect, the fixed learning rate for afferent projections ($\alpha_A = 0.137$) is additionally scaled in order to compensate for changes in the pre-synaptic afferent activity:

$$\alpha_{jA}(t) = \frac{\alpha_A}{\tau_{jA}} \left(\frac{\kappa_{lr}}{\overline{X_{jA}(t)}} \right) \quad (2.13)$$

This automatic mechanism allows comparisons to be made easily between different simulations by compensating for differences in pre-synaptic activity without having to manually adjust the afferent learning rate. For simulations using Gaussian patterns $\kappa_{lr} = 0.045$ and for noisy disks and natural images $\kappa_{lr} = 0.025$. This scaling amounts to decreasing the learning rate in the case where retinal activity is more broadly distributed to avoid learning short term correlations due to noise. If afferent activity was instead controlled by a pre-synaptic mechanism (either in the retina or LGN) this mechanism would not be required. Removing this mechanism does not have any significant effect on the results.

The lateral inhibitory projection learning rate (α_I) is constant throughout the simulation, and the lateral excitatory weights are fixed ($\alpha_E = 0.0$; as mentioned above, enabling plasticity for these short-range connections has little effect on the simulation).

Unlike ALISSOM, short range lateral excitatory connections in the LISSOM model have plasticity enabled and have an initial projection learning rate of $\alpha_E = 2.56$. The projection learning rate is also adjusted during development to account for the decrease in the number of individual connections which results from the change in the lateral excitatory bounds. The initial learning rate for afferent connections in the LISSOM model also starts off at a larger value ($\alpha_A = 0.4795$) and decreases exponentially throughout development to a final value of $\alpha_A = 0.137$. Again, this contributes to the instability in receptive fields, i.e., in LISSOM, at the start of development, thresholds are low and learning rates are high at the start of development, leading to receptive fields being overwritten at each iteration.

2.2.6 Model comparisons

The LISSOM model presented in section 2.3.1 for comparison is described in detail in Miikkulainen et al. (2005) and is available in the examples directory of the topographica simulator (*examples/lissom_oo_or.ty*). The main differences between ALISSOM and LISSOM have been highlighted throughout the rest of section 2.2.

The SOM model used in section 2.3.1 for comparison is based on the model presented in Obermayer et al. (1990). This model is also available in the examples directory of the topographica simulator (*examples/obermayer_pnas90.ty*). The model is a simplified network which contains only retina and V1 sheets. In this model, lateral interactions between neurons are replaced by a computationally less expensive mechanism. The algorithm finds the maximally active neuron in the cortex and adapts the afferent weights only in a specified circular (Gaussian) neighborhood around it. The area of this neighborhood is decreased throughout learning. Learning rates in this model also decrease linearly during development. Input patterns are elongated and oriented Gaussian patterns, similar to those in the LISSOM and ALISSOM simulations shown in comparison.

The simulation differs slightly from the published version in order to match the LISSOM and ALISSOM models more closely. Periodic boundary conditions are not taken into account, therefore the results presented here will contain edge effects not present in the original simulation. However, these should not affect the general purpose of this comparison which was to highlight the instability of orientation preference during development.

2.2.7 Analysis.

Orientation maps are calculated based on the vector sum method (Blasdel and Salama, 1986) (*topo.command.analysis.measure_or_pref*) Sine grating inputs that cover the full range of parameter values (combinations of all orientations, frequencies, and phases) are presented, and for each orientation, the peak response of the neuron is recorded. The orientation preference is calculated by constructing a vector for each orientation θ (between 0° and 180°) with the peak response as the length and 2θ as its orientation. These vectors are summed and the preferred orientation is calculated as half of the orientation of the summed vector. The selectivity is given by the magnitude of the

summed vector.

To ensure consistent map measurement across all conditions, afferent scaling is removed and a single fixed activation function is used for all V1 neurons. Individual differences in the activation function of each model neuron at the beginning of development manifest as an artificial increase or decrease in orientation selectivity. Hypothetically, this could be avoided if random noise was added to the orientation map measurement (as in the real data), but this is a more computationally intensive approach (as it requires making many repeated measurements) and therefore I have used this alternative mechanism for simplicity. Average orientation selectivity values are the median value of orientation selectivity across the map.

The orientation similarity index (Si) is calculated as in Chapman et al. (1996). The difference between preferred orientations for corresponding units in the two maps is averaged and subtracted from 1.0, giving a value of 1.0 for identical maps and approximately 0.5 for maps that are uncorrelated. Specifically when comparing two orientation preference maps ($ORpref$) at $t = 1$ and $t = 2$:

$$Si = 1.0 - \frac{\sum_i |(ORpref_{t=1} - ORpref_{t=2})_{Cyclic}|}{N_d^2} \quad (2.14)$$

Where N_d is the cortical density (i.e. N_d^2 is the total number of neurons in the V1 sheet) and i iterates over all neurons in the sheet. *Cyclic* indicates that the orientation difference must be calculated taking into account that orientation is a cyclic quantity. To account for noise in experimental measurement, ferret similarity indices are linearly extrapolated to obtain an estimate of the similarity that would be expected between two maps measured on the same day (0.76). Similarity indices are then normalized to this maximum value to obtain the values plotted in figure 2.15b.

Receptive fields are calculated using reverse correlation (Jones and Palmer, 1987) (*topo.command.analysis.measure_rfs*). Each unit on the model retina is activated by a light and dark square. Each time the unit in the V1 sheet responds, the stimulus that elicited the response is added to a retina-sized accumulator array, scaled by the unit's response level. The accumulated value of that array after all retinal patterns have been presented is then plotted as the receptive field.

2.2.8 Parameter values

In general, the ALISSOM algorithm is sensitive to changes in the frequency parameter v that determines the target mean V1 activity μ . The ratio $\mu = 0.0045v$ remains constant throughout all of the simulations shown and is initially determined by judging the stability of an initial simulation using 2 Gaussians at each iteration. The system is also sensitive to the ratio between μ and κ (desired mean afferent contribution to each V1 unit). We find that the relationship between κ and μ that gives the best map organization and selectivity is $\kappa = 4\mu$. This value remains constant in all the simulations presented in this chapter. Other parameters can be varied over a wide range without significant changes to the results.

Parameter	Value	Description
N_d	100	Cortical density, i.e., width and height of a unit area of cortex, section 2.2.1 and equation 2.14
L_d	50	LGN density, i.e., width and height of a unit area of LGN, section 2.2.1
R_d	50	Retina density, i.e., width and height of a unit area of Retina, section 2.2.1
r_A	0.271	Maximum radius of the afferent connections from LGN (On and Off) to V1, section 2.2.5
r_E	0.02	Maximum radius of the lateral excitatory connections, section 2.2.5
r_I	0.229	Maximum radius of the lateral inhibitory connections, section 2.2.5
σ_{center}	0.037	Standard deviation of LGN difference of Gaussian center Gaussian, equation 2.2
$\sigma_{surround}$	0.148	Standard deviation of LGN difference of Gaussian surround Gaussian, equation 2.2
σ_A	0.542	Standard deviation of initial Gaussian random afferent connection weights, equation 2.11
σ_E	0.04	Standard deviation of Gaussian lateral excitatory connection weights, equation ??
σ_I	0.458	Standard deviation of initial Gaussian random lateral inhibitory connection weights, equation 2.11
γ_L	2.33	Strength scaling factor for the LGN's afferent weights, equation 2.1
Continued on next page		

Table 2.1 – continued from previous page

Parameter	Value	Description
γ_A	1.0	Strength scaling factor for the V1 afferent weights, equation 3.4
γ_E	1.0	Strength scaling factor for the lateral excitatory weights, equation 3.4
$\gamma_{I,Gauss}$	-1.0	Strength scaling factor for the lateral inhibitory weights when input patterns are Gaussian, equation 3.4
$\gamma_{I,nd}$	-2.0	Strength scaling factor for the lateral inhibitory weights for noisy disk input patterns, equation 3.4
$\gamma_{I,nat}$	-2.0	Strength scaling factor for the lateral inhibitory weights for natural image input patterns, equation 3.4
Ψ_{nd}	0.85	Activation value for disks in noisy disks pattern, equation 2.1
$Noise_{nd}$	[-0.5, 0.5]	Bounds of uniform zero-mean random noise added to disk in noisy disks pattern, section 2.2.2
r_{nd}	1.0	Radius of the disk in noisy disks pattern, section 2.2.2
σ_{nd}	0.1	Radius of the Gaussian falloff in brightness at the disk edges in noisy disks pattern, section 2.2.2
Ψ_{Gauss}	1.0	Maximum activation value for Gaussian inputs, equation 2.1
σ_{Gauss}	0.044	Radius of the major axis of oriented Gaussian inputs, section 2.2.2
ar_{Gauss}	4.67	Aspect ratio of oriented Gaussian inputs (used to calculate radius of minor axis), section 2.2.2
α_A	0.137	Afferent projection learning rate, section 2.2.5, equation 3.2
α_E	0.0	Lateral excitatory learning rate, section 2.2.5, equation 3.2
α_I	1.808	Lateral inhibitory learning rate, section 2.2.5, equation 3.2
$\kappa_{lr,Gauss}$	0.045	Target afferent activity for use in learning rate scaling for Gaussian input patterns, equation 2.13
$\kappa_{lr,nd}$	0.025	Target afferent activity for use in learning rate scaling for noisy disk input patterns, equation 2.13
Continued on next page		

Table 2.1 – continued from previous page

Parameter	Value	Description
$\kappa_{lr,nat}$	0.025	Target afferent activity for use in learning rate scaling for natural image input, equation 2.13
β	0.999	Smoothing parameter in exponential average of afferent activity, equation 2.5
v_{2Gauss}	2.0	Expected frequency of V1 responses for 2 Gaussians, section 2.2.8
v_{4Gauss}	3.56	Expected frequency of V1 responses for 4 Gaussians, section 2.2.8
v_{nd}	7.5	Expected frequency of V1 responses for noisy disk patterns, section 2.2.8
v_{nat}	5.0	Expected frequency of V1 responses for natural images, section 2.2.8
μ	$0.045v$	Target average V1 activity (mean of target exponential distribution), section 2.2.8
κ	4μ	Target average contribution from the afferent projections to the V1 activity, equation 2.3
ϵ	0.016	Homeostatic learning rate, equations 2.9,2.10
a_{init}	[10, 20]	Bounds of initial slope parameter in sigmoid transfer function chosen randomly from a uniform random distribution, equation 2.9
b_{init}	[-8.0, -4.0]	Bounds of initial threshold parameter in sigmoid transfer function chosen randomly from a uniform random distribution, equation 2.10
t_s	9	Number of settling iterations, equation 3.1

Table 2.1: **ALISSOM model parameters** Model parameters used in the ALISSOM simulations presented in this chapter. Lengths shown are in sheet dimensions, where the size of the V1 sheet is 1.0×1.0 nominal area in sheet dimensions. ALISSOM is based on the hierarchical LISSOM architecture, more details of which are available in Miikkulainen et al. 2005.

Many of the parameters (particularly connection radii and input pattern parameters) for this model have been inherited from the original LISSOM model (Miikkulainen et al., 2005). Although most these parameters including connection radii (r_p), strength scaling factors (γ_p), Hebbian learning rates (α_p) and frequencies (v) can be varied over a substantial range (at least within 50% of the values given here) without affecting the

main findings, more substantial changes may require additional tuning. For example, changing the connection radii (r_p) by more than 50% is likely to require subsequent changes to the strength scaling factors (γ_p) or the homeostatic parameters to compensate for the effects of weight normalization (equation 3.2). Homeostatic parameters in the model have been discussed in section (for the smoothing parameter β) and section . The homeostatic learning rate η can also be varied within a wide range (approximately 100% of the value given here).

Initial connection weight parameters (σ_p) can be varied dramatically without changing the results, since these weights learn during the course of development. The parameters chosen here are to give a realistic approximation to what might be the initial state of these connection weights in the immature animal. Parameters defining the input pattern types can also be changed substantially, since this is one of the main benefits of the ALISSOM model, although dramatic changes in the statistics will require tuning of the frequency parameter as discussed in section 2.3.3.

Although many of the parameters in table 2.1 are given to several decimal places, this is not a reflection of how specific these parameters need to be. Instead this reflects the inability in such a large computationally intensive developmental modeling project to do an exhaustive parameter search. In this case many of the parameters are fixed to some reasonable arbitrary historic value which gives a good result, and only a small number of the parameters are tuned. The simulations presented in this chapter, and indeed in the rest of this thesis are intended to show the feasibility of concepts (such as using adaptive homeostatic mechanisms to ensure robust and stable development), rather than providing an exhaustive analysis of the parameter space. Exploring the space of these parameters and their interactions would be an extremely interesting and worthwhile future project.

2.3 Results

In order to highlight clearly the different homeostatic mechanisms required in ALISSOM, first results are presented from examples in which map development is driven by simple, abstract visual input (elongated Gaussian patterns). These mechanisms are then later used to reproduce the experimentally observed orientation map selectivity, stability and organization in a model that simulates spontaneous activity before eye

opening, followed by experience with natural images.

2.3.1 Orientation map stability in ALISSOM in comparison with previous SOM models

In a basic self-organizing map algorithm, when V1 neurons receive input from the LGN, lateral interactions lead to ‘bubbles’ of activity in V1 (figure 2.3). This drives the self-organization of the orientation map by ensuring that nearby neurons have similar orientation preferences. As the weights in the model change, receptive fields of V1 neurons become orientation selective and a closer match to the patterns in the input, leading to an increase in the responses of V1 neurons. Neurons across the map become orientation selective at different times, since each neuron receives a different sequence of activations.

In reality, there may also be differences in the initial thresholds for activation or the initial strength of connections between neurons. Therefore, some neurons in V1 will be activated more often at the start of development, be more excitable, or have stronger lateral connections, and will thus become orientation selective before others, increasing their responses more quickly. In general, (as discussed in section 2.1) for self organizing map models without homeostatic mechanisms, common but ad-hoc methods for controlling V1 activity (such as exponentially decaying learning rates, high initial excitability across the map, and changes in the lateral interaction radius during development) lead to unstable orientation map development.

For example, in the LISSOM model (Miikkulainen et al., 2005) the level of stability of map development is not consistent with the results from chronic optical imaging studies (Chapman and Bonhoeffer (1998); Godecke et al. (1997); see figure 1.9). There are several factors that contribute to this instability. Firstly, differences in the activation of individual neurons are minimized by ensuring that LGN input results in larger ‘bubbles’ of activation in V1 at the start of development (by starting with a larger radius of lateral excitatory connections). This increases the likelihood that every neuron will be activated at the start of development but also results in instability in the size of orientation domains. Secondly, a global change in the excitability of V1 neurons is manually implemented in the model during development. This ensures that neurons will all have a low enough threshold to be activated at the start of development, and that excitability is reduced as neurons become orientation selective. The rate of Hebbian learning in




Figure 2.6: Comparison of orientation map development in three self-organizing models.

Model development in all three cases is driven by elongated Gaussian input (one per iteration in (a) and two per iteration in (b) and (c)). Orientation preference maps from the beginning of development to the point at which the final map structure is formed are shown. In order to produce an orientation map responses of neurons/units are recorded while sine gratings of different orientations are presented to the model retina, as in experimental map measurement (see methods). Each neuron/unit in V1 is then color coded according to the orientation it prefers (as shown in the color key above). (a) Development in a simple SOM model (based on Obermayer et al. 1990; see methods). Various ad-hoc methods in this model, such as a shrinking of the radius of excitability in V1, and a decaying learning rate, ensure map formation but also lead to massive reorganization during development. (b) The LISSOM model (Miikkulainen et al., 2005) also contains similar ad-hoc methods which lead to instability. For example, the shrinking lateral excitatory radius leads to larger domains of similar orientation preference at the beginning of development. (c) The ALISSOM model remains very stable throughout the course of development. Once established, orientation preferences do not change and domain sizes remain constant. This model is a much better match to the experimental recordings of orientation map development (see figure 1.9). Maps are scaled to ensure an approximate match in orientation domain size at the end of development. This also results in differences in the number of units in each case.

afferent projections is also higher at the start of development to ensure that learning and reorganization occurs. The combination of an initial period of high learning rate and an initial period of high V1 responses means that neurons in the LISSOM model are initially very unstable in their preferred orientations.

In the SOM model presented in Obermayer et al. (1990), the algorithm finds the maximally active neuron in the cortex and adapts the afferent weights only in a specified circular (Gaussian) neighborhood around it. Similar to LISSOM, the area of this neigh-




Figure 2.7: **Stability of map development.** (a) Average (median) orientation selectivity for the simulations shown in figure 2.6. Selectivity across the simulated maps during development increases gradually in both LISSOM and ALISSOM as observed in optical imaging of animal maps, however, in the simple SOM network the selectivity increase is not as gradual. (Selectivity values are normalized to allow better comparison between models.) (b) Orientation similarity index during development for the simulations shown in figure 2.6. Correlation of orientation preference with the final map consistently increases during development for both LISSOM and ALISSOM, unlike the SOM simulation where the maps remain uncorrelated with the final map until nearly the end of the simulation. The ALISSOM simulation is much more stable than LISSOM, as the similarity index curve in this case increases more quickly at the start of development, as soon as the average selectivity begins to increase.

borhood is decreased throughout learning, and initial learning rates are higher, leading to a decrease in the size of orientation domains and a dramatic reorganization of the map during development. Orientation map development in the LISSOM model and in a simple SOM model based on Obermayer et al. (1990) is shown in figure 2.6. It is clear to see that there is a lot of reorganization during development in the SOM model (figure 2.6a), and early in development in LISSOM (figure 2.6b).

A main feature of the SOM model which may be key to explaining why this model requires such a large reorganization is that each V1 neuron begins with initial connections to the whole of the retina. This requires the SOM model to organize both retinotopic and orientation selective receptive fields simultaneously. The process of lateral neighborhood decrease and learning rate decrease in this model ensure that self-organization does not fall into a local minimum where the mapping is not retinotopic (i.e. they provide a kind of annealing of the map). However, in biology, maps are actually formed in a more sequential manner, with a coarse retinotopic map already in place before the orientation map. It is possible to speculate that the biological system may have made a trade-off, taking longer to form each map in turn, in order to ensure map stability.

The stability of the map can be quantified by comparing corresponding units in orientation maps measured during development with the final organized map (Chapman et al., 1996). The orientation preference differences for individual units are averaged across the map to obtain a mean similarity index. Identical maps have a mean orientation similarity index of 1, completely anti-correlated maps have a mean similarity index of 0, and uncorrelated maps have an intermediate similarity index of around 0.5. The SOM model (figure 2.6a) reorganizes massively, which is reflected in the similarity index measurement that remains close to 0.5 until very late in development (figure 2.7b). The LISSOM model, although reasonably stable in the latter stages of development does not resemble the final map as closely in the early stages, and therefore the similarity index does not increase quickly at the start of development (figures 2.6b, 2.7b). The quicker the similarity index reaches a value close to 1, the more stable the development. It is also important to consider the average selectivity during development (figure 2.7a) together with the stability curve, as in some cases a low stability can indicate a lack of selectivity rather than an unstable map. Stability is therefore characterized by a similarity index curve which increases quickly as soon as the average selectivity begins to increase.

In the ALISSOM model, the ad-hoc mechanisms which lead to instability in LISSOM have been removed, instead neurons in V1 automatically adjust their activation function in order to bring their activity (i.e., firing rate) into an exponential distribution with a fixed mean, using the gradient rule proposed by Triesch (2005). In this results section, a simple case where the input to V1 does not change throughout development is considered. In this case, at each iteration, two elongated Gaussian patterns are presented to the model retinal photoreceptors (figure 2.8). The automatic adjustment of the slope and threshold parameters of the activation function during development for several different neurons is shown in figure 2.8c,d. The final map organization, typical receptive fields, and the gradual increase in orientation selectivity across the map throughout development are shown in figures 2.7a and 2.8a. The ALISSOM model orientation maps contain all of the features observed in animal maps and develop stably as observed experimentally (figures 2.6c and 2.7b). Unlike the simple SOM and LISSOM examples, the map similarity index in ALISSOM increases quickly from 0.5 to 1 throughout development as selectivity increases.

In the ALISSOM model, the increase in V1 activity due to Hebbian learning is compensated by automatically reducing the intrinsic excitability of neurons as they become




Figure 2.8: **Automatic adjustment of intrinsic excitability in the ALISSOM model.** (a) Orientation preference map recorded using optical imaging at postnatal day 42 in ferret (left). Scale bar 2mm. Reprinted from Chapman et al., 1996. Simulated orientation preference map (right) for a simple simulation in which two randomly positioned and oriented elongated Gaussians are presented to the retina on each iteration (example shown in inset). Also shown (bottom) are receptive fields for five randomly selected simulated neurons. (b) Average activity during development for three example simulated model neurons. As each neuron becomes orientation selective, there is a transient increase in its average activity. In each case the target average activity ($\mu = 0.009$) of the neuron is maintained before and after the neuron becomes selective, despite differences in the timing of development of orientation selectivity for different cells. (c) Automatic adjustment of the ‘a’ parameter for three example neurons. This parameter controls the slope of the sigmoidal activation function. (d) Automatic adjustment of the ‘b’ parameter for three example neurons. This parameter controls the threshold of the sigmoidal activation function.

selective. The ‘modifiability’ of synapses in ALISSOM during Hebbian learning depends on V1 activity, and if activity is reduced as soon as the neuron becomes orientation selective, neurons are more likely to be activated only when their receptive field matches the stimulus well, strengthening the current connection weights rather than overwriting the previous orientation preference. This promotes orientation map stability. The adaptation of intrinsic excitability is also purely local, and so each neuron in the map can maintain the target average activity despite differences in the average activation across the map over time (figure 2.8b). Unlike SOM and LISSOM, broader

activation patterns are not required at the start of development to ensure that every neuron gets activated before excitability is reduced by the global changes in threshold. This also improves the stability of the orientation map in ALISSOM in comparison with these other SOM models.

There are two parameters in this initial simulation which require tuning. (1) The target average contribution from the afferent projections to the V1 activity (κ) and, (2) The target average V1 activity (μ ; mean of target exponential distribution). However, once the relationship between these two values is established, it does not require tuning throughout the rest of the simulations. All other parameters can be changed within reasonable ranges without affecting the results. Tuning can be achieved by setting κ to be the actual average contribution from the afferent projections and setting this equal to the target, i.e., no scaling will be required. The target average activity of a V1 neuron μ is then tuned in order to ensure good map organization and stable development. Organization is judged by observing how close the Fourier transform of the orientation map is to the ring shape that has been observed in animal maps (see figure 1.5); (Erwin et al., 1995)). Stability is judged by observing how quickly the similarity index curve increases during development during the period of development where average selectivity is increasing.

At this point, I then define a frequency parameter (v), which for this simulation (using 2 Gaussian inputs at each iteration) is arbitrarily chosen to be $v = 2$. From this point on μ is now defined in terms of v multiplied by a constant which is determined by the tuning of this simulation (i.e., for this simulation, the best value of $\mu = 0.09 = 0.045v$). The ratio between μ and κ is also fixed to the value obtained in this initial simulation, since their relationship defines the balance between afferent and lateral input required for good map organization. I find that the best relationship between these two parameters is $\kappa = 4\mu$. The only parameter changed in all subsequent simulations in this chapter (and only when the input pattern is changed) is the frequency parameter v .

2.3.2 Homeostatic regulation of balance.

In contrast to the previous simple case, the strength of retinogeniculate or thalamocortical drive is unlikely to remain fixed during development, particularly at the point of eye opening where driving activity changes from spontaneously generated patterns to natural visual input (Beckmann and Albus, 1982; Cragg, 1975; Hooks and Chen, 2006;




Figure 2.9: **Homeostatic regulation of balance rescues map selectivity and organization.**

Final orientation selectivity maps and representative receptive fields for three example simulations with and without afferent scaling. In the orientation selectivity map, each unit is color coded according to the orientation it prefers (as shown in the color key), and the saturation of the color indicates the level of orientation selectivity (how closely the input must match the unit's preferred orientation for it to respond). Two randomly positioned and oriented elongated Gaussians are presented to the retina on each training iteration (as in figure 2.8). In the 'strong' simulation (a), the strength of retinal photoreceptor output is 150%, and in the 'weak' case (b) 50%, of that in the balanced case (figure 2.8). In the 'changing' case (c) the strength of retinal output is raised to 200% at iteration 7500. Without afferent scaling, receptive fields are largely unoriented and orientation maps are unselective and/or poorly organized. With afferent scaling, selectivity and map organization are rescued.

Huberman et al., 2006; Tavazoie and Reid, 2000). This section considers how scaling of afferent inputs can ensure that this balance is maintained even if the average afferent input is changing.

The balance between afferent and lateral interactions has not been addressed in previous network implementations of intrinsic excitability adjustment (Butko and Triesch, 2007). In these simpler SOM networks there is an artificial 'winner take all' mechanism where the activity of the most active neuron in V1 is updated regardless of the driving activity from the LGN. In ALISSOM, it is crucial to maintain this balance.




Figure 2.10: **Afferent scaling in ALISSOM.** Average afferent activity for three different example V1 units, before and after scaling, for the simulations shown in figure 2.9, i.e., (a) strong inputs, (b) weak inputs, (c) changing inputs. The target average afferent activity in the models with scaling is $\kappa = 0.36$ and is marked by a straight line. The average activity before scaling is marked by a dotted line. The average activity after scaling is marked by a solid black line. The ALISSOM algorithm calculates a smoothed exponential average of the afferent activity in each case and subsequently scales these inputs to meet the target average afferent activity.

For example, if the average activity in the LGN is increased, intrinsic plasticity regulation will maintain a target mean activity in V1 and therefore a target lateral interaction strength. However, the afferent input strength is still increased. This imbalance leads to unselective receptive fields, because lateral interactions are no longer strong enough, in comparison with afferent input, to form ‘bubbles’ of activity in V1 that drive organization (figures 2.9a and 2.11). Similarly, if the average LGN activity is decreased, map organization is unstable because the pattern of V1 responses is determined by lateral rather than afferent interactions (figures 2.9c and 2.11).

In ALISSOM, balance is maintained by ensuring that for each neuron, the activity from afferent connections remains constant on average (figure 2.10). If the afferent activity increases or decreases on average, the afferent contribution will then be scaled post-synaptically to reach the target average afferent activity (figure 2.10). Maintaining the balance between afferent and lateral inputs rescues orientation map organization (figure 2.9b,d,f), selectivity (figure 2.11a), and stability (figure 2.11b).




Figure 2.11: **Homeostatic regulation of balance rescues stability of map development.** (a) Average orientation selectivity and (b) average orientation similarity index during development for the simulations shown in figures 2.9 and 2.10. Scaling of afferent input based on average afferent activity ensures that the balance between lateral and afferent input is maintained, rescuing the stability and selectivity of orientation maps despite changes in the strength of inputs. Stability is indicated by a stability curve which increases quickly as the average selectivity is increasing.

2.3.3 Targets for homeostatic regulation depend on input frequency.

In the previous examples, although the strength of the afferent input may have been changing, the likelihood that a V1 neuron would be activated at all was the same for every iteration during development, i.e., the patterns of activation were similar. However, neurons in V1 will be receiving various patterns of intrinsically generated activity before eye opening, and visual input after eye opening. In the ALISSOM model the stability of orientation map development is based on preserving the responses of V1 neurons, and the balance between afferent and lateral contributions to that response. This section considers how this stability is maintained despite the fact that, in reality, the patterns of visual input during development are unlikely to be constant.

The average activity of a neuron is a combination of the *magnitude* of the response when activated by a stimulus, and how *frequently* the neuron is activated. Therefore, a measure of average activity cannot distinguish between being activated more often, or having larger responses each time. Consequently, when viewing two very different types of visual scene (such as a starry sky versus dense underbrush), neurons in V1 maintaining an average activity, would also need to have some indication of how often they are likely to be activated in order to continue to maintain meaningful responses.

To illustrate, in the previous simple cases where two Gaussian patterns were presented




Figure 2.12: **Map development in ALISSOM for different types of input.** Examples of retinal activity (red square indicates the 1.0×1.0 sheet area to which V1 responds directly), subsequent V1 activity patterns, representative receptive fields, and orientation selectivity maps. Three simulations are shown (a) 4 Gaussian patterns per iteration, (b) randomly placed noisy disk patterns and (c) natural images. Selectivity values also depend on the likelihood of fine, high contrast edges appearing in the input, and so simulations using more abstract Gaussian patterns lead to higher average selectivity values, and therefore more saturated selectivity plots.

to the retinal photoreceptors at each iteration, intrinsic excitability adjustment preserves the responses of V1 neurons because the frequency of activation is kept constant. If the number of Gaussians presented at each iteration is increased to four, for example, neurons are twice as likely to be activated at each iteration. Without any change in μ , intrinsic plasticity adjustment would maintain a constant average V1 activity by decreasing the *magnitude* of the responses to individual activations. Therefore, in order to preserve average *responses* (rather than just average activity levels) neurons in the model must have some indication of how frequently they will be activated in order to choose the correct targets. For instance, if the number of Gaussians doubles, the target average activity μ should double to compensate and ensure normal development. In the ALISSOM model, this is achieved by tuning the frequency parameter (ν). The target average afferent activity κ is automatically regulated based on the parameters tuned in




Figure 2.13: **Adjustment of the frequency parameter (v) for different types of input results in stable map development.** (a) Average orientation selectivity across the map and (b) average orientation similarity index during development for the simulations shown in figure 2.12. Despite the vast difference in input types, development is stable and receptive fields are orientation selective. Three simulations are shown: 4 Gaussian patterns per iteration ($v = 3.56$), randomly placed noisy disk patterns ($v = 7.5$), and natural images ($v = 5$); $v = 2$ for the reference simulation that used 2 Gaussian patterns (figure 2.8). Selectivity values also depend on the likelihood of fine, high contrast edges appearing in the input, and so simulations using more abstract Gaussian patterns organize more quickly, and lead to higher average selectivity values, as expected.

the original simulation (as discussed in section 2.3.1).

For the simple case of Gaussian inputs, the frequency of V1 responses is approximately related to the number of Gaussian patterns in the input. However, for input images where activity is more broadly distributed across the retinal photoreceptors, the frequency of V1 responses can also depend on the contrast of the image, the strength of input, the amount of noise in the image, and the strength of lateral inhibition and excitation. For instance, for structured patterns on a background of noise, the frequency of V1 responses depends on whether just the structured edges result in V1 responses, or if noise in the LGN is also strong enough to activate V1 neurons. Similarly, for natural images, the frequency of V1 responses will depend on the exact nature of the image, the contrast of edges, and the strength of LGN afferents.

Orientation maps developed using ALISSOM for these different input types are shown in Figure 2.12. Examples of V1 responses in each case show that in images with more broadly distributed activity (noisy disks and natural images) it is not clear which features in the input are likely to result in activation in V1. For each simulation, the frequency parameter v is manually chosen to give the best results (see section 2.2). This manual adjustment leads to well organized orientation maps (figure 2.12), orientation selective receptive fields (figures 2.12 and 2.13a), and stable map development

(figure 2.13).

