Topographic and laminar models for the development and organisation of spatial frequency and orientation in V1

Chris M Palmer

Doctor of Philosophy

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

2009

Abstract

Over the past several decades, experimental studies of the organisation of spatial frequency (SF) preference in mammalian visual cortex (V1) have reported a wide variety of conflicting results. A consensus now appears to be emerging that in the superficial layers SF is mapped continuously across the cortical surface. However, other evidence suggests that SF may differ systematically with cortical depth, at least in layer 4, where the magnocellular (M) and parvocellular (P) pathway afferents terminate in different sublaminae. It is not yet clear whether the topographic organisation for SF observed in the superficial layers is maintained throughout the input layers as well, or whether there is a switch from a laminar to a topographic organisation along the vertical dimension in V1.

I present results from two alternative self-organising computational models of V1 that receive natural image inputs through multiple SF channels in the LGN, differing in whether they develop laminar or topographic organisation in layer 4. Both models lead to topographic organisation for orientation (OR) and SF preference in upper layers, consistent with current experimental evidence. The results suggest that in either case separate sub-populations of neurons are required to obtain a wide range of SF preference from Hebbian learning of natural images. These models show that a laminar organisation for SF preference can coexist with a topographic, columnar organisation for orientation, and that the columnar organisation for orientation is dependent upon inter-laminar feedback. These results help clarify and explain the wide range of SF results reported in previous studies.

Acknowledgements

I am extremely grateful to my supervisor Jim Bednar for his excellent guidance and enthusiasm over the past four years. His patience and good humour have been inspiring and appreciated throughout. I also wish to thank my close colleagues Jan Antolik, Chris Ball, and Jude Law, all of whom have been immeasurably helpful throughout my research. Jude Law has been particularly amazing in her motivation and support - we have shared many dark times together, staring despairingly into the abyss of the PhD, but she has helped me find the light! There are many other students in the DTC programme who have been of great encouragement, and have provided kind words; I would particularly like to thank Sim Bamford, John Davey, Nicolas Heess, Adrianna Teriakidis, and Lysimachos Zografos. The DTC and IANC staff have been very supportive thoughout the programme, and I am very grateful to Mark Van Rossum, David Willshaw, Peggy Seriès, Matthias Hennig, David Sterratt, and of course Pat Ferguson.

Thoughout this research I have fought with technology, battling against the limits of the patience of computers while trying to run week-long simulations, and struggling with computers that crash five times a day. I am extremely grateful to the computer support personnel, particularly Alison Downie, Ewan Grant, and Jennifer Oxley, who always met my frustration with a friendly smile and proceeded to show me what I had done wrong.

Thanks to forest for the all night dance parties, the eternal chaos, and all the good times. Outside of University my friends have kept me alive. Particular love goes to Yasmin Fedda, Jane Flett, James Gloyn, Dan Gorman, Nicholas Holdstock, Dai Jones, Faith Nicholson, and Ryan Van Winkle. These people understand what lazers are for. Enormous thanks goes to my loving family who have always been there for me, and are deeply appreciated.

I am also grateful for financial support of the EPSRC and MRC, which made this work possible by funding me through the Edinburgh Doctoral Training Centre (DTC) in Neuroinformatics and Computational Neuroscience.

This work has made use of the resources provided by the Edinburgh Compute and

Data Facility (ECDF). (http://www.ecdf.ed.ac.uk/). The ECDF is partially supported by the eDIKT initiative (http://www.edikt.org).

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.

(*Chris M Palmer*)

“Start at infinity. Then go backwards until it’s possible.” - Jim Bednar

Dedicated to my parents, Christine and John,

and to Ellie Maxwell

whose memory will always be cherished.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Introduction

Vision is perhaps the most extraordinary and powerful of all our senses, allowing us to acquire richly varied and precisely detailed information about our surroundings. It is also the most accessible for study, due to the relative ease and accuracy of controlling stimuli to drive it. It seems likely that cortical regions that process other senses operate in a fashion grossly similar to those for visual processing, and that the actual function is determined by the driving input to a region (Sur et al., 1988). Thus success in understanding visual processing may provide insight into other aspects of brain function. It is for these reasons that vision is the most commonly researched of the senses.

Through experiments in monkeys and cats, it is known that neurons in the primary visual cortex (V1) have preferences for various localised components of images, such as orientation (OR), colour, spatial frequency (SF), and direction of motion (see figure 1.1, and chapter 2). In higher mammals these properties are typically mapped systematically across the cortical surface, forming a topographic mapping that covers the range of possible feature values (Hubel et al., 1978). This mapping is also typically columnar, in that neurons under a certain location on the surface tend to have similar preferences, leading to aligned topographic maps in each of the cortical layers (Vanduffel et al., 2002). Figure 1.1D and E show an example set of aligned columnar topographic maps for orientation.

Although it has been known for decades that V1 neurons have a wide range of spatial frequency (SF) preferences (Campbell et al., 1969; Maffei and Fiorentini, 1977), there has been much debate about how SF preference might be organised in cortex: whether SF preference is topographic and continuous (Everson et al., 1998, Issa et al., 2000), topographic but with binary values (Hübener et al., 1997; Shoham et al., 1997), laid out laminarly (with cortical depth) (Maffei and Fiorentini, 1977), or in no apparent order (Sirovich and Uglesich, 2004). Recently, experimental techniques have improved, and a number of studies have published continuous topographic SF preference maps for a variety of species (Xu et al., 2007; Zhang et al., 2007; Farley et al., 2007; Mallik et al., 2008; Purushothaman et al., 2009). However, due to limitations of the imaging techniques, the maps from these studies apply only to the uppermost (superficial) layers in V1. Thus we know that in these upper layers SF is organised topographically, but we do not know what form the organisation takes in the deeper layers.

For instance, there is some evidence for a laminar organisation in layer 4 (Tootell et al., 1988). V1 receives its input from the lateral geniculate nucleus (LGN). The LGN afferents projecting to V1 form more connections in layer 4 than any other layer (Hubel and Wiesel, 1972). Hendrickson et al. (1978) report that in primate the segregation of pathways found in LGN is preserved in their projections to layer 4. Magnocellular neurons (with lower SF preference) project to layer 4Cα, while parvocellular neurons (with higher SF preference) project to the 4Cβ lamina. The amount of segregation between these pathways appears to vary widely across species. The pathways are strongly separated in macaque (Tootell et al., 1988), and while separation has also been observed in cat (Ferster and LeVay, 1978; Lund et al., 1979), it may be weaker or only partial (Humphrey et al., 1985).

This aspect of the cortical architecture raises some interesting questions. If different laminae in V1 receive separate SF information, one would reasonably expect a laminar organisation for SF preference in layer 4. If this laminar organisation exists, is it compatible with columnar organisation for OR preference? Can a topographic organisation for SF preference arise in the upper layers from a laminar SF organisation in layer 4? If the SF pathways from the LGN are not segregated and neurons in

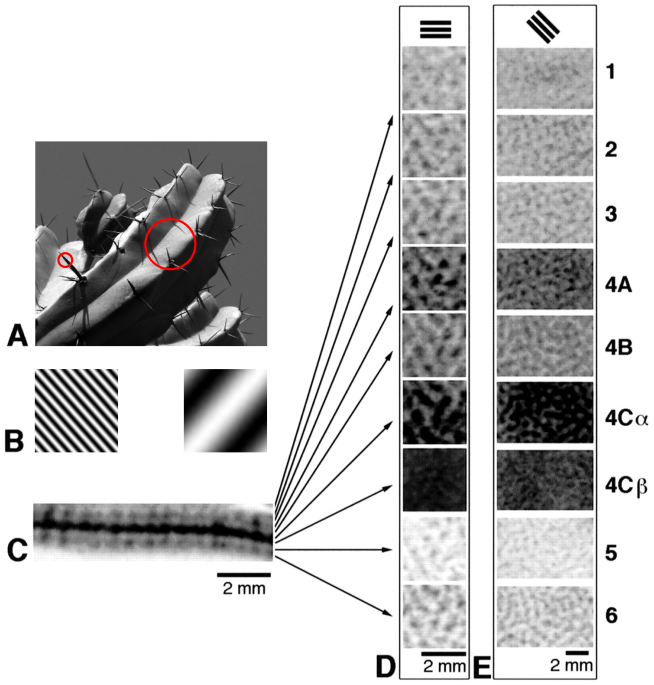

Figure 1.1: V1 neuron feature preferences and columnar topographic organisation. (A) Neurons in the visual cortex can have preferences for different local features of an image, such as position, orientation, and spatial frequency. Red circled regions show examples of two different features that would excite different neurons, because they differ in position, orientation and spatial frequency. (B) In an experimental laboratory, the responses of neurons can be tested using artificial stimuli, such as sine gratings of different orientation and frequency, to determine their feature preferences. The two patterns shown would be likely to excite neurons that respond to one of the two circled areas in A, if presented at the correct location in the visual field. (C) A cross section through V1 of macaque, perpendicular to the cortical surface. Dark regions indicate laminae with high activity as a consequence of neurons being activated by a particular stimulus. (D and E) Slices taken tangentially to the cortical surface, at 9 different depths corresponding to different laminae of cortex, showing cortical activity in response to a horizontal or oblique stimulus, respectively. Note that in each column the activated patches are spatially coincident in different laminae, which demonstrates the columnar organisation for OR preference. Note the variation in uptake between the different laminae, which may be due to laminar differences in OR selectivity. C, D, and E reprinted from Vanduffel et al. (2002).

layer 4 receive afferents with many different SF preferences, can an SF map develop to represent the entire range of SF? These are questions that are currently very difficult to investigate experimentally, due to the difficulty in determining organisation from electrode studies, and because the regions in question are too deep for current imaging technology.

Computer modelling is a powerful tool in this context. One can devise a model to explore these issues, based upon what are considered to be the pertinent aspects of the system. Indeed, there has been a great amount of success in reproducing topographic feature maps using computational models (for reviews see Swindale, 1996 and Goodhill, 2007). These models often succeed in producing OR maps that are qualitatively similar to those observed experimentally, but typically use high-level algorithms such as the elastic net (Durbin and Willshaw, 1987; Carreira-Perpiñán et al., 2005). While this approach serves well to illustrate what appears to be one of the overarching aims of cortex organisation – dimension reduction, whereby a space with multiple feature dimensions is mapped onto a two dimensional sheet – it does not capture the underlying organisational or developmental mechanisms. Organisational models that bear a closer resemblance to the physiological detail of brain architecture, while still being as simple as possible, may serve better for this purpose.