In the future, this adjustment of target activities could be achieved using a suitable homeostatic mechanism. For example, neurons could directly maintain their average responses (rather than just activity levels) independent of how often they are activated, i.e., the response of a neuron could be directly determined on a short timescale by the strength of its afferent input. Possibilities for automating this adjustment are discussed further in chapter 6.

2.3.4 Stable orientation map development before and after eye opening.

In this last set of results, the ALISSOM model reproduces the experimentally observed orientation map selectivity, stability, and organization in a simulation with retinal input that closely resembles realistic inputs during development (figures 2.14 and 2.15). In this example the initial map development is driven before eye opening using intrinsically generated activity (modeled after retinal waves; Wong 1999, and represented as noisy disk-shaped patterns randomly placed on the retina at each iteration; Bednar et al. 2004). After eye opening, the retinal inputs are image patches extracted from a data set of natural images (Shouval et al., 1996).

Results from this model are compared with experimental results from chronic optical imaging in ferret in figures 2.14 and 2.15. When the input type changes from noisy disks to natural images (at iteration 6000), the frequency parameter is manually set to the value used for that input type alone (as shown in figure 2.13) to ensure that responses are maintained. Average selectivity values (figure 2.15a) and similarity index values (figure 2.15b) increase after eye opening, as observed experimentally. The gradient of increase of the similarity index is similar in both the model and experimental data. The mechanisms described above ensure that the change in input type does not affect average V1 responses or the balance of lateral and afferent input, and therefore orientation preferences and the orientation map remain stable. The results demonstrate for the first time how realistic maps can develop in a stable way.




Figure 2.14: **A realistic model of stable map development before and after eye opening.**

(a) Simulated orientation selectivity maps. The hue indicates the preferred orientation and the brightness of the color indicates the selectivity; backgrounds are changed to black in this case and selectivity is normalized to the peak value for better comparison with the experimental results. Noisy disk patterns drive the map development until 6000 iterations, after which natural images are presented to the model retina (b) Polar orientation maps recorded using chronic optical imaging at different ages in one ferret (beginning at eye opening). Reprinted from Chapman et al., 1996. The hue indicates the preferred orientation and the brightness of the color indicates orientation selectivity. As the maps develop, the optical signal becomes stronger, and iso-orientation domains become visible as colored regions in the map. Scale bar 2mm. In both the model and the ferret data, similar orientation map features are maintained before and after eye opening.




Figure 2.15: **Comparison of selectivity and stability in ALISSOM and ferret.** (a) Average orientation selectivity across the map (normalized to peak) and (b) average orientation similarity index throughout development, for the simulation shown in figure 2.14a and the ferret shown in figure 2.14b. The red dotted line indicates the point of eye-opening in the ferret and the start of natural image input in the simulation. Ferret similarity indices are re-scaled to account for the reduced similarity due to the lower signal to noise ratio in experimental data in comparison with simulation results (see section 2.2).

2.4 Discussion

This chapter has presented a computational model of stable orientation map development that is robust to changes in the strength or pattern of visual activity. The model suggests that it is important for neurons in V1 to maintain the magnitude of their typical *responses* to an input during development, not simply their average activity levels as previously proposed. The model also highlights the importance of the balance between afferent input and lateral interactions for map development. Although each of the particular mechanisms implemented have a biological basis, it is necessary to draw attention to the fact that the same aims may be achieved by the combination of many different underlying mechanisms. For example, homeostatic regulation in the LGN or retina has not been included in this model. However, it is likely that similar mechanisms are occurring in these earlier regions of the visual pathway and may also contribute to stability in V1. By clarifying the higher level homeostatic requirements of V1 as it develops, this model should direct experimental efforts toward the investigation of the underlying mechanisms.

Homeostatic regulation is often described as maintaining the average activity of a neuron. However, these results show that for stable map development, it is important that the average *response* to a stimulus, rather than average activity, is maintained. If only an average activity is maintained, sustained changes in the strength or pattern of incoming activity will affect the post-synaptic response of neurons in V1. Together with ensuring stable and selective map development, maintaining the response to a stimulus in these ways is also likely to be important in the adult network to ensure the propagation of signals without loss of information (figure 2.1; Turrigiano and Nelson 2004).

This work has highlighted the necessity for a balance between afferent and lateral input in the cortex in order to maintain selectivity, and in turn stability. This aspect has been previously neglected in models of homeostatic plasticity in self-organizing maps (Butko and Triesch, 2007; Sullivan and de Sa, 2006) because these models do not require specific lateral interaction strengths, unlike the LISSOM based architecture used in ALISSOM.

Since the rate of Hebbian learning in the ALISSOM network depends primarily on the post-synaptic response, maintaining this response ensures that neurons will respond to pre-synaptic activation and therefore develop orientation selective receptive fields.

Once they become orientation selective, neurons will also compensate for their increase in responses, by reducing their intrinsic excitability and ensuring that they do not overwrite this initial orientation preference. This is consistent with experimental results which show that binocularly matched orientation preference maps can be formed despite a lack of common visual experience between the two eyes (Crair et al., 1998; Godecke and Bonhoeffer, 1996). In this case, homeostatic maintenance of V1 responses ensures that once orientation preference is established by spontaneous activity before eye opening, only those neurons that are strongly correlated have their connections strengthened further. Therefore, as soon as the initial binocular orientation preference is formed before eye opening, the map will maintain the same structure even without common visual experience.

For computational models, avoiding the need to search a large space of parameters when changing the input type makes the models much more robust and easier to use. The same is likely true for developing cortical neurons, which will robustly organize and develop regardless of the differences in afferent input (within reasonable bounds). In ALISSOM, the number of model parameters specific to each input type have been greatly reduced in comparison with LISSOM; typically only the frequency (ν) of the V1 responses needs to be specified. This model is now a much simpler starting point for subsequent models of map development. The model also shows that after eye opening, despite the massive reorganization occurring at this time, it is possible for cortical neurons to maintain a stable orientation map, matching that developed before eye opening, using plausible homeostatic mechanisms.

Chapter 3

A model of orientation map development consistent with the known wiring of V1

3.1 Introduction

Previous self-organizing map models, such as SOM (Obermayer et al., 1990) or LIS-SOM (Miikkulainen et al., 2005) have required ‘Mexican-hat’ connectivity in V1 (figure 1.13) in order to restrict cortical activity to a spatially localized ‘bubble’ of excitation. Spatially restricting cortical activity ensures that nearby neurons are simultaneously activated and will therefore learn to prefer similar features in the input. At longer distances, neurons will inhibit each other and learn to prefer different features. As described in more detail in chapter 2, if the input to the network consists of images that contain oriented lines, over the course of development these interactions will result in a topographic mapping of orientation preference. Maps for other features such as ocular dominance, and direction preference can develop in a similar way (Miikkulainen et al., 2005).

In the cortex, long-range lateral connections between V1 neurons originate from excitatory pyramidal cells and connect to primarily excitatory neurons, which comprise approximately 80% of cortical cells. These connections also contact a smaller percentage of inhibitory neurons, which comprise approximately 20% of cortical cells (Gilbert and Wiesel, 1989; Hirsch and Gilbert, 1991; Weliky et al., 1995). Long-range

connections generally show a periodic, patchy pattern of termination zones that have been found to link regions of similar orientation or direction selectivity (Buzás et al., 2006; Roerig and Chen, 2002) and can spread on average 6-7mm in macaque monkey (Angelucci and Bullier, 2003; Angelucci et al., 2002a,b). In some cases, patchy connections are extended along the retinotopic dimension that reflects the orientation that they represent (Bosking et al. 1997; figure 1.11).

The anatomy of local connections (figure 3.1) have been studied through labeling of neurons with tracer injections. It appears that neurons contact other neurons within $\approx 300 - 500\mu\text{m}$ of the injection site in rodents (Burkhalter, 1989; Hellwig, 2000; Holmgren et al., 2003; Matsuzaki et al., 2008), tree shrew (Bosking et al., 1997), cat (Gilbert and Wiesel, 1983; Kisvarday et al., 1997), ferret (Roerig and Chen, 2002; Roerig and Kao, 1999) and monkey (Katz et al., 1989; Malach et al., 1993). At the anatomical level (in carnivores and primates) there does not appear to be any anisotropy in these local connections and the spatial distributions of excitatory and inhibitory inputs are similar (although see chapter 4.6 for a discussion of the non-specificity of local connectivity in rodents).

In general less is known about inhibitory neuron connectivity. Despite composing only approximately 20% of cortical neurons, there appear to be a wealth of highly diverse inhibitory cell types (Gupta et al., 2000; Kisvarday et al., 2002; Markram et al., 2004; Yoshimura and Callaway, 2005). Large GABAergic basket cells in V1 have been shown to provide lateral projections up to 2mm in cat (Kisvarday and Eysel, 1993) and there is evidence that individual inhibitory basket cells may inhibit locally in an isotropic fashion while simultaneously providing anisotropic lateral inhibition at a slightly longer scale (Budd and Kisvarday, 2001; Buzas et al., 2001). In general however, in layers 2/3 of V1, the majority of longest range connections are excitatory (Gilbert and Wiesel, 1989).

There is no doubt that the biological evidence points to a more complex connectivity than just the long-range inhibitory and short range excitatory connections used in self organizing map models such as SOM or LISSOM. Regardless of the finer scale details, experimental results indicate that in general, in layer 2/3 of V1 (where orientation maps are measured), long-range connections are excitatory and shorter range local connections can be both excitatory and inhibitory. This appears to be at odds with the Mexican hat connectivity used to ensure development of topographic maps in self-organizing models. However, it may be possible to reconcile this more complex




Figure 3.1: Local connections between V1 neurons. (a) Tracer injection site in layers 2/3 of tree shrew V1. Axons are labeled and contact a local region around the injection site. Scale bar, $100\mu\text{m}$. Reprinted from Bosking et al. (1997) (b)(c) Examples of the distribution of cells within the orientation map in cat following an injection of tracer in an orientation domain (b) and a pinwheel center (c). Black: excitatory cells; white: inhibitory cells. The two outer concentric circles are at $250\mu\text{m}$ intervals from the injection site (inner circle). Scale bar: $250\mu\text{m}$. The connections between excitatory and inhibitory cells appear to be locally isotropic. Reprinted from Marino et al. (2005).

wiring with self-organization if we consider the effect of multiple connections between neurons, i.e., their overall effective interactions.

In this chapter I have reproduced orientation map organization using a model which contains long-range connections that are exclusively excitatory. Short range connections in the model are both excitatory and inhibitory, consistent with the experimental data. The LESI (**L**ong-range **E**xcitatory, **S**hort-range **I**nhibitory) model works by using di-synaptic inhibition to ensure that long-range excitatory connections actually have an overall inhibitory effect by acting on local inhibitory neurons.

There are a number of anatomical, functional, computational and behavioural observations that support this idea that long range connections in V1 can have an overall inhibitory effect despite being mainly excitatory. Many of these studies suggest that the net effect of long range connections (i.e., whether they are overall excitatory or inhibitory) depend on the properties of the stimulus (such as contrast). The topographic organization of V1 may also affect the dynamics of long-range interactions; another motivation for reconciling the underlying connectivity in V1 with orientation map development. Motivation for the LESI model is discussed in more detail in section 3.1.1.

3.1.1 Motivation

It has previously been suggested that long-range excitatory connections can have an overall inhibitory effect (Martin, 2002; Somers et al., 1998; Weliky et al., 1995). This




Figure 3.2: Di-synaptic Inhibition. This circuit can potentially explain how lateral interactions can depend on the input contrast. A long-range lateral connection from an excitatory cell contacts two pyramidal excitatory cells (large black triangles) and one inhibitory cell (large circle). The inhibitory cell has a high threshold for activation, but strongly inhibits the pyramidal cells when activated. Weak excitation activates the pyramidal cells monosynaptically, and does not activate the inhibitory cell. However, strong excitation activates the inhibitory cell as well, causing a net inhibitory effect. In this manner, a single incoming excitatory long-range lateral connection could have inhibitory effects for strong stimuli (e.g. high-contrast patterns), and excitatory effects for weak stimuli. Reprinted from Weliky et al. (1995).

idea rests on the fact that long range excitatory connections connect to local inhibitory neurons (Hirsch and Gilbert, 1991; Ren et al., 2007; Silberberg and Markram, 2007). The suggestion is that when long-range excitatory connections are sufficiently activated, their local connections to inhibitory cells will result in net inhibition. When more weakly activated, long-range connections will not sufficiently activate local inhibitory neurons and will therefore act to excite rather than inhibit their targets (figure 3.2).

A variable effect of long-range connections also gains support from studies of surround modulation in the adult V1. For example, one simple phenomena that is thought to arise from this underlying di-synaptic inhibition is contrast dependent receptive-field size tuning. Size tuning of a cell in V1 is experimentally measured by stimulation of a cell with a moving, high-contrast sinusoidal grating that has the optimal orientation, spatial and temporal frequency for the cell (figure 3.3). The size of the stimulus is increased until the response of the cell reaches a maximum, defined as the high contrast summation field (hsRF). The underlying di-synaptic inhibition is revealed by the dependence of this receptive field size measurement on the contrast of the stimulus. When a low-contrast, grating is used to define the low contrast summation-field (lsRF), the size obtained is around twice that of the hsRF obtained with a high-contrast grating.




Figure 3.3: Contrast dependent receptive field measurement. (a)(b) Contrast dependent receptive field measurement in macaque monkey. Curves show the response of a single cell in V1 to presentation of a circular patch of sine-wave grating at the cell's preferred position, orientation and spatial and temporal frequency. Stimuli are shown at low and high contrasts and at different stimulus radii. The vertical arrows indicate the radius of the peak response for low and high contrast. Reprinted from Sceniak et al. (1999). (c,d) Cartoon version of high contrast and low contrast receptive field measurement respectively of which (a) and (b) are examples. In general, for each cell, the size giving the peak response (and therefore preferred receptive field size) increases as contrast decreases. Reprinted from Schwabe et al. (2006).

Several experimental and computational modeling studies (Angelucci et al., 2002a,b; Dragoi and Sur, 2000; Somers et al., 1998; Stemmler et al., 1995; Sullivan and de Sa, 2006; Wielaard and Sajda, 2006) have shown that a combination of local and lateral (or feedback) connections in V1 are capable of reproducing this type of cellular level surround modulation effect using disynaptic inhibition. In particular, the recent model of Schwabe et al. (2006) suggests that using a low contrast stimulus, a larger area of the surround needs to be activated before inhibitory neurons are activated, leading to the observed contrast dependent size tuning. This idea rests on local inhibitory neurons having a higher threshold and gain than excitatory neurons, therefore being activated by long-range excitatory connections only when a stimulus reaches a particular size, and thus a level of excitation sufficient to make the inhibitory neuron respond.




Figure 3.4: **Perceptual level surround modulation.** Although the physical contrast of the ringed target is 40%, surround modulation makes it appear much lower. This type of perceptual level effect is thought to arise from underlying surround modulation at the cellular level in early visual areas. Reprinted from Dakin et al. (2005).

Other surround modulation phenomena measured in adult visual cortex depend not only on the contrast but also on the orientation, spatial frequency or position of the experimental stimulus (reviewed in Angelucci et al. 2002a,b; Series et al. 2003; see figure 1.8). These cellular level effects are thought to underlie many perceptual phenomena. For example, when looking at the central patch in figure 3.4 and matching the contrast of this patch with the examples on the left, subjects often estimate the contrast to be similar to either the 20% or 30% patch. However, the actual contrast of this central patch is 40%. The perceived contrast of the central stimulus is reduced by the high contrast stimulus in the surround. This strong link between cellular effects and perception makes the origin of surround modulation phenomena in V1 an interesting and important question for neuroscience research. Unfortunately, due to the number and interrelationship of stimulus parameters, the experimental results regarding surround modulation are difficult to interpret and often inconsistent. Sensitivity to the type of anesthetic, species, attentional factors, laminar position of the recorded cells and differences in the type of stimuli used in different studies all add to the diversity of behaviours reported.

It is also possible that discrepancies in the experimental results may be explained by the unaccounted for underlying map structure. The types of neuronal response properties, on which these behaviours appear to depend, such as orientation preference, are smoothly topographically organized across the cortical surface. As described above, this topographic organization is intimately tied with the intracortical connectivity i.e. each individual neuron will have a very specific connectivity depending on its position within the orientation map (figure 1.11). Therefore, in order fully describe these surround modulation effects, it appears to be crucial to take into account the structure of the orientation map. Interestingly, studies have suggested that extra-classical receptive field effects are indeed clustered within V1 (Yao and Li, 2002) and that particular features such as corners and T-junctions are preferentially detected by neurons within different map structures, due to the differential responses to flanking objects in the receptive field surround (Das and Gilbert, 1997, 1999). However, this may not be the case in all species; see Fitzpatrick (2000).

There are two obvious questions here to address. First, are models of self organizing maps consistent with a lack of long range-inhibitory connections in V1? The above discussion suggests that it is possible that these models are actually providing a simpler way of implementing the more complex circuitry that is actually present in layer 2/3 of V1. The long range inhibitory connections in the LISSOM model, for example, could be replaced with long-range excitatory connections and di-synaptic inhibition.

The phenomena of di-synaptic inhibition has received support from experimental and modeling studies of adult function, particularly surround modulation. Long range connections could have either excitatory or inhibitory effects depending on the ability of these connections to activate local inhibition. The current version of LISSOM would be operating in a 'high contrast regime' where long range inputs have only an overall inhibitory effect.

Second, does the structure of the orientation map and its underlying connectivity influence surround modulation at the cellular level? Ideally a full description of the orientation map and its underlying connectivity would be required to answer this question, together with the ability to test the model with many different experimental paradigms. The orientation map is a complex structure that arises during development, and unfortunately a complete experimentally measured description of the orientation map and its underlying connectivity are not yet available. However, a developmental self-organizing map model is the simplest method for providing a mathematical description

of this 2-dimensional orientation map structure together with its underlying connectivity. The LESI model will therefore contribute to answering the question of how connectivity influences surround modulation at the cellular level.

In this chapter I introduce and describe the LESI model of orientation map development. LESI is consistent with di-synaptic inhibition via long-range excitatory connections in V1, showing for the first time that it is possible to reproduce orientation map development using this more realistic architecture. LESI is also a starting point for the future understanding of surround modulation and its dependence on lateral interactions in both the adult and developing visual cortex.

3.2 Methods

The LESI model is based on the ALISSOM model presented in chapter 2. The differences between this architecture and the LESI architecture are described below, however, the majority of the details are as outlined in section 2.2.

The architecture of the LESI model is shown in figure 3.8. It is important to discuss why the LESI model has this particular architecture, since there are several different model architectures that are consistent with long-range excitation and di-synaptic short-range inhibition. A detailed investigation of these architectures leads to several subtle constraints that must be taken into consideration. Di-synaptic inhibition alone is not enough to reproduce successful orientation map development.

In order to motivate the LESI architecture I will first describe a simpler implementation of di-synaptic inhibition that fails to develop an orientation map (section 3.2.1). Section 3.2.2 then describes the slightly more complex LESI architecture used in the rest of the chapter.

3.2.1 Simple di-synaptic inhibition

The first important adaptation made to the ALISSOM model is to separate the V1 sheet into populations of inhibitory and excitatory units. This separation is more biologically realistic, since the majority of neurons produce a single type of neurotransmitter (either excitatory or inhibitory), and also allows for different connectivity between different neuron types. The model architecture from the photoreceptors to the LGN output

is identical to the ALISSOM model as shown in figure 2.3, and therefore only the structure of the V1 architecture is illustrated in this chapter.

In LISSOM, long-range inhibitory connections create a ring of inhibition that during settling sharpens the input activity, whereas short-range excitation strengthens the activity in the maximally active areas. In order to reproduce spatially restricted patterns with only short-range disynaptic inhibition, inhibitory neurons would need to be active over a larger area than excitatory neurons, ensuring that inhibition is stronger than excitation in a ring around the maximally active units.

The first simple model architecture that implements these requirements is shown in figure 3.5a. This architecture replaces long-range inhibition and short-range excitation in the ALISSOM V1 sheet with long-range excitation, short-range excitation and short range inhibition. The initial activity patterns in the V1 excitatory sheet (V1Exc) are spatially restricted bubbles, whereas the activity patterns in V1 inhibitory neurons (in the inhibitory sheet - V1Inh) have a broader activity pattern (figure 3.5b).

The combined interaction of the connections in this architecture is illustrated in figure 3.6. This illustration is a simplified view of the complex recurrent architecture, but gives an indication of the interactions at the beginning of the settling process that lead to the observed behaviour. Despite containing no long-range inhibitory connections, the long range excitatory connections to V1Inh and the short range excitatory connections in V1Exc lead to an overall combined interaction in V1Exc that has a Mexican hat profile. This architecture is sufficient to produce spatially restricted activation patterns in V1Exc at iteration one in the model. At each timestep the pattern of activity in V1Inh is broader because long-range excitatory connections to V1Inh neurons are stronger than long-range excitatory connections to V1Exc neurons. Although the result is a Mexican hat, these broader activity patterns in V1Inh begin to cause problems during learning and self-organization.

In this architecture, the activity of V1Inh neurons is determined only by long range excitatory connections. There are no local inhibitory to inhibitory self-connections or local excitatory to inhibitory connections. During learning in the model, initially random long-range lateral excitatory connections within the V1Exc sheet and between the V1Inh and V1Exc sheet are strengthened based on correlated activity. As these connections begin to learn, the new pattern of long-range connections to V1Inh neurons will determine their activity patterns. Because the patterns of activity in V1Inh and V1Exc




Figure 3.5: Simplest disynaptic architecture showing center surround interactions (a) This model architecture contains long-range excitatory connections from V1Exc units to both V1Exc and V1Inh units, short-range excitatory connections from V1Exc to V1Exc units and short range inhibitory connections from V1Inh to V1Exc units. Maximum connection radii are indicated by colored circles. Excitatory connections are shown in red and inhibitory in blue. Solid lines indicate connections between units in different sheets, and dashed lines indicate connections between units in the same sheet. (b) Initial activity patterns in response to a single oriented Gaussian on the retina (see inset) show that in V1Exc activity is spatially restricted in a 'bubble' as in LISSOM. However, activity in the V1Inh sheet is slightly broader because it is determined by long-range excitatory connections from V1Exc and therefore does not match the structure of the V1Exc activity exactly, i.e., there is no distinct border around the pattern in the V1Inh sheet as there is in V1Exc, and the bubble is slightly larger. The density of the V1Inh sheet is reduced in comparison with V1Exc in order to reproduce the 80:20 ratio of excitatory to inhibitory neurons in V1.

are initially different (although only subtly at iteration one; figure 3.5b) the pattern of connections to V1Exc and V1Inh will diverge during development. Therefore, during learning in this model, activity patterns in V1Exc and V1Inh become increasingly dissimilar (figure 3.7a).

The Mexican hat interaction in this simple architecture is highly dependent on the relationship between the activity patterns in V1Exc and V1Inh. As described above for the first iteration, at each timestep the pattern of activity in V1Inh needs to be broader than in V1Exc, in order to form blobs of activity in V1Exc and maintain the Mexican hat interaction. As the patterns of activity in V1Inh and V1Exc become in-




Figure 3.6: **Combined interaction of connections in simple disynaptic architecture.** Two-dimensional slice through V1Exc and V1Inh showing the profile of interactions in the architecture shown in figure 3.5. Line color and style reflects connection type as indicated in figure 3.5. Combined interactions are shown in solid green. (a) V1Inh units receive long range excitatory connections from V1Exc units. (b) V1Exc units receive both long and short range excitatory connections from other V1Exc units. (c) The net pattern of excitation and inhibition in V1Exc is shaped like a Mexican hat as in LISSOM. However, the activity patterns in V1Inh are determined only by long-range excitatory input.

creasingly dissimilar, this destroys the Mexican hat interaction in V1, and therefore causes a break-down in map organization. Orientation maps and orientation selective receptive fields do not form in this simple model architecture.

In the ALISSOM model presented in chapter 2, the organization of the orientation map depends on maintaining the pattern of activity in V1 during development. This is achieved in ALISSOM by regulating the strength of lateral interactions in response to a stimulus. In ALISSOM the patterns of inhibitory and excitatory activity are automatically similar because there is no separation of these neuron types and therefore no differences in their connectivity. However, in LESI, the separation of excitatory




Figure 3.7: **Effects of learning in simple disynaptic architecture.** (a) After 100 iterations in the simple disynaptic architecture, patterns of activity in V1Exc and V1Inh are not similar. (b) Long-range lateral excitatory connection weights for one unit in V1Exc and the corresponding unit in V1Inh. Lateral connection weights organize differently because of the differences in activity patterns (and subsequent correlations) in the two sheets. Initial small differences in activity patterns lead to increasingly different patterns in lateral connectivity. This eventually leads to a breakdown in the Mexican hat interactions in V1.

and inhibitory neurons means that additional mechanisms are required to constrain the activity patterns in V1.

Constraining the patterns of activity to be similar can be achieved by driving V1Inh neurons in a similar way to V1Exc neurons, i.e., ensuring that both sets of neurons receive the same set of connections. In this way, regardless of the changes that occur during learning, the overall interaction profile will still maintain a Mexican hat profile i.e. any changes in connection weights between neurons will affect V1Exc and V1Inh similarly. As long-range connections learn, the change in the pattern of lateral connections will no longer differentially effect V1Inh and V1Exc activity patterns. This will ensure that Mexican hat interactions will be maintained throughout development.

The LESI model additionally includes both local excitatory connections from V1Exc to V1Inh and inhibitory self-connections in V1Inh. This ensures that patterns of activity in V1Exc and V1Inh are similar throughout development, allowing orientation maps to form in V1. The LESI model is described in detail in the next section.




Figure 3.8: **LESI model architecture.** (a) In the LESI architecture there are short range excitatory connections between all units. Long-range connections arise from the excitatory neurons only and contact both excitatory and inhibitory units. Inhibitory units contact excitatory and other inhibitory units only over a short range. Maximum connection radii are indicated by colored circles. Excitatory connections are shown in red, inhibitory in blue, solid lines indicate connections between units in different sheets, and dashed lines indicate connections between units in the same sheet. (b) Initial activity patterns in response to an oriented Gaussian pattern (inset) indicate that patterns are similar in both V1Exc and V1Inh. Although there is a small ‘halo’ of activity in V1Inh, there is a distinct edge to the activity blob in V1Inh corresponding to the edge of the blob in V1Exc that is not present in the architecture shown in figure 3.5. The density of the V1Inh sheet is reduced in comparison with V1Exc in order to reproduce the 80:20 ratio of excitatory to inhibitory neurons in V1.

3.2.2 LESI model architecture

The architecture of the LESI model is shown in figure 3.8a. In this architecture there are recurrent, short range excitatory connections between all units. Long-range connections arise from the excitatory neurons only and contact both excitatory and inhibitory units. Inhibitory units contact excitatory and other inhibitory units only over a short range. Having the same set of connections to every V1 neuron ensures that all excitatory and inhibitory units receive similar connections, leading to similar activity patterns (figure 3.8b).

As discussed above, the ability of the LESI model to self-organize and develop an orientation map is dependent on the ability of the interactions between neurons to produce

spatially localized ‘bubbles’ of excitation in V1 in response to images presented on the retina. The ‘halo’ of activity in the V1Inh sheet (figure 3.8b) gives an indication of why this occurs. The relative strengths of connections within V1 have been chosen such that the overall effect of long-range excitatory inputs is inhibitory, i.e., recreating the Mexican hat interaction via di-synaptic inhibition. Mexican hat interactions in this model are illustrated in figure 3.9, this illustration is a simplification of the complex recurrent activation that occurs during the process of settling, but gives an indication of how the relative connection strengths can combine at the beginning of the settling process.

The choice of relative projection strengths (see table 3.1) that achieve this Mexican hat interaction was reached through a time-consuming method of reasoning and trial and error parameter searching. Unfortunately the nature of this type of non-linear, developmental model, together with the large space of possible parameters, makes an exhaustive parameter search computationally infeasible. Ultimately, the ability of the model to self-organize and form a realistic orientation map is dependent on the activity patterns in V1Exc and V1Inh. These are determined by both the strength of lateral projections and the intrinsic excitability of the neuron units (determined by the sigmoidal transfer function). Consequently there are likely to be other sets of strength and transfer function parameters that also achieve similar results. In the future we hope to further constrain these parameters using experimental evidence and/or automatic homeostatic mechanisms (see chapter 6).

The LESI model highlights the requirements that V1 connections must satisfy in order to reproduce orientation map development. Excitatory and inhibitory neurons must receive similar connections in order to maintain similar patterns of activity and thus the relationships between interaction strengths during development. Inhibition must be strongly activated in a spatial range that exceeds excitation, allowing Mexican hat interactions to be maintained, and activity patterns to become spatially restricted during settling. In LESI this is achieved by having strong long range excitatory connections to V1Inh neurons. It is possible that this could also be achieved by regulating the gain of inhibitory neurons, i.e., long range connections could be relatively weak but could activate inhibitory neurons strongly.




Figure 3.9: **Combined interaction of lateral connections in the LESI model.** Two dimensional slice through V1Exc and V1Inh showing the profile of lateral interactions in the architecture shown in figure 3.8. Line color and style reflects connection type as indicated in figure 3.8. Combined interactions are shown in solid green. (a) V1Inh units receive both short range and long range excitatory connections from V1Exc units. V1Inh units also receive short range connections from other V1Inh units (not shown). (b) V1Exc units receive both long and short range excitatory connections from other V1Exc units. (c) The combined interactions of excitation (red) and inhibition (blue) in V1Exc yield a Mexican hat shaped interaction pattern (green). The combined interaction in V1Inh is not shown since this is a non-linear function of the recurrent excitatory connections to V1Inh and cannot be illustrated accurately. However, in this architecture, as a result of receiving similar connections, V1Inh units also have a similar pattern of activity as V1Exc units.

3.2.3 Model parameters

The LESI model consists of five sheets that represent the input (retinal photoreceptors), RGC/LGN On and Off cells, and V1 excitatory (V1Exc) and V1 inhibitory (V1Inh) cells. The density of units per 1.0×1.0 area are as in section 2.2. All parameters are identical to those in table 2.1 except those indicated in table 3.1. The equations are as defined in chapter 2 for the ALISSOM model except that the V1 sheet is now split

into V1Exc and V1Inh. Therefore the main equations relating to V1 activation are reproduced here for clarity.

Recurrent lateral connections are updated within both V1Exc and V1Inh at every settling step t_s . The number of settling steps in V1Inh is one less than for V1Exc in order for all projections to arrive in V1Exc and be updated before settling finishes. The contribution X_{jp} to the activation of unit j in the V1 (either V1Exc or V1Inh) sheets at time t from each projection type p , is updated in each sheet at every settling step t_s :

$$X_{jp}(t_s) = \sum_{i \in F_{jp}} \eta_{i,V1}(t_s - 1) \omega_{ij} \quad (3.1)$$

Where $\eta_{i,V1}$ is the activation of unit i taken from the set of neurons in either V1Exc or V1Inh to which unit j is connected (its connection field F_j) and ω_{ij} is the connection weight from unit i in V1 (either V1Exc or V1Inh) to unit j in V1 (either V1Exc or V1Inh) for the projection p .