Hubel and Wiesel's model for the formation of a V1 simple cell receptive field (RF) through the convergence of aligned geniculate input RFs suggests that the simple cell would inherit the SF preference of the geniculate input. If this is the case, then the range of optimum SF preference found in V1 would be a consequence of the SF preference found earlier in the visual pathway, specifically the preferences of the centre/surround retinal ganglion and LGN cells. As individual LGN cells tend to receive afferent input principally from a single retinal ganglion cell in fovea (McMahon et al., 2000), or from a small number of retinal ganglion cells with similar size receptive fields and spatial location (Usrey et al., 1999), the range of SF preference found in the retina is reflected in the LGN. At a given eccentricity on the retina the range of RF sizes (and hence SF preference) of LGN cells (or, similarly, retinal ganglion cells) can cover a broad range whose precise values differ between species. For example, at 10° eccentricity, grey squirrel LGN cell RF centres may vary from about 1° to 12° (Van

Hooser et al., 2003), while marmoset may vary from 0.03° to 0.7° (Kremers and Weiss, 1997). LGN cells are typically classified into different cell types (magno, parvo and konio in primate; X, Y, and W in cat), according to differences in receptive field properties (size, temporal response, etc.). If, however, we solely consider the SF preference of LGN cells' RFs at a given eccentricity, we often find a large and essentially continuous range (for instance see Levitt et al., 2001) - typically spanning 3 to 4 octaves. (An octave is an interval between two frequencies that has a ratio of 2 to 1; e.g. the range between 1 cycle per degree (cpd) and 2 cpd is one octave, as is the range between 20 cpd and 40 cpd.) One of the questions I try to address in this work is how a large range of SF preference centre/surround cells in the LGN, all making afferent connections to cells in V1, can result in continuous OR and SF preference maps.

To examine these questions, I have created computational models that can develop cortical maps from natural images, when given input from a model LGN with centre/surround cells covering a range of sizes comparable to that in animals. I think it is important to base the models upon the physiology of the visual system, as this is likely to offer deeper understanding of the mechanisms at play in cortical functioning and organisation, and generate further questions, which may be pursued in later research. The models are driven by natural images, while their architecture captures the form of the flow of information found in the visual pathway, allowing development of cortical organisation to arise through Hebbian learning. I present results demonstrating that a laminar organisation for SF preference can coexist with columnar topographic mapping for orientation preference in layer 4C, and that from this a continuous topographic organisation of spatial frequency preference can then arise in the superficial layers.

1.2 Thesis outline

The following chapters outline my investigation into SF organisation:

Chapter 2 begins by providing a broad functional overview of the early visual pathway in cat and macaque. It then focuses upon specific aspects of visual processing that are particularly relevant for SF. The second half of the chapter is a literature review

of published experimental work relating to SF maps in V1. Despite a long history of controversy, there is now a broad consensus about the existence and form of SF maps. In this chapter the biological foundations for the models described in later chapters are established.

Chapter 3 presents a review of previous models pertaining to SF organisation. We see that for many of these models SF mapping is not the central aim. Most of the maps produced by these models bear little resemblance to observed animal maps, and of those that do the models have considerable limitations. The LISSOM model is introduced and its suitability for forming the basis of a model to investigate SF organisation is considered.

Chapter 4 introduces a model with a single cortical sheet, which receives input from 4 LGN channels, each tuned for a different SF. The results from this model show that topographic SF maps can form from such a model, but the SF preference range is less than the input space. Further results demonstrate that by segregating the input from the SF channels to each neuron, the SF range can be improved so that it reflects the input space.

Chapter 5 introduces a model with a hierarchical V1. Layer 4C comprises 4 interconnected cortical sheets, each receiving separate input from one of 4 LGN channels. These 4 cortical sheets in turn project to layer 2/3, represented by a single cortical sheet. The results demonstrate that a laminar organisation for SF preference can coexist with a topographic, columnar organisation for orientation. The results also suggest that establishing a columnar organisation for orientation is dependent upon interlaminar feedback, even when each lamina gets inputs from LGN channels that have overlapping response properties. Furthermore, a topographic organisation for SF can develop in layer 2/3, but the neurons are poorly tuned for SF preference. Segregation of the channels allows neurons to become more tuned for SF preference.

Chapter 6 discusses the implications of these findings and how they may affect experimental work. Possible extensions to these models and other future work are then discussed.

Chapter 7 summarises the outcomes of this study.

Appendix A lists the parameters used in the simulations.

Chapter 2

Background and literature review

2.1 Mammalian visual system - Overview

This thesis presents detailed computational models for the organisation of SF preference in cortex. These models are based upon the known physiology of the early visual system, and so a thorough grounding in the mammalian visual system is required. Here I provide a basic overview of the mammalian early visual pathway, before focusing upon specific features of the primate and cat visual systems in further depth that relate to my modelling work. Primate and cat are the most commonly studied species for investigation of cortical SF preference representation, so knowledge of the visual mechanics of these species forms the foundation on which the computational models described later (in chapters 4 and 5) are based. There are some important differences between the visual pathways of these species, and it is often the case that an experiment conducted upon one species may not have been repeated on the other, leading to difficulties in comparing the visual systems of these species. This factor also makes it very difficult to gain a fully detailed picture of the visual system, as it is difficult to know if information only known for one species may generalise to others, or if it is a unique feature.

The mammalian early visual system is the pathway from eye to neocortex (see figure 2.1), and it can be broken down into a number of functional sections, which are briefly introduced here before considering aspects relevant to subsequent work and discussion in more detail. Sight begins with the eye, and the natural world provides

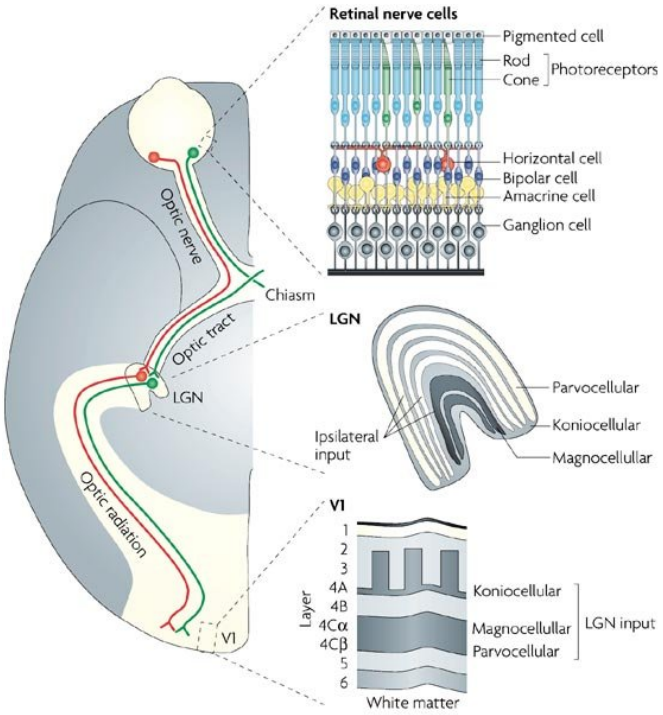

Figure 2.1: Diagram showing gross anatomy of the early visual pathway, and the organisation of the retina, LGN and V1 of the macaque.

Illustrations on the right of the figure provide detail of regions of the pathway concerned with visual processing. An image is formed upon the retina, whereupon photoreceptors produce signals inversely proportional to the incident light intensity. These signals are processed by the retinal circuitry, and then encoded by the ganglion cells as spike rates. Ganglion cells provide the output from the retina, their axons forming the optic nerve, which projects to the LGN. The LGN is a folded laminate structure, with each lamina consisting of specific cell types. The LGN output projects mainly to V1, with the different cell types in LGN providing segregated input to specific V1 sublaminae. The termination site in layer 4 depends on the layer in which the LGN neuron is found: layer 4Cβ receives input predominantly from parvocellular cells, 4Cα receives input predominantly from magnocellular cells, and layer 4A and lower layer 3 receive koniocellular cell input. Reprinted from Solomon and Lennie (2007).

us with a wealth of different eye types. The very simplest pit eyes, found in some multicellular animals some 540 million years ago, can do little more than tell light from dark (for review see Land and Fernald, 1992). Such eyes have diameters of less than 100μm, and comprise only 1 to 100 photoreceptors, but these form the evolutionary beginnings of vision. Through time the eye has evolved, becoming more optically sophisticated, resulting in an increase in its resolving power. Land vertebrates utilise corneal refraction for image formation upon the retina, with adjustment of focus performed by changing the focal length of the lens through mechanical deformation by contraction or relaxation of the ciliary muscle. The iris, which forms an adjustable aperture, can dilate or constrict to try to maintain a constant flux of light that enters the pupil, allowing the eye to adapt to available light levels. The eye is an extraordinary device, carefully controlled both consciously and unconsciously by the brain through an elaborate feedback system, all to try to produce the best possible image for subsequent processing.

The primate retina is a thin (100-200μm) laminar structure that covers the back of the eye (see figure 2.2). The outermost layer of retina consists of a loosely hexagonal mosaic of photoreceptor cells. The layer of photoreceptors provides the interface between the visual world and the brain. The photoreceptors are optoelectical transducers; absorption of incident photons leads to hyperpolarisation of their electrical membrane potential (for review see Sterling and Demb, 2004). This hyperpolarisation leads to a proportionate release of neurotransmitter. It is these neurotransmitter signals that are then processed by the other retinal circuitry, resulting in the final coded signal being transmitted as action potentials to the LGN. There are two classes of retinal photoreceptors: cones, which are sensitive to bright light levels (and thus used for daylight vision), and rods, which are highly sensitive to low light, responding even to single photons, but saturating at daylight levels. The number of cone types determines whether an animal has colour vision in daylight or not. Colour vision is a consequence of slightly differing types of cone cells having different frequency response profiles. Cones and rods are mixed on the retina, although in the foveal region there are no rods, and peripheral areas have more rods than cones. The fovea is a small pit in the centre of the retina, created by the separation of the retinal layers, allowing light to fall directly upon the photoreceptors, resulting in the greatest resolution. Near the fovea, cones are

much smaller than rods, so higher resolution vision is a result of cones.