Hebbian connection weight adjustment at each iteration is dependent on the pre-synaptic activity, the post-synaptic response, and the Hebbian learning rate:

$$\omega_{ij}(t) = \frac{\omega_{ij}(t-1) + \alpha \eta_j \eta_i}{\sum (\omega_{ij}(t-1) + \alpha \eta_j \eta_i)} \quad (3.2)$$

where for unit j , α is the Hebbian learning rate for the connection field F_{jp} in projection p .

For the V1Exc sheet the projections p are: afferent connections from the LGN (On and Off combined, A), short range lateral excitatory from other V1Exc neurons (*EtoEshort*), long range lateral excitatory from other V1Exc neurons (*EtoElong*) and short range lateral inhibitory from V1Inh neurons (*ItoEshort*). For the V1Inh sheet the projections p are: short range lateral excitatory from V1Exc neurons (*EtoIshort*), long range lateral excitatory from V1Exc neurons (*EtoIlong*) and short range lateral inhibitory from other V1Inh neurons (*ItoIshort*).

The contributions from all projections p described above are combined in each sheet (either V1Exc or V1Inh) using equations 3.3 and 3.4 to calculate the activation of a neuron j in V1 (either V1Exc or V1Inh) at time t ($\eta_j(t)$):

$$\eta_j(t) = \sigma(t) = \frac{1.0}{1.0 + e^{(-a(t)x(t) - b(t))}} \quad (3.3)$$

where σ is a sigmoid transfer function. In this case only the a and b parameters in the V1Exc sheet are varying with time, a and b in V1Inh are fixed. The input activity $x(t)$

in each sheet is a sum of the contributions from all projection types in that sheet:

$$x(t) = \sum_p \gamma_p X_{jp}(t) \quad (3.4)$$

The projection strength scaling factors for each projection type γ_p are given in table 3.1.

Parameter	Value	Description
r_A	0.271	Maximum radius of the afferent connections from LGN (On and Off) to V1Exc
r_E	0.02	Maximum radius of all short-range connections, section 3.2.2
r_I	0.229	Maximum radius of all long-range connections, section 3.2.2
σ_A	0.542	Standard deviation of initial Gaussian random afferent connection weights, equation 2.11
σ_E	0.04	Standard deviation of Gaussian short-range connection weights, equation ??
σ_I	0.458	Standard deviation of initial Gaussian random long-range connection weights, equation 2.11
γ_L	2.33	Strength scaling factor for the LGN's afferent weights, equation 2.1
γ_A	1.0	Strength scaling factor for the V1 afferent weights to V1Exc, equation 3.4
$\gamma_{EtoEshort}$	0.7	Strength scaling factor for short-range lateral excitatory weights to excitatory neurons, equation 3.4
$\gamma_{EtoElong}$	0.1	Strength scaling factor for the long-range lateral excitatory weights to excitatory neurons, equation 3.4
$\gamma_{EtoIshort}$	0.6	Strength scaling factor for short-range lateral excitatory weights to inhibitory neurons, equation 3.4
$\gamma_{EtoIlong}$	1.5	Strength scaling factor for long-range lateral excitatory weights to inhibitory neurons, equation 3.4
$\gamma_{ItoIshort}$	-0.35	Strength scaling factor for short-range lateral inhibitory weights to inhibitory neurons, equation 3.4
$\gamma_{ItoEshort}$	-0.35	Strength scaling factor for short-range lateral inhibitory weights to excitatory neurons, equation 3.4
α_A	0.137	Afferent projection learning rate, equation 3.2
Continued on next page		

Table 3.1 – continued from previous page

Parameter	Value	Description
$\alpha_{EtoEshort}$	0.0	Learning rate for short-range lateral excitatory projection to excitatory neurons, equation 3.2
$\alpha_{EtoElong}$	1.808	Learning rate for the long-range lateral excitatory projection to excitatory neurons, equation 3.2
$\alpha_{EtoIshort}$	0.0	Learning rate for short-range lateral excitatory projection to inhibitory neurons, equation 3.2
$\alpha_{EtoIlong}$	1.808	Learning rate for long-range lateral excitatory projection to inhibitory neurons, equation 3.2
$\alpha_{ItoIshort}$	0.0	Learning rate for short-range lateral inhibitory projection to inhibitory neurons, equation 3.2
$\alpha_{ItoEshort}$	0.0	Learning rate for short-range lateral inhibitory projection to excitatory neurons, equation 3.2
$a_{initV1Exc}$	14.5	Initial slope parameter in V1Exc sigmoid transfer function (single value rather than randomly chosen as in ALISSOM), equation 2.9
$b_{initV1Inh}$	-4.0	Initial threshold parameter in V1Exc sigmoid transfer function (single value rather than randomly chosen as in ALISSOM), equation 2.10
a_{V1Inh}	12.0	Initial slope parameter in V1Inh sigmoid transfer function (single fixed value rather than randomly chosen as in ALISSOM, equation 3.3
b_{V1Exc}	-5.5	Initial threshold parameter in V1Inh sigmoid transfer function (single fixed value rather than randomly chosen as in ALISSOM), equation 3.3
t_{sV1Inh}	50	Number of settling iterations in V1Exc, equation 3.1
t_{sV1Exc}	49	Number of settling iterations in V1Inh, equation 3.1

Table 3.1: LESI model parameters Model parameters used in the LESI simulations presented in this chapter. Lengths shown are in sheet dimensions, where the size of the V1 sheet is 1.0×1.0 nominal area in sheet dimensions. All other parameters are as in ALISSOM model (table 2.1).

Afferent connections from the LGN On and Off cells contact only neurons in the V1Exc sheet. These connections are scaled similarly to the afferent connections in the ALISSOM model. For each lateral projection p between neurons in LESI, the

multiplicative scaling factor for their overall strength γ_P is shown in table 3.1. The number of settling steps is increased in LESI in comparison with ALISSOM, in order to allow the more complex set of connections to form spatially restricted blobs of activity during settling. Recurrent lateral connections from V1Exc are updated 50 times per image presentation (one training iteration) and recurrent lateral connections from V1Inh are updated 49 times per image presentation (at each settling step t_s).

The average strength of each lateral projection also depends on the average responses in the V1Exc and V1Inh sheets. As discussed in chapter 2, average responses are also determined in part by the sigmoidal transfer function of the neurons. The parameters of the sigmoidal transfer function $\sigma(x)$ for each unit in V1Exc are adapted using the homeostatic intrinsic plasticity rule proposed by Jochen Triesch (Triesch, 2005) as described in section 2.2. The initial slope and threshold parameters for all neurons in V1Exc are $a = 14.5$ and $b = -4.0$. In ALISSOM these values are initially randomly chosen from a uniform distribution. However, the LESI model in its current form is not as robust as ALISSOM to randomness in this initial parameter. This is discussed in more detail in section 3.4. The sigmoidal output function in V1Inh remains fixed throughout development ($a = 12.0$ and $b = -5.5$). As V1 inhibitory neurons do not receive direct afferent input, the adaptation of intrinsic excitability in the V1Exc sheet is sufficient to control V1Inh average activity levels.

The target average activity for the V1Exc sheet μ_{exc} is chosen, as before, as a function of the frequency parameter v that reflects how frequently a V1 neuron is activated. As for the ALISSOM model, $\mu_{exc} = 0.0045v$ where $v = 2$ for simulations with 2 oriented Gaussians presented per iteration, $v = 7.5$ for noisy disks, and $v = 5$ for simulations using natural image input.

As discussed in section 3.1, only 20% of cortical cells are inhibitory. In the model the relative number of excitatory and inhibitory units is determined by the density of the V1Exc sheet relative to the density of the V1Inh sheet (as in figure 3.8b). However, it should be noted that this distinction between neuron numbers is not necessary for the model's success. The crucial parameters in the model are the overall strength of the connections between the different neuron populations. Using a larger number of more weakly connected inhibitory neurons is equivalent to using fewer more strongly connected neurons, because divisive weight normalization in the model ensures that any changes in sheet density are compensated by increasing connection weight values. Reducing the density of V1Inh to lower values in the model can introduce practical

problems due to the constraint of using a regular grid of neuron units, necessitating the use of high, computationally intensive densities for the V1Exc units. Therefore, because the results are the same, below I use simulations in which similar numbers of excitatory and inhibitory units are used in each sheet.

3.2.4 Learning

As in section 2.3.4, initial map development in LESI is driven before eye opening using intrinsically generated activity (modeled after retinal waves; Wong 1999, and represented as noisy disk-shaped patterns randomly placed on the retina at each iteration; Bednar et al. 2004). After eye opening, the retinal inputs are image patches extracted from a data set of natural images (Shouval et al., 1996), see section 2.2.2.

As discussed in section 2.2.5, the lateral projection learning rates are constant throughout the simulation and short-range connection weights are fixed (enabling plasticity for these short-range connections has little effect on the simulation). Learning rates for each projection are shown in table 3.1.

All other parameters are as described in section 2.2. As discussed in detail in section 2.2.8 many of the ALISSOM parameters can be varied within a wide range (50-100% of the values given here), this also applies to the new LESI model parameters such as the connection radii (r_p), initial weights, learning rates (α_p), and settling times (t_s). However, the LESI model is not as robust to variation in the initial sigmoid parameters (a and b) or to changes in the strength scaling factors γ_p . This is because it is these parameters which determine the ability to produce mexican-hat interactions. Although one strength scaling factor could be varied within approximately 25% of its value, any greater variation would require compensatory changes in the other projection strengths.

3.3 Results

The main result of the LESI model is that it is possible to reproduce orientation map development in a model that contains long-range excitatory connections only. Orientation maps developed in the LESI model are described in section 3.3.1 and receptive fields and final connectivity are described in section 3.3.2.




Figure 3.10: **Activity patterns in the LESI model.** Activity patterns in the LESI model during development in response to a small natural image patch on the retina (red square indicates the 1.0×1.0 sheet area to which V1 responds directly). Activity patterns in V1Exc and V1Inh are similar during development (unlike those in the simpler disinaptic model, see figure 3.7).

3.3.1 Excitatory and inhibitory orientation maps in LESI model

The LESI model develops realistic orientation preference maps that are similar for both excitatory and inhibitory neurons (figure 3.11). The orientation maps contain all of the observed features of animal maps, i.e., linear zones, saddle points, pinwheel-centers and fractures, and have ring-shaped Fourier transforms. The Fourier transform of these maps is not as clearly ring-shaped as those observed experimentally, suggesting that there are map features at a wider range of spatial scales. Figures 3.8b and 3.10 show that activity bubbles in V1Exc are accompanied by a faint extra ‘ring’ of activity. This ringing is likely to be affecting the map organization but, as discussed in section 3.2.2, may be improved by using additional homeostatic mechanisms, or by performing a more exhaustive parameter search (which is currently computationally infeasible).

Final orientation maps in the V1Exc and V1Inh sheets are similar because activity pat-





Figure 3.11: **Matching maps in the LESI model.** (a) Orientation preference maps in V1Exc and V1Inh. Noisy disk patterns drive the map development until 6000 iterations, after which natural images are presented to the model retina. The hue indicates the preferred orientation, as shown in the color key. Preference maps in both V1Exc and V1Inh contain all of the features of experimentally measured animal maps (see figure 1.5). (b) Fourier transform of the central portion of the maps shown in (a), edges are excluded from the analysis in order to ignore the edge effects due to lateral connections in the model. Fourier transforms are approximately ring shaped, as in animal maps (see figure 1.5).

terns during development are similar, as shown in figure 3.11, i.e., both V1Exc and V1Inh neurons will learn similar correlations between the input patterns and V1 activity.

3.3.2 Receptive fields and connectivity in LESI model

The receptive fields of excitatory and inhibitory neurons in the LESI model were measured using reverse correlation at the end of development. Example pairs of excitatory and inhibitory neurons measured at the same retinotopic location are shown in figure 3.12. Inhibitory and excitatory neurons in LESI are similarly tuned for orientation, but inhibitory receptive fields are also slightly larger than excitatory receptive fields in

V1 Excitatory



V1 Inhibitory




Figure 3.12: **LESI receptive fields.** Examples of excitatory and inhibitory receptive fields of pairs (at the same retinotopic location) of LESI model neurons measured using reverse correlation at the end of development. Receptive fields are orientation selective and resemble receptive fields measured in real V1 neurons. Inhibitory cells, although similarly tuned, appear to have slightly larger receptive fields.

general (qualitative observation). This is again the result of the slightly broader activity patterns or ‘halo’ of activity observed in the V1Inh sheet in comparison with the V1Exc sheet.

Examples of the long-range lateral excitatory connection weights between excitatory neurons and from excitatory neurons to inhibitory neurons are shown in figure 3.13. Connections are strong locally, but there are also longer range patchy connections that contact other neurons with similar orientation preferences, (figure 3.13a) as observed experimentally (figure 1.11). Patterns of lateral connection weights between excitatory neurons and from excitatory to inhibitory neurons are similar (figure 3.13b). The range of these long-range connections will be calibrated in future implementations of LESI in order to model the specific lateral interaction radii in specific species.

3.4 Discussion

The LESI model presented in this chapter is more consistent with the known wiring of V1 layer 2/3 than previous self-organizing map models. The LESI model, rather




Figure 3.13: **Long range lateral excitatory connections in LESI.** Examples of long range excitatory connection weights from excitatory neurons to other excitatory neurons (V1Exc to V1Exc) and to inhibitory neurons (V1Exc to V1Inh) at the end of development. Nine examples that correspond to excitatory and inhibitory neurons at the same retinotopic location in the two sheets are shown. (a) Strong connections are shown as white on a black background and the orientation preference of neurons is indicated by the color of connections (as shown in the color key). Connections are strongest locally where neurons prefer similar positions and orientations, but also contact patches of neurons with similar orientation preference (although not strictly) further away in the orientation map. (b) The same neurons as in (a) with color coding for orientation preference removed and strong connections shown in black in order to highlight the spatial structure of connections. Long range connections contact both excitatory and inhibitory neurons in similar locations because they learn similar correlations in V1 activity. This is due to the similar patterns of activity in excitatory and inhibitory neurons throughout development.

than undermining previous models, suggests that they may be accurate in simplifying the connectivity between neurons to long-range inhibitory connections and short-range excitatory connections to model development.

In the LESI model architecture, if the interaction strengths between neurons satisfy the requirements outlined in section 3.2.2, i.e., that inhibition is strong enough to form spatially restricted bubbles of excitation, and that a similar pattern of activity is present in the inhibitory and excitatory neurons, then orientation maps will develop. LESI shows that these requirements can be satisfied by using the di-synaptic inhibition also thought to underlie many surround modulation phenomena in the adult cortex. This property makes LESI a good starting point for the investigation of the relationship between two dimensional map organization and surround modulation function.

There may also be other possible architectures that will reproduce the desired behaviour. For example, there is still a question over whether or not there should be afferent connections from the LGN to inhibitory neurons. The current LESI model does not contain these connections but they should not affect the ability of LESI to develop maps. It may be necessary to include these connections when the model is constrained further by adult function, however, the current model highlights only those connections required for successful map development.

Although previous models are simpler to implement and to understand, and have fewer parameters, there are benefits to a model that more closely resembles the underlying connectivity. In particular, the long-range patchy connections observed in the LISSOM model are retained in LESI. LISSOM is among the few models (see also Burger and Lang 2001; Weber 2001) that have specific modifiable lateral connections that learn to represent the correlations in V1 activity, thus reproducing the properties of long-range lateral connections observed in V1 (figure 1.11). However, LISSOM cannot reproduce many of the observed effects in V1 that arise from long-range excitation when stimuli are at low contrast. With LESI, it is now possible to investigate this low-contrast behavior.

Future work using the LESI model to investigate the effect of topographic organization on surround-modulation requires several additional steps outside the scope of this thesis, discussed in detail in chapter 6. For example, a first important extension of the LESI model would be to include surround modulation arising earlier in the visual pathway. The LESI model should also be tuned to represent a specific species for which

surround modulation data is available. The ultimate aim would be to test this model with a very specific set of experimental paradigms in order to ensure that the model can reproduce the known effects. At this point, the model could then be used to make predictions about how topographic organization may affect function, or if indeed the cortex somehow compensates for differences in local topography.

Another important benefit of the LESI model, in comparison to LISSOM and other simpler developmental models, is the separation of excitatory and inhibitory neurons. There is much to be gained by being able to separately manipulate excitation and inhibition, and the specific connections between inhibitory and excitatory neurons. For example, experiments in which the balance between lateral excitation and inhibition is manipulated during development have shown that this balance can control the width of ocular dominance columns (Hensch and Stryker, 2004) and may be responsible for determining the scale of topographic maps in general. The LESI model may therefore provide further predictions for the importance of excitatory/inhibitory balance in determining the overall structure of topographic maps.

Separating excitatory and inhibitory neurons is also intimately related to homeostatic plasticity. In the previous chapter, the higher level requirements for homeostatic plasticity in the developing orientation map were highlighted. Individual neurons in the ALISSOM model made both excitatory and inhibitory connections, requiring both these populations of neurons to be regulated together. In the LESI model, inhibitory and excitatory neurons can separately regulate their intrinsic excitability to maintain their responses and/or average activity (although currently the level of inhibition is determined by the level of excitation). This separation adds a layer of complexity that can be further investigated and exploited. Although the previous chapter highlighted the requirement for a balance between lateral and afferent interactions in V1, the LESI model can further explore the homeostatic regulation of the balance between excitation and inhibition during development.

The LESI model has successfully shown how topographic maps, in particular the map of orientation preference, can develop in a model that has no long range inhibitory connections. This model shows that inhibitory connections that are longer range than excitatory connections are not required for development, provided that the combined interaction of long-range excitation with local inhibition leads to Mexican hat interactions between V1 neurons. Future work to extend this model and investigate the low-contrast regime in V1, together with the effect that topographic organization may

have on functional properties such as extra-classical receptive field effects is now possible in this new architecture.

Chapter 4

Organization of orientation preference in mouse V1

4.1 Introduction

Rodents have been used as a model system in neuroscience for many decades, and have been essential for major scientific breakthroughs in the basic understanding of brain function and the study of neurological disease. However, it is only in recent years that rodents have been used widely for visual neuroscience research. Rodents are not thought to be highly visual mammals, and rely primarily on their olfactory and somatosensory systems (especially whiskers) for navigation and object recognition (Diamond et al., 2008). The visual systems of larger mammals such as monkeys, cats, and ferrets are also much less fragile and easier to study experimentally than those of rodents. Despite this, the emergence of powerful genetic techniques for manipulating structures in the rodent brain (particularly the mouse), together with advances in imaging and recording techniques, have led to a shift towards using rodents to study the visual system.

Electrophysiological studies of the receptive field properties of cells in mouse V1 (Dräger, 1975; Mangini and Pearlman, 1980; Métin et al., 1988), and in particular a recent comprehensive study by Niell and Stryker (2008), have determined that the basic properties of visual processing that have been studied in other mammalian species are also present in mice. Cells in mouse V1 are tuned for orientation and spatial frequency and have contrast-invariant tuning properties similar to those observed in higher mam-




Figure 4.1: **Receptive fields of mouse V1 neurons.** Six example receptive fields measured using spike triggered average. Example cells with two (A1, A2), three (B1, B2) or one sub-field (C1, C2) with various orientations, ON/OFF centers and spacing of subfields are shown. Reprinted from Niell and Stryker (2008).

mals. Both simple and complex cells are also observed (see figure 1.2 in chapter 1.1). Typical mouse receptive fields measured using reverse correlation/ spike triggered average are shown in figure 4.1. The majority of excitatory cells in layer 2/3 of mouse V1 are orientation selective (71%) and most of these are simple cells (75%). The total percentage of orientation selective neurons in mouse V1 (74%) is lower than the value for cat (>95%, Van Hooser, 2007), but the tuning width (median half width at half height for neurons showing some orientation selectivity) is comparable between the two species (Niell and Stryker, 2008).

Another recent study has also confirmed that, as in higher mammals, there are multiple distinct visual areas in mouse visual cortex (Wang and Burkhalter, 2007). In this study, three different fluorescent tracers were injected into three distinct areas of mouse V1. The subsequent distribution of these tracers in other areas connected to V1 showed that there are at least 9 different areas in the mouse cortex that have a topographic representation of the visual field (figure 4.2).

Therefore, it seems that both at the single cell level, and in large scale organization, there are many similarities between rodents and higher mammals. However, at the intermediate level of individual cortical areas, there do appear to be significant dif-




Figure 4.2: **Distinct visual areas in mouse cortex.** Maps are generated by making three simultaneous injections of fluororuby (red), fluoroemerald (green) and biotinylated dextran amine (yellow) into different adjacent locations in V1. These tracers are transported via anterograde transport in intracortical connections. Fluorescently labelled axons that respect the retinotopic relationship between the 3 different injections can be observed in at least 9 different cortical locations as indicated by the dashed lines. Scale bar 0.1mm in inset. Reprinted from Wang and Burkhalter (2007).

ferences in functional organization. As discussed in section 1.3, two-photon calcium imaging has recently confirmed that the organization of orientation preference in V1 in rodents is very different from that of higher mammals (figure 1.6). The next section discusses the functional organization of rodent primary visual cortex at the level of topographic maps in more detail.

4.1.1 Imaging of organization in rodent visual cortex

Traditional optical imaging methods for measuring topographic maps (as described in figure 1.4) measure optical signals that result from activity-related changes in the cortex, either generated intrinsically by changes in reflectance properties (Zepeda et al., 2004) or by voltage sensitive dyes (Grinvald and Hildesheim, 2004). Optical imaging of this type can be used to measure the averaged preferences of neurons occupying a particular area in a map at the scale of millimeters (under favourable conditions, optical imaging can provide activity maps with a spatial resolution of 80 – 100 μ m) but cannot resolve the preferences of individual neurons (which are approximately 10 μ m in




Figure 4.3: In vivo two-photon calcium imaging. (a) Schematic of two-photon imaging. A calcium-sensitive fluorescent dye is ejected from a micropipette into the extracellular space of the cortex, from which it diffuses into cells. Intracellular mechanisms render the indicator dye membrane-impermeable once it has entered the cell. Activity-dependent changes in intracellular calcium concentration can be monitored using two-photon fluorescence microscopy in a depth of up to $500\mu\text{m}$. Additional labeling with genetically encoded fluorescent proteins enables identification of specific cell types (e.g. interneurons). (b) Reconstruction of a 3D volume in mouse V1. Scale bars $50\mu\text{m}$. (c) Fluorescence image of one slice through the imaging volume, parallel to the V1 cortical surface. (d) Response of three neurons in the imaged volume to the direction of a drifting grating stimuli. Traces of $\Delta F/F$ (where F is fluorescence signal) versus time reflect changes in calcium concentration. Reproduced with permission from Tom Mrsic-Flogel.

diameter). Optical imaging of large scale populations of neurons can therefore not rule out the possibility of differences in organization of neurons at a smaller scale. Two-photon calcium imaging, however, provides a method that can measure topographic maps at the scale of hundreds of individually resolvable neurons in a small portion of the brain, and can therefore clarify smaller-scale organization.

Two-photon calcium imaging (Gbel and Helmchen, 2007; Stosiek et al., 2003) is described in figure 4.3. An infrared laser beam is focused through an objective lens and used to excite a fluorescent dye. If the fluorescent dye absorbs two infrared photons it will emit a single high-energy photon (the wavelength of which depends on the type of fluorescent dye used), and these photons can be collected by a high sensitivity detector. Because the probability that two photons will be absorbed increases with the laser intensity, a larger fluorescence signal is generated when the laser beam is tightly focused. Using a scanning protocol to move the beam focus through a volume of neuronal tissue can build up a light-intensity image from the individual focal point (pixel) values. Neuronal activity is associated with the influx of calcium ions into the neuron, therefore, if cells in the imaged volume are filled with calcium sensitive fluorescent




Figure 4.4: **Organization for orientation in mouse and cat, measured using two-photon imaging.** (a) Cell-based orientation maps from five different depths in cat visual cortex (210, 230, 250, 270 and 290 μm , as indicated). Selective cells are coloured according to their preferred orientation as shown in the color key. The last panel shows an overlay of images from nine different depths. Reprinted from Ohki et al. (2006). (b) Similar cell-based orientation maps from five different depths in mouse visual cortex (217, 239, 260, 286 and 322 μm , as indicated). Selective cells are coloured according to their preferred orientation as shown in the color key. The last panel shows the overlay of images from all five depths. Scale bars 100 μm . Unpublished results, reproduced with permission from Tom Mrsic-Flogel.

dyes, two-photon imaging can also provide real-time analysis of the activity in individual cells. Individual cell types can also be identified using specific labels that each emit light of a different wavelength.

Figure 4.4 shows orientation preference maps measured at different depths (together with the corresponding combined volume map) in both mouse and cat, measured using two-photon imaging. It is clear from this data that at single-cell resolution there is a vastly different organization, or lack of organization, for orientation in the mouse visual cortex in comparison with cat.

Despite the lack of organization for orientation preference in rodent visual cortex, many studies, using several different optical imaging techniques, have suggested that mouse V1 is well organized with respect to retinotopic position preference (Husson et al., 2007; Kalatsky and Stryker, 2003; Schuett et al., 2002). Retinotopic map mea-




Figure 4.5: Optical imaging of the retinotopic organisation of mouse visual cortex. (a) For retinotopic stimulation, square-shaped grating stimuli are presented at 21 adjacent positions within the visual field. The colour code representing stimulus position is used to generate colour-coded retinotopic maps. (b) Single condition maps from one animal, with each map corresponding to a stimulus in (a), according to the indices. Dark regions in each map were activated by the respective stimulus. Note the gradual shift in the position of the activated patch relative to the red cross, which was placed at identical positions in the maps. The smaller dark patches visible in some maps correspond to activity in higher visual areas. (c) Colour-coded map of the overall retinotopic organisation of the left primary visual cortex, on the basis of averaged single condition maps from six mice. The colour of each pixel corresponds to the stimulus position that elicited the strongest response at this pixel. To mask out regions without cortical response, colour saturation equals the maximum intensity projection of all single condition maps. Abbreviations: a, anterior; p, posterior; m, medial; l, lateral. Scale bars: 2 mm. Reprinted from Hübener (2003).

surement in mouse using optical imaging is illustrated in figure 4.5. However, as discussed above, optical imaging provides resolution that is too low to really test whether the retinotopic map is smooth at a finer spatial scale. The observed disorder in orientation preference at the single cell level indicates that it is important to establish organization at high resolution. In order to investigate the local smoothness of the retinotopic map in mouse, my collaborator Tom Mrsic-Flogel at University College London has performed two-photon calcium imaging. My analysis of this two-photon imaging data (presented in section 4.2) shows that, in fact, at this smaller spatial scale, the retinotopic preference of individual neurons is also highly disordered.

The next section reviews previous models that address the difference in organization for orientation preference between rodents and higher mammals and also outlines my approach to this question.

4.1.2 Modeling species differences in organization

There are, as yet, no mechanistic or developmental models of organization for orientation preference in rodents. However, previous theoretical models have addressed the possible evolutionary optimizations that could lead to a lack of organization in mouse visual cortex. The model of Koulakov and Chklovskii (2001) is based on the suggestion of Hubel et al. (1977) that a limitation on the size of the brain will result in evolutionary pressure to minimize the length of cortical wiring. Intra-cortical wiring length can be optimized by both smoothly varying orientation maps (as observed in higher mammals) and the ‘salt and pepper’ organization observed in rodents.

The critical parameters that underlie the differences between the two organizations in this theoretical model are the connectivity rules between neurons of different orientation preferences. A ‘salt and pepper’ organization is optimal for a uniform connection function, where each neuron connects with equal probability to neurons that prefer any other orientation. In contrast, smoothly varying maps are optimal for a ‘Mexican-hat’ connection function where neurons connect primarily to similar orientations. This predicts that rodents and higher mammals may both be minimizing the wiring length of neurons in V1, but that neurons in rodent visual cortex connect to all orientations equally. It is not clear from this conceptual model what the developmental mechanisms or physical constraints are that could lead to this difference, nor how individual neurons in mouse visual cortex can have similar functional properties as higher mammals without a similar functional organization. The key prediction of this model, i.e., that rodent neurons connect in an isotropic fashion to other neurons with different orientation preferences (unlike in, for example, tree shrew visual cortex, where neurons preferentially connect to similar orientations; see figure 1.11b), has not yet been tested. Note that in testing this prediction care should be taken to ensure that both long and short range connections are taken into account, since considering short range connections alone may give a false confirmation.

Other possible constraints for map development include coverage uniformity, i.e., the necessity for neurons at each retinotopic location in V1 to be selective for a full range

of orientations, spatial frequencies, directions, etc. This idea has received support from a study in which individual feature maps in cat visual cortex were perturbed by various transformations, with the remaining maps remaining optimal for coverage uniformity (Swindale et al., 2000). It is possible that given the smaller size of rodent visual cortex and therefore the small cortical magnification factor (i.e. only a small area of cortex is representative of a large region in visual space), that a ‘salt and pepper’ organization would optimize coverage uniformity in rodents. However, this issue has not yet been addressed experimentally or theoretically.

The lack of orientation maps in rodents may also be due to a lack of evolutionary pressure to improve visual function. However, the function of smooth mapping of feature preferences is still largely unknown, and it may not necessarily be of any evolutionary advantage. This hypothesis is also at odds with the preliminary observation of a lack of orientation maps in squirrel, which is a highly visual rodent (Van Hooser et al., 2005). The constraints that lead to species diversity in organization may become clearer if the mechanisms that lead to the development of both ‘salt and pepper’ and smooth orientation map organizations are known. It may be either the case that both rodent and higher mammals have similar developmental paradigms that are parameterized differently, or that there are inherent differences in the mechanisms of development in these different species.

Although he has not implemented this theory directly, Ringach (2007) suggests that salt and pepper maps may arise if there is more overlap in LGN ON and OFF cells in rodents and if rodent V1 cells sample from only a small number of LGN cells. This speculative model has not previously been implemented computationally, but results to be presented in chapter 5 (section 5.3.3) suggest that sparse sampling of LGN cells may indeed contribute to the disorder observed in mouse.

In any case, there is a vast array of differences between rodent species and higher mammals, many of which will be completely unrelated to the differences in topographic organization in V1. It is impossible to experimentally test every parameter in order to isolate those responsible for this specific difference. However, establishing the relevant parameters *is* possible in a computational model. Simplifying the system and modeling only the features necessary for reproducing the observed behaviour can allow us to convert one species into another merely by adjusting the parameters, and therefore isolate those which are functionally relevant.

My approach to investigating the differences between functional organization in rodents and higher mammals is to parameterize the ALISSOM model presented in chapter 2 using information about the specific receptive field sizes and connection lengths in both mouse and cat. In this chapter I have initialized the model with the amount of retinotopic scatter observed in the adult mouse (as calculated from two-photon calcium imaging data in section 4.2). These simulations show that the disorder in topographic organization in rodents could arise from spatially sparse activation of neurons in V1 during activity driven development. However, ‘building-in’ this scatter does not address how this scatter may arise in earlier stages of development. Several studies of mechanisms that could lead to the development of both retinotopic and orientation organization are therefore explored in chapter 5. These preliminary findings suggest an alternative hypothesis that salt and pepper orientation maps could arise alongside retinotopic scatter, even without activity-driven development, provided that the LGN is sparsely connected to V1. I will also outline some proposed experiments for testing these two hypothesis.

4.2 Analysis of two-photon calcium imaging of retinotopy in mouse visual cortex

The unpublished data on which the following analysis has been performed was provided by Tom Mrsic-Flogel at University College London. Two-photon calcium imaging of mouse visual cortex was carried out using a method similar to that used for optical imaging of the full retinotopic map in V1 (figure 4.5), but the activity was recorded using the equipment illustrated in figure 4.3 rather than by optical imaging. The two-photon imaging protocol is illustrated in figures 4.6 and 4.7.

Mice are anesthetized and restricted in movement, and the exposed skull above the primary visual cortex is positioned below the two photon microscope. For retinotopic stimulation, square-shaped windows of drifting gratings (in the illustration in figure 4.6a these have a side length of 12°) are placed at adjacent positions within the visual field (presented on a video display). Fluorescence signals from an approximately $340 \times 340 \mu\text{m}$ square area of V1 are measured as the position and orientation of the drifting square grating is changed. The time course of the signal (the change in fluorescence, $\Delta F/F$) corresponding to each visual field position for each of the cells within the




Figure 4.6: **Two photon calcium imaging of retinotopy in mouse.** (a) Square drifting sine gratings of side length (in this case) 12° are positioned at different adjacent points in the visual field. (b) Recordings of the change in fluorescence signal over time in each of the imaged cells at each of the visual field positions are made. (c) The position preference of each pixel within the recorded area gives an indication of the average preference and is likely to correspond to the optical imaging data. (d) Individual cells in the imaged area are color coded according to the position that maximally activated the cell. It is clear that the retinotopic map is not smooth at this single-cell resolution. Scale bars $50\mu m$. Reproduced with permission from Tom Mrsic-Flogel.

square area is recorded (figure 4.6a,b).

The individual cells in figure 4.6d are color coded based on the position that elicited the maximum (peak) response for that cell. Although the data used for the analysis is based on the individual cells, the position preference of individual pixel values (figure 4.6c) can also be recorded. It is generally accepted that the values for pixels in the area between cells (neuropil) give a good indication of the mean preference value in that region (due to the overlapping connections of nearby neurons), and thus correspond to the optical imaging data.

It is clear from the data in figure 4.6d that at single-cell resolution, retinotopic preference is not smoothly organized, unlike the retinotopic maps observed in larger scale optical imaging studies. However, in order to parameterize the mouse ALISSOM model,




Figure 4.7: **Stimulus configuration and two-photon imaging results in mouse.** (a) Stimulus configuration used in the data set on which the analysis in this section is performed. Square drifting sine gratings of side length 15° are positioned at different adjacent points in the visual field as a $340 \times 340 \mu\text{m}$ square area of cortex is imaged. (b) Individual cells in the recorded area. (c) The peak retinotopy preference of all pixels in the volume. (c) Individual cells color coded according to their (peak) position preference. Reproduced with permission from Tom Mrsic-Flogel.

I have quantified this organization more precisely. Figure 4.7 shows the protocol used for the dataset on which this analysis was performed. In this case, gratings of side length 15° are used and a total visual field area of 75° azimuth and 60° elevation is activated during the stimulus presentation (figure 4.7a). The individual cells are shown in figure 4.7b and the pixel based and cell based retinotopy preference maps (again based on the peak response) are shown in 4.7c and d.

The data set obtained from Tom Mrsic-Flogel contains the averaged fluorescence time-courses for each cell and the baseline response for each cell in response to gratings placed at each of the stimulus grid positions (as in figure 4.6b). To obtain the averaged time-courses, all pixels within a cell's outline are averaged to give a single time-course and these single time-courses are averaged over several stimulus repetitions. The baseline is obtained from the averaged time-courses when there is no grating stimulus. In order to quantify the scatter more exactly the following protocol was used:

1. All fluorescence time-courses not significantly different from the baseline response at the $p=0.05$ level (ANOVA, $F>4$) were discarded.
2. For each cell, the weighted average of the response to each stimulus location l (using the peak to base value of the fluorescence time-course signal, F^l) was calculated separately for the elevation and azimuth directions. The azimuth (X_{pref}) and elevation (Y_{pref}) preferences for each cell are shown in figures 4.8a and c, respectively. Specifically:

$$X_{pref} = \frac{\sum_l F^l r_x^l}{\sum_l F^l}$$

$$Y_{pref} = \frac{\sum_l F^l r_y^l}{\sum_l F^l}$$

Where (r_x, r_y) for a location l are the retinal coordinates of the midpoint of the square on the stimulus grid and l iterates over all the locations in the stimulus grid.

3. The mean preferred elevation ($\bar{Y}_{pref}$) and azimuth ($\bar{X}_{pref}$) for the imaged area was calculated:

$$\bar{X}_{pref} = \frac{\sum_i X_{pref}^i}{N}$$

$$\bar{Y}_{pref} = \frac{\sum_i Y_{pref}^i}{N}$$

Where i iterates over the N cells in the dataset.

4. A ‘perfect’ retinotopy preference was calculated by assuming that the mean preferred elevation and azimuth preferences are accurate but that the range is determined by the magnification factor recorded using optical imaging (because the actual range observed is dependent on the scatter). Azimuth and elevation preference values for the individual cells obtained using this method are shown in figures 4.8b and d respectively. For each cell, the perfect azimuth ($perf_x$) and perfect elevation ($perf_y$) are calculated using:

$$perf_x = \bar{X}_{pref} + (c_x - 170)/m_x$$

$$perf_y = \bar{Y}_{pref} + (c_y - 170)/m_y$$

Where m_x and m_y are the cortical magnification factor measured using optical imaging for the azimuth and elevation directions respectively and (c_x, c_y) is the cortical location of the cell. The midpoint of the two-photon patch is assumed to

correspond to the mean retinotopy, therefore the value $(c_x - 170)$ is the distance from the midpoint (of the $340\mu\text{m}$ square patch) and when transformed into retinal space using the magnification factor gives the distance from the mean in retinal space.

5. For each individual cell, the difference between the actual preferred retinal position and the perfect position preference is calculated (separately for azimuth, D_x , and elevation, D_y). For each cell:

$$D_x = X_{pref} - perf_x$$

$$D_y = Y_{pref} - perf_y$$

6. These differences in retinal position are binned in a histogram which is then fitted with a Gaussian distribution to obtain a mean and standard deviation for the observed retinotopic scatter in both directions.

The calculated ‘perfect retinotopy’ used in the above process corresponds to the retinotopic preference that would be likely to arise if this same cortical area were imaged using optical imaging. This measure also corresponds well to the retinotopic preferences of the individual pixel values (figure 4.7c), i.e., the underlying neuropil preference. In future, a method could be developed that would separate the neuropil pixel values from the cell pixel values and extrapolate these values to obtain an alternative measure of perfect retinotopy for each cell.

The actual retinotopic position preference and the calculated perfect preference are illustrated again in figure 4.9. The preferred locations of each cell are plotted on the corresponding visual field position on the stimulus grid (black diamonds). The locations of the calculated preferred positions if the retinotopic map were perfect are also plotted (white dots). The distribution of scatter in retinopic position is calculated using the distances between each black diamond and white dot for each cell.

The plot in figure 4.9 also highlights a small potential error in the retinotopy estimation, introduced by the biased coverage of the visual field by the stimulus grid. Neurons whose preferences are in the top left portion of the grid tend to lie along the central points in the square positions. This is due to an under representation of the upper left portion of the visual field by the visual stimulus. These neurons are only significantly activated when the gratings are placed in one grid position, leading to the observed pattern. Expanding the stimulus grid in the upper left direction would likely lead to a




Figure 4.8: **Illustration of retinotopic scatter in azimuth and elevation in mouse.** (a) Individual cells in the cortical position (x,y) are color coded according to their azimuth preferences. (b) Azimuth preferences of each cell if the retinotopic map was perfect, generated using the average azimuth preference value and the magnification factor from optical imaging data. (c),(d) As for (a),(b) for elevation preference. The color key in the top right corner of each plot indicates the retinal position with respect to the full stimulus grid.

greater number of contributions to the weighted average and a more accurate measurement of receptive field position. However, this error is a conservative one, as it leads to an underestimation of the scatter. In future two-photon imaging experiments, it will be necessary to ensure that the visual field representation is unbiased, by measuring the entire area of the visual field to which the imaged neurons may respond.

Histograms of the differences between the preferred elevation and azimuth and the calculated perfect preferences for each cell are shown in figure 4.10. Each histogram is well fitted with a Gaussian distribution, the mean and standard deviation of which




Figure 4.9: **Scatter in receptive field position.** The preferred location of each cell is plotted on the corresponding visual field position (see figure 4.7) on the stimulus grid as a black diamond. The preferred position in a 'perfect' retinotopic map (calculated using the mean position preference and the magnification factor from optical imaging data) is plotted as a white dot. The difference between the real and 'perfect' position for each cell is used to calculate the retinotopic scatter.




Figure 4.10: **Analysis of retinotopic scatter in mouse.** The differences between the azimuth (a) and elevation (b) preference and the calculated perfect preference for each cell are binned in a histogram of bin width 5° . These histograms are fit with Gaussian distributions with mean μ and standard deviation σ as shown. The distribution of scatter appears to be similar in the two directions.

appear to be similar for scatter in both the azimuth and the elevation directions. The average of these mean and standard deviation values are used in a parameterized ALISSOM model of mouse visual cortex that is described in the next section.

4.3 Methods

The models presented in this chapter are based on the ALISSOM model, described in chapter 2. The model architectures, specific parameterization and analysis methods, where they differ from those in the previous ALISSOM model, are described here. All other methods are as described in section 2.2.

4.3.1 Mouse model architecture

A high-level diagram of the early visual system in the mouse is shown in figure 4.11. The mouse system has the same basic features as higher mammals (as shown in figure 1.1), i.e., retinal ganglion cells contact LGN cells, which then make connections to primary visual cortex. The mouse LGN is less laminated than higher mammals, but information from the two eyes does remain separate. Mouse V1 has a small binocular region of cortex that receives input from both eyes and, as discussed above, the retinotopic relationship between connections is largely preserved, although it appears to be scattered at the level of individual cells. Mouse receptive field sizes and the cortical magnification factor (CMF) are relatively consistent across the visual field (CMF varies by a factor of ≈ 2), unlike higher mammals where there are more significant changes in RGC morphology and RGC density between the central and peripheral regions of the retina (Chalupa and Williams, 2008).

Parameters from the experimental literature are used to constrain the size of connection radii in the pathway from the retinal photoreceptors to V1, making the ALISSOM model specific to the mouse. The simulated area of V1 in both cases represents a number of two-photon imaging ‘patches’ (either 3×3 or 4×4), each representing $340 \times 340 \mu\text{m}$ of cortex. More than one patch is used in order to simulate more accurately the horizontal connections in V1, that in reality will extend well beyond the imaged area.

The V1 sheet dimensions are chosen such that the edge of one two-photon patch cor-




Figure 4.11: **Mouse early visual system.** The mouse visual system has the same basic features as higher mammals. Signals from the retinal ganglion cells project to the LGN that, although less laminated than in higher mammals, has separate layers for input arising from each eye. Axons projecting from the LGN into V1 have arbors that terminate mainly in layer 4. The mouse has only a small binocular region receiving input from both eyes, due to the lateral placement of the eyes. Reproduced with permission from Tom Mrsic-Flogel.

responds to $340\mu\text{m}$ in cortical distance and 1.0 in sheet coordinates. The density of V1 units per 1.0×1.0 area is chosen such that the number of V1 neuron units in one patch (100 neurons on a 10×10 grid, i.e., density 10) is approximately the same number of neurons in an imaged area at one particular depth. Thus the model V1 sheet corresponds to one depth of the two-photon imaging volume (see figure 4.3c). In order to accurately represent the type of data observed in two-photon imaging experiments, the exact cortical location of each neuron unit within the V1 grid is randomly scattered by a small amount.

The parameters taken from the literature and used for the subcortical areas in the mouse model are shown in table 4.1. In some cases these are unknown or there are discrepancies in the literature; the values chosen for this model are highlighted in bold. The basic mouse model architecture (to scale) is shown in figure 4.12. Where parameters are




Figure 4.12: **Mouse model architecture.** The model of the mouse visual pathway is based on the ALISSOM architecture, described in figure 2.3. Connections of one unit in each sheet are shown. The array of V1 units shown corresponds to 3×3 two-photon imaging patches, each $340\mu\text{m} \times 340\mu\text{m}$. V1 units receive lateral excitatory (small red circle) and lateral inhibitory (large blue circle) connections from nearby V1 units. V1 units receive afferent input from two sheets representing the ON and OFF channels in the retinal ganglion cells (RGC) and lateral geniculate nucleus (LGN) that in turn receive input from the retinal photoreceptors (only one sheet is shown here). The smaller square in the Retina and RGC/LGN sheets corresponds to the area to which V1 responds directly. In the initial models presented in this chapter, retinotopic jitter is included in the model by randomly jittering the connections from RGC/LGN to V1 by an amount corresponding on average to the experimentally measured scatter (as indicated by the dashed arrow and circle).

changed in subsequent simulations, the changes are highlighted in the relevant section. The size of the Retina and RGC/LGN sheets is calculated using the mouse cortical magnification factor ($15\mu\text{m}/^\circ$) i.e. $1020\mu\text{m}$ of V1 corresponds to $1020/15 = 68^\circ$ in the Retina and LGN. As in section 2, the overall size of the retinal photoreceptor and LGN/RGC sheets is then increased in order to minimize edge effects in V1.

Retinotopic scatter can be ‘built-in’ to this model by randomly adjusting (‘jittering’) the central point of the LGN to V1 afferent connection fields (referred to as *inconnections* from here onwards), as illustrated in figure 4.12. A similar experimental protocol and analysis as used in section 4.2 is used in the model to ensure that similar mean and standard deviation values (shown in table 4.2) are obtained for the retinotopic scatter. This scatter in the model is illustrated in figure 4.13.

In each simulation, retinotopy is measured in two ways (see section 4.3.5). The first takes only the afferent input to V1 into account, allowing the measurement of retinotopic scatter without making assumptions about the form of the lateral connectivity. The second takes both afferent and lateral connections into account, which may be

Parameter	Values from literature and/or chosen value
Size of one two-photon patch	$340\mu m \times 340\mu m$ ¹
Cortical magnification factor	• Vertical: $24.8 \pm 5.9\mu m/^\circ$² • Horizontal: $14.4 \pm 5.4\mu m/^\circ$² • Value used for simplicity: $15\mu m/^\circ$
Average LGN Receptive field center size (diameter)	6° ³
Average LGN Receptive field surround size (diameter)	17° ³
Average V1 Receptive field size (diameter)	• 6°⁴ • 14°⁵ • Initial value unknown; value chosen to allow receptive fields in the above range to develop: 28°
Extent of long range horizontal connections (diameter)	• Up to $1800\mu m$⁶ • Up to $2000\mu m$ but typically within $800\mu m$¹ • Arbitrary value $1020\mu m$ (chosen to cover 3 two-photon patches)
Extent of local horizontal connections (diameter)	Arbitrary value $200\mu m$ (see chapter 3)
Mean of Gaussian retinotopic scatter	-0.4° ⁷
Standard deviation of Gaussian retinotopic scatter	8.5° ⁷

Table 4.1: **Parameters for mouse ALISSOM model.** 1. Tom Mrsic-Flogel, unpublished observations. 2. Schuett et al. (2002), from optical imaging data. 3. Grubb et al. (2003). 4. Niell and Stryker (2008), measured with small sweeping bars, half width, half maximum of Gaussian fit. 5. Métin et al. (1988) layer 2/3, hand mapped, average of large and small diameter. 6. Burkhalter (1989). 7. Average of azimuth and elevation scatter from section 4.2.




Figure 4.13: **Analysis of retinotopic scatter in the mouse model.** Retinotopic scatter is included in the model by ‘jittering’ the center of LGN afferent connection fields (similarly for both ON and OFF) with mean $\mu = -0.4^\circ$ and standard deviation $\sigma = 8.5^\circ$. Retinotopy is measured using two different protocols, before settling (afferent only) and after settling (afferent and lateral connections). The inset in each plot shows the color key assigned to the full area of the Retina to which stimuli were presented. (a) Results from the before settling measure of retinotopy in the central 10×10 units of the V1 sheet. The retinotopy in the azimuth (or x) direction is scattered to a similar degree before (iteration 0), and after development (iteration 20000). (b) As for (a) for the elevation (or y) direction. (c) Histogram of scatter for azimuth (red) and elevation (blue) preference before (solid line) and after (dashed line) development. (d,e,f) as for (a,b,c) using the after settling protocol. Including the lateral connections in retinotopy measurement decreases the spread in scatter in comparison with the afferent only measure before development. The spread in scatter also decreases during development when using the after settling protocol.

	Azimuth (x)		Elevation (y)	
Parameters of Gaussian distribution fit	μ	σ	μ	σ
Afferent only				
Iteration 0	-3.4	9.1	1.5	8.2
Iteration 20000	-3.1	8.4	1.6	7.2
Afferent and lateral				
Iteration 0	-2.1	6.2	2.5	5.4
Iteration 20000	-2.6	-3.8	1.1	3.4

Table 4.2: **Scatter properties for mouse model with ‘built-in’ scatter.** Table indicates the mean μ and standard deviation σ values for all Gaussian fits for the histograms shown in figure 4.13c and f. The parameters obtained using the afferent only measure are a good match to the experimentally measured values ($\mu = -0.4^\circ$, $\sigma = 8.5^\circ$). Measurement of retinotopy also depends on the lateral interactions in the model, which in this case reduce the spread of the scatter.

particularly important in simulations where the lateral connectivity is specifically being changed to alter the course of development. When the measured retinotopic scatter is added to the model, including lateral connections results in a smaller spread of retinotopy values (figure 4.13; table 4.2) that decreases during development. Lateral connections in this case are driving the network to make the most active units more active, effectively sharpening the receptive field of each unit. Refinement of lateral connections during development will result in increased sharpening of retinotopy preference during development. By measuring retinotopy preference in these two different regimes, it is possible to separate the relative contributions of afferent and lateral connections.

Unfortunately, there is very little information available on the extent of connections in young animals. The initial extents of lateral connections are chosen arbitrarily in the model, since it is the balance of excitatory and inhibitory interaction strengths that is the crucial parameter (as discussed in chapter 2). The size of afferent connection fields from the RGC/LGN sheets is chosen based on values of the average receptive field size in degrees (although the initial value is chosen to be larger than the values obtained in the literature for adult mice to allow for refinement during development). The size of difference-of-Gaussian afferent connections from the photoreceptors to the RGC/LGN units are chosen based on LGN receptive field center and surround values in the literature. Connection field weights are fixed, isotropic 2D Gaussians for the lateral

excitatory projection, and initially random within a Gaussian envelope for afferent and lateral inhibitory projections.

For comparison with the mouse model, I have also parameterized an ALISSOM model using values obtained from the literature on the cat visual pathway. This model is described in more detail in the next section.

4.3.2 Cat model architecture

The early visual system in the cat is more complex than the mouse, but with the same basic features. The main issue with parameterizing a cat ALISSOM model is that receptive field properties are much more variable across the visual field (unlike the mouse). Since this model is primarily for qualitative comparison with the mouse model I have chosen parameter values that I feel reasonably approximate the average and are within the reasonable bounds of experimental values; there is unfortunately no complete, consistent set of values from a single eccentricity for the cat.

The experimental values used in parameterizing the cat model are shown in table 4.3; again the chosen parameters are highlighted in bold. Rather than choosing values which reflect only either the area centralis (where orientation maps are not usually measured), the cortical magnification factor chosen from the range of experimental values is an average over the central 10° of the visual field. Similarly, other values which best approximate the average properties in this region are chosen from the range of possible values. The LGN receptive field center size is chosen from a study which includes only X-cells (thought to be the homologue of macaque P cells, see section 1.2), since it is these cells which have the most sustained response to a stimulus and therefore likely to be those contributing most to map formation. The LGN receptive field surround size is chosen based on the chosen center size (rather than on the single value available in the literature), so that the relative values combine to give a receptive field with reasonable edge detection properties.

Choosing parameters for the model also suffers from a lack of available data for immature animals, therefore the chosen initial value for the V1 receptive field size is outside of the range of adult values. This is to allow receptive fields within the adult range to develop and is based on the assumption that larger initial receptive field refine over the course of development. The chosen extent of lateral connections is also outwith




Figure 4.14: **Cat model architecture.** The model of the cat visual pathway is based on the ALISSOM architecture, described in figure 2.3. Connections of one unit in each sheet are shown. The array of V1 units shown corresponds to 3×3 two-photon imaging patches, each $340\mu\text{m} \times 340\mu\text{m}$. V1 units receive lateral excitatory (small red circle) and lateral inhibitory (large blue circle) connections from nearby V1 units. V1 units receive afferent input from two sheets representing the ON and OFF channels in the retinal ganglion cells (RGC) and lateral geniculate nucleus (LGN) that in turn receive input from the retinal photoreceptors (only one sheet is shown here).

the experimentally observed range. This is to reduce the computational resources required, and does not significantly influence the results, since it is the relative strengths of interactions rather than the extent of connections which is the key parameter in the model.

The basic cat model architecture (to scale) is shown in figure 4.14. The size of the Retina and RGC/LGN sheets is calculated using the cat cortical magnification factor ($667\mu\text{m}/^\circ$) i.e. $1020\mu\text{m}$ of V1 corresponds to $1020/667 = 1.5^\circ$ in the Retina and LGN. The overall size of the retinal photoreceptor and LGN/RGC sheets is then increased in order to minimize edge effects in V1. The initial weights and sizes of connection fields are as discussed in section 4.3.1 for the mouse.

Comparison of the mouse architecture (figure 4.12) and cat architecture (figure 4.14) highlights the differences between these two species. The cortical area is the same in both models, but in cat this cortical area corresponds to a much smaller visual field area (because of the large cortical magnification factor in cat). The relative afferent connection field sizes also indicate that in the cat visual pathway, receptive fields of V1 neurons overlap more extensively.

Parameter	Values from literature and/or chosen value
Size of one two-photon patch	$340\mu\text{m} \times 340\mu\text{m}$ ¹
Cortical magnification factor	• $2300\mu\text{m}/^\circ$ centrally, decreasing paracentrally to $600\mu\text{m}/^\circ$² • $1870\mu\text{m}/^\circ$³ at area centralis • $667\mu\text{m}/^\circ$⁴ average over central 10°
Average LGN Receptive field center size (diameter)	• $1.22$⁵ • $0.515 \pm 0.181$⁶ • Value used for simplicity 0.5°
Average LGN Receptive field surround size (diameter)	• $2.95$⁵ • Value used for simplicity: 2°
Average V1 Receptive field size (diameter)	• $1.6 \pm 0.44$⁷ • 0.7 increasing to 2.6° at 10° eccentricity⁸ • Initial value unknown; value chosen to allow receptive fields in the above range to develop: 3°
Extent of long range horizontal connections (diameter)	• $2500\mu\text{m}$⁹ • Up to $4000\mu\text{m}$¹⁰ • Arbitrary value $1020\mu\text{m}$ (chosen to cover 3 two-photon patches)
Extent of local horizontal connections (diameter)	Arbitrary value $200\mu\text{m}$ (see chapter 3)
Mean of Gaussian retinotopic scatter	0.008 ¹¹
Standard deviation of Gaussian retinotopic scatter	0.19 ¹¹

Table 4.3: **Parameters for cat ALISSOM model.** 1. Tom Mrsic-Flogel. 2. Albus (1975). 3. Tusa et al. (1978). 4. Van Hooser (2007). 5. Cai et al. (1997), within 20° of area centralis, includes all cell types. 6. Bullier and Norton (1979), X-cells, within 20° of area centralis. 7. Warren et al. (2001). 8. Buzás et al. (2003), Average width assuming square shape. 9. Kisvarday et al. (1997). 10. Gilbert and Wiesel (1983). 11. Mouse scatter values converted into cat values as described in the text.

Single electrode electrophysiological recordings suggest that receptive field positions in the cat are also scattered (Albus, 1975). The exact scale of this scatter is difficult to conclude using single electrode penetrations, since these are likely to suffer from difficulty in quantifying electrode position, eye-drift, and other methodological problems. Unfortunately, two-photon data measuring the retinotopic scatter in cat primary visual cortex is not yet available. However, multi-electrode arrays have recently been used to quantify scatter in cat V1 receptive fields (Warren et al., 2001) and have also found that there is a local scatter in receptive field positions.

In multi-electrode recordings, a similar method of analysis as that used above for the mouse two-photon data can be applied to the data to quantify the scatter. In Warren et al. (2001), the recorded receptive field positions for each isolated electrode were compared with a linear map of the electrode array onto its visual space representation (by a least-mean-squared-error fit of the difference in receptive field positions). It is difficult to compare two-photon and multi-electrode analyses directly since, by isolating single units, the multi-electrode data does not result in a complete sampling of the cells in the recorded area. However, a rough comparison of the standard deviation in scatter for cat calculated from the multi-electrode data (0.59° and 0.45° for the horizontal and vertical axis) in comparison with calculated values for the mouse (9.9° and 7.1°) suggests that there is a much smaller overall amount of scatter in receptive field positions in cat than in mouse.

In order to calculate a value for the scatter in cat to include in the model, I have transformed the scatter in degrees in mouse (σ_{mouse}) into cortical distance -using the mouse magnification factor (m_{mouse}) and calculated a similar cortical distance scatter in the cat model (σ_{cat}) using the cat magnification factor m_{cat} :

$$\sigma_{cat} = \frac{\sigma_{mouse} m_{mouse}}{m_{cat}} \quad (4.1)$$

This results in a smaller overall scatter in visual field distance (mean $\mu = 0.008^\circ$ and standard deviation $\sigma = 0.19^\circ$), that is an approximate order of magnitude match to the value of retinotopic scatter in cat measured using multi-electrode recordings. The fact that these values are an order of magnitude match suggests that similar physical mechanisms that give rise to the scatter observed in mouse are also present to a similar degree in cat.

4.3.3 Input patterns and model parameters

In order to simplify the models as much as possible for most of the results in this chapter, simple input patterns are used for each simulation. In a final section (section 4.5) natural images are to drive development and the implications and future work to improve these simulations is discussed. Again, natural images are retina-sized patches from images of natural objects and landscapes from a dataset by Shouval et al. (1996) as described in section 2.2.2.

In order to make the mechanisms clear in the initial simulations, at each iteration two elongated Gaussian patterns are presented to the model retinal photoreceptors. The center coordinates and orientation of each Gaussian pattern are chosen from a uniform random distribution and cover the full extent of the retina. The size of the Gaussian patterns are based on the center size of RGC/LGN units for each species, chosen to maximally activate the RGC/LGN in each case.

Because of the different extents of connections in each species, values for the homeostatic parameters differ from the generic ALISSOM model presented in chapter 2. The initial model parameters (as defined in section 2.2) for each species model are shown in table 4.4. Those parameters not listed in table 4.4 are identical to those in the original ALISSOM model (table 2.1). Where subsequent simulations use parameters different from these values this is highlighted in the text. The equations which describe the model are identical to those used in the original ALISSOM model.

4.3.4 Retinotopic scatter

The ‘built-in’ scatter described in section 4.3.1 is a way of including retinotopic scatter without requiring that it arise during development in the model. The models presented in this chapter are used to illustrate how retinotopic disorder may affect the development of organization for orientation without considering how retinotopic disorder arises. However, the mechanisms underlying retinotopic scatter development are explored in detail in chapter 5. In this chapter, scatter is implemented using a Gaussian random number generator (with specified mean and standard deviation) for the x and y coordinates of afferent *in-connection* field centers (F_{jA}).

Parameter	Mouse value	Cat value	Description
N_d	20	20	Cortical density, i.e., width and height of a unit area of cortex, section 2.2.1 and equation 2.14
L_d	10	10	LGN density, i.e., width and height of a unit area of LGN, section 2.2.1
R_d	10	10	Retina density, i.e., width and height of a unit area of Retina, section 2.2.1
γ_L	6.0	4.0	Strength scaling factor for the LGN's afferent weights ¹ , equation 2.1
γ_A	1.0	1.0	Strength scaling factor for the V1 afferent weights ² , equation 3.4
γ_E	1.0	1.0	Strength scaling factor for the lateral excitatory weights ³ , equation 3.4
γ_I	-1.0	1.0	Strength scaling factor for the lateral inhibitory weights ³ , equation 3.4
κ_{lr}	0.017	0.06	Target afferent activity for use in learning rate scaling ⁴ , equation 2.13
v_{2Gauss}	2.0	2.0	Expected frequency of V1 responses for 2 Gaussians ² , section 2.2.8
v_{nat}	5.0	5.0	Expected frequency of V1 responses for natural images ² , section 2.2.8
μ	0.05v	0.05v	Target average V1 activity (mean of target exponential distribution) ⁴ , section 2.2.8
κ	1.7 μ	6 μ	Target average contribution from the afferent projections to the V1 activity ⁴ , equation 2.3
a_{init}	12.0	12.0	Initial slope parameter in sigmoid transfer function ⁵ , equation 2.9
b_{init}	-4.0	-4.0	Initial threshold parameter in sigmoid transfer function ⁵ , equation 2.10

Table 4.4: **ALISSOM model parameters for mouse and cat models.** 1. Increased in comparison with generic ALISSOM and different for each species due to differences in size and activation of LGN difference-of-Gaussian receptive fields. 2. Identical to generic ALISSOM values. 3. Identical to generic LISSOM values except where highlighted in the text. 4. Different from generic ALISSOM and also different for each species because of the differences in size and therefore activation of afferent receptive fields. 5. Unlike the generic ALISSOM, these initial values are fixed and identical for each unit in order to make the differences in initial activation patterns due to retinotopic scatter only.

4.3.5 Analysis

Orientation preference and selectivity maps are measured as described in section 2.2. Selectivity values (as indicated by the saturation in each plot) are usually normalized across simulations in order to allow for better comparison. Where sets of plots have been normalized, this is indicated in the text. The ‘zoomed in’ two-photon orientation plots are constructed by plotting a circle representing each unit with a position determined by the individual coordinate values of the central 10×10 units in V1 (i.e. one two-photon patch). Each unit is then colored according to its preferred orientation.

Receptive fields are calculated using reverse correlation as described in section 2.2. In most cases, the stimuli used for reverse correlation are circular 2D Gaussian patterns with $\sigma_x = \sigma_y$ equal to the model RGC/LGN center size. The effect of using different sized stimuli on the receptive field measurement is discussed in more detail in section 5.3.3. Receptive field plots are displayed as a 3×3 patch of V1 showing adjacent units that have been jointly normalized across the full V1 sheet. Differences in the background color indicate the halfway point between the maximum and minimum value across the whole sheet.