The retinal processing (for review see Field and Chichilnisky, 2007) serves to compress the encoding of the visual image, which is elegantly performed by essentially discarding regions of uniform illumination, while edges of visual features and regions changing in illumination with time are retained and enhanced. The retinal circuitry accomplishes this task by pooling the output from regions of photoreceptors and subtracting this collective activity from central sub-population activity, which results in the annular centre/surround receptive fields possessed by the retinal ganglion cells. The receptive field for a specific neuron is a map that illustrates the regions of retina that, when illuminated by light (or darkness), result in a maximal change in membrane potential. These retinal ganglion cells annular RFs are of two main varieties, ON cells, with an excitatory centre and inhibitory surround, and OFF cells, with an inhibitory centre and an excitatory surround (see figure 2.3). When uniformly illuminated, the net activity of these cells is nearly zero, as the inhibitory and excitatory contributions mostly cancel each other out. Thus they respond most strongly to points or edges. These responses provide the fundamental building blocks the brain uses to construct an understanding of the world. Ganglion cells generate action potentials at a rate proportional to how closely the light pattern matches their RF. These spikes are mediated along their axons, to the LGN.

It is indicative of the compression that the retinal circuitry performs that in adult humans there are about 100 million photoreceptors, but there are only around 1 million retinal ganglion cells. The axons of these retinal ganglion cells collectively form the optic nerve, which primarily terminate by innervating the LGN of the dorsal thalamus. The optic nerves from each eye meet at the optic chiasm, where the optic nerves split in two, such that ganglion cells in the temporal retina project to the ipsilateral hemisphere, while ganglion cells in the nasal retina project to the contralateral hemisphere. Thus the left half of the visual image is processed in the right brain hemisphere and vice versa.

The LGN is a folded laminar structure, where each lamina retains a retinotopic mapping. The number of laminae and their cellular organisation is quite species spe-

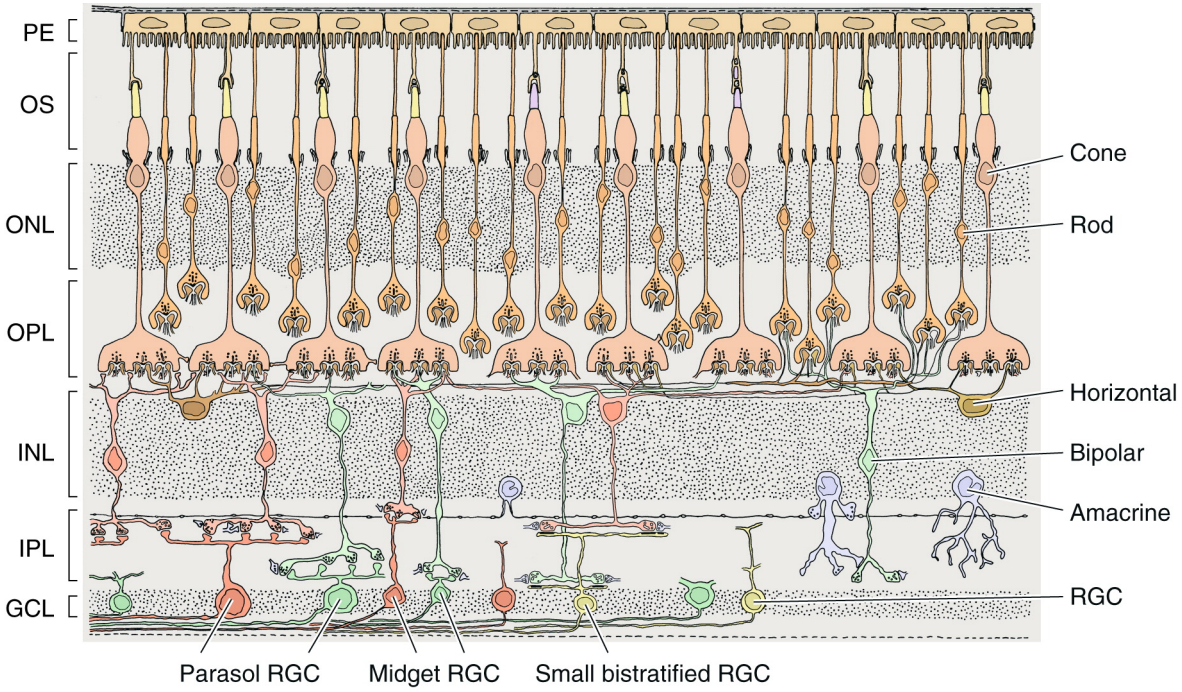

Figure 2.2: Cross section illustration of macaque retina.

The rod and cone photoreceptors transduce photons into electrical signals, which are conveyed through synaptic connections to bipolar and horizontal cells. The bipolar, amacrine, horizontal and retinal ganglion cells form specialised circuitry that creates the centre/surround RF output. Retinal ganglion cells of different types form the output layer, sending the visual signal to LGN and other brain regions. Abbreviations are as follows: Retinal ganglion cell (RGC), pigment epithelium (PE), outer segments of photoreceptors (OS), outer nuclear layer (ONL), outer plexiform layer (OPL), inner nuclear layer (INL), inner plexiform layer (IPL), ganglion cell layer (GCL). Reprinted from Field and Chichilnisky, 2007.

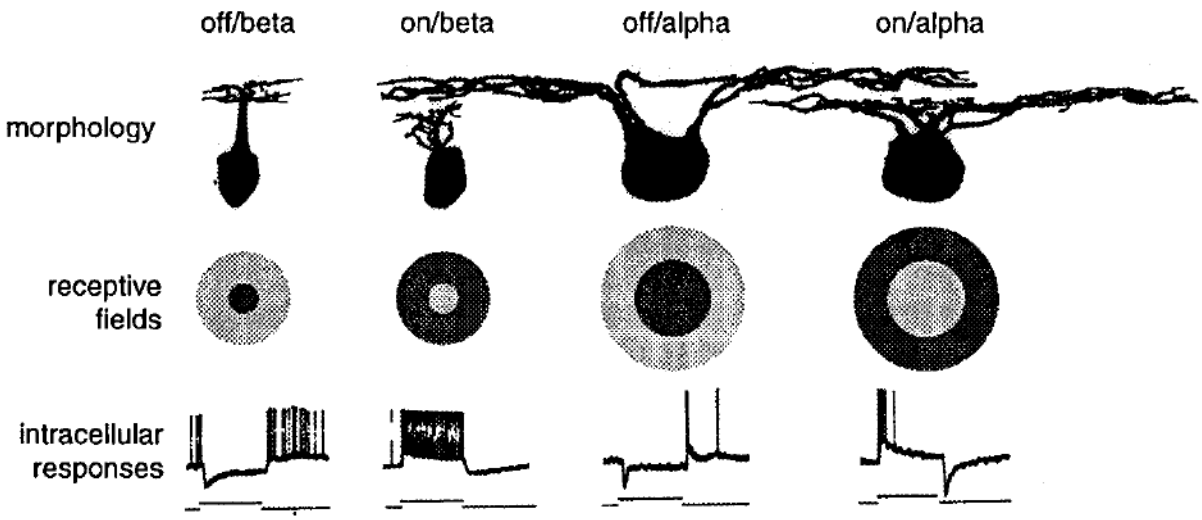

Figure 2.3: Cat retinal ganglion cells.

First row shows the different morphology and size of alpha (also termed X) and beta (also termed Y) cells in cat. The alpha cells are larger and their dendritic field is more extensive. The second row shows the corresponding receptive fields for these ON and OFF cells. The ON cells are maximally excited when light covers the whole dendritic field, which corresponds to the receptive field centre shown (RFs and neuron morphology are at differing spatial scales). They are maximally suppressed when light fills the outer receptive field surround. For OFF cells it is the other way around. The third row shows responses to onset and offset of light. Beta cells have a mainly sustained response, while alpha cells have a transient response. Reprinted from Sterling and Demb (2004).

cific. The exact target of retinal ganglion cells tends to depend upon the cell properties and the eye of origin. The specifics of LGN organisation are examined for primate and cat in subsections 2.2.2 and 2.2.5 below. LGN cells receive afferent input from a small number of retinal ganglion cells. The connected ganglion and geniculate cells tend to have very similar RFs, and the geniculate response is often dominated by a single afferent ganglion contribution (Usrey et al., 1999). Due to this, the spatial form of the annular ON and OFF RFs is essentially preserved from the retinal ganglion cells to the LGN cells. As well as receiving afferent projections from the retina, the LGN also receives extensive backprojections from brainstem, V1 and higher visual processing areas (Sherman and Koch, 1986).

The primary visual cortex is at the back of the occipital lobes of the cortex. It is also a laminar structure, historically described as consisting of 6 laminae, within which there are a number of sublaminae. The laminae are distinguished by the different distributions of cell types that are found in them. Thalamocortical projections terminate predominantly in layer 4 of V1. As will be described in detail below, there is some lamina specificity of the projections, depending upon the LGN cell type. In cat, neurons receiving thalamic input may have connections from around 30 LGN neurons (Alonso et al., 2001).

There are a variety of neurons that make up the visual cortex, but most of them fall into two basic classes. Pyramidal cells are excitatory, projection neurons, and have axons that form connections to local neurons, but also to more distant brain areas. Their dendritic trees are long and spiny. Non-pyramidal cells are small local interneurons. They can be excitatory, with spiny denrites, or inhibitory with smooth dendrites. The nonpyramidal neurons' axons are confined to V1.

In terms of response properties, two main classes of cell have been categorised: Simple cells, which are orientation selective and have a preference for a specific spatial phase (e.g. the precise location of a sine grating), and complex cells, which are largely phase invariant (Hubel and Wiesel, 1962). Simple cells are found more frequently in the layers of cortex that receive direct thalamic input, while complex cells are found more frequently in layers that are somewhat further along the processing

stream, such as layer 2/3 (Gilbert, 1977).

Understanding of cortical circuitry is still in its infancy. It appears that within the V1 cortical circuitry, the initial segregation between input pathways from the LGN is not maintained throughout the circuit; information that was segregated in LGN is mixed at different points in the circuit as processing develops. This mixing of signal types perhaps allows for the emergence of new, more complex, RF types in higher parts of the circuit. There is a great deal of feedback both within the V1 circuit and from higher cortical areas. It has been argued that the circuit is essentially feedforward, but activity is modulated by weak input from lateral and higher regions (Sherman and Guillery, 1998).