Retinotopy preference is measured in the mouse model in a similar way as in the experimental protocol (figure 4.6, Schuett et al., 2002). The area on the retina to which V1 responds directly (see figure 4.12) is divided into squares of approximately 18° for a 4×4 stimulus grid (as used in sections 4.4.1 and 5.3.1) or 17° for a 3×3 stimulus grid (as used in sections 5.3.2 and 5.3.3). These squares are then activated by presenting sine gratings of various phases, frequencies and orientations within each square. For each unit in V1, a weighted average of its responses to stimuli at each x and y position (using the center point of the grid) at the preferred orientation, frequency and phase is used to calculate a preferred retinotopic position (as described in section 4.2).

In retinotopic position preference plots the retina is color coded along the x and y directions and each V1 unit is assigned a color according to its preference. As discussed above, retinotopy is usually measured in the model using the afferent input to V1 only (i.e. before settling) and taking into account both afferent and lateral interactions (after settling). This retinotopic analysis is not applied to the cat model because the small size of the retina in degrees makes this analysis impractical. However, it is not necessary to know the exact scatter in cat, since at present we are comparing the effects of similar amounts of scatter in the mouse and cat models.

Retinotopic scatter in both the model and the experimental data is measured by defining a perfect retinotopic location for each individual cell/unit and calculating the difference in degrees of visual angle between the actual preferred position and the perfect position (as described in section 4.2). The differences are then binned in a histogram of bin width 5° and the Matlab statistics toolbox (*dftool*) used to fit a Gaussian distribution to the data. As described above, for the experimental data a ‘perfect’ retinotopy preference is defined. Calculation of this ‘perfect’ preference uses the mean preferred retinal position of all units in the specified area (one two-photon patch in both the experimental and model data). The central cortical coordinate of the two-photon imaging patch is then assumed to prefer this mean retinal position. The cortical positions of each cell/unit are then mapped to preferred retinal positions using their cortical distance from the central point and the cortical magnification factor. For the experimental data the cortical magnification factor measured from optical imaging data is used (Schuett et al., 2002), and for the model data the simplified value (15°) as used in the model is used (see table 4.1).

4.4 Results

The first part of this results section investigates how the retinotopic scatter measured above can contribute to a ‘salt and pepper’ mapping of orientation preference in the mouse. The second section shows how a similar amount of scatter in the cat model does not necessarily result in similar disorder in organization for orientation preference. These results show that the underlying differences in parameterization between species, rather than any fundamental difference in developmental algorithm, can result in the observed differences in topographic organization.

The final results section investigates how using natural images rather than simple Gaussian patterns affects these results.

4.4.1 Development of ‘salt and pepper’ maps in mouse driven by sparse activation of V1

The first set of results (figure 4.15) are from the parameterized mouse model without retinotopic scatter in the V1 afferent *in-connections*. In this simulation, presentation

of an oriented Gaussian to the model retinal photoreceptor sheet will result in a ‘blob’ of activation, as in previous models (figure 4.15a), both at iteration 0 and after map development (iteration 20000).

This simulation, which does not include any retinotopic scatter, will develop an orientation map (figure 4.15b), with orientation selective receptive fields that are clustered in their orientation preference (figure 4.15c, d). The scatter histograms illustrate that an initial perfect retinotopy in this simulation is also maintained as the orientation map forms during development (figure 4.15e,f). Measuring the retinotopic preference after settling (taking into account lateral interactions) increases the spread of the retinotopic scatter in comparison with the before settling measure (afferent input only). This increase in scatter could be due to under-sampling of the retina, i.e, placing the stimulus grid in non-overlapping positions will not necessarily activate each V1 unit during the measurement protocol if V1 activity (as in this case) is spatially restricted ‘blobs’.

Figure 4.16 illustrates a simulation in which *in-connection* jitter is added to the mouse model. Activity patterns in V1 (in response to oriented Gaussians) are more scattered in comparison to the patterns observed without retinotopic jitter (figure 4.16a). This scattered activity, over the course of development, leads to a more scattered map of orientation preference (figure 4.16b,c). This is unsurprising, since orientation maps arise because lateral connections result in ‘blobs’ of activity, ensuring that nearby neurons learn to prefer similar orientations. Replacing ‘blobs’ of activity with more sparse activation will cause scattered units in V1 to be correlated and therefore learn similar orientation preferences, resulting in a more scattered orientation map.

However, including the measured amount of retinotopic scatter by itself does not appear to result in a fully ‘salt and pepper’ map since there are clusters of similar receptive field preferences remaining (figure 4.16b,c). Scatter in retinotopy measured before settling does not change remarkably during development, with the final scatter (at iteration 20000) comparable to the scatter observed in the two-photon imaging experiments (figure 4.16e). Measurement of scatter after settling does reduce the spread of scatter both before and during development (as discussed in section 4.3.1). Although the form of lateral interactions and their subsequent contribution to retinotopic scatter in the mouse V1 is not yet known, this may indicate that in order to reproduce the measured amount of scatter in the adult mouse, a larger initial *in-connection* scatter should be added to the model. This result also suggests that lateral connections play an important role in determining the amount of retinotopic scatter.




Figure 4.15: **Mouse model without retinotopic scatter.** (a) The response to an elongated and oriented Gaussian pattern (as shown in inset) at iteration 0 and at iteration 20000 is a ‘blob’ of activity similar to those observed in previous implementations of the ALISSOM model. Activity plots (in this, and in subsequent figures in this chapter) are normalized, and the initial activation function of model units is the same for all units in the sheet to allow for clear comparison between the activity patterns in different simulations. (b) Orientation preference and selectivity maps show that there is clearly development of an organized orientation map structure in this simulation at iteration 20000. Selectivity is normalized across all plots in section 4.4.1 for ease of comparison. (c) Two-photon like orientation preference plots at iteration 0 and 20000. (d) Receptive fields for a 3×3 patch of V1 model units at iteration 0 and 20000. (e) Initial (solid line) and final (dashed line) retinotopic scatter histograms for the azimuth (red) and elevation (blue) directions respectively using the before settling protocol (using afferent activity only). (f) as for (e) using after settling protocol (takes into account lateral interactions).




Figure 4.16: **Mouse model with 'built-in' retinotopic scatter.** As presented in figure 4.15 except this simulation includes a scatter in V1 afferent *in-connections* equivalent to the amount of scatter measured from the experimental data. (a) The response to an elongated and oriented Gaussian pattern (as shown in inset) at iteration 0 is slightly more spread out or scattered 'blob' of activity than in the simulation without any scatter. At iteration 20000, once receptive fields have become more selective, this blob of activity has become more sparse (as only the units preferring similar orientations will respond) (b) Orientation preference and selectivity maps show that there is development of some map structure in this simulation at iteration 20000 but that orientation preference is more scattered. (c) Two-photon like orientation preference plots at iteration 0 and 20000 show that there is less clustering of orientation preference in this simulation than in the simulation presented in figure 4.15. This result is more similar to the organization of orientation measured experimentally. (d) Receptive fields for a 3×3 patch of V1 model units at iteration 0 and 20000 show that receptive fields are slightly less selective but also more random and less clustered according to orientation. (e) Initial (solid line) and final (dashed line) retinotopic scatter histograms for the azimuth (red) and elevation (blue) directions respectively for the before settling protocol using afferent activity only. (f) as for (e) for the after settling protocol that also takes into account lateral interactions.




Figure 4.17: **Mouse model with larger spread in 'built-in' retinotopic scatter.** As presented in figure 4.16 except this simulation includes a greater amount of scatter in V1 afferent *in-connections*. (a) The response to an elongated and oriented Gaussian pattern (as shown in inset) at iteration 0 and at iteration 20000 is a more spread out or scattered 'blob' of activity than in the simulation without any scatter. (b) Orientation preference and selectivity maps show that there is less map structure in this simulation at iteration 20000 than in the simulation presented in figure 4.15. (c) Two-photon like orientation preference plots at iteration 0 and 20000 show that there is less clustering of orientation preference in this simulation than in the simulation presented in figure 4.15. This result is more similar to the organization of orientation measured experimentally. (d) Receptive fields for a 3×3 patch of V1 model units at iteration 0 and 20000 show that receptive fields are slightly less selective but also more random and less clustered according to orientation. (e) Initial (solid line) and final (dashed line) retinotopic scatter histograms for the azimuth (red) and elevation (blue) directions respectively using the before settling protocol. (f) as for (e) using the after settling protocol. Taking into account lateral connections in this case more closely resembles scatter at the beginning of development but is still less than the measured scatter at the end of development.

There are therefore a number of possible changes to this initial model that may increase the disorder in the orientation preference map. As described in section 4.2, the experimental measure of retinotopic scatter is likely to be biased by undersampling of the visual field. The above results also suggest that lateral connections may sharpen the retinotopy preference and decrease the amount of measured scatter during development. It is therefore possible that including a greater initial spread in retinotopic scatter in the model may result in a more disordered map. Including a greater spread in the *in-connection* jitter (standard deviation $\sigma = 11^\circ$) is illustrated in figure 4.17. This greater amount of jitter does indeed increase the disorder in the orientation map and reduce the amount of clustering.

Lateral interactions in the model may also play a part in increasing the disorder in the map. The previous simulations suggest that the level of clustering in the orientation map is determined by the sparsity of activity patterns in V1. Increasing the strength of lateral inhibitory interactions ($\gamma_I = -3.0$) and decreasing the radius of lateral excitatory interactions ($r_E = 0.1$) alone, without including any jitter, can result in smaller activation patterns in V1 in response to a stimulus (figure 4.18a). This can also increase the disorder of the map by decreasing the extent of correlations between the activation of neighboring units. However, these changes in lateral connections do not result in any changes in the measured retinotopy preference in comparison with the simulation in which there is no jitter (see figures 4.18e,f and 4.15e,f). These lateral interaction changes alone also do not completely eliminate clustering of orientation preference (figure 4.18b,c,d).

Results from a simulation in which changes in lateral interactions are combined with the experimentally measured retinotopic scatter are shown in figure 4.19. This simulation reproduces both the experimentally observed scatter in retinotopy preference (figure 4.19e,f) and a much more ‘salt and pepper’ like organization for orientation preference (figure 4.19b,c). The activity plots in figure 4.19a highlight that in this case, the response to an oriented stimulus on the retina is a set of sparse activations of nearby individual model units. The scatter in retinotopy ensures that the response area is spread out over the cortical surface, and the changes in lateral interactions ensure that activation is sparse, i.e., neighboring neurons do not necessarily respond.

A more detailed investigation of the receptive fields in this ‘salt and pepper’ model shows that there are examples of cells with two (figure 4.20a,b), three (figure 4.20c,d) or one subfield (figure 4.20e,f), with various orientations, ON/OFF centers and spac-




Figure 4.18: Mouse model with increased inhibition and decreased excitatory radius. As presented in figure 4.15 except this simulation has $\gamma_I = -3.0$ and $r_E = 0.1$. (a) The response to an elongated and oriented Gaussian pattern (as shown in inset) at iteration 0 and at iteration 20000 is a much smaller ‘blob’ of activity than in the simulation without any scatter. (b) Orientation preference and selectivity maps show that there is less map structure in this simulation at iteration 20000 than in the simulation presented in figure 4.15. (c) Two-photon like orientation preference plots at iteration 0 and 20000 show that there is less than in the simulation presented in figure 4.15, but still some clustering of orientation preference in this simulation after development. (d) Receptive fields for a 3×3 patch of V1 model units at iteration 0 and 20000 show that receptive fields are more random with smaller clusters in orientation preference. (e) Initial (solid line) and final (dashed line) retinotopic scatter histograms for the azimuth (red) and elevation (blue) directions respectively using the before settling protocol. (f) as for (e) using the after settling protocol. Comparison with figure 4.15 shows that there is very little difference between this simulation and a simulation without any *in connection* scatter, suggesting that the lateral connections have maintained the perfect retinotopy in this case.




Figure 4.19: **'Salt and pepper' orientation maps with sparse activation.** As presented in figure 4.16 except this simulation has $\gamma_I = -3.0$ and $r_E = 0.1$. (a) The response to an elongated and oriented Gaussian pattern (as shown in inset) at iteration 0 and at iteration 20000 is a cloud of sparse activation of individual model V1 units. (b) Orientation preference and selectivity maps show that although the selectivity of individual units has increased during development there is very little structure in the orientation map. (c) Two-photon like orientation preference plots at iteration 0 and 20000 show that it is very difficult to determine any clustering of orientation preference in this simulation after development. It is difficult to distinguish any differences in organization between the plots at 0 and 20000 iterations. (d) Receptive fields for a 3×3 patch of V1 model units at iteration 0 and 20000 show that receptive fields are randomly oriented and not clustered according to orientation. (e) Initial (solid line) and final (dashed line) retinotopic scatter histograms for the azimuth (red) and elevation (blue) directions respectively using the before settling protocol. (f) as for (e) using the after settling protocol. In this case there is very little difference between the two different protocols and the amount of scatter in the model is similar to the experimentally measured scatter. The amount of scatter also remains fairly constant throughout development.




Figure 4.20: **Receptive fields in ‘salt and pepper’ mouse model.** Six example model receptive fields measured using reverse correlation, chosen to illustrate the different types present in the model. Example cells with two (a,b), three (c,d) or one subfield (e,f), with various orientations, ON/OFF centers and spacing of subfields are measured in the model as in the experimental data (figure 4.1).

ing of subfields, just as in the experimental data (figure 4.1). There do seem to be a large proportion of receptive fields in the model that have only one subfield (although the exact proportions in the experimental data have not been reported). As with all of the simulations presented here, further refinement of the model to improve the development of selective receptive fields (by adjustment of homeostatic parameters for example) may result in a larger proportion of more orientation selective model V1 units.

Long-range lateral connections in the mouse model are initially random within a Gaussian envelope but also adapt during development (as in all the models presented). A comparison of lateral weights in the ‘salt and pepper model’ (figure 4.19) with those observed experimentally in rat is shown in figure 4.21. Longer range connections in the model appear to be mostly sparse rather than patchy, similar to those observed in rodent visual cortex (rat, figure 4.21a; unpublished observations in mouse, Chalupa and Williams, 2008; grey squirrel, Van Hooser et al., 2006).

Lateral connections in the model adapt based on the correlations in V1 activity and are therefore mainly between neurons with similar orientation preferences (figure 4.21b).

In species with maps, regions of similar orientation preference in the map are preferentially connected. Here, due to the ‘salt and pepper’ arrangement in the rodent orientation map, having connections between neurons that prefer similar orientations is consistent with the lack of patchy connections in these species. It is not yet known if connections in rodents are between neurons with similar orientation preferences and therefore this result constitutes a prediction of this model. Note that this prediction is a direct contradiction of the prediction made by the wiring length minimization model of Koulakov and Chklovskii (2001).

Interestingly, in some rodent studies, patchy lateral connections *have* been observed in V1 (rat; Rumberger et al., 2001, grey squirrel; Kaas et al., 1989). The models presented above suggest that this discrepancy in experimental results could be due to the fact that a number of different parameters may contribute to the organization for orientation, and therefore orientation preference in rodents may not be truly random. If this is the case, clustered regions of orientation may still occur. It is unlikely that small clusters of this size would be observed using optical imaging methods due to its low signal to noise ratio. It is also possible that due to the small size of the recording area in two-photon imaging experiments (and small amount of data that has so far been recorded) that these clusters have not yet been observed. Therefore, the model also predicts that regions of clustered lateral connections may be observed, and that these would also be regions in which orientation preference is clustered. Interestingly in this model, sparse lateral connections maintain the scatter in retinotopy rather than decrease it during development. That is, the scatter in afferent receptive fields is reflected in the lateral connections during development.

These results suggest that mechanisms that result in activation ‘clouds’ (sparse activation of cells in V1) during development, together with activity-driven self-organization, could lead to the ‘salt and pepper’ organization for orientation observed in rodents. However, in order to answer the question ‘is the organization of rodent cortex fundamentally different from that of higher mammals’ it is also necessary to consider how similar mechanisms might affect the cortex of higher mammals. If similar amounts of scatter or changes in lateral connections can also result in a ‘salt and pepper’ map in cat, such a result would suggest that these parameters constitute the crucial differences between these species. If these changes do not simply result in a salt and pepper map, this implies that other differences, such as the magnification factor or receptive field overlap also contribute to the species differences in organization.




Figure 4.21: **Lateral connections in ‘salt and pepper’ mouse model** (a) Lateral connections in a slice of adult rat visual cortex labelled using a transsynaptic tracer based on the rabies virus that labels connections made between single cells. The red cell in B is the initially infected cell and the green cells label those that are connected. Connections between cells appear to be sparse rather than patchy and more dense around the region closest to the cell. Scale bars (in both A and B) $200\mu\text{m}$. Reprinted from Wickersham et al., 2007 (b) Examples of lateral connection weights for individual units in the model V1 for the simulation shown in figure 4.19. Lateral weights in the model units are also more dense in the region closest to the unit but are sparse at longer distances.

4.4.2 Orientation maps persist despite retinotopic scatter in cat model

In this section, the different simulation paradigms as tested in the mouse model are carried out for the parameterized ALISSOM cat model. The mouse results presented in the last section are collated and presented for ease of comparison in figure 4.22.

As for the mouse model above, without including the jitter in connection fields in the parameterized ALISSOM cat model, the simulation will develop an orientation map (figure 4.23a). The scale and features in this map are a reasonable match to experimentally measured orientation maps in cat area 17 (V1) as shown in figure 4.24. As discussed above (section 4.3), implementing a similar amount of cortical scatter in the cat model results in a smaller overall scatter in visual field position. This scatter is, in general, not enough to result in sparse activation in V1 and therefore the map structure in this case is conserved (figure 4.23b). If the standard deviation of connection field jitter in the cat model is increased to $\sigma = 0.25^\circ$ (figure 4.23c) the orientation map still remains salient, unlike in the mouse case (figure 4.22c). In future, even greater amounts of scatter could be included in the model in order to determine if there is a point at which scatter would influence the map structure to a similar extent as in the mouse model.

Changes in lateral connections in the cat model identical to those made in the mouse model (figures 4.23d,e and 4.22d,e) result in more spread out and sparse activation, but the responses of nearby neurons are still correlated. As a result, these changes (although affecting the structure of the map) do not destroy the clustering of similar orientation preferences as they do in the mouse model.

These results suggest that due to the differences in magnification factor and receptive field overlap, a much larger amount of ‘physical’ scatter or more extreme changes in the lateral interactions would be required in the cat in order to observe disorder in the organization of orientation preference. In order to achieve a ‘salt and pepper’ organization, it seems necessary for oriented stimuli on the retina to activate only a subset of cortical neurons that are not neighbours, i.e. ‘blobs’ of activation will lead to maps whereas ‘sparse clouds’ of activation will lead to ‘salt and pepper’ organization.

It is possible to speculate from these results, although they are by no means conclusive, that rodents and higher mammals merely differ in the parameterization of the same




Figure 4.22: **Summary of mouse model results.** From left to right, at iteration 20000: V1 activity (retinal stimulus is a single oriented and elongated Gaussian as displayed in inset), orientation preference, two-photon like plot of central 10×10 patch, receptive fields in a 3×3 patch of V1. (a) Simulation without any *in-connection* scatter. (b) Simulation with experimentally measured amount of retinotopic scatter 'built in' via scatter in V1 afferent *in-connections*. (c) Simulation with an increased amount of *in-connection* scatter. (d) Simulation with no retinotopic scatter but with increased lateral inhibitory strength and decreased lateral excitatory radius. (e) Simulation with experimentally measured amount of retinotopic scatter, increased lateral inhibitory strength and decreased lateral excitatory radius. These results suggest that an increase in the scatter and sparsity of activity patterns during development can lead to the observed 'salt and pepper organization' in mouse V1.




Figure 4.23: **Cat model results.** From left to right, at iteration 20000: V1 activity (retinal stimulus is a single oriented and elongated Gaussian as displayed in inset), orientation preference and selectivity, two-photon like plot of central 10×10 patch, receptive fields in a 3×3 patch. (a) Simulation without any *in-connection* scatter. (b) Simulation with experimentally measured amount of retinotopic scatter 'built in' via scatter in V1 afferent *in-connections*. (c) Simulation with an increased amount of *in-connection* scatter. (d) Simulation with no retinotopic scatter with increased lateral inhibitory strength and decreased lateral excitatory radius. (e) Simulation with experimentally measured amount of retinotopic scatter, increased lateral inhibitory strength and decreased lateral excitatory radius. Similar changes to retinotopy and lateral interactions as those made in the mouse model do not result in similar sparse clouds of activity and therefore do not lead to salt and pepper organization in the cat model.




Figure 4.24: **Comparison of cat model and experimentally measured orientation map.** (a) 1mm $\times$ 1mm of cat primary visual cortex (area 17) color coded for orientation preference as shown in the color key. Adapted from Engelmann et al. (2002). (b) 1mm $\times$ 1mm of the ALISSOM cat model V1 sheet orientation map. The map features appear to be of a qualitatively similar scale in both the model and the animal data.

underlying developmental process. The key differences between these species are then likely to be those that determine the difference between ‘blobs’ or ‘sparse clouds’ of activation in V1 during development, that is, in rodents, the parameterization may make it more likely that activity patterns fall into the class of ‘sparse clouds’ rather than blobs of activity.

4.5 Development using natural images

In order to determine whether these results hold when using more realistic images, the simulations described above were repeated using images from a natural image dataset instead of elongated Gaussian patterns. The results of these simulations for mouse and cat are shown in figures 4.25 and 4.26 respectively. There are a number of considerations to make when using natural images. As described in section 2.3.3, it is currently necessary to tune the ALISSOM model depending on how frequently V1 neurons are activated by the stimulus. In the mouse and cat models, changes have been made to the LGN receptive field sizes (and hence spatial frequency preferences), and extent and strength of lateral connections. This means that the frequency parameter in these models requires tuning in order to achieve good receptive field development (and maps in the no jitter case). The results presented here have not be subject to extensive tuning, the same frequency parameter ($\nu = 5$) as used in the simulations presented in




Figure 4.25: **Summary of mouse model results using natural images.** From left to right, at iteration 20000: Orientation preference and selectivity (plots in the figure are normalized to highlight differences accross the different simulations), orientation preference, two-photon like plot of central 10×10 patch. (a) Simulation without any *in-connection* scatter. (b) Simulation with experimentally measured amount of retinotopic scatter 'built in' via scatter in V1 afferent *in-connections*. (c) Simulation with an increased amount of *in-connection* scatter. (d) Simulation with no retinotopic scatter but with increased lateral inhibitory strength and decreased lateral excitatory radius. (e) Simulation with experimentally measured amount of retinotopic scatter, increased lateral inhibitory strength and decreased lateral excitatory radius.




Figure 4.26: **Cat model results using natural images.** From left to right, at iteration 20000: Orientation preference and selectivity (plots in the figure are normalized to highlight differences across the different simulations), orientation preference, two-photon like plot of central 10×10 patch. (a) Simulation without any *in-connection* scatter. (b) Simulation with experimentally measured amount of retinotopic scatter ‘built in’ via scatter in V1 afferent *in-connections*. (c) Simulation with an increased amount of *in-connection* scatter. (d) Simulation with no retinotopic scatter but with increased lateral inhibitory strength and decreased lateral excitatory radius. (e) Simulation with experimentally measured amount of retinotopic scatter, increased lateral inhibitory strength and decreased lateral excitatory radius.

chapter 2 has been used. Future work to improve the parameter tuning should lead to clearer results.

In any case, the results for the mouse model (figure 4.25) appear to be similar to those obtained for Gaussian images. Maps are present in the mouse simulation without any jitter in afferent connections (figure 4.25a), these maps become much more scattered as increasing amounts of retinotopic jitter are added to the simulation (figure 4.25b,c). Lateral changes similar to those made for the Gaussian input simulations in section 4.4.1 (increasing the strength of lateral inhibitory interactions ($\gamma_I = -3.0$) and decreasing the radius of lateral excitatory interactions ($r_E = 0.1$)) alone, again decreases the map structure by making the size of orientation domains smaller (figure 4.25d). Combining afferent jitter with lateral changes (figure 4.25e) increases both the map scatter and the orientation selectivity of receptive fields. It is expected that a more careful tuning of lateral connection changes to ensure *sparse clouds* of activity in response to the stimulus would further improve these results.

For the cat model (figure 4.26) the results aren't as clear, however maps which are present without any jitter in afferent connections (figure 4.26a) do appear to be maintained despite increasing amounts of scatter, in contrast to the mouse model (figure 4.26b,c). This is consistent with the results presented above for Gaussian inputs. In this case the effect of changing the extent of connections has a much more dramatic effect on the changes in lateral connections shown in figure 4.26d,e). Again, this model will require better tuning of the frequency parameter in order to compare the results of the mouse and cat simulations.

Overall, changing the input to natural images, although requiring more tuning, does not significantly effect the main results. However, this chapter is not supposed to be an exact model of the cortex, to some extent the parameters are chosen in order to achieve the results and provide a possible theory for how the mouse cortex works. It is harder to tune the model to achieve the required organization when using natural images because the activation of V1 is much more variable during development. It is expected that there is a set of parameters which will achieve a clearer demonstration of salt and pepper organization in mouse provided oriented lines within natural images lead to *sparse clouds* of activity in V1.

4.6 Discussion

The first results in this chapter (section 4.2) provide a comprehensive analysis and quantification of retinotopic scatter using two-photon imaging data. This data analysis has been successful in providing an approximate measure of the type and extent of the retinotopic scatter in mouse at the single-cell level. It has also highlighted some important methodological considerations for future experiments, namely that it is important to ensure good coverage of the visual field by the stimulus grid.

The simulations presented in this chapter suggest that ‘salt and pepper’ orientation maps could arise in the mouse because of sparse activation of V1 during development. In this case, ‘built in’ retinotopic scatter helps to decorrelate the activity of neighboring model V1 units and strong lateral connections help to ensure sparsity. Results from the introduction of similar mechanisms in the cat model suggest that a larger amount of scatter or more drastic changes would be required in higher mammals to result in salt and pepper organization due to the larger magnification factor and greater receptive field overlap in these species.

A crucial test of this hypothesis would be to perform a similar analysis on the gray squirrel, which is a highly visual rodent with no apparent orientation map (Van Hooser et al., 2005). The squirrel has a magnification factor of approximately $30\mu\text{m}/^\circ$ which is intermediate between the value for mouse ($15\mu\text{m}/^\circ$) and cat ($667\mu\text{m}/^\circ$). The simulations presented here suggest that there are several factors that could contribute to the formation or lack of map structure, and therefore that the switch between mapping and no mapping may be a continuum depending on the combination of important parameters. The tree shrew (which is a member of the order scandentia) also has a similar magnification factor as the gray squirrel, yet has an exquisitely well organized V1 (Bosking et al., 1997). It would also be interesting to compare these more closely matched species and establish whether the same conclusions can be drawn regarding the important parameters.

As discussed previously, several suggestions have been made as to the possible optimal nature of a ‘salt and pepper’ organization in mouse. The wiring length minimization model of Koulakov and Chklovskii (2001) suggests that ‘salt and pepper’ organization is optimal when each neuron connects an equal amount to other neurons that prefer other orientations. The lateral connection development in the simulations presented here suggest that just as in other species, lateral connections in rodents will be stronger

between neurons that prefer similar orientations, reflecting the correlations in the visual input. The observation of long and short range sparse connections in rodent V1 (Van Hooser et al., 2006; Wickersham et al., 2007) also indicates that connections in rodents are unlikely to be uniform. Evidence from cross correlation analysis of synaptic currents in rodents has also shown that at a fine scale ($0.5 - 2\mu\text{m}$) there are connected subnetworks of excitatory cells, i.e., functional local connections between neurons in rodent V1 may in fact be very specific (Ohki and Reid, 2007; Shepherd et al., 2005; Yoshimura and Callaway, 2005; Yoshimura et al., 2005). Some studies also suggest that each neuron has the potential for extensive rewiring without the need for extensive axonal or dendritic growth (Bé and Markram, 2006; Kalisman et al., 2005).

Optimization based on coverage of the visual field for all orientations seems a more likely driving force behind ‘salt and pepper’ organization given the small magnification factor in mouse. Inspection of the orientation map that develops without any scatter in mouse (figure 4.15) indicates that if there was an orientation map in mouse, all orientations would be represented within a cortical region corresponding to approximately the size of a two-photon patch, i.e. within a sizable $15 \times 15^\circ$ portion of the visual field. In the salt and pepper map, all orientations may be represented within a smaller area. Whether the random selection of orientation preferences in the salt and pepper mouse maps cover the visual field optimally is a crucial topic for future investigation.

The ‘built-in’ scatter used in the simulations in this chapter is a way of including the measured amount of retinotopic scatter observed in the adult mouse V1 without requiring consideration of how this scatter might arise in a realistic way during development. Chapter 5 investigates in detail how this scatter may arise. Consideration of these underlying mechanisms also introduces a possible alternative to the activity-driven development presented here. The implications of these different models are discussed in more detail in chapter 6.

Chapter 5

Mechanisms underlying development of disorder in mouse V1

5.1 Introduction

The models presented in chapter 4 suggest that retinotopic scatter in mouse visual cortex contributes to the development of ‘salt and pepper’ orientation maps by increasing the likelihood of sparse clouds of activity during development. Although this suggests that there is a link between retinotopic scatter and organization of orientation preference, the actual implementation of retinotopic scatter in the models presented in chapter 4 reflects the final retinotopic scatter in the adult cortex, without describing how this type of scatter could arise. If this implementation was to arise naturally in the cortex it would require that LGN cells make very specific initial connections with V1 neurons, which is contrary to the observed course of development. Initially, LGN connection arbors grow into cortex in a retinotopically defined manner that is further refined (O’Leary and McLaughlin, 2005), as will be illustrated in figure 5.1.

There are several possible mechanisms by which retinotopic scatter may arise during development from an initial unrefined set of retinotopic connections. These are summarized as follows and illustrated in figure 5.1 (although this may not be an exhaustive list):

- Sparse activation patterns in V1. As described in chapter 4, sparse clouds of activity can lead to the development of salt and pepper maps and sparse lateral connections. Without directly ‘building-in’ retinotopic scatter, it may be possible

to drive the development of this scatter if activity patterns are already sparse.

- Scatter in RGC/LGN receptive fields on the retinal photoreceptors (figure 5.1d). A recent experimental study of mouse LGN (Grubb et al., 2003) suggests that although the organization of mouse LGN may be ordered with respect to representation of the visual field, individual LGN neuron receptive field positions may be quite highly scattered.
- Scatter in LGN to V1 arbors (figure 5.1e). Although there is, as yet, little evidence to confirm or deny this possibility, this type of scatter would contribute to the scatter in V1 receptive fields in a similar way to scatter in RGC/LGN receptive fields (figure 5.1f).
- Sparse afferent connections from LGN to V1 (figure 5.1g). Again there is little experimental data available regarding the sparsity of afferent connections to V1, however, this sparsity could also contribute to an increase in the scatter of V1 receptive fields.

The possibility that these mechanisms, or a combination of them, could lead to the observed differences in organization between species is explored in this chapter.

5.2 Methods

The models presented in this chapter are based on the general mouse and cat model architectures described in chapter 4. The specific parameterization and analysis methods, where they differ from those in the previous chapter, are explained here. All other methods are as described in sections 2.2 or 4.3.

5.2.1 Mechanisms of retinotopic scatter

The ‘built-in’ scatter described in section 4.3.1 is a direct but artificial way to include retinotopic scatter without requiring that it arise during development in the model. This case was used in chapter 4 to illustrate how retinotopic disorder may affect the development of organization for orientation (section 4.4.1). Figure 5.1 illustrates all of the underlying mechanisms explored in this chapter that may lead to the retinotopic scatter presented in section 4.2.




Figure 5.1: Different types of scatter. Connections of one unit in the model can be defined as either *in connections* (a), the units from which one unit receives connections, or *out connections* (b), the units to which that unit connects. The different scatter types can be illustrated, with in some cases one or both of these definitions being used. (c) In-connection scatter of V1 units is used in the initial models from chapter 4 to ‘build-in’ the retinotopic scatter as measured experimentally. This type of scatter is also illustrated in figure 4.12. (d) A similar type of scatter as in (c) applied instead to RGC/LGN cells, i.e., the preferred retinal location of each LGN cell is randomly scattered. (e) Scatter in the *out connections* of RGC/LGN units corresponds to scatter in the retinotopic location of afferent arbors in V1; this type of scatter is equivalent to RGC/LGN receptive field scatter as illustrated in (f). (g) Increasing the sparsity of afferent connections to V1 can also result in retinotopic scatter. (h) This type of scatter combined with scatter in RGC/LGN receptive fields is investigated further in section 5.3.2. (i) The effect of increasing the sparsity of lateral connections on development is investigated in section 5.3.1.

In most cases, scatter is implemented in the model using a Gaussian random number generator for the x and y coordinates of the centre of *in-connection* fields (see figure 5.1). I.e. if (x_{new}, y_{new}) are the new center coordinates of the connection fields (F_i) and (x_{old}, y_{old}) are the original (retinotopic) center coordinates, the scattered projection is specified using:

$$\text{For } i \in F_i, (x_{new}, y_{new})_i = (x_{old} + x, y_{old} + y) \quad (5.1)$$

where x and y are values chosen randomly from a Gaussian distribution with mean=0 and standard deviation= σ . The standard deviation (σ) is usually specified from the experimental data (the values for each simulation are indicated in the text). In the case of LGN receptive field scatter (section 5.3.2) a uniform random distribution within a specified range is used (i.e. values chosen from a uniform distribution around zero rather than a Gaussian distribution).

Sparse Gaussian connections are implemented by applying a binary threshold to Gaussian random noise.

$$\begin{aligned} P &= 1 \text{ if } V \geq \text{sparsity} \\ &= 0 \text{ if } V < \text{sparsity} \end{aligned}$$

where P is the connection probability of a weight $\omega_{i,j}$ in the connection field F_i . V is a scalar value chosen from a Gaussian random distribution with mean zero and unit variance. The threshold level controls the probability of a connection existing or not, and is controlled by a *sparsity* value that varies between 0 and 1. Connection weights ($\omega_{i,j}$) are unrelated to the existence of a connection (P), i.e., initial weights are still random within a Gaussian envelope but only apply to the set of available sparse connections. The set of sparse connections is also fixed and does not change throughout development; the implications of this assumption are discussed in section 4.6.

5.3 Results

This results section is split into three sub-sections. In the first I have investigated a preliminary hypothesis that both retinotopic scatter of the type observed experimentally and ‘salt and pepper’ orientation maps could arise through activity-driven development. In the second section I explore the possibility that sub-cortical scatter could underlie the scatter observed in V1 receptive fields. In the third and final section I

investigate how sparsity of connections could also influence retinotopic scatter and orientation selectivity. Both the implementation of sub-cortical scatter and sparse connections lead to orientation selective receptive fields *without* activity driven development. Establishing the relative merits of these two alternative hypothesis, i.e., activity driven or non-activity driven development will require future experimental work, and is discussed in more detail in section 5.4 and chapter 6.

5.3.1 Development of retinotopic scatter driven by sparse clouds of activity

In the simulation described in this section there is no ‘built in’ scatter at the start of development (as in chapter 4). Instead, patterns of sparse, spread out ‘clouds’ of initial activity are encouraged in the model by making lateral excitatory and inhibitory connections sparse ($\text{sparsity} = 0.95$) and increasing their strength ($\gamma_e = 1.5, \gamma_i = -2.0$). The size of afferent *in-connection* fields at the start of development is also increased ($r_A = 1.0$) in order to allow scattered receptive fields to develop, and lateral excitatory connections are longer than in previous simulations ($r_E = 1.0$).

These sparse activity patterns, over the course of development, result in an increase in the retinotopic scatter (shown in figure 5.2). The increase in scatter does not result purely from the action of lateral connections, since it is present even when the retinotopy preference is measured in the absence of lateral interactions (figure 5.2a). That is, retinotopic scatter is being embedded in the afferent connection weights during development.

Orientation preference maps and receptive fields in this model (figures 5.3a,b,c) are less selective than in previous simulations, but this is not necessarily a feature of the model, i.e., this may be improved by changes in other model parameters. The purpose of this simulation is to illustrate the possibility that the type of retinotopic scatter ‘built-in’ to the model in section 4.4.1 could arise during development, provided activity patterns in V1 are sparse and distributed. This possibility requires further investigation and is discussed in more detail in section 6.




Figure 5.2: **Development of retinotopic scatter driven by sparse lateral connections.** Activity patterns in this simulation are forced to be more sparse during development. (a) Initial (solid line) and final (dashed line) retinotopic scatter histograms for the azimuth (red) and elevation (blue) directions respectively before settling (using afferent activity only). (b) As for (a) for the after settling protocol (also takes into account lateral interactions). In both cases there is an increase in the spread of scatter during development, with final scatter similar to that observed experimentally (figure 4.10).

5.3.2 Subcortical scatter

In the previous models presented, the mapping of receptive field position in LGN units is perfectly retinotopic. However, experimental measurement of retinotopic mapping in mouse LGN suggests that this mapping may only be quite rough (figure 5.4). In Grubb et al. (2003) retinotopic mapping in the dorsal LGN is assessed using pairs of cells. For each cell pair, the separation in cortical distance and the separation in preferred receptive field position is measured. When plotted on the x and y axis there is a correlation in both the dorso-ventral (elevation) direction and in the naso-temporal (azimuth) direction for measured pairs (figure 5.4c), but the retinotopic mapping is clearly not perfect. Inspection of these measurements gives an approximate amount of scatter of $200\mu\text{m}$ (using the maximum distance observed between cells with overlapping receptive fields). This measurement can be transformed into a receptive field scatter of approximately 15° using the magnification factor in the mouse model. This value also corresponds well to the amount of scatter in receptive field distance for near overlapping cells in figure 5.4. More precise quantitative measurement should be carried out in the future, but at this preliminary stage such a rough measure is appropriate.




Figure 5.3: **Orientation preference and receptive field development driven by sparse lateral connections.** (a) Orientation preference and selectivity maps at iteration 0 and 20000. Selectivity is normalized across iteration 0 and 20000 showing that there is little increase in selectivity. (c) Two-photon orientation preference plots at iteration 0 and 20000. (d) Receptive fields for a 3×3 patch of V1 model units at iteration 0 and 20000. Although receptive fields have refined during development, they are not very orientation selective.

Another possible source of scatter in V1 retinotopy could arise from the non-perfect mapping of LGN to V1 arbors. Although we have previously considered this mapping to be perfectly retinotopic, it is also possible that there is some scatter from a perfect mapping. This type of scatter is likely to have an equivalent effect as the scatter in LGN receptive field position, with the retinotopic scatter merely relayed to V1 by a different mechanism (see figure 5.1d,e,f for an illustration). I am not aware of any studies measuring the amount of scatter in afferent arbors (although some studies have shown that the size of arbors in mouse V1 is highly variable; Antonini et al. 1999),




Figure 5.4: Measurement of retinotopic scatter in mouse LGN. (a) The Dorso-Ventral axis of the visual field is mapped onto the rostromedial-caudolateral dorsal LGN (dLGN) while Nasal-Temporal visual space is mapped to the caudomedial-rostromedial dLGN. (b) On a given penetration of the recording electrode, receptive fields with the similar positions can be mapped for cells with different cortical locations. The separation of cells in the dLGN Δ_{cell} and the separation of visual field positions in the azimuth and elevation directions of the corresponding receptive fields $\Delta\text{RF}(\text{NT})$ and $\Delta\text{RF}(\text{DV})$ can be recorded for each pair. Plotting these pairs (separately for the DV (c) and NT (d) directions) shows that there is a correlation, and therefore systematic variation in both directions, but that retinotopic organization is not perfect. In each plot each dot represents one cell pair. Solid lines show the best fitting linear regression of the data and dotted lines show the 95% confidence interval of this fit. Reprinted from Grubb et al. (2003).

but it is possible to speculate that this might introduce an additional similar amount of scatter as that estimated from the LGN scatter results.

The hypothesis that these types of sub-cortical scatter could result in the scatter observed in adult mouse V1 is tested by including this scatter in the mouse model (without any other type of scatter). Figures 5.5a and b illustrate these different types of sub-cortical scatter. Figure 5.5c illustrates the effect of this type of scatter on the activity patterns in the LGN and V1, i.e., this type of scatter results in sparse activity in the LGN but a ‘blurred’ blob of activity in V1.

The effect of this scatter on the V1 receptive field position is illustrated in the plots of azimuth preference in figure 5.5d and the histograms of measured scatter in figure 5.5e. Adding an increasing amount of sub-cortical scatter spreads out the receptive field and increases the amount of measured scatter, but not in the same way as measured in the experimental data. Comparison of the azimuth preference plots in figure 5.5 with

those in figure 4.8 shows that rather than preferring scattered and different portions of the receptive field, the relative mapping of nearby cells is maintained.

The relationship between different amounts of subcortical scatter and measurement of scatter using the retinotopy protocol in the model is shown in figure 5.6. As the amount of added scatter is increased, the measured scatter also increases, however the graph also appears to show a non-linear relationship (for both x and y directions). Measured scatter in V1 increases less as an increased amount of subcortical scatter is added. The overall values of measured retinotopic scatter in V1 do not come close to those measured experimentally (figure 4.10, mean 8.5°) until very large (and unrealistic) subcortical scatter values.

It is much clearer to interpret what is happening when this type of scatter is added to the model by looking at the receptive field plots in figure 5.7. Adding LGN receptive field scatter increases the size of the V1 receptive field but does not change the mean preferred retinotopic position, nor the relationship between nearby cells. In fact, it increases the amount of overlap in receptive fields, making them more similar overall. Therefore, regardless of the amount of this type of scatter in the mouse architecture, it is unlikely to (alone) be able to result in the scatter observed experimentally.

Interestingly, orientation maps measured in this model (figure 5.8) show that even before development there is initial orientation selectivity and initial clustering of orientation preference. An initial orientation selectivity (at iteration 0) arises in this model because V1 units are no longer sampling from overlapping LGN On and Off cells, but are instead sampling from randomly offset On and Off cells. Initial clustering then arises because nearby V1 units are connected to overlapping sets of LGN units. This is reminiscent of the model presented by Ringach (2007) for the initial development of orientation maps in higher mammals, although in Ringach (2007) model LGN units were arranged on a hexagonal grid rather than being randomly arranged. As described in section 1.6, it is not clear whether the maps produced by this algorithm are realistic in structure when compared to maps from higher mammals, since they do not appear to be as regular as measured animal maps. For instance, they do not have a ring-shaped 2D Fourier transform (Blasdel, 1992). In any case, this type of scatter alone is not capable of reproducing either the retinotopic scatter, or the ‘salt and pepper’ orientation maps observed in mouse V1.




Figure 5.5: Effect of subcortical scatter mechanisms on mouse model retinotopy. (a) Illustration of LGN receptive field scatter in the model, implemented by jittering the *in connections* to LGN/RGC units. (b) Illustration of afferent LGN to V1 arbor scatter in the model by jittering the *out connections* of RGC/LGN units, this type of scatter can be implemented in an equivalent way to (a). (c) This type of scatter results in sparse activity patterns in the RGC/LGN sheet and a larger ‘blurred’ blobs of activity in V1 (d) Azimuth preference at iteration 20000 for increasing amounts of uniform random scatter with the corresponding maximum value shown. Measurements in this figure are made using the before settling protocol, as the protocol has little effect on the results except to slightly increase the spread in a similar way as in figure 4.15e,f. Preferences are not significantly different before or after development or between azimuth and elevation preferences, therefore these plots are eliminated for clarity. Increasing amounts of scatter appear to preserve the retinotopic relationship between nearby cells and the overall mean preference. (e) Histograms of scatter measurement and Gaussian fits with the parameters shown for azimuth (red) and elevation (blue) at iteration 20000. Increasing amounts of scatter also increase the standard deviation of the scatter but not by an equivalent amount, nor in a similar way, as that observed experimentally (figure 4.10).




Figure 5.6: **Relationship between subcortical scatter and retinotopic scatter** Different amounts of subcortical scatter added to the model are plotted against the mean value of the measured retinotopic scatter of V1 cells in each case. Mean values are calculated from five different simulations in which different seed values are used for the random scatter in LGN connection fields. Error bars show the standard deviation around the mean. The value of V1 retinotopic scatter increases as subcortical scatter is added, for both the x and y directions.

5.3.3 Salt and pepper maps arising from sparse connections to V1

Ringach (2007) also suggests that in mouse there may be fewer connections between the LGN and V1 and this may contribute to the salt and pepper organization. As discussed in section 3.1 there is some evidence to suggest that, in general, connections in rodent V1 may be sparse (Ohki and Reid, 2007; Shepherd et al., 2005; Yoshimura and Callaway, 2005; Yoshimura et al., 2005). Sparseness of connectivity between the LGN and V1 may also contribute to the observed scatter in retinotopy. This hypothesis is tested in the mouse model by increasing the sparsity of afferent *in-connections* from the LGN to V1.

Figure 5.9a,b and c illustrate the effect of increasingly sparse afferent *in-connections* (without any other types of scatter). Although this has little effect on V1 activity at the start of development, at the end of development activity is more sparse (figure 5.9b). Retinotopy analysis in this model shows that although increasing amounts of scatter lead to a slightly increasing spread in preferred position, the scatter is not as dramatic




Figure 5.7: **Effect of subcortical scatter mechanisms on mouse model receptive fields.** (a) Receptive fields at iteration 0 and 20000 for the mouse model without any scatter. (b) Adding sub-cortical scatter spreads out and increases the size of the receptive field area but preserves the mean position preference. Nearby units also have more similar receptive fields because they receive connections from more overlapping LGN cells.

as the levels of scatter observed experimentally.

The relationship between different values for the sparsity of LGN to V1 connections and measurement of scatter using the retinotopy protocol in the model is shown in figure 5.10. As sparsity increases the amount of V1 retinotopic scatter increases in a non-linear way. The variance in the scatter values also increases with increasing sparsity. Large sparsity values indicate a low probability of connection, however in order to reach values approximating the scatter measured experimentally (figure 4.10, mean 8.5°) it seems that sparsity values greater than 0.9 would be required, which would indicate around zero probability of connection, which is clearly unrealistic.

Orientation maps for increasingly sparse connections are shown in figure 5.11. Even before any activity driven development (at iteration 0), in the simulations with more sparse connections, orientation selectivity is increased. The maps also maintain a random salt and pepper organization during development. In this case, orientation selec-




Figure 5.8: Effect of subcortical scatter on mouse model orientation maps. Orientation preference and selectivity maps and two-photon orientation maps for increasing amounts of subcortical scatter at 0 and 20000 iterations respectively. (a) No scatter. (b) 15° scatter (c) 30° scatter. Orientation preference and selectivity maps are jointly normalized across all the maps in the figure to highlight the differences in selectivity.

tivity is again arising from the sampling of non-overlapping LGN On and Off cells, but because of the sparse connectivity, nearby V1 units are no longer necessarily getting inputs from nearby LGN cells. Therefore, orientation preference is not clustered in V1. This model is very similar to the rodent model suggested (but not implemented) by Ringach (2007), with fewer constraints on the positioning of LGN receptive fields.

Consideration of the receptive fields in this model (figure 5.12) confirms that as sparsity increases, cells sample from fewer overlapping LGN cells. It is difficult to appreciate in these receptive field plots (particularly figure 5.12d) how this arrangement of On and Off regions can lead to orientation selective receptive fields, since these plots do not look like the plots of receptive fields measured experimentally (figure 4.1). Interpreting receptive field measurements is highly dependent on the type of stimulus that is used. The receptive fields of units in the model presented previously are measured using a small Gaussian stimulus the same size as the center of an LGN unit receptive




Figure 5.9: Effect of afferent connection sparsity on retinotopy in the mouse model. (a) Illustration of afferent *in connection* receptive field sparsity in the model. (b) Sparse afferent connections have little effect on the activity patterns at the start of development but do result in more sparse activation at iteration 20000. (c) Examples of initial afferent connection weights for one V1 unit for increasing values of the sparsity measure. (d) Azimuth preference at iteration 20000 for increasing sparsity. Measurements in this figure are made using the before settling protocol, as the protocol has little effect on the results except to slightly increase the spread in a similar way as in figure 4.15e,f. Preferences are not significantly different before or after development or between azimuth and elevation preferences, therefore these plots are eliminated for clarity. Increasing sparsity increases the scatter in retinotopy but not very significantly. (e) Histograms of scatter measurement and Gaussian fits for azimuth (red) and elevation (blue) at iteration 20000. Increasing sparsity increases the spread in retinotopy preference but not by an amount equivalent to that observed experimentally (figure 4.10).




Figure 5.10: Relationship between sparsity and retinotopic scatter Different values for the sparsity of LGN to V1 connections are plotted against the mean value of the measured retinotopic scatter of V1 cells in each case. Mean values are calculated from five different simulations in which different seed values are used to determine the existing connections. Error bars show the standard deviation around the mean. The value of V1 retinotopic scatter increases as sparsity is increased, for both the x and y directions.

field (smallest Gaussian in figure 5.13a). In the experimental measurements of receptive fields the stimulus is a contrast modulated white noise movie (shown in figure 5.13a) that contains a range of spatial frequencies. Measuring the receptive field of the same 3×3 units in the model V1 with different spatial frequency stimuli (figures 5.13b,c,d) can lead to very different interpretations of the same underlying connectivity. The experimentally measured receptive field plots are obtained using a stimulus that contains features that are a similar size as the larger of the Gaussian stimuli used in figure 5.13c or d. Therefore the experimental measurement may not reveal the full underlying complexity of the receptive field.

Although this simulation does not reproduce the observed scatter in retinotopy, the model has confirmed the possibility that increasing the sparsity of connections in the model may be a mechanism for the development of salt and pepper orientation maps in rodents. A simulation in which the sub-cortical scatter in LGN receptive fields is combined with sparse connectivity is illustrated in figure 5.14. The combination of scatter and sparse afferent connectivity leads to an increased spread in the amount of scatter. Again, orientation preference maps are initially selective and have a salt and




Figure 5.11: Effect of afferent connection sparsity on orientation preference maps in the mouse model. Orientation preference and selectivity maps and two-photon like orientation maps for increasingly sparse connections at 0 and 20000 iterations. (a) Full connection fields. (b) 0.5, (c) 0.7, and (d) 0.9 sparsity (sparsity measure corresponds approximately to the probability of a connection existing). Increasing sparsity leads to increased orientation selectivity both before and after development. Orientation preference and selectivity maps are jointly normalized across all the maps in the figure to highlight the differences in selectivity.

pepper organization. Figures 5.14e and f show the receptive fields measured before and after activity driven development. During development the orientation preference is maintained, but the receptive fields are refined.

This model is a plausible mechanistic model for the formation of ‘salt and pepper’ maps in the mouse. Further experimental studies will be required to evaluate the relative merit of this model and the model presented in chapter 4 in which salt and pepper map development is driven by sparse activity. The sparse connectivity model shown




Figure 5.12: **Effect of afferent connection sparsity on mouse model receptive fields.** (a) Receptive fields at iteration 0 and 20000 for the mouse model with full connections. (b),(c),(d) Increasing the sparsity of connections reduces the number of overlapping On/Off connections and increases the orientation selectivity.




Figure 5.13: **Reverse correlation of receptive fields with different stimuli.** (a) Three different sizes of Gaussian stimuli and an example frame from the contrast modulated white noise movie (containing a specific range of spatial frequencies) that is used to perform spike triggered averaging in the experimental data. Stimuli are matched in scale for comparison. (b),(c),(d) Receptive fields in a 3×3 patch of model V1 measured using the stimuli shown in the inset (corresponding to the sizes shown in (a)). Using larger stimuli (or lower spatial frequency) can ‘blur’ the underlying complexity of the receptive field, making the plots smoother, but making it more difficult to interpret the receptive field plots.

in figure 5.14 does not reproduce the amount of scatter observed in the data, but only a small amount of sub-cortical scatter is used in this simulation (15°). In the future it will be necessary to experimentally quantify the amount of scatter and the sparsity of each set of connections more precisely, in order to establish whether these mechanisms are realistic.

5.3.4 Sparse afferent connections also lead to ‘salt and pepper’ maps in cat model

As in chapter 4, it is also necessary to investigate how these same changes affect the parameterized ALISSOM cat model. Results for these simulations are illustrated in figure 5.15. Unlike in the simpler simulations presented in section 4.4.1, the same changes (LGN receptive field scatter and sparse afferent *in connections*) have similar




Figure 5.14: Combination of sparse afferent connections and LGN receptive field scatter in the mouse model. (a) Illustration of afferent *in connection* receptive field sparsity and LGN receptive field scatter in the model. (b) Azimuth preference for units in V1 at iteration 20000. Retinotopy measurements in this figure are made using the before-settling protocol, as the protocol has little effect on the results except to slightly increase the spread in a similar way as in figure 4.15e,f. (c) Histogram of scatter values for azimuth (red) and elevation (blue) at iteration 20000. Standard deviation values (σ) are larger than in the simulations with the mechanisms included separately, but still do not reproduce the full spread in values measured experimentally (figure 4.10). (d) Orientation preference and selectivity map (normalized together with all previous maps in this section) and two-photon orientation preference map at iteration 20000. (e) Receptive fields measured at iteration 0 consist of randomly placed but non-overlapping separated On and Off regions; these are further refined but remain selective at iteration 20000, as shown in (f).




Figure 5.15: **Cat model results.** From left to right: Orientation preference and selectivity, two-photon plot of central 10×10 patch and receptive fields in a 3×3 patch for iteration 0 and iteration 20000. (a) Simulation without any scatter. (b) Simulation with scatter in LGN receptive fields (the equivalent scatter for cat is $\sigma = 0.34^\circ$). (c) Simulation with sparse afferent connections (sparsity=0.9) (d) Simulation with both LGN receptive field scatter and sparse afferent connections. Changes in both subcortical scatter and afferent sparsity in the cat and mouse model have a similar effect.

effects in both cat and mouse. This suggests that afferent connections in rodents are more sparse than in higher mammals.

The inclusion of LGN scatter alone in the cat model again results in a clustering of orientation preferences (figure 5.15b), as in Ringach (2007). It will be interesting in future to establish whether this mechanism may be an alternative to the spontaneous activity-driven development discussed in chapter 2.

5.4 Discussion

This chapter has investigated the mechanisms that could lead to retinotopic disorder in rodents. Both chapter 4 and this chapter have been successful in highlighting the mechanisms that could lead to ‘salt and pepper’ orientation maps in rodents and have therefore isolated the possible species differences that underlie the differences in or-

ganization. However, a number of possible explanations have emerged out of this investigation, the relative merits of which will require both further experimental and computational research.

In section 5.3.2, scatter in LGN receptive fields or in afferent arborizations in V1 (possibly due to errors in axon pathfinding and/or molecular gradient cues) are explored. These types of sub-cortical scatter do not reproduce the same type of scatter measured experimentally in mouse. However, combining this type of scatter with sparse afferent connections to V1 can result in scatter similar to the type measured experimentally. Sparse connections alone can reproduce a ‘salt and pepper’ organization for orientation in the model mouse V1, even without any activity driven development (section 5.3.3). This type of model, where orientation selectivity results purely from the initial wiring of connections in the early visual system has also been explored previously in higher mammals by Ringach (2007). My comparison of these mechanisms in the mouse and cat model suggest that if this is the case, connections in rodent species would be likely to be significantly more sparse than in higher mammals.

Even though the type of orientation maps formed by initial random wiring of V1 receptive fields do not correspond well to experimentally measured maps in higher mammals, it is possible that these mechanisms do underlie orientation organization in rodent species. One further consideration, not included in these models, is that connections are likely to be both forming and retracting in the developing cortex. In my implementation of sparse afferent connections, the existence of connections is fixed during development (although the weights of existing connections change). Including more dynamic connection patterns in the model may result in reorganization of the initial wiring that I have not taken into consideration here.

The overriding conclusion of the investigation, both in this chapter and in chapter 4, is that a program of comparative experimental work is required to distinguish between the possible models. Measurement of retinotopy, orientation preference and selectivity, detailed receptive fields, and activity patterns in the same animal, at different stages of development would clarify the most likely mechanisms underlying the differences between cat and mouse. These experiments should be particularly targeted at the following questions:

- What is the form of activation in mouse V1 during early development, i.e., is activity sparse in mouse and more ‘blob-like’ in cat?

- How much do retinotopic scatter and orientation preference change during development? Less change in the amount of scatter or in orientation tuning width of V1 neurons would indicate that receptive fields are fixed at the beginning of development by initial wiring and that activity does not play an important role in establishing the map. Ideally these measurements would somehow dissociate the effects of afferent and lateral connections, as lateral connections that are also adapting during development may sharpen orientation tuning and/or decrease the amount of scatter.
- What do receptive fields look like during development? As observed in the models above it is often informative to look at the detailed features of the receptive fields in order to establish how they are constructed from LGN On and Off afferent connections. Receptive fields should be reverse correlated using a stimulus with a spatial frequency corresponding to the approximate size of the LGN receptive field centre in order to fully appreciate the receptive field features.
- How sparse are afferent connections in rodents in comparison with higher mammals?
- How scattered are LGN receptive fields in rodents in comparison with higher mammals?

The models presented here, and in chapter 4, support the hypothesis that differences in organization between rodents and higher mammals are likely to arise purely through a different parameterization of the same underlying architecture. Each species may also be constrained by the physical size of the system and respond differently to similar underlying mechanisms. Although the models presented cannot conclusively determine the mechanisms, they have been valuable in highlighting the important parameters for investigation in future experimental and computational work.

Chapter 6

Overview and future work

6.1 Introduction

The aims of this thesis are summarized in section 1.7 and in this chapter I will discuss both whether these aims have been achieved, and how these aims could be further extended. The four results chapters in the thesis have covered a breadth of topics, each of which have opened up many avenues for future research. Each individual topic is discussed here in a separate section (chapter 2 in section 6.2, chapter 3 in section 6.3 and chapters 4 and 5 in section 6.4), though there is some overlap between each of these topics and therefore between each section. Some general points which apply to all of the work presented in this thesis are discussed in section 6.5.

6.2 Homeostatic requirements for stable development

As discussed in section 1.4, experimental evidence from ferrets (Chapman and Bonhoeffer, 1998; Chapman et al., 1996) and cats (Crair et al., 1998; Godecke et al., 1997) indicates that orientation preference maps in the primary visual cortex (V1) develop in a stable way, with the initial layout determined before eye opening. This stability is remarkable given the massive circuit reconstruction and change in visual inputs occurring at this time.

Homeostatic mechanisms such as adjustment of intrinsic excitability (Schulz, 2006), scaling of synaptic strengths (Turrigiano and Nelson, 2004), or sliding thresholds for

plasticity (Abraham, 2008; Bienenstock et al., 1982) have been shown to allow individual neurons to maintain stability of overall activity levels, whether during learning or, over longer timescales, during development (Davis and Bezprozvanny, 2001; Desai, 2003). In chapter 2, I explored the hypothesis that homeostatic mechanisms could also underlie the stability of organization of orientation preference maps.

I have proposed a model of stable orientation map development based on realistic homeostatic mechanisms (Adaptive LISSOM; ALISSOM). The ALISSOM model shows that adjustment of intrinsic excitability and synaptic normalization can achieve a desired average activity in a V1 model neuron (as previously proposed), but that this would not be enough to stably develop an orientation map. Neuron units in the model V1 must also ensure a balance between afferent inputs and lateral interactions which is required for good map organization. Neurons must also maintain their average response to an input, rather than simply average activity levels, and therefore should adjust to changes in the statistics of visual input.

This model is the first to reproduce stable orientation map development, using realistic homeostatic mechanisms, despite changes in the strength and pattern of inputs before and after eye opening. This work has thus achieved the initial aims of investigating how stable development of orientation preference maps could occur using homeostatic mechanisms, and also to highlight the requirements for homeostatic plasticity in the developing V1. There are a number of immediate future aims which have arisen from this work which should now be addressed. These are discussed in the next section.

6.2.1 Immediate future aims

Automating parameter tuning

The ALISSOM model is not only a useful tool for understanding the homeostatic requirements for development, but is also a simplified starting point for future research. The model has completely eliminated the scheduled changes required throughout development in the LISSOM and SOM models (such as learning rate decay and shrinking of the lateral excitatory radius). This has greatly reduced the number of parameters to be ‘hand tuned’ and makes it much easier to use different types of retinal input and maintain robust orientation map development. Unfortunately, the model has not completely eliminated the requirement for parameter tuning. In order to maintain the

balance between afferent and lateral inputs to V1 in the face of changing retinal input, a hand tuned ‘frequency’ parameter has been introduced. The ‘balance’ parameter (the ratio of afferent to lateral input strength) is also a hand tuned parameter. In future work it will be important and useful to eliminate these parameters.

Contrast gain control

Another consideration for future work is the reconciliation of this model with the contrast invariant properties of V1 neurons. The threshold and gain of the sigmoidal output function regulate the ‘intrinsic excitability’ of each model neuron. However, this function, together with the dynamics of lateral connectivity, is also important in determining the neuron’s contrast response. As a result, the ALISSOM model does not have a realistic contrast response (similar to LISSOM), but preliminary work (discussed below) suggests that it will be possible in future to reconcile ALISSOM with contrast invariant orientation tuning.

As discussed in section 1.3, saturation of a V1 cell response depends on stimulus contrast but is also independent of stimulus properties such as orientation. That is, responses to different orientations are lower, yet saturate at the same contrast (contrast gain control). This ensures that the orientation tuning curve of a V1 cell is invariant with contrast (Sclar and Freeman, 1982; Skottun et al., 1987). As contrast increases, the tuning width of the orientation tuning curve does not change (see figure 6.1a).

The orientation tuning curve and contrast response of an example model neuron from the ALISSOM simulation (section 2.3.1) is shown in figure 6.1b. The response to all orientations, even those not preferred by the neuron, increase with increasing contrast, leading to a widening of the orientation tuning curve and an overall increase in response as contrast increases. In order to achieve contrast invariance, several other mechanisms need to be included in the ALISSOM model.