In addition to connections extending between laminae and with other cortical regions, neurons often make lateral connections that can form an extensive network within a lamina. Short-range local connections link neighbouring neurons in a small region of cortex. Typically these interconnections are between neurons that share feature preferences, and it has been suggested that this local interlinking may serve as a sort of amplifier, strengthening the signal of the geniculate input (Douglas et al., 1995).

Long range horizontal connections are the other main form of lateral connectivity and can span as far as a several millimetres, connecting neurons that share the same OR preference (Gilbert and Wiesel, 1989; Bosking et al., 1997). These connections are patchy and tend to extend in the direction of the preferred OR. This patchy network of lateral connections is thought to be important for the development of feature preferences, and for cortical organisation (Ruthazer and Stryker, 1996). Functionally, the long-range connections may account for some of the responses observed outside of the classical RF, such as contextual modulation (Angelucci et al., 2002).

Different laminae tend to have cells with differing distributions of properties (Gur et al., 2005). Certain feature preferences appear to arise at certain points in the cortical circuit, which in turn correspond to particular laminae. For example, there are more neurons with phase invariance in layer 2/3 than layer 4 (Gilbert, 1977).

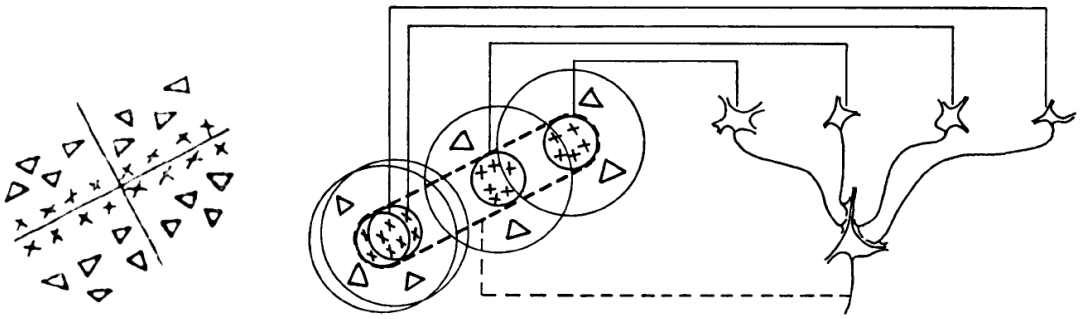

Figure 2.4: Simple cell circuitry.

(Left) Example receptive field of a simple cell in the cat visual cortex. Onset of light in the ON domain (x), or offset of light in an OFF domain (△) excites the cell, while light onset in an OFF domain or offset in an ON domain inhibits the cell. (Right) Suggested thalamocortical circuit for how the receptive field of a simple cell might be constructed from LGN afferent input (Hubel and Wiesel, 1962). Simple cell receives afferent input from LGN cells (top right). The LGN receptive fields are superposed with their centres in alignment, creating the simple cell receptive field. OFF cells are not shown, but their centres would be in alignment with the OFF domains of the simple cell. Modified from Hubel and Wiesel (1962).

It was suggested by Hubel and Wiesel (1962) that a simple hierarchical architecture allows for the construction of receptive fields that can become increasingly selective for complex stimuli further along the visual pathway. There is much evidence to support the emergence of simple cell properties in a cortical cell from afferent LGN input (for example see Reid and Alonso, 1995). The superposition of centre/surround LGN cells RFs is able to produce new Gabor-patch shaped RFs consisting of alternating bars preferring light or dark. Thus, a cortical neuron integrating input from LGN cells aligned correctly would produce such a receptive field (figure 2.4). These simple cells are maximally responsive to a light-dark boundary, orientated identically to the RF, and as such these neurons can be viewed as edge selective. They have a clear orientation preference, and as the orientation is altered the response decreases, until at the orthogonal orientation the neuron's activity drops to the background level. These simple cells are also selective for spatial position, spatial frequency, phase and eye. Similarly, it is thought that complex cells are constructed from superposition of simple cells with similar SF and OR but slightly different positions.

The retinotopic mapping found in LGN in preserved in V1. V1 is columnar, and neurons in different laminae that share the same position tangentially along the cortex surface have RFs that are principally at the same position on retina. A topographic columnar organisation is also found for orientation preference (see figure 2.5 and figure 1.1 in chapter 1). Neighbouring neurons tend to have the same orientation preference, forming regions called iso-OR domains, and these preferences are maintained perpendicular to the cortical surfaces, resulting in a columnar organisation. As well as iso-domains there are a number of other characteristic features found in OR maps: pinwheels, i.e. points around which OR changes continuously, covering the entire range of OR preference; linear zones, regions where along a straight line OR preference changes continuously; and fractures, sudden jumps from one OR preference to another. Other established feature preferences known to be topographically organised in visual cortex are direction selectivity and ocular preference. The emerging picture from recent work is that SF preference is also organised topographically (see section 2.4 for review).

The functional organisation appears to affect how often each neuron fires. When the visual cortex is stained for cytochrome oxidase, a metabolic enzyme indicative of high metabolic activity, a regular pattern of repeating blobs is revealed spanning layer 2/3. Livingstone and Hubel (1984) observed that cytochrome oxidase blobs contain neurons that are selective for colour, are monocular, and have little orientation selectivity.

Despite investigators' best efforts, complete understanding of the functional organisation of V1 remains elusive (for review see Sincich and Horton, 2005). While more complex RF properties clearly arise along the V1 cortical circuit, how this takes place is poorly understood. For instance, Livingstone and Hubel (1988) suggested that the image components of form, colour, and motion were segregated, such that cells in each pathway were not selective for properties outwith their pathway (for example a colour selective cell would not be selective for orientation), and were spatially compartmentalised both within V1 and at the output stage in projections to V2. However, the evidence supporting this view has been reinterpreted in light of new studies demonstrating that most cells have selectivity for many of the features purported to be

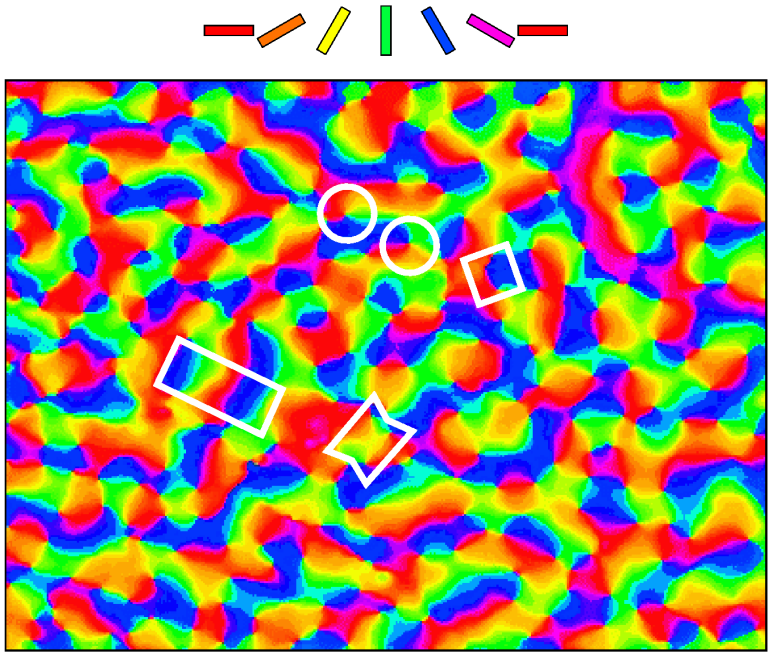

Figure 2.5: Adult macaque orientation map.

Neurons are colour coded according to orientation preference, as indicated in colour key at top. Adjacent neurons typically share orientation preference, forming iso-orientation domains. There are a number of characteristic OR map features picked out in this map. Pinwheels are points around which all orientations are represented continuously. Pinwheels often occur in pairs, joined by iso-domains. Two example pinwheels are indicated by white circles. Fractures occur when there is a non-continuous jump in OR preference; an example is shown in the white square. Linear zones are regions where OR preference changes continuously along a straight line; an example is shown in the white oblong. Saddle points occur when a iso-zone is nearly bisected by another OR iso-domain; an example is shown in the white bow tie. Cortex area shown is a 7.5mm by 5.5mm region. Reprinted from Blasdel (1992).

segregated, and that compartmentalisation of specialised neurons remains uncertain. It is now thought that two distinct projections from the cytochrome oxidase blobs and interblobs are made to V2 rather than three from specialised compartments, and that within these two projections the magno, parvo, and koniocellular content is mixed. What the functional difference between these pathways may be is unknown.

If one considers these methods of representation of features in visual space in V1, it becomes apparent that for a point in visual space, it is essential that all feature possibilities for the corresponding point in cortex are available. This is the requirement for good coverage (Swindale, 1991), and leads to the idea of the repeating hypercolumn (Hubel and Wiesel, 1974). The hypercolumn is considered to be a functional module of cortex that corresponds to one small region of visual space and within it represents every possible combination of stimuli. This is achieved by having cells that have preferences spanning the entire range of possible stimuli at this point. The hypercolumn is also considered to be orderly structured (see figure 2.6), comprising a full range of OR iso-zones, a pair of ocular dominance stripes, a cytochrome oxidase blob, and other feature preferences required for visual coverage. Some debate exists as to the possible arrangements of these feature preferences in a hypercolumn (see Bressloff and Cowan, 2003 for discussion). As the hypercolumn is repeated across cortex, it ensures full coverage and continuity in an orderly fashion. While this notion of hypercolumn is a simplification that overstates the modularity and uniformity of cortex, there is a lot of evidence to support this idea of a loosely co-ordinating tiling of feature preferences across cortex to achieve good coverage (Yu et al., 2005).

Although this thesis only considers V1, it is important to realise that V1 is only one of many visual processing regions in the cortex. After V1, there are around 30 further regions identified to be involved in visual processing. The visual pathway is thought to split into two functionally distinct pathways, the dorsal or parietal stream, and the ventral or temporal stream (Mishkin et al., 1983). The dorsal stream is thought to process information concerning motion and spatial location to guide motor actions. The dorsal stream runs up from the occipital lobe to the parietal lobe encompassing such areas as MT, MST, VIP, and LIP. The ventral stream is thought to relate to object identity and recognition. The ventral stream runs through the lower portion of cortex from occipi-

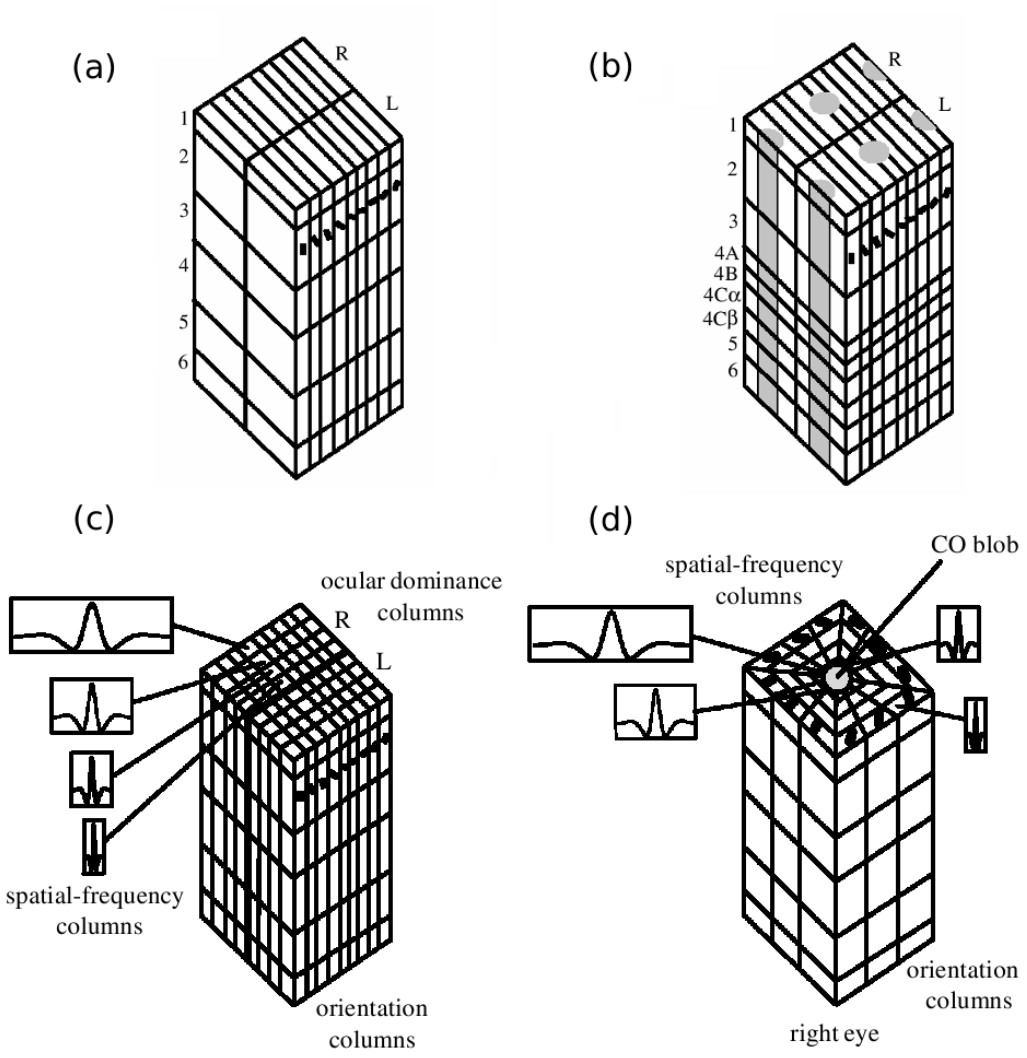

Figure 2.6: Different suggested hypercolumn organisations.

(a) The original ice cube V1 hypercolumn model, modified for cat. Ocular preference (R, right eye; L, left eye) is represented as rows orthogonal to orientation preference. Different cortical laminae are shown as horizontal slabs. All possible eye and orientation preferences are represented in the hypercolumn and in all layers. (b) The ice cube model hypercolumn model for macaque, including cytochrome oxidase blobs. (c) The modified ice cube model for cat, extended to incorporate a possible linear SF organisation. (d) Modified ice cube model for macaque, including cytochrome oxidase blobs, and showing an SF organisation where SF preference increases away from the cytochrome oxidase blobs. Reprinted from Bressloff and Cowan (2003).

tal lobe to the inferior temporal cortex, including areas V2, V3 and V4. The areas in these two pathways are highly interconnected, and there are also connections between the dorsal and ventral pathways themselves. A functional hierarchy exists amongst the regions in each of these pathways, with neurons further along a pathway becoming increasingly specialised for larger and higher order features of the visual scene.

This section has provided a brief overview of the visual system. The visual system is complex in detail, and both hierarchical and parallel in its organisation. It offers precision with reliability, which in day-to-day use is both essential and remarkable. The sections that follow consider in more detail aspects of the early visual system that relate to SF processing, without which much of the richness of the visual world would be lost.

2.2 SF processing in the visual system

There are many further details of the early visual system that are important for the generation of neurons with SF preference, and for understanding their possible organisation. This section aims to identify these elements, with a view to understanding the role that they play in vision so that we may judge whether and how they should be incorporated into a model of the visual system concerned with SF representation. This is done separately for primate and then for cat, because many of these details are species specific.

2.2.1 Primate retina

Retinal ganglion cells can be differentiated by a number of factors, both anatomical and physiological. There are three broad classes of primate retinal ganglion cell that are known to have axons that terminate in LGN. Midget cells, with narrow dendritic trees and sustained firing, project to the parvocellular layer in LGN. Most midget cells have red/green colour opponent receptive fields, e.g., a red ON centre with green OFF surround. However, because the spectral sensitivities of red and green cells overlap significantly, these cells also respond well to luminance edges. Parasol cells, which have broader dendritic arbors, are not colour selective, as their centre and surround

receive input from a mixture of cone types. Parasol cells project to the magnocellular layers of LGN. Bistratified cells have large RFs and project to the koniocellular layers. All these cell types can be further classified as ON or OFF cells. In and near the fovea, midget cells are tightly packed in the ganglion cell layer; here they typically receive afferent input from a single midget bipolar cell, which in turn receives input from a single cone photoreceptor in the centre (McMahon et al., 2000). The cone spacing sets the acuity limit of the eye. Other retinal ganglion cell types are rarer and not well characterised (for review see Field and Chichilnisky, 2007).

Apart from these large differences, there is a large variation in retinal ganglion cell RF size with eccentricity. Figure 2.7 shows a species comparison for dendritic field area of retinal ganglion cells. The dendritic field is the area over which retinal ganglion cells pool responses to form the centre region of their RFs. It has been shown (Croner and Kaplan, 1995) that this anatomical measure of dendritic tree size closely matches the functional RF centre size (see figure 2.8). This diversity in RF sizes corresponds to a large range in SF preference of these centre/surround cells at a given eccentricity. The range of preference typically spans 3 to 4 octaves in most species, and provides the SF input to the visual system. It is this range of SF input to the LGN and V1 that is referred to later as the "SF input space". If this information were not to be represented further along the processing stream, then its existence at this level would presumably not be necessary.

2.2.2 Primate LGN

The primate LGN (for review see Callaway, 2005) is composed of three cell types: magnocellular (M), parvocellular (P), and koniocellular (K). M cells receive input from parasol cells, and have large RFs and fast latencies and conduction velocities. The spatial summation characteristics of most of these cells are linear, ie. the response activity of the cell is proportional to the amount of light falling on a subunit within its RF. A smaller subpopulation with nonlinear properties is also observed (Kaplan and Shapley, 1982). P cells receive input from colour opponent midget cells, and have small RFs, slower and more sustained responses, slower conduction velocities, low contrast sensitivity, linear spatial summation, and have red/green colour opponency. K cell

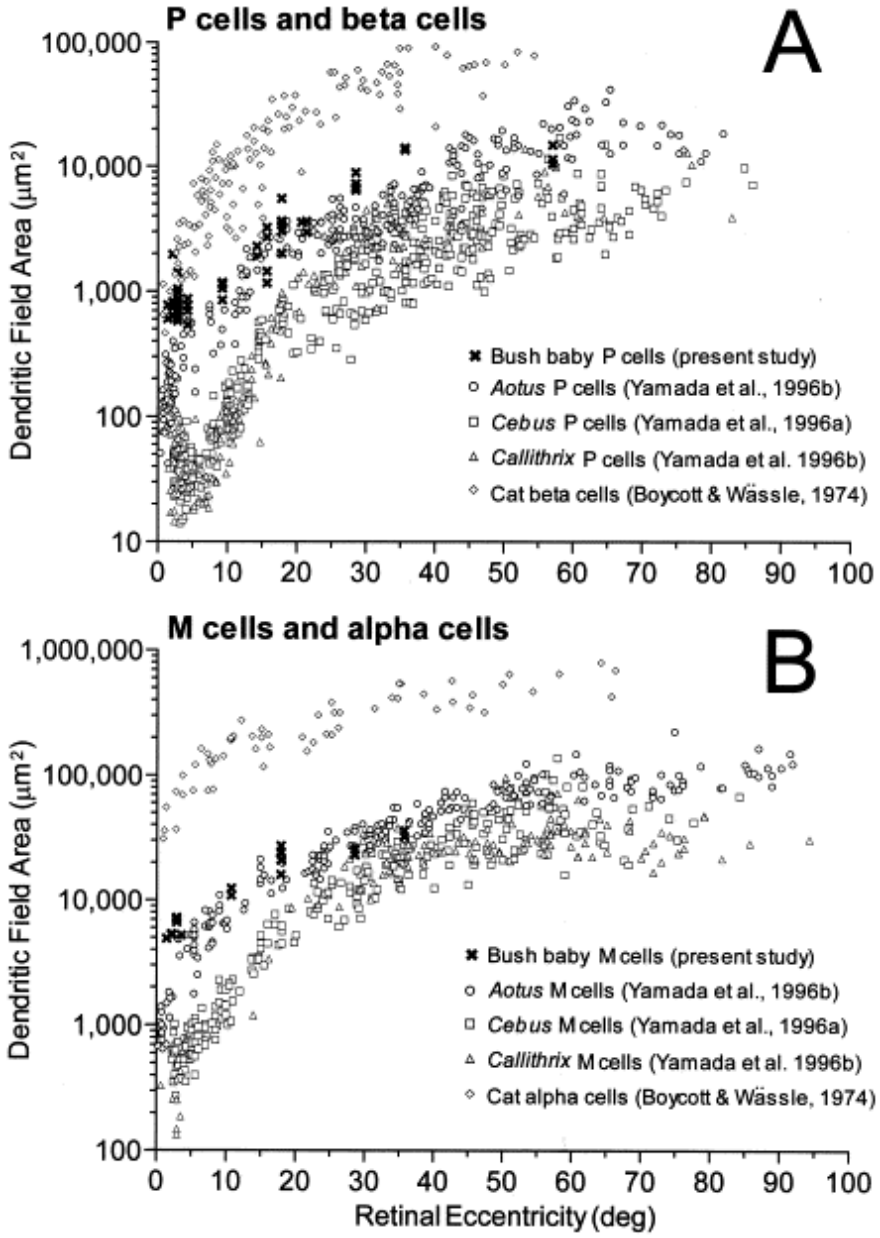

Figure 2.7: Dendritic field area of retinal ganglion cells.