The topic of contrast invariance has been the subject of a large amount of research. The models of Albrecht and Geisler (1991); Ben-Yishai et al. (1995); Carandini et al. (1997); Kayser et al. (2001); Lauritzen et al. (2001); Somers et al. (1995), and the ‘normalization’ models: Carandini and Heeger (1994); Carandini and Ringach (1997); Heeger (1992) use varying mechanisms of intracortical inhibition to explain contrast-invariant orientation selectivity in the adult V1. In such models, intracortical inhibition that is very broadly tuned causes suppression of the response to all non-preferred orien-




Figure 6.1: **Contrast responses in real neurons in comparison with ALISSOM model.**

(a) The left hand graph shows the response of one V1 cell to stimuli of different orientations at 4 different contrasts (10, 20, 40, 80%). Reprinted from Sclar and Freeman (1982). The cell is tuned to a preferred orientation of approximately 220° . On the right, the same data is replotted as contrast tuning curves at different orientations. The response to the preferred orientation is always greater than non-preferred, even with increasing contrast. (b) Orientation tuning curve as in (a) for a model neuron in the ALISSOM simulation described in section 2.3.1. The width of the orientation tuning curve increases with contrast, unlike real neurons. The inset shows a magnification of the lower contrast curves. Re-plotting the data as a contrast response curve reveals that as contrast increases, the response to all orientations increases, eventually reaching a saturation point which is the same regardless of orientation. (c) Adjusting the model by replacing the saturating non-linear output function of model V1 neurons with a non-saturating function (half-wave rectifier) and adjusting the strength of connections (from $\gamma_E = 1.0$, $\gamma_I = -1.0$ and $\gamma_A = 1.0$ to $\gamma_E = 2.0$, $\gamma_I = -3.6$ and $\gamma_A = 3.0$ with 16 settling steps t_s) improves the contrast invariance of orientation selectivity. Contrast responses at non-preferred orientations in the improved model now have consistently lower values than the preferred orientation, as in real V1 neurons.

tations at all contrasts. However, experiments based on cross-orientation suppression, together with a lack of evidence showing cross-orientation inhibition, argue against this type of model (Carandini et al., 2002).

In ‘feed-forward’ models of contrast gain control (Anderson et al., 2000; Carandini et al., 2002; Finn et al., 2007; Geisler and Albrecht, 1997), contrast gain control in V1 is due to various possible mechanisms, including feedforward gain control from the LGN, or synaptic depression at geniculocortical synapses. This feed-forward contrast-gain control, together with an expansive nonlinearity in V1 neuron outputs, could in theory produce contrast invariant orientation selectivity. This model is also supported by evidence that gain control occurs before convergence of information from the two eyes (Truchard et al., 2000).

Preliminary work in the ALISSOM model has shown that implementing a gain control mechanism in the LGN can produce contrast invariant orientation selectivity in V1 model units. If the sigmoidal activation function (σ) is replaced by an expansive non-linearity (e.g. half-rectifying or squaring function), and projection strengths are adjusted to compensate for a decrease in responses, then contrast invariant orientation tuning curves in the fully developed model can be obtained (see figure 6.1c). This method however, is not yet consistent with the homeostatic mechanisms implemented in ALISSOM, where the adjustment of intrinsic excitability is based on a sigmoidal non-linearity. It is also not yet clear how gain control should be implemented throughout development (or indeed whether in reality this property is maintained in the developing cortex).

Reconciling the current version of ALISSOM with contrast invariant orientation selectivity may also eliminate some of the hand tuned parameters discussed above, if the model is changed as follows:

- Rather than maintaining a target average activity in V1, the intrinsic excitability of V1 neurons could be adjusted on a shorter timescale to produce an output activity which is directly related to the afferent input each time the neuron is activated. This would be a more direct implementation of a response homeostasis which also takes into account the strength of afferent activity and therefore the balance between lateral and afferent input types. This would eliminate the requirement for the ‘frequency’ parameter.
- This intrinsic excitability adjustment in ALISSOM could be achieved using a

much simpler threshold adjustment of a non-linearity to replace the adjustment of a sigmoidal activation function.

- Adding contrast gain control to the LGN would mean that changes in strength (from changes in input contrast, for example) would be automatically compensated at the level of the LGN. This would eliminate the requirement for afferent scaling in the model.
- Together, contrast gain control in the LGN and an expansive non-linearity in V1 would also result in contrast invariant orientation selectivity.

These mechanisms would eliminate the ‘frequency’ parameter and reproduce contrast invariant orientation selectivity but would still not eliminate the ‘balance’ parameter. However, the ‘balance’ parameter would be constant for all simulations which have the same model architecture, i.e., the balance would only require tuning if the sizes of connection fields in the model are changed. In order to eliminate this parameter, it would be necessary to also include a response homeostasis which is based on the sizes of receptive fields. This type of response homeostasis has been observed experimentally in the afferent projection from the retina to the superior colliculus in mouse (Chandrasekaran et al., 2007). In this case, genetic manipulations which result in larger receptive fields also result in lower peak visual responses. Receptive field size homeostasis could be included in the model by regulating the balance parameter using the size of connection fields (or the total number of connections).

Implementation of these immediate future aims would transform the current ALISSOM model into an even simpler and more robust starting point for future research. The possible future avenues of this research are described in the next section.

6.2.2 Future outlook

Comparison with experiment

The ALISSOM model has so far shown that it is consistent with the small amount of experimental data available addressing the development of orientation maps in V1. However, in order to contribute to our future understanding of V1 development, it will be necessary to confirm the predictions made by this model experimentally i.e. that both responses, and the balance between afferent and lateral interactions in V1, are

maintained during development despite changes in the statistics or strength of visual inputs. A program of future work to establish how the model would react to changes in key parameters could provide concrete predictions for how the real system would react under experimental manipulations.

Further homeostatic mechanisms

Although the ALISSOM model has introduced the higher level requirements for homeostasis in the developing orientation map, it is clear that underlying these requirements are many complex interactions between many different types of homeostatic regulation. For instance, it will be interesting in future to understand more clearly how adding a detailed mechanism for synaptic weight normalization (such as synaptic scaling) would affect the developing network and interact with the homeostatic regulation of responses. As discussed in section 2.4, there may also be regulatory processes earlier in the visual pathway which contribute to the stability of the network, such as gain control in the retina or LGN.

Applying ALISSOM to other systems

The main computational benefit of a simpler model such as ALISSOM is its ability to be applied more readily to other research questions, both in the visual system and for the investigation of topographic maps in other areas. Making the model robust to changes in the type of inputs allows many more features of topographic organization to be investigated. Although in this thesis I have mainly considered the map of orientation preference, there are also many other features mapped in V1, some of which have yet to be fully experimentally investigated i.e. direction, spatial frequency, color, disparity and speed. Both the forms of these maps, and their interactions could now be investigated in a combined model based on the simpler ALISSOM starting point which has far fewer parameters than LISSOM while being more biologically realistic.

This ability to use many different types of inputs to ALISSOM without making many parameter changes also makes it much easier to begin to use much more realistic inputs to the model retina. For example, input from a simulated virtual reality environment has been used in the ALISSOM model to develop an orientation map in V1 (preliminary work by James Adwick). Extension of this work to using real world video inputs

to drive development should also be possible in the near future.

The ALISSOM model is also an important first step towards building a visual system model which includes higher visual areas. Regulating the responses of V1 neurons allows for much more consistent signal propagation in this hierarchical network (as described in figure 2.1). For example, it would now be possible to add a V2 layer to this model that is receiving consistent inputs from V1 during development. Using this model we could make predictions about the type of features which V2 neurons will become selective for during development. Since most neurons in layer 2/3 of V1 from which V2 receives connections are phase invariant, this would also require modeling complex cells in V1; such a project has been carried out by Jan Antolik in parallel with the work in this thesis.

As well as using ALISSOM to investigate other visual areas, it is also possible to use this simpler model to investigate the emergent properties that result from self-organization of other cortical areas. For example, models of other primary sensory areas such as the auditory or somato-sensory cortex could use the ALISSOM algorithm, provided the inputs to the model could be well defined. In fact, work carried out with other colleagues in the lab and at Sheffield University (Stuart Wilson, James Bednar, Ben Mitchison, and Tony Prescott) reproducing the self-organization of direction maps in rat somato-sensory (or ‘barrel’ cortex) has shown that a specific correlational structure of the inputs to this system is crucial in order to reproduce the observed behavior. Therefore, applying the self-organizing algorithm (using the simpler ALISSOM model) to other cortical areas could make testable predictions about how these systems will behave.

6.3 A model of orientation map development consistent with the known wiring of V1

As discussed in chapter 3, experimental, theoretical and computational studies have shown that there is a consensus that the longest range connections between V1 neurons are excitatory rather than inhibitory (Gilbert and Wiesel, 1989; Hirsch and Gilbert, 1991; Weliky et al., 1995). Moreover, long-range excitatory connections are thought to account for a wide range of experimental data from adult V1, such as surround modulation or contrast dependent receptive field size (Angelucci and Bullier, 2003; Angelucci

et al., 2002a,b). However, previous models of orientation map development have been based on long-range inhibitory connections and short-range excitatory connections, which is at odds with the general consensus.

In order to address this issue I aimed in chapter 3 to implement a model that is consistent with the known wiring of primary visual cortex, yet also reproduces the features of successful developmental models of topographic map formation. The LESI model is a successful demonstration of this aim, with short-range inhibitory connections and long-range excitatory connections between neurons in a model V1 which develops a realistic orientation map.

Although the LESI model has been successful in this respect, and demonstrates that the approach in the later chapters (using ALISSOM) is valid, there are a number of immediate extensions and improvements which would make the LESI model a more practical model for use in future research. These are discussed in the next section.

6.3.1 Immediate future aims

Reconciling with homeostasis

The first things to consider in the LESI model would be those improvements that are related to the homeostatic mechanisms in the ALISSOM model. As discussed in section 6.2, simplification of the ALISSOM model could lead to a model in which no parameter tuning is required. However, in the LESI model there is an added layer of complexity arising from the separation of excitatory and inhibitory neurons. In fact, the LESI model introduces at least four more free parameters which require tuning in order to reproduce successful map development.

Additionally, as illustrated in figure 6.2, experiments suggest that levels of inhibition may control levels of excitation in V1. Increased levels of activity in excitatory cells could be reduced by reducing levels of excitation but also by increasing levels of inhibition. Similarly, reducing excitatory activity could be compensated by increasing excitation and decreasing inhibition. Therefore it is possible that responses of inhibitory neurons is regulated directly by levels of excitatory neural activity in V1. It may be possible to ‘couple’ the homeostatic plasticity of these two populations in order to further simplify the LESI model.




Figure 6.2: Coupled excitatory and inhibitory homeostatic regulation. Activity in recurrent cortical networks is strongly affected by feedback excitation and inhibition. Pyramidal neurons (blue) make excitatory outputs (triangles) onto other pyramidal neurons, and also onto inhibitory interneurons (red). These inhibitory neurons in turn feed inhibition (red circles) back onto the pyramidal neurons. In cortical cultures, raising activity for two days produces a coordinated set of changes in synaptic strength that result in reduced feedback excitation and increased feedback inhibition onto pyramidal neurons (lower left). Conversely, blocking activity for two days increases the gain of excitatory feedback and decreases inhibitory feedback. Similar changes in the cortical excitation-inhibition balance are induced by sensory deprivation. Reprinted from Turrigiano and Nelson (2004).

Contrast gain control and surround modulation

As discussed in section 6.2, homeostatic plasticity is also related to contrast behavior. In order to make the LESI model a practical model for investigating functional behaviours such as surround modulation, it is crucial that the contrast behaviour of the model is consistent with the experimental data. Adding separate inhibitory and excitatory populations as in LESI is an extra complication, since it has been shown experimentally that excitatory and inhibitory neurons differ in their response to con-

trast. In fact, computational models suggest that it is this difference which underlies surround-modulation effects such as the contrast dependence of receptive field size. The recent model of Schwabe et al. (2006) showed that if inhibitory neurons have a higher threshold and gain than excitatory neurons, then this could (in part) lead to the observed contrast-dependent size tuning.

Results from surround modulation experiments also show some significant species differences. For example, surround suppression may be stronger in monkey than in cat (Jones et al., 2001), which could feasibly be related to the different extents of lateral connections in these species. Therefore, in order to fully reproduce results from surround modulation experiments in LESI, it will also be necessary to more closely tie the LESI model to a specific species. It may be possible to combine the parameterized ALISSOM cat model presented in section 4 with the LESI model in order to specifically represent the extents of connections in the visual pathway. This closely tied model could then be compared directly with recent surround modulation results in the cat (Wang et al., 2009).

Given all of these changes, it will then be possible to use the LESI model to investigate the relationship between orientation map organization and functional properties arising from specific lateral connections in V1. As discussed in section 1.5, the relationship between orientation maps and modulatory effects is an important consideration in answering the question ‘what are the functional consequences of topographic organization?’ There are also some additional mechanisms that should be taken into account when considering this question in V1, and these are described in the next section.

6.3.2 Future outlook

Dependence of surround modulation on map position

A full investigation of surround modulation in V1 should also take into account that effects may be arising not just from lateral connectivity in V1, but also from connections from other visual areas. For example, surround modulation effects in V1 may arise purely from feed-forward connections from the retina via the LGN (Alitto and Usrey, 2008; Naito et al., 2007; Solomon et al., 2002). Further complication also arises from the reciprocal feedback between V1 and the LGN (Angelucci and Sainsbury, 2006; Webb et al., 2002). Feed-back connections from higher visual areas (such as V2, V3,

MT etc.) have also been suggested to account for surround modulation in V1 (Angelucci and Bressloff, 2006; Angelucci and Bullier, 2003).

This is obviously a complex system, which the LESI model is not yet well equipped to deal with. The added complexity of considering development of the orientation map means that including all of these connections is not computationally viable at present. However, it should be possible to use the fully developed LESI model as a starting point. The fully developed LESI V1 sheet contains realistic lateral connectivity embedded within an orientation map structure. Feed-forward and feed-back connections could then be added to this adult V1 model (without considering development) and the functional consequences of each connection type could be investigated. This type of model would also make it much easier to break the complex system down into its constituent parts, and force the modeler to consider all possible mechanisms contributing to the observed behaviors in V1.

6.4 Mechanisms underlying topographic organization of mouse V1

As discussed in detail in chapter 4, two-photon calcium imaging of networks in the visual cortex has confirmed that in rodents such as rats and mice, neurons selective for orientation are organised randomly in a “salt and pepper” arrangement, unlike the smooth topographic organization found in higher mammals.

Finding such dramatic species differences is exciting, because the differences may allow us to determine whether and how topographic maps are important for cortical function. Unfortunately, directly comparing experimental findings across species is difficult, due to the many unconstrained variables that might account for the observed differences. Computational modelling allows specific hypotheses to be tested rigorously, with complete control, to establish what is and is not crucial to reproduce a set of experimental results. The aim of chapters 4 and 5 was to address the difference in functional organization between rodents (particularly the mouse) and higher mammals (particularly the cat) using computational models specifically parameterized to represent the mouse and cat visual pathway. The models presented in these chapters have been very successful in highlighting the possible differences between these species that could lead to the lack of orientation maps in mouse, and generated testable predictions

for future experiments, thus achieving the aims set out in section 1.7.

Above all, these chapters have demonstrated that there are several possible mechanisms or combinations of mechanisms that may underlie the differences in topographic organization. This work therefore forms the starting point for a large program of future research as described in the next two sections.

6.4.1 Immediate future aims

The first step, on which I have already made significant progress, is highlighting the likely candidates for developmental mechanisms that could underlie the observed differences in topographic organization. There are now two avenues for immediate future work. The first is a program of experimental investigation, both to improve the currently available data and to collect more data. The second is a series of improvements to the current models to produce more tightly constrained comparisons with the two-photon data.

6.4.1.1 Experimental work

Determining whether or not organization is activity-driven

Correlational structure of the preferences of mouse V1 neurons may provide some clues as to whether or not receptive field development is activity-driven. In order to highlight any correlations there may be between, for example, retinotopic scatter and orientation preference or selectivity, it would be necessary to measure orientation preference and retinotopy in the same cortical area in the same animal. The model presented in chapter 5 might suggest that if orientation preference arises from random wiring of the LGN to V1 that there would be no correlation between orientation preference and scatter. However, if orientation selectivity develops from an initially un-selective receptive field, neurons with neighbors which prefer similar orientations may also correspond to areas of low scatter between receptive field positions.

Sparsity appears to have some role in the development of ‘salt and pepper’ orientation maps, whether it is sparsity in connectivity, or sparsity of activity patterns. In order to distinguish these possibilities it will be interesting to determine how receptive fields and activity patterns in the mouse V1 neurons change during development. It may

also be possible to make genetic or chemical manipulations in the mouse, to determine whether lateral connections play an important role in development. We could, for instance, speculate that if maps in the mouse are driven by activity, that adaptation of lateral interactions (such as increased levels of inhibition) could decrease the sparsity of activity patterns and lead to more clustered orientation preferences in mouse.

Increasing the volume of experimental data

Firstly, as discussed in section 4.2, the dataset on which the analysis was performed suffers from a bias in the sampling of the visual field. It is therefore necessary both to improve the experimental paradigm and to increase the number of animals in the dataset to ensure that the results are statistically significant. It would also be extremely interesting in these experiments to dissociate the effects of feed-forward and lateral contributions to the retinotopic scatter. This would allow us to distinguish whether the retinotopic scatter is truly embedded within the feed-forward connections to V1.

It would also be useful to observe directly the form and extent of lateral connections in the mouse visual cortex, for example by using a transsynaptic tracer based on the rabies virus (as used in Wickersham et al. 2007 for the rat). It should also be noted that knowing the extents of lateral connections does not necessarily tell us anything about the extent or nature of the interactions between neurons, since the extent of overall interactions are not necessarily defined by the extent of connections (as is demonstrated by chapter 3).

The ideal experimental data would be two-photon imaging of all maps (retinotopy, orientation and possibly spatial frequency) in the same animal, throughout development, in both mouse and cat. However, since achieving this data set is unlikely, a more realistic option is to measure maps (in the same animal) at a series of developmental stages in the mouse and compare these to the model. In this way we can produce a mouse model which is tightly constrained by the experimental data and then make predictions about how manipulations may affect development in both mouse and cat.

6.4.1.2 Computational work

Predicting the outcome of genetic manipulations

Recent experimental work has shown that the nasotemporal (azimuth) axis of the retinotopic map in mouse V1 is distorted in mice deficient in both molecular cues (ephrin-As) and spontaneous activity (by disruption of retinal waves using $\beta 2$ knock-out; Cang et al. 2008). Using a correlate for this type of genetic manipulation in the model, we could predict what might happen to the mapping of orientation preference (either in the sparse wiring model or during activity-driven development of orientation selectivity). In the knockout mice it would appear that both LGN receptive fields and V1 receptive fields are distorted, with LGN receptive fields abnormally scattered along the azimuthal axis and V1 receptive fields enlarged in the azimuth direction. We might expect that in the sparse wiring model (section 5.3.3) an increased amount of LGN receptive field scatter and arbourization scatter in the azimuth direction would indeed result in a corresponding increase in V1 receptive field sizes (as discussed in Cang et al. 2008, for the case of a non-sparse sampling of LGN cells). We may be able to predict from the ALISSOM mouse model how this anisotropy in scatter would affect orientation tuning and then measure whether this prediction is consistent with the experimental data in ephrin-A/ $\beta 2$ knockout mice.

It is less clear how to reconcile the model of activity-driven salt and pepper map development with this experimental data, but we could test this model by ‘building-in’ the type of scatter observed in the knockout mouse, i.e., increasing the scatter in LGN receptive fields and manually elongating V1 receptive fields in the azimuth direction. The original model with ‘built-in’ scatter (section 4.4.1) assumes that V1 cells sample from overlapping LGN ON and OFF cells. In this case, however, overlap is unlikely because pairs of ON and OFF cells are unlikely to be scattered in exactly the same way in the knockout. If connections between the LGN and V1 are not sparse, this type of increase in scatter may lead to an initial clustering of orientation preferences as observed in the models presented in section 5.3.2.

Improvements to the model

As well as further work to test predictions of the current models and to improve the model parameterization, there are some improvements in the model architecture which

may make it easier to closely tie the experimental and model results. In the models presented in chapter 4 and 5, I have considered only one slice through the two-photon imaging volume. Extending the models to the full 3D volume would allow modelling of the full connectivity between neurons in all directions, and may be functionally important. Since actual cortical location of neurons in the 3D volume is important for the exact measurement of retinotopic scatter, another improvement of the model would be to give each neuron a cortical location which is not constrained to lie within a grid. Another extension of the model would be to include more dynamic connection patterns. As discussed in section 5.4, when sparse connections are included in the model these are fixed (although the weights of connections can change). In reality connections are likely to be both forming and retracting, and including this in the model may also be an important consideration for fully investigating the course of development.

Improvements could also be made to the type of stimuli used to drive development in these models. The current models use elongated Gaussian inputs, however it would be interesting to use realistic natural images with longer range correlations that would be reflected in the lateral connections during development. Improving the ALISSOM model with the changes mentioned above in section 6.2, i.e., implementing contrast gain control and eliminating the frequency parameter, would also make changing the type of input pattern much easier.

6.4.2 Future outlook

Understanding the functional consequences of map organization

Provided the future work mentioned above results in a well defined model for mouse visual cortical development, this model could be extremely useful in understanding the functional consequences of a lack of topographic mapping. As described in section 1.5, the connectivity between neurons in V1 is highly dependent on their particular topographic organization. In the modelling work thus far, it appears that connectivity patterns in mouse visual cortex differ substantially from higher mammals, and that similar developmental paradigms can lead to sparse connections in the mouse and patchy connections in the cat. These different connectivity patterns may have a measurable effect on properties such as surround modulation or orientation tuning. It may there-

fore be possible to predict which types of behaviourally relevant visual functions are not present in rodents.

As discussed in section 4.6, it would also be interesting to determine whether, given the architecture of the mouse visual system, ‘salt and pepper’ organization for orientation preference in the mouse is optimal for coverage of the visual field by all orientations.

Knowing the specific ways in which rodents and higher mammals are different is crucial if we are to extrapolate results from mouse experiments to higher mammals. Tightly constrained models of different species are extremely useful in this case because we could use these to predict how similar manipulations in higher mammals might affect function or behaviour without the need to perform experiments. This would increase the impact of model driven, specifically targeted experiments in higher mammals.

6.5 General points

A general issue with modelling of topographic organization is that it is currently difficult to quantitatively evaluate the results. For example, the large variation in map (or receptive field) structure, both across and within species makes it difficult to use some sort of distance measure between model and experimental maps (or receptive fields). Some measures have been suggested, such as the ring-shaped fourier transform (Erwin et al., 1995) (often used to qualitatively evaluate maps in this thesis) and it may be possible to create a more quantitative measure based on this. This would also involve evaluating the fourier transform of a large set of animal maps in order to get a quantitative measure of the experimental variability in this measure. A measure based on pinwheel density has also recently been suggested (Kaschube et al., 2008), however this is again likely to require a large amount of future analysis of experimentally measured maps.

It is often easier to evaluate the success of a model when looking at the relationship between maps of multiple features. For example, joint models of orientation maps and ocular dominance columns can be evaluated based on whether or not orientation pinwheels fall preferentially away from ocular dominance boundaries (Erwin et al., 1995). The LISSOM model, on which the models in this thesis are based, has been shown to reproduce many of these observed relationships (Miikkulainen et al., 2005).

Again however, these measures are often qualitative rather than quantitative.

For now, methods of estimating the goodness of a model or parameter settings in map models are mainly qualitative, however this is an excellent topic for future research. Advances in technology will also make more exhaustive parameter searches in the large computationally intensive models presented in this thesis possible. Therefore, automating this process using a measure of model goodness of fit will be extremely useful. Because of the current inability to exhaustively search parameters, the models presented here are primarily feasibility studies of the proposed hypotheses but it is hoped that in future these can be extended to give a more complete understanding of the parameter space.

6.6 Conclusion

Each chapter of this thesis has achieved the initial aims as set out in section 1.7 and has made a significant and original contribution to scientific research. However, as this chapter has described, there are many possible extensions and improvements which can be made in the future. Overall, the work has covered a breadth of topics for which a lot of the depth remains to be explored. This is an exciting consequence since the future research avenues that this work has uncovered should make significant contributions to our understanding of V1 and the computational function of the mammalian cortex in general.

Chapter 7

Conclusions

The predominant driving force behind this research is to understand the brain from a computational perspective, using detailed biological data to constrain and inspire computational models. Rather than working at a low level in order to fully describe the behaviour of individual neurons, or at a high level, abstracting away the cortical circuitry and modelling only aspects of cortical behaviour, I have used the simplest possible models required to realistically explain the observed phenomena that I am interested in.

Chapter 2 highlights the homeostatic requirements for stable development of the orientation map in V1 of higher mammals, and also provides a much simpler and more robust starting point for future models of topographic map development with the ALIS-SOM model.

Chapter 3 shows that orientation map development can be achieved in a model (LESI) that is consistent with the experimentally observed cortical circuitry. This work suggests that previous simpler models of map development (e.g. SOM, LISSOM) are justified in abstracting away from the more complex cortical architecture in order to highlight the required interactions (rather than connections) between neurons. However, the more complex LESI model is a crucial starting point for investigating properties of V1 neurons which depend on both the connectivity and the topographic organization of V1.

Chapters 4 and 5 have established the crucial parameters that could determine the differences between topographic organization in rodents and higher mammals. A program of future experimental and computational work based on the predictions made in

this thesis can begin to establish the functional consequences of species differences in topographic organization.

The original contribution of this thesis consists of three feasibility studies. The ALIS-SOM model is the first to show that it is possible to self-organize the orientation map in V1 in a stable way using biologically realistic homeostatic mechanisms. The LESI model is the first to show that it is possible to self-organize the orientation map in V1 using an architecture which contains no inhibitory connections which are longer range than excitatory connections. Chapters 4 and 5 are the first mechanistic modelling approach to understanding the underlying mechanisms of organization for orientation in the rodent visual pathway.

The work presented here has therefore made significant progress in understanding how the experimentally observed cortical circuitry can lead to the functional properties and topographic organization of individual neurons.

Bibliography

- Abraham, W. (2008). Metaplasticity: tuning synapses and networks for plasticity. *Nat. Rev. Neurosci.*, 9:387.
- Albrecht, D. G. and Geisler, W. S. (1991). Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Vis Neurosci*, 7(6):531–546.
- Albus, K. (1975). A quantitative study of the projection area of the central and the paracentral visual field in area 17 of the cat. I. The precision of the topography. *Exp Brain Res*, 24:159–179.
- Alitto, H. J. and Usrey, W. M. (2008). Origin and dynamics of extraclassical suppression in the lateral geniculate nucleus of the macaque monkey. *Neuron*, 57(1):135–146.
- Alonso, J. and Swadlow, H. (2005). Thalamocortical specificity and the synthesis of sensory cortical receptive fields. *J. Neurophysiol.*, 94:26–32.
- Alonso, J. M., Usrey, W. M., and Reid, R. C. (2001). Rules of connectivity between geniculate cells and simple cells in cat primary visual cortex. *J Neurosci*, 21(11):4002–4015.
- Anderson, J. S., Carandini, M., and Ferster, D. (2000). Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *J Neurophysiol*, 84(2):909–926.
- Angelucci, A. and Bressloff, P. C. (2006). Contribution of feedforward, lateral and feedback connections to the classical receptive field center and extra-classical receptive field surround of primate v1 neurons. *Prog Brain Res*, 154:93–120.
- Angelucci, A. and Bullier, J. (2003). Reaching beyond the classical receptive field of v1 neurons: horizontal or feedback axons? *J Physiol Paris*, 97(2-3):141–154.

- Angelucci, A., Levitt, J. B., and Lund, J. S. (2002a). Anatomical origins of the classical receptive field and modulatory surround field of single neurons in macaque visual cortical area V1. *Progress in Brain Research*, 136:373–388.
- Angelucci, A., Levitt, J. B., Walton, E. J., Hupe, J. M., Bullier, J., and Lund, J. S. (2002b). Circuits for local and global signal integration in primary visual cortex. *Journal of Neuroscience*, 22:8633–8646.
- Angelucci, A. and Sainsbury, K. (2006). Contribution of feedforward thalamic afferents and corticogeniculate feedback to the spatial summation area of macaque v1 and lgn. *J Comp Neurol*, 498(3):330–351.
- Antonini, A., Fagiolini, M., and Stryker, M. (1999). Anatomical correlates of functional plasticity in mouse visual cortex. *J. Neurosci.*, 19:4388–4406.
- Bé, J.-V. L. and Markram, H. (2006). Spontaneous and evoked synaptic rewiring in the neonatal neocortex. *Proc Natl Acad Sci U S A*, 103(35):13214–13219.
- Beckmann, R. and Albus, K. (1982). The geniculocortical system in the early postnatal kitten: an electrophysiological investigation. *Exp. Brain Res.*, 47(1):49–56.
- Bednar, J. A., Choe, Y., De Paula, J., Miikkulainen, R., Provost, J., and Tversky, T. (2004). Modeling cortical maps with Topographica. *Neurocomputing*, 58–60:1129–1135.
- Ben-Yishai, R., Bar-Or, R. L., and Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proc Natl Acad Sci U S A*, 92(9):3844–3848.
- Bienenstock, E., Cooper, L., and Munro, P. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J. Neurosci.*, 2:32–48.
- Binzegger, T., Douglas, R. J., and Martin, K. A. C. (2004). A quantitative map of the circuit of cat primary visual cortex. *J Neurosci*, 24(39):8441–8453.
- Blakemore, C. and Cooper, G. F. (1970). Development of the brain depends on the visual environment. *Nature*, 228(5270):477–478.
- Blasdel, G. G. (1992). Orientation selectivity, preference, and continuity in monkey striate cortex. *Journal of Neuroscience*, 12:3139–3161.