Comparison for bush baby, owl monkey (Aotus), capuchin monkey (Cebus), marmoset (Callithrix), and cat. (A) Dendritic field area for primate P cells and cat beta cells. (B) Dendritic field area for primate M cells and cat alpha cells. There is a large range in size of dendritic field area for any species at any given eccentricity (note the different logarithmic y-axis scales for the two graphs). Reprinted from Yamada et al., 1998.

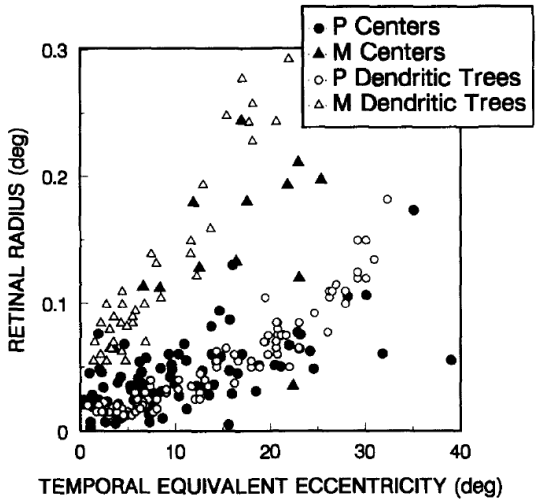

Figure 2.8: RF centre size variation compared with dendritic tree size with eccentricity in macaque.

Midget cells (parvocellularly projecting, hence P cells), and parasol cells (magnocellularly projecting, and so M cells) centre RF size radii, plotted against eccentricity. Also shown are data from a different study of midget and parasol cells dendritic tree radii, with eccentricity. The dendritic tree size of a specific cell type matches the corresponding RF centre size at a given eccentricity. Reprinted from Croner and Kaplan (1995).

properties vary widely; they receive input from bistratified retinal ganglion cells, and typically have large RFs, slow response latencies and conduction velocities, and also have blue/yellow colour opponency.

These cells are segregated into six layers in macaque LGN, the dorsal four layers primarily containing P cells, and the ventral two layers M cells. These layers are alternately contralateral (receiving input from the eye on the opposite side of the head) and ipsilateral (from the same side). In between each of these main layers are much thinner K layers, whose eye preference matches that of the main layer above it.

As LGN cell RFs are essentially determined by their retinal ganglion cell input, the range of receptive field sizes of macaque LGN cells closely matches that found amongst ganglion cells at the same eccentricity, also spanning around 3-4 octaves of SF preference (see figure 2.9). Although LGN cells are typically classified into different cell types, according to differences in receptive field properties, if we solely consider the size of LGN cell RFs at a given eccentricity we find a large, essentially continuous range (see figure 2.9), with a large overlap between different cell types (Levitt et al., 2001). The overlap is large enough that size alone may not be sufficient to distinguish between them.

Using difference-of-Gaussians (DoG) models to estimate centre/surround LGN cell (and retinal ganglion cell) spatial responses to stimuli has proven to be highly effective (Rodieck, 1965; for a more recent LGN example see Xu et al., 2001). The DoG remains the typical model for representation of linear LGN spatial receptive fields. Figure 2.10 shows a comparison of fitted difference-of-Gaussians models with M, P and K cell responses in owl monkey.

It is well established that the M, P and K pathways respectively project to layers 4Cα, 4Cβ and the cytochrome oxidase blobs in the upper layers of macaque V1 (Hubel and Wiesel, 1972; Lachica and Casagrande, 1992). Further weaker projections have been found, with the M and P pathways projecting to layer 6, the P pathway to layer 4A, and the K pathway to layer 1, but the projections to layer 4C provide the bulk of the LGN input to V1.

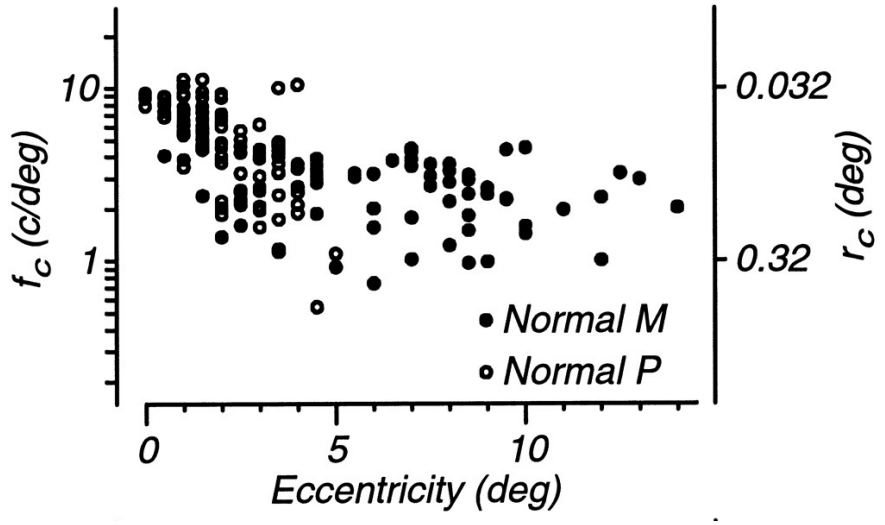

Figure 2.9: Spatial frequency preference for macaque LGN cells versus eccentricity. RF centre radius (right hand vertical axis)/ SF preference (left hand vertical axis) reduces as eccentricity increases. Note the large spread in values at a given eccentricity, and the large overlap between the cell classes. Reprinted from Levitt et al. (2001).

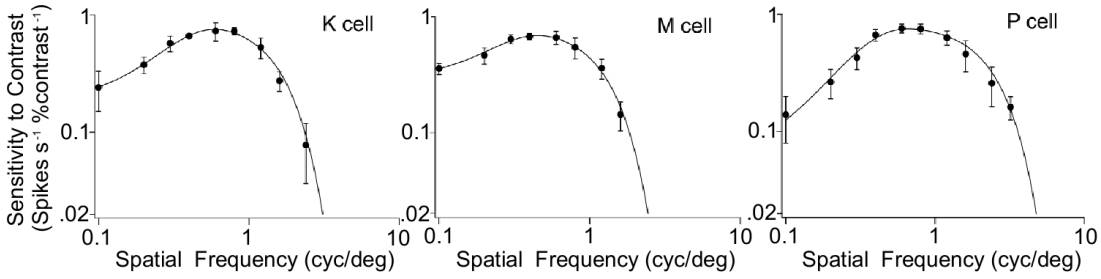

Figure 2.10: Difference of Gaussian fits for M, P, and K LGN cells in owl monkey. The responses of each LGN cell to achromatic drifting gratings at 28% contrast and different SFs. Cell data points are shown with one standard deviation error bars. Difference of Gaussians best fit to cells responses curves are shown in solid line. The DoG curves show a good fit to the data. Reprinted and modified from Xu et al., 2001.

Note that the retinal ganglion cell input provides only 10% of the synaptic contact to the LGN, while interneurons and feedback connections from V1 and other brain regions account for the other 90% (Sherman and Koch, 1986). The other input serves to modulate the LGN response to the driving input from retina. This modulation is not well understood, but is known to subtly influence the properties of LGN cells, such as whether their firing mode is burst or tonic (Sherman and Guillery, 1998).

2.2.3 Primate V1

V1 is a laminar structure that receives a variety of specific afferent inputs to distinct laminae, sublaminae and compartments. As mentioned in the preceding subsection, the bulk of the LGN input terminates in layer 4C, with a segregation between the magno and parvocellular streams into two sub laminae, 4Cα and 4Cβ respectively.

The size of simple-cell RF sub regions (and hence their SF selectivity) is thought to arise from correspondingly sized thalamic input RFs. In macaque, Angelucci and Sainsbury (2006) compared the RFs of simple cells in V1 with the visuotopic extent of LGN neurons whose axonal arbors provide thalamic input. They conclude that these regions are spatially coextensive, which provides further evidence to support Hubel and Wiesel's hypothesis concerning the formation of simple cell RFs from aligned thalamic input.

The SF preference of primate V1 neurons can vary across a range of as much as 4 octaves (see review in section 2.4), which matches the range measured earlier in the visual pathway. The bandwidth of SF tuning can vary greatly, from half an octave to more than 2 octaves (Valois et al., 1982). Functional SF segregation is thought to parallel anatomical structures to some extent, as neurons with lower SF preference tend to be confined to cytochrome oxidase blobs, and neurons with higher SF preference tend to be in the inter blob areas (Silverman et al., 1989). The spatial organisation for SF will be analysed in section 2.4, where it will be treated in detail.

As described in chapter 1 (e.g. figure 1.1), orientation selectivity is known to have

a columnar organisation. However, the degree of selectivity for features like OR can vary greatly depending upon lamina position (Gur et al., 2005). Consensus amongst investigators suggests that OR selectivity in layer 4Cα is high. However, the selectivity in 4Cβ is contested. Livingstone and Hubel (1984) found no OR selective cells amongst 124 cells measured in 4Cβ. However, Ringach et al. (2002) report that 25% of cells measured in layers 4Cα and 4Cβ are highly selective for OR, having bandwidths of less than 20°. Gur et al. (2005) also find high OR selectivity amongst cells in sublaminae 4Cα and 4C*m* (a suggested sublaminae between 4Cα and 4Cβ). In layer 4Cβ, all the cells had bandwidths greater than 30°, but Gur et al.'s sample size was small, only consisting of 9 cells. Vanduffel et al. (2002) suggest that columnar organisation for OR is maintained amongst cells that are OR selective in layer 4C, as they found iso-OR columns throughout the cortical layers. These columns are particularly strong for sublaminae 4Cα, and present but faint in 4Cβ.

Lateral connections in V1 extend primarily within a single layer, but are not strictly confined to that layer, and can extend into other layers (see figure 2.11). Axonal connections from neurons receiving thalamic input within layer 4C may extend to either or both sublaminae, as well as other layers of V1 (Yabuta and Callaway, 1998). As will be investigated in chapter 5, such interconnections may partially serve to reinforce feature preference mapping in different layers. For instance, this connectivity may lead to columnar mapping for OR preference. Patchy clusters of axon terminals have also been observed to extend not just horizontally but into multiple layers (Yabuta and Callaway, 1998). Although the input segregation to the layer 4C sublaminae is quite strict, deviations from this have been observed, with axons principally terminating in 4Cβ making small forays into 4Cα (Blasdel and Lund, 1983).

In addition to the LGN input, V1 receives a great deal of input from V2 and from higher visual regions. The targeting of these feedback streams is not well understood, but it is thought that they serve to modulate the activity of V1 cells, as is thought to occur in the LGN. Indeed, inactivation of V2 input to V1 has been seen to cause only a small change in the response properties of V1 cells and their RFs (Sandell and Schiller, 1982; Hupé et al., 2001).

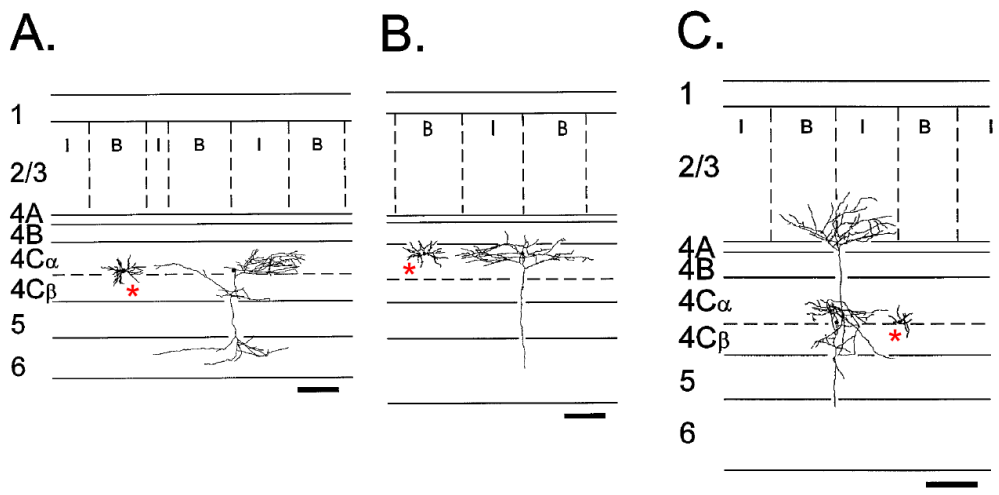

Figure 2.11: Lateral connections of neurons in layer 4 of macaque.

(A) Spiny stellate neuron with a cell body located in lower 4Cα. This neuron makes extensive axonal connections to 4Cα, but also some connections to 4Cβ and layer 6. (B) For this neuron, somata, dendritic, and axonal trees are all confined to 4Cα, with a little spillage to layer 4B and one long projection to layer 6. (C) Cell with soma based in lower 4Cα. This neuron makes many local projections to 4Cα, 4Cβ, and layer 2/3, and single projections to layers 5 and 6. A and C illustrate that even if there were a strict segregation in afferent input between the layer 4C sublaminae, these sublaminae have direct connections to each other. This could aid neurons with the same retinotopic position in different sublaminae in developing the same OR preference, leading to a columnar OR organisation. Dendritic trees are shown (marked with red asterix) horizontally displaced from the soma body, to aid distinction from axons, shown in thinner lines. I: interblob, B: blob. Scale bars are 200μm. Reprinted from Yabuta and Callaway (1998).

Of the primate visual system features listed above, the most important aspects that form the basis of the models in chapters 4 and 5 are the linear centre/surround RFs of the RGCs and LGN, the large variation in SF preferences of RGC and LGN cells, the feed-forward architecture to V1, and the separation of LGN channels targeting different populations of neurons in V1. Without each of these features, the models would be unable to account for the experimental results.

2.2.4 Cat retina

The domestic cat has much lower visual acuity, with far fewer (and larger) retinal ganglion cells than macaque. It has dichromatic daylight vision, with two types of cone photoreceptors. Aside from these factors, the essential circuitry of the cat retina (for review see Sterling, 1983) is functionally similar to that found in primate.

In cat, the principal types of retinal ganglion cell that project to LGN are called X, Y, and W cells. X cells have a narrow, bushy dendritic tree and correspondingly small RFs, with linear spatial summation, and have a sustained spike output to a maintained stimulus. Y cells have larger, sparser dendritic trees, also with larger RFs, with non linear spatial summation, and have a more transient spike output. W cells have a wide range of different properties, typically with large RFs that may be linear or nonlinear, and may include edge and direction-sensitive cells. These classifications can be somewhat blurry, and many other retinal ganglion cell types exist.

There are many similarities between the properties of cat and primate retinal ganglion cell types. In particular, there are notable response similarities between X and P, Y and M, and W and K cells. These similarities suggest that similar functional pathways may exist in these different species, although it is important to note there are also clear differences (such as the wavelength selectivity of macaque P cells, and that M cells mostly have linear spatial summation).

Figure 2.12 shows the great variation in dendritic diameter of retinal ganglion cells in cat. Figure 2.13 shows the RF centre diameter for retinal ganglion cells in cat, at differing retinal eccentricities. The RF centre size can be derived from the photoreceptor sampling circuitry, and is linked to the ganglion cell dendritic tree size. As for

primate, there is a large range in RF centre at a given eccentricity. This range may be somewhat lower in cat, at around 3 octaves (Peichl and Wässle, 1979).

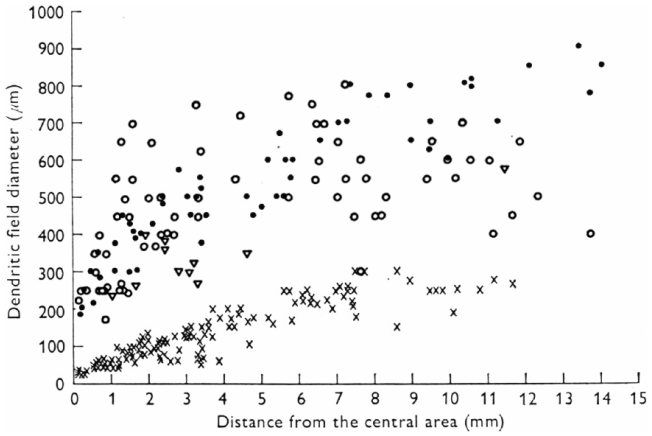

Figure 2.12: Dendritic field diameters of X, Y, and W cells with distance from area centralis. Key is as follows: $\times$, X cells; $\bullet$, Y cells; $\circ$ and $\bigtriangledown$, W cells. Modified and reprinted from Boycott and Wässle (1974).

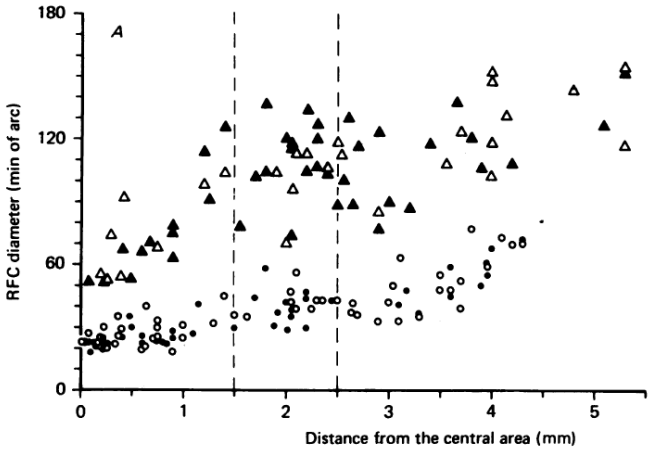

Figure 2.13: Receptive field centre sizes of retinal ganglion cells with eccentricity. Key: $\circ$, ON X cell; $\bullet$, OFF X cell; $\triangle$, ON Y cell; $\blacktriangle$, OFF Y cell. Note that RF centre size correlates with the dendritic field size (shown in figure 2.12), for a given cell type at a specific eccentricity. Also note the large range of receptive field centre sizes present at a given eccentricity. Reprinted from Peichl and Wässle (1979).

2.2.5 Cat LGN

The cat LGN is considered to be composed primarily of three neuron types, also called X, Y and W cells as their response properties are largely inherited from their retinal afferents. The population of cat LGN cells at a given eccentricity spans a large range of SF preferences, closely matching the range found amongst retinal ganglion cells. Cat LGN cells within 20° of area centralis are found to vary in centre RF size by up to around 3 octaves (for example see Cai et al., 1997).

The six LGN layers in cat are distinguished by the cell types predominant in them, and by the eye input they receive. Layers A and A1 receive contralateral and ipsilateral input respectively, and consist of a mixture of X and Y cells, with X cells making up around 80% of those in the area centralis (LeVay and Ferster, 1977). Layer C receives contralateral input, and is mainly composed of Y cells. Layers C1, C2, and C3 receive contra-, ipsi-, and contralateral input respectively, and primarily W cells are found in these thinner layers. All the layers are in retinotopic registration with each other, and as such there is a columnar organisation for retinotopy.

These different cell types project to different regions of V1. X and Y cells axons terminate in layer 4 and less commonly in layer 6. There is a bias towards X cells projecting primarily into the bottom of layer 4, and Y cells to the top (Ferster and LeVay, 1978; Lund et al., 1979), but X cell axon terminals have been found throughout layer 4 (Humphrey et al., 1985). W cells project either side of layer 4, and also less commonly to layer 1. Similar to primate segregation between the M and P pathways, X and Y pathways are partially segregated between separate sublaminae in cat V1, but the segregation may not be as pronounced in this species. In addition to these projections, Y cells also project to V2, and the bulk of W cell projections are to extrastriate areas.

Studies have found correlations between RF properties of simple cells in V1 and the RF properties of geniculate afferent cells that are connected to them. It is well established that retinotopic position of the RFs of these connected cells is shared. It has also been shown that if the sub regions of a cortical simple cell and the centre region of a geniculate cell are spatially coextensive and of the same sign (ON or OFF) than there is a high likelihood that the two cells will be connected (Reid and Alonso,

1995). Other response properties have also been shown to be important in determining connectivity, including temporal properties and the RF size (Alonso et al., 2001).

Weng et al. (2005) found a correlation between RF size and response latency in cat. They observe that the larger the cell, the faster the response time, and that this generalises for all three LGN cell types. If indeed this is the case, then it means that V1 may receive visual information relating to low SFs before high SFs. Thus a low resolution image may be available quickly in cortex, and then it rapidly sharpens with time.

2.2.6 Cat V1

What is known of the connectivity of cat V1 has many parallels with the macaque circuitry (for anatomical comparison see Lund et al., 1979; for review see Alonso, 2002). The principal thalamic recipient, layer 4, sends strong connections to layers 2/3. Most other cortical layers also project to layer 2/3. Although the specificity of targets is known to vary greatly, layer 5 projections generally target inhibitory neurons, while layer 4 targets excitatory pyramidal cells. Layer 2/3 projects to layer 5, and layer 5 projects to layer 6. Layers 5 and 6 both project to the superficial layers (1-3), and layer 6 completes the cycle by projecting to layer 4. How these vertical feedforward and feedback loops operate is not well understood, but they appear to allow for the emergence of complex cells in the superficial layers, and presumably help govern cell response properties along with the input from other regions of cortex.

In cat, selectivity for OR is high in all V1 laminae, including the input layers (Hubel and Wiesel, 1962), while macaque layer 4Cβ has fewer OR-selective cells (see section 2.2.3). It is not yet clear how strongly segregated the LGN input to V1 sublaminae may be in cat. Ferster and LeVay (1978) found a clear division in input in layer 4, but Humphrey et al. (1985) found only a weak and partial segregation. Whether weak or strong, the segregation may lead to a laminar distribution for SF preference. Figure 2.14 demonstrates a tendency for neurons deeper within layer 4 (which receives predominantly X cell input) to have small RFs, with RF size increasing while ascending the layer to the upper region (which receives predominantly Y cell input). Martinez et al. (2005) reservedly draw attention to this observation, as their sample size was not

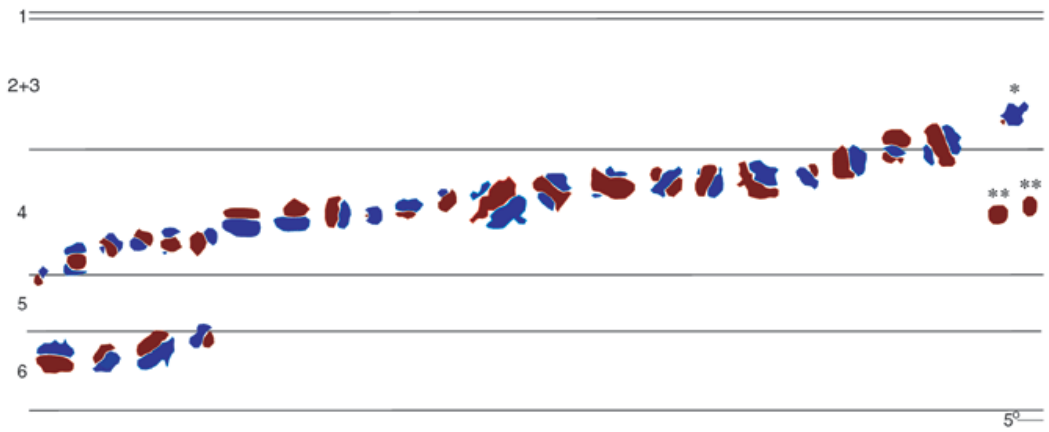

Figure 2.14: Laminar distribution of simple cell RFs.
Simple cell receptive fields shown with ON regions in red and OFF regions in blue. Note the overall trend from small receptive field size at the bottom of layer 4, to large RF size at the top of layer 4. The RFs are ordered left to right according to the depth of the soma. Reprinted from Martinez et al. (2005).

sufficient to draw strong conclusions.

Complex cells predominate in layer 2/3. A study temporarily blocking activity in a small region of LGN (by a local injection of GABA) blocked activity amongst simple cells in layer 4, and also blocked activity in layer 2/3 complex cells that had correlated spiking activity, matching OR preference, and matching vertical registration (Martinez and Alonso, 2001). This evidence provides support in cat for Hubel and Wiesel's suggestion of the generation of complex cells through the convergence of simple cells.

Spatial organisation for SF will be considered in detail in section 2.4.

2.2.7 Species comparison in relation to SF processing

As argued above, primates and cats have similar visual systems, particularly for processing of SF information. Both have a wide range of sizes of ON/OFF centre/surround LGN cells, which can be well represented by the difference of Gaussians model. Spatial summation of RFs is linear for most P and M cells in primate; for cat this is also

true for the X cells, which dominate the A layers of LGN around the area centralis. In the thalamic recipient layers of V1 in cat and primate, simple cells form the main population of neurons. Gabor functions have been found to accurately describe the spatial RFs of simple cells in cat (Jones and Palmer, 1987a) and macaque (Ringach, 2002). Observations suggest that the variation in RF properties found amongst simple cells may be described by a simple relationship between the parameters of the Gabor function (Ringach, 2002). Cortical neurons in both species reflect the range in SF preference found in LGN, and there is a lot of evidence to suggest that their RFs are formed from the superposition of the afferent LGN RFs that connect to them. However, the exact circuitry that may accomplish this remains elusive in both species. LGN cell RF properties and V1 cell RF properties are correlated in both species, which supports the idea that thalamic feedforward connections may primarily determine the preferences of simple cells (despite the more numerous feedback connections).

In both, there is a large range of SF preference amongst LGN cells at a given eccentricity, and evidence suggests that V1 cells are more likely to be connected to LGN cells with a similar SF preference to them. Moreover, the range of SF preference available to RFs if they were constructed from only a fixed size DoG cell may be limited. The response from a single sub-region of a simple cell is approximately linearly proportional to the area illuminated within it (this is a tenet of reverse correlation; Jones and Palmer, 1987b). Were a simple cell to be constructed from many small ON/OFF cells, then the response in a subregion would not be linear (see figure 2.15), but would be similar for any size bar stimulus falling within that subregion. This consideration in itself suggests that simple cells are constructed from ON/OFF cells with similar SF preference to the simple cell. Figure 2.16 illustrates an observed correspondence between an LGN cell's RF size and position, and the simple cell RF to which they contribute. These considerations suggest (but do not yet definitely establish) that LGN cell RFs are broadly similar in size to the width of the simple cell sub-regions of their target cortical cells.

Segregation of LGN input type into V1 sublaminae is also observed in both species. In fact, the nature of the separation at the laminar level is very similar, although evidence suggests it is stricter in primate, while a more complex partial laminar segre-

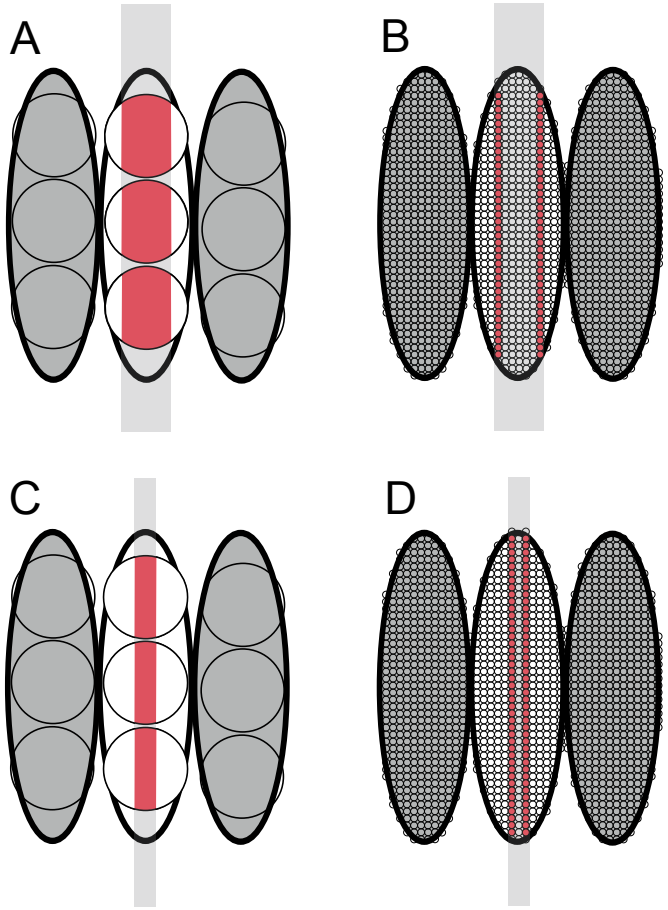

Figure 2.15: V1 simple cell construction from either large or small ON/OFF cells.
3 sub-regions (OFF, ON, OFF) of a sample V1 simple cell are shown as large ovals. Circles inside sub-regions are the centre regions of ON (in the inner sub-region) or OFF (the outer sub-regions) cells (surrounds are not shown). Grey rectangles are even illumination falling on the centre sub-region. Red regions indicate the proportional excitation of that ON/OFF cell. (A) An incident light bar produces a strong response in this simple cell, as a result of heavily activating its three constituent ON cells. (B) The same light source as A produces a much weaker response in this simple cell, due to the surround contributions from all the ON cells within the illuminated region cancelling their corresponding centre regions' positive contribution. As a result, only cells near the edge of the bar stimulus are activated. (C) A narrow bar of light elicits a much reduced response for the same cell as in A. (D) A narrow bar of light produces a very similar response for this cell constructed from many small ON/OFF cells as the response from a broad bar of light in B. Thus constructing V1 cells from LGN cells of similar SF preference can provide greater SF selectivity than if much smaller SF LGN cells are used.

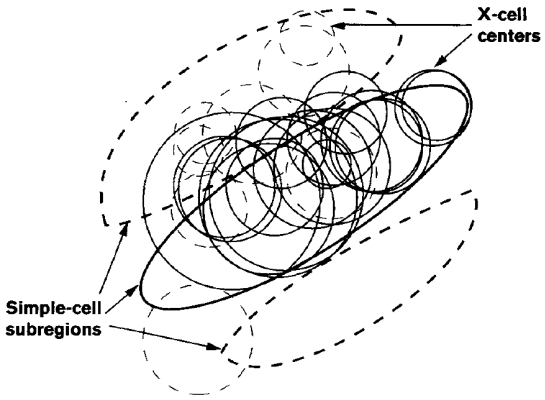

Figure 2.16: Contributing LGN cells and simple cell spatial RF plot in cat.
Simplified