- Blasdel, G. G. and Salama, G. (1986). Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321(6070):579–585.
- Bonhoeffer, T. and Grinvald, A. (1993). The layout of iso-orientation domains in area 18 of cat visual cortex: optical imaging reveals a pinwheel-like organization. *J Neurosci*, 13(10):4157–4180.
- Bosking, W. H., Zhang, Y., Schofield, B., and Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *J Neurosci*, 17(6):2112–2127.
- Budd, J. M. and Kisvarday, Z. F. (2001). Local lateral connectivity of inhibitory clutch cells in layer 4 of cat visual cortex (area 17). *Exp Brain Res*, 140(2):245–250.
- Bullier, J. and Norton, T. (1979). X and Y relay cells in cat lateral geniculate nucleus: quantitative analysis of receptive-field properties and classification. *J. Neurophysiol.*, 42:244–273.
- Burger, T. and Lang, E. W. (1999). An incremental Hebbian learning model of the primary visual cortex with lateral plasticity and real input patterns. *Zeitschrift für Naturforschung C — A Journal of Biosciences*, 54:128–140.
- Burger, T. and Lang, E. W. (2001). Self-organization of local cortical circuits and cortical orientation maps: a nonlinear hebbian model of the visual cortex with adaptive lateral couplings. *Z Naturforsch [C]*, 56(5-6):464–478.
- Burkhalter, A. (1989). Intrinsic connections of rat primary visual cortex: laminar organization of axonal projections. *J Comp Neurol*, 279(2):171–186.
- Butko, N. and Triesch, J. (2007). Exploring the role of intrinsic plasticity for the learning of sensory representations. *Neurocomputing*, 70(7-9):1130–1138.
- Buzas, P., Eysel, U. T., Adorjan, P., and Kisvarday, Z. F. (2001). Axonal topography of cortical basket cells in relation to orientation, direction, and ocular dominance maps. *J Comp Neurol*, 437(3):259–285.
- Buzás, P., Kovcs, K., Ferecsk, A., Budd, J., Eysel, U., and Kisvrday, Z. (2006). Model-based analysis of excitatory lateral connections in the visual cortex. *J. Comp. Neurol.*, 499:861–881.
- Buzás, P., Volgushev, M., Eysel, U., and Kisvrday, Z. (2003). Independence of vi-

- suotopic representation and orientation map in the visual cortex of the cat. *Eur. J. Neurosci.*, 18:957–968.
- Cai, D., DeAngelis, G., and Freeman, R. (1997). Spatiotemporal receptive field organization in the lateral geniculate nucleus of cats and kittens. *J. Neurophysiol.*, 78:1045–1061.
- Cang, J., Niell, C. M., Liu, X., Pfeiffenberger, C., Feldheim, D. A., and Stryker, M. P. (2008). Selective disruption of one cartesian axis of cortical maps and receptive fields by deficiency in ephrin-as and structured activity. *Neuron*, 57(4):511–523.
- Carandini, M. and Heeger, D. J. (1994). Summation and division by neurons in primate visual cortex. *Science*, 264(5163):1333–1336.
- Carandini, M., Heeger, D. J., and Movshon, J. A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci*, 17(21):8621–8644. Comparative Study.
- Carandini, M., Heeger, D. J., and Senn, W. (2002). A synaptic explanation of suppression in visual cortex. *J Neurosci*, 22(22):10053–10065.
- Carandini, M. and Ringach, D. L. (1997). Predictions of a recurrent model of orientation selectivity. *Vision Res*, 37(21):3061–3071.
- Carreira-Perpin, M. A., Lister, R. J., and Goodhill, G. J. (2005). A computational model for the development of multiple maps in primary visual cortex. *Cereb Cortex*, 15(8):1222–1233.
- Chalupa, L. M. and Williams, R. W., editors (2008). *Eye, retina, and visual system of the mouse*. Cambridge, Mass: MIT Press.
- Chandrasekaran, A., Shah, R., and Crair, M. (2007). Developmental homeostasis of mouse retinocollicular synapses. *J. Neurosci.*, 27:1746–1755.
- Chapman, B. and Bonhoeffer, T. (1998). Overrepresentation of horizontal and vertical orientation preferences in developing ferret area 17. *Proc Natl Acad Sci U S A*, 95(5):2609–2614.
- Chapman, B., Stryker, M. P., and Bonhoeffer, T. (1996). Development of orientation preference maps in ferret primary visual cortex. *J Neurosci*, 16(20):6443–6453.

- Cragg, B. (1975). The development of synapses in the visual system of the cat. *J. Comp. Neurol.*, 160:147–166.
- Crair, M., Gillespie, D., and Stryker, M. (1998). The role of visual experience in the development of columns in cat visual cortex. *Science*, 279:566–570.
- Crair, M. C., Ruthazer, E. S., Gillespie, D. C., and Stryker, M. P. (1997). Ocular dominance peaks at pinwheel center singularities of the orientation map in cat visual cortex. *J Neurophysiol*, 77(6):3381–3385.
- Dakin, S., Carlin, P., and Hemsley, D. (2005). Weak suppression of visual context in chronic schizophrenia. *Curr Biol*, 15(20):822–824. Letter.
- Daoudal, G. and Debanne, D. (2003). Long-term plasticity of intrinsic excitability: learning rules and mechanisms. *Learn. Mem.*, 10(6):456–465.
- Das, A. and Gilbert, C. D. (1997). Distortions of visuotopic map match orientation singularities in primary visual cortex. *Nature*, 387(6633):594–598.
- Das, A. and Gilbert, C. D. (1999). Topography of contextual modulations mediated by short-range interactions in primary visual cortex. *Nature*, 399(6737):655–661.
- Daugman, J. G. (1980). Two-dimensional spectral analysis of cortical receptive field profiles. *Vision Res*, 20(10):847–856.
- Davis, G. and Bezprozvanny, I. (2001). Maintaining the stability of neural function: a homeostatic hypothesis. *Annu. Rev. Physiol.*, 63:847–869.
- Desai, N. (2003). Homeostatic plasticity in the CNS: synaptic and intrinsic forms. *J. Physiol. Paris*, 97:391–402.
- Diamond, M. E., von Heimendahl, M., Knutsen, P. M., Kleinfeld, D., and Ahissar, E. (2008). 'where' and 'what' in the whisker sensorimotor system. *Nat Rev Neurosci*, 9(8):601–612.
- Dougherty, R. F., Koch, V. M., Brewer, A. A., Fischer, B., Modersitzki, J., and Wandell, B. A. (2003). Visual field representations and locations of visual areas v1/2/3 in human visual cortex. *J Vis*, 3(10):586–598.
- Douglas, R. J. and Martin, K. A. C. (2007). Recurrent neuronal circuits in the neocortex. *Curr Biol*, 17(13):R496–R500.

- Dräger, U. (1975). Receptive fields of single cells and topography in mouse visual cortex. *J. Comp. Neurol.*, 160:269–290.
- Dragoi, V., Rivadulla, C., and Sur, M. (2001). Foci of orientation plasticity in visual cortex. *Nature*, 411(6833):80–86.
- Dragoi, V. and Sur, M. (2000). Dynamic properties of recurrent inhibition in primary visual cortex: contrast and orientation dependence of contextual effects. *J Neurophysiol*, 83(2):1019–1030.
- Durbin, R. and Willshaw, D. (1987). An analogue approach to the travelling salesman problem using an elastic net method. *Nature*, 326(6114):689–691.
- Engelmann, R., Crook, J. M., and Lwel, S. (2002). Optical imaging of orientation and ocular dominance maps in area 17 of cats with convergent strabismus. *Vis Neurosci*, 19(1):39–49.
- Ernst, U. A., Pawelzik, K. R., Sahar-Pikielny, C., and Tsodyks, M. V. (2001). Intracortical origin of visual maps. *Nat Neurosci*, 4(4):431–436.
- Erwin, E. and Miller, K. (1998). Correlation-based development of ocularly matched orientation and ocular dominance maps: determination of required input activities. *J. Neurosci.*, 18:9870–9895.
- Erwin, E., Obermayer, K., and Schulten, K. (1995). Models of orientation and ocular dominance columns in the visual cortex: a critical comparison. *Neural Comput*, 7(3):425–468.
- Felsen, G., Shen, Y.-s., Yao, H., Spor, G., Li, C., and Dan, Y. (2002). Dynamic modification of cortical orientation tuning mediated by recurrent connections. *Neuron*, 36(5):945–954.
- Finn, I. M., Priebe, N. J., and Ferster, D. (2007). The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron*, 54(1):137–152.
- Fitzpatrick, D. (2000). Seeing beyond the receptive field in primary visual cortex. *Curr Opin Neurobiol*, 10(4):438–443.
- Geisler, W. S. and Albrecht, D. G. (1997). Visual cortex neurons in monkeys and cats: detection, discrimination, and identification. *Vis Neurosci*, 14(5):897–919.

- Gilbert, C. and Wiesel, T. (1983). Clustered intrinsic connections in cat visual cortex. *J. Neurosci.*, 3:1116–1133.
- Gilbert, C. D. and Wiesel, T. N. (1989). Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *J Neurosci*, 9(7):2432–2442.
- Godecke, I. and Bonhoeffer, T. (1996). Development of identical orientation maps for two eyes without common visual experience. *Nature*, 379:251–254.
- Godecke, I., Kim, D. S., Bonhoeffer, T., and Singer, W. (1997). Development of orientation preference maps in area 18 of kitten visual cortex. *Eur J Neurosci*, 9(8):1754–1762.
- Goodhill, G. J. (2007). Contributions of theoretical modeling to the understanding of neural map development. *Neuron*, 56(2):301–311.
- Goodhill, G. J. and Cimoneriu, A. (2000). Analysis of the elastic net model applied to the formation of ocular dominance and orientation columns. *Network: Computation in Neural Systems*, 11:153–168.
- Grabska-Barwinska, A. and von der Malsburg, C. (2008). Establishment of a scaffold for orientation maps in primary visual cortex of higher mammals. *J. Neurosci.*, 28:249–257.
- Grinvald, A. and Hildesheim, R. (2004). VSDI: a new era in functional imaging of cortical dynamics. *Nat Rev Neurosci*, 5(11):874–885.
- Grossberg, S. and Olson, S. J. (1994). Rules for the cortical map of ocular dominance and orientation columns. *Neural Networks*, 7:883–894.
- Grubb, M., Rossi, F., Changeux, J., and Thompson, I. (2003). Abnormal functional organization in the dorsal lateral geniculate nucleus of mice lacking the beta 2 subunit of the nicotinic acetylcholine receptor. *Neuron*, 40:1161–1172.
- Gupta, A., Wang, Y., and Markram, H. (2000). Organizing principles for a diversity of GABAergic interneurons and synapses in the neocortex. *Science*, 287(5451):273–278.
- Gbel, W. and Helmchen, F. (2007). In vivo calcium imaging of neural network function. *Physiology (Bethesda)*, 22:358–365.

- Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Vis Neurosci*, 9(2):181–197.
- Hellwig, B. (2000). A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biol Cybern*, 82(2):111–121.
- Hensch, T. K. and Stryker, M. P. (2004). Columnar architecture sculpted by gaba circuits in developing cat visual cortex. *Science*, 303(5664):1678–1681.
- Hirsch, J. A. and Gilbert, C. D. (1991). Synaptic physiology of horizontal connections in the cat's visual cortex. *J Neurosci*, 11(6):1800–1809.
- Holmgren, C., Harkany, T., Svennenfors, B., and Zilberter, Y. (2003). Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *J Physiol*, 551(Pt 1):139–153.
- Hooks, B. M. and Chen, C. (2006). Distinct roles for spontaneous and visual activity in remodeling of the retinogeniculate synapse. *Neuron*, 52(2):281–291.
- Hubel, D. H. and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol*, 160:106–154.
- Hubel, D. H., Wiesel, T. N., and LeVay, S. (1977). Plasticity of ocular dominance columns in monkey striate cortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 278:377–409.
- Hübener, M. (2003). Mouse visual cortex. *Curr. Opin. Neurobiol.*, 13:413–420.
- Huberman, A. D., Feller, M. B., and Chapman, B. (2008). Mechanisms underlying development of visual maps and receptive fields. *Annu Rev Neurosci*, 31:479–509.
- Huberman, A. D., Speer, C. M., and Chapman, B. (2006). Spontaneous retinal activity mediates development of ocular dominance columns and binocular receptive fields in V1. *Neuron*, 52(2):247–254.
- Husson, T. R., Mallik, A. K., Zhang, J. X., and Issa, N. P. (2007). Functional imaging of primary visual cortex using flavoprotein autofluorescence. *J Neurosci*, 27(32):8665–8675.
- Jones, H. E., Grieve, K. L., Wang, W., and Sillito, A. M. (2001). Surround suppression in primate V1. *J Neurophysiol*, 86(4):2011–2028.

- Jones, J. P. and Palmer, L. A. (1987). The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *J. Neurophysiol.*, 58(6):1187–1211.
- Kaas, J. H., Krubitzer, L. A., and Johanson, K. L. (1989). Cortical connections of areas 17 (V-I) and 18 (V-II) of squirrels. *J Comp Neurol*, 281(3):426–446.
- Kalatsky, V. A. and Stryker, M. P. (2003). New paradigm for optical imaging: temporally encoded maps of intrinsic signal. *Neuron*, 38(4):529–545.
- Kalisman, N., Silberberg, G., and Markram, H. (2005). The neocortical microcircuit as a tabula rasa. *Proc Natl Acad Sci U S A*, 102(3):880–885.
- Kaschube, M., Schnabel, M., and Wolf, F. (2008). Self-organization and the selection of pinwheel density in visual cortical development. *New Journal of Physics*, 10(1):015009 (20pp).
- Katz, L. C., Gilbert, C. D., and Wiesel, T. N. (1989). Local circuits and ocular dominance columns in monkey striate cortex. *J Neurosci*, 9(4):1389–1399.
- Kayser, A., Priebe, N. J., and Miller, K. D. (2001). Contrast-dependent nonlinearities arise locally in a model of contrast-invariant orientation tuning. *J Neurophysiol*, 85(5):2130–2149.
- Kisvarday, Z. F. and Eysel, U. T. (1993). Functional and structural topography of horizontal inhibitory connections in cat visual cortex. *Eur J Neurosci*, 5(12):1558–1572.
- Kisvarday, Z. F., Ferecsko, A. S., Kovacs, K., Buzas, P., Budd, J. M. L., and Eysel, U. T. (2002). One axon-multiple functions: specificity of lateral inhibitory connections by large basket cells. *J Neurocytol*, 31(3-5):255–264.
- Kisvarday, Z. F., Toth, E., Rausch, M., and Eysel, U. T. (1997). Orientation-specific relationship between populations of excitatory and inhibitory lateral connections in the visual cortex of the cat. *Cereb Cortex*, 7(7):605–618.
- Koulakov, A. A. and Chklovskii, D. B. (2001). Orientation preference patterns in mammalian visual cortex: A wire length minimization approach. *Neuron*, 29:519–527.
- Kritzer, M. F., Cowey, A., and Somogyi, P. (1992). Patterns of inter- and intralaminar gabaergic connections distinguish striate (V1) and extrastriate (V2, V4) visual

- cortices and their functionally specialized subdivisions in the rhesus monkey. *J Neurosci*, 12(11):4545–4564.
- Lauritzen, T. Z., Krukowski, A. E., and Miller, K. D. (2001). Local correlation-based circuitry can account for responses to multi-grating stimuli in a model of cat V1. *J Neurophysiol*, 86(4):1803–1815.
- LeVay, S., Connolly, M., Houde, J., and Essen, D. C. V. (1985). The complete pattern of ocular dominance stripes in the striate cortex and visual field of the macaque monkey. *J Neurosci*, 5(2):486–501.
- Linsker, R. (1986). From basic network principles to neural architecture: Emergence of spatial-opponent cells. *Proceedings of the National Academy of Sciences, USA*, 83:7508–7512.
- Lund, J. S. (1988). Anatomical organization of macaque monkey striate visual cortex. *Annu Rev Neurosci*, 11:253–288.
- Malach, R., Amir, Y., Harel, M., and Grinvald, A. (1993). Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo targeted biocytin injections in primate striate cortex. *Proc Natl Acad Sci U S A*, 90(22):10469–10473.
- Mangini, N. J. and Pearlman, A. L. (1980). Laminar distribution of receptive field properties in the primary visual cortex of the mouse. *J Comp Neurol*, 193(1):203–222.
- Marino, J., Schummers, J., Lyon, D. C., Schwabe, L., Beck, O., Wiesing, P., Obermayer, K., and Sur, M. (2005). Invariant computations in local cortical networks with balanced excitation and inhibition. *Nat Neurosci*, 8(2):194–201.
- Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nat Rev Neurosci*, 5(10):793–807.
- Martin, K. A. C. (2002). Microcircuits in visual cortex. *Curr Opin Neurobiol*, 12(4):418–425.
- Matsuzaki, M., Ellis-Davies, G. C. R., and Kasai, H. (2008). Three-dimensional mapping of unitary synaptic connections by two-photon macro photolysis of caged glutamate. *J Neurophysiol*, 99(3):1535–1544.

- Métin, C., Godement, P., and Imbert, M. (1988). The primary visual cortex in the mouse: receptive field properties and functional organization. *Exp Brain Res*, 69:594–612.
- Miikkulainen, R., Bednar, J., Choe, Y., and Sirosh, J. (2005). *Computational Maps in the Visual Cortex*. Springer.
- Miller, K. (1994). A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between ON- and OFF-center inputs. *J. Neurosci.*, 14:409–441.
- Miller, K. D., Erwin, E., and Kayser, A. (1999). Is the development of orientation selectivity instructed by activity? *J Neurobiol*, 41(1):44–57.
- Murphy, E. H. and Berman, N. (1979). The rabbit and the cat: a comparison of some features of response properties of single cells in the primary visual cortex. *J Comp Neurol*, 188(3):401–427.
- Naito, T., Sadakane, O., Okamoto, M., and Sato, H. (2007). Orientation tuning of surround suppression in lateral geniculate nucleus and primary visual cortex of cat. *Neuroscience*, 149(4):962–975.
- Nauhaus, I., Benucci, A., Carandini, M., and Ringach, D. L. (2008). Neuronal selectivity and local map structure in visual cortex. *Neuron*, 57(5):673–679.
- Niebur, E. and Wörgötter, F. (1993). Orientation columns from first principles classical visual receptive field. *Biological Cybernetics*, 70:1–13.
- Niell, C. and Stryker, M. (2008). Highly selective receptive fields in mouse visual cortex. *J. Neurosci.*, 28:7520–7536.
- Obermayer, K., Ritter, H., and Schulten, K. (1990). A principle for the formation of the spatial structure of cortical feature maps. *Proc Natl Acad Sci U S A*, 87(21):8345–8349.
- Ohki, K., Chung, S., Ch'ng, Y. H., Kara, P., and Reid, R. C. (2005). Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature*, 433(7026):597–603.
- Ohki, K., Chung, S., Kara, P., Hbener, M., Bonhoeffer, T., and Reid, R. C. (2006). Highly ordered arrangement of single neurons in orientation pinwheels. *Nature*, 442(7105):925–928.

- Ohki, K. and Reid, R. C. (2007). Specificity and randomness in the visual cortex. *Curr Opin Neurobiol*, 17(4):401–407.
- O’Leary, D. D. M. and McLaughlin, T. (2005). Mechanisms of retinotopic map development: Ephs, ephrins, and spontaneous correlated retinal activity. *Prog Brain Res*, 147:43–65.
- Pettigrew, J. D. (1974). The effect of visual experience on the development of stimulus specificity by kitten cortical neurones. *J Physiol*, 237(1):49–74.
- Raizada, R. D. S. and Grossberg, S. (2003). Towards a theory of the laminar architecture of cerebral cortex: computational clues from the visual system. *Cereb Cortex*, 13(1):100–113.
- Rao, S., Toth, L., and Sur, M. (1997). Optically imaged maps of orientation preference in primary visual cortex of cats and ferrets. *J. Comp. Neurol.*, 387:358–370.
- Reid, R. C. and Alonso, J. M. (1995). Specificity of monosynaptic connections from thalamus to visual cortex. *Nature*, 378(6554):281–284.
- Ren, M., Yoshimura, Y., Takada, N., Horibe, S., and Komatsu, Y. (2007). Specialized inhibitory synaptic actions between nearby neocortical pyramidal neurons. *Science*, 316(5825):758–761.
- Ringach, D. (2004). Haphazard wiring of simple receptive fields and orientation columns in visual cortex. *J. Neurophysiol.*, 92:468–476.
- Ringach, D. (2007). On the origin of the functional architecture of the cortex. *PLoS ONE*, 2:e251.
- Rochester, N., Holland, J. H., Haibt, L. H., and Duda, W. L. (1956). Tests on a cell assembly theory of the action of the brain, using a large digital computer. *IRE Transactions on Information Theory*, 2:80–93.
- Roerig, B. and Chen, B. (2002). Relationships of local inhibitory and excitatory circuits to orientation preference maps in ferret visual cortex. *Cereb Cortex*, 12(2):187–198.
- Roerig, B. and Kao, J. P. (1999). Organization of intracortical circuits in relation to direction preference maps in ferret visual cortex. *J Neurosci*, 19(24):RC44.

- Roger, A. S. and Schwartz, E. L. (1990). Cat and monkey cortical columnar patterns modeled by bandpass-filtered 2D white noise. *Biological Cybernetics*, 62:381–391.
- Rumberger, A., Tyler, C. J., and Lund, J. S. (2001). Intra- and inter-areal connections between the primary visual cortex V1 and the area immediately surrounding V1 in the rat. *Neuroscience*, 102(1):35–52.
- Sceniak, M. P., Ringach, D. L., Hawken, M. J., and Shapley, R. (1999). Contrast's effect on spatial summation by macaque V1 neurons. *Nat Neurosci*, 2(8):733–739.
- Schuett, S., Bonhoeffer, T., and Hubener, M. (2002). Mapping retinotopic structure in mouse visual cortex with optical imaging. *J Neurosci*, 22(15):6549–6559.
- Schulz, D. (2006). Plasticity and stability in neuronal output via changes in intrinsic excitability: it's what's inside that counts. *J. Exp. Biol.*, 209:4821–4827.
- Schummers, J., Marino, J., and Sur, M. (2002). Synaptic integration by V1 neurons depends on location within the orientation map. *Neuron*, 36(5):969–978.
- Schummers, J., Marino, J., and Sur, M. (2004). Local networks in visual cortex and their influence on neuronal responses and dynamics. *J Physiol Paris*, 98(4-6):429–441.
- Schwabe, L., Obermayer, K., Angelucci, A., and Bressloff, P. C. (2006). The role of feedback in shaping the extra-classical receptive field of cortical neurons: a recurrent network model. *J Neurosci*, 26(36):9117–9129.
- Sclar, G. and Freeman, R. D. (1982). Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp Brain Res*, 46(3):457–461.
- Sengpiel, F. and Bonhoeffer, T. (2002). Orientation specificity of contrast adaptation in visual cortical pinwheel centres and iso-orientation domains. *Eur J Neurosci*, 15(5):876–886.
- Sengpiel, F., Stawinski, P., and Bonhoeffer, T. (1999). Influence of experience on orientation maps in cat visual cortex. *Nat Neurosci*, 2(8):727–732.
- Series, P., Lorenceau, J., and Fregnac, Y. (2003). The "silent" surround of V1 receptive fields: theory and experiments. *J Physiol Paris*, 97(4-6):453–474.
- Shepherd, G. M. G., Stepanyants, A., Bureau, I., Chklovskii, D., and Svoboda, K.

- (2005). Geometric and functional organization of cortical circuits. *Nat Neurosci*, 8(6):782–790.
- Shouval, H., Goldberg, D., Jones, J., Beckerman, M., and Cooper, L. (2000). Structured long-range connections can provide a scaffold for orientation maps. *J. Neurosci.*, 20:1119–1128.
- Shouval, H., Intrator, N., Law, C. C., and Cooper, L. N. (1996). Effect of binocular cortical misalignment on ocular dominance and orientation selectivity. *Neural Comput.*, 8(5):1021–1040.
- Silberberg, G. and Markram, H. (2007). Disynaptic inhibition between neocortical pyramidal cells mediated by martinotti cells. *Neuron*, 53(5):735–746.
- Skottun, B. C., Bradley, A., Sclar, G., Ohzawa, I., and Freeman, R. D. (1987). The effects of contrast on visual orientation and spatial frequency discrimination: a comparison of single cells and behavior. *J Neurophysiol*, 57(3):773–786.
- Solomon, S. G. and Lennie, P. (2007). The machinery of colour vision. *Nat Rev Neurosci*, 8(4):276–286.
- Solomon, S. G., White, A. J. R., and Martin, P. R. (2002). Extraclassical receptive field properties of parvocellular, magnocellular, and koniocellular cells in the primate lateral geniculate nucleus. *J Neurosci*, 22(1):338–349.
- Somers, D. C., Nelson, S. B., and Sur, M. (1995). An emergent model of orientation selectivity in cat visual cortical simple cells. *J Neurosci*, 15(8):5448–5465.
- Somers, D. C., Todorov, E. V., Siapas, A. G., Toth, L. J., Kim, D. S., and Sur, M. (1998). A local circuit approach to understanding integration of long-range inputs in primary visual cortex. *Cereb Cortex*, 8(3):204–217.
- Stemmler, M., Usher, M., and Niebur, E. (1995). Lateral interactions in primary visual cortex: a model bridging physiology and psychophysics. *Science*, 269(5232):1877–1880.
- Stosiek, C., Garaschuk, O., Holthoff, K., and Konnerth, A. (2003). In vivo two-photon calcium imaging of neuronal networks. *Proc Natl Acad Sci U S A*, 100(12):7319–7324.
- Sullivan, T. and de Sa, V. (2006). Homeostatic synaptic scaling in self-organizing maps. *Neural Netw*, 19:734–743.

- Swindale, N. V. (1996). The development of topography in the visual cortex: a review of models. *Network*, 7(2):161–247.
- Swindale, N. V., Shoham, D., Grinvald, A., Bonhoeffer, T., and Hübner, M. (2000). Visual cortex maps are optimized for uniform coverage. *Nat Neurosci*, 3(8):822–826.
- Tanaka, K. (1983). Cross-correlation analysis of geniculostriate neuronal relationships in cats. *J Neurophysiol*, 49(6):1303–1318.
- Tanaka, S. (1990). Theory of self-organization of cortical maps: Mathematical framework. *Neural Networks*, 3:625–640.
- Tavazoie, S. F. and Reid, R. C. (2000). Diverse receptive fields in the lateral geniculate nucleus during thalamocortical development. *Nat. Neurosci.*, 3(6):608–616.
- Tiao, Y. C. and Blakemore, C. (1976). Functional organization in the visual cortex of the golden hamster. *J Comp Neurol*, 168(4):459–481.
- Triesch, J. (2005). A gradient rule for the plasticity of a neurons intrinsic excitability. *ICANN - Lecture Notes in Computer Science*, 3696(2):65–70.
- Truchard, A. M., Ohzawa, I., and Freeman, R. D. (2000). Contrast gain control in the visual cortex: monocular versus binocular mechanisms. *J Neurosci*, 20(8):3017–3032.
- Turrigiano, G. (1999). Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends Neurosci.*, 22:221–227.
- Turrigiano, G. and Nelson, S. (2004). Homeostatic plasticity in the developing nervous system. *Nat. Rev. Neurosci.*, 5:97–107.
- Turrigiano, G. G. and Nelson, S. B. (2000). Hebb and homeostasis in neuronal plasticity. *Curr Opin Neurobiol*, 10(3):358–364.
- Tusa, R., Palmer, L., and Rosenquist, A. (1978). The retinotopic organization of area 17 (striate cortex) in the cat. *J. Comp. Neurol.*, 177:213–235.
- Van Hooser, S. (2007). Similarity and diversity in visual cortex: is there a unifying theory of cortical computation? *Neuroscientist*, 13:639–656.
- Van Hooser, S. D., Heimel, J. A., Chung, S., and Nelson, S. B. (2006). Lack of patchy

- horizontal connectivity in primary visual cortex of a mammal without orientation maps. *J Neurosci*, 26(29):7680–7692.
- Van Hooser, S. D., Heimel, J. A. F., Chung, S., Nelson, S. B., and Toth, L. J. (2005). Orientation selectivity without orientation maps in visual cortex of a highly visual mammal. *J Neurosci*, 25(1):19–28.
- von der Malsburg, C. (1973). Self-organization of orientation-sensitive cells in the striate cortex. *Kybernetik*, 15:85–100.
- Walker, G. A., Ohzawa, I., and Freeman, R. D. (1999). Asymmetric suppression outside the classical receptive field of the visual cortex. *J Neurosci*, 19(23):10536–10553.
- Wang, C., Bardy, C., Huang, J. Y., FitzGibbon, T., and Dreher, B. (2009). Contrast dependence of center and surround integration in primary visual cortex of the cat. *J Vis*, 9(1):20.1–2015.
- Wang, Q. and Burkhalter, A. (2007). Area map of mouse visual cortex. *J. Comp. Neurol.*, 502:339–357.
- Warren, D., Fernandez, E., and Normann, R. (2001). High-resolution two-dimensional spatial mapping of cat striate cortex using a 100-microelectrode array. *Neuroscience*, 105:19–31.
- Webb, B. S., Tinsley, C. J., Barraclough, N. E., Easton, A., Parker, A., and Derrington, A. M. (2002). Feedback from V1 and inhibition from beyond the classical receptive field modulates the responses of neurons in the primate lateral geniculate nucleus. *Vis Neurosci*, 19(5):583–592.
- Weber, C. (2001). Self-organization of orientation maps, lateral connections, and dynamic receptive fields in the primary visual cortex. In *Proceedings of the International Conference on Artificial Neural Networks*, Lecture Notes in Computer Science 2130, pages 1147–1152, Berlin. Springer.
- Weliky, M., Bosking, W. H., and Fitzpatrick, D. (1996). A systematic map of direction preference in primary visual cortex. *Nature*, 379(6567):725–728.
- Weliky, M., Kandler, K., Fitzpatrick, D., and Katz, L. C. (1995). Patterns of excitation and inhibition evoked by horizontal connections in visual cortex share a common relationship to orientation columns. *Neuron*, 15(3):541–552.

- Weliky, M. and Katz, L. C. (1997). Disruption of orientation tuning in visual cortex by artificially correlated neuronal activity. *Nature*, 386(6626):680–685.
- White, L. and Fitzpatrick, D. (2007). Vision and cortical map development. *Neuron*, 56:327–338.
- White, L. E., Coppola, D. M., and Fitzpatrick, D. (2001). The contribution of sensory experience to the maturation of orientation selectivity in ferret visual cortex. *Nature*, 411(6841):1049–1052.
- Wickersham, I. R., Lyon, D. C., Barnard, R. J. O., Mori, T., Finke, S., Conzelmann, K.-K., Young, J. A. T., and Callaway, E. M. (2007). Monosynaptic restriction of transsynaptic tracing from single, genetically targeted neurons. *Neuron*, 53(5):639–647.
- Wieland, J. and Sajda, P. (2006). Extraclassical receptive field phenomena and short-range connectivity in V1. *Cereb Cortex*, 16(11):1531–1545.
- Wiesel, T. and Hubel, D. (1974). Ordered arrangement of orientation columns in monkeys lacking visual experience. *J. Comp. Neurol.*, 158:307–318.
- Wong, R. O. (1999). Retinal waves and visual system development. *Annu. Rev. Neurosci.*, 22:29–47.
- Yao, H. and Li, C.-Y. (2002). Clustered organization of neurons with similar extrareceptive field properties in the primary visual cortex. *Neuron*, 35(3):547–553.
- Yoshimura, Y. and Callaway, E. M. (2005). Fine-scale specificity of cortical networks depends on inhibitory cell type and connectivity. *Nat Neurosci*, 8(11):1552–1559.
- Yoshimura, Y., Dantzker, J. L. M., and Callaway, E. M. (2005). Excitatory cortical neurons form fine-scale functional networks. *Nature*, 433(7028):868–873.
- Yuille, A. L., Kammen, D. M., and Cohen, D. S. (1989). Quadrature and the development of orientation selective cortical cells by Hebb rules. *Biological Cybernetics*, 61(3):183–194.
- Zepeda, A., Arias, C., and Sengpiel, F. (2004). Optical imaging of intrinsic signals: recent developments in the methodology and its applications. *J Neurosci Methods*, 136(1):1–21.

- Zhang, W. and Linden, D. J. (2003). The other side of the engram: experience-driven changes in neuronal intrinsic excitability. *Nat. Rev. Neurosci.*, 4(11):885–900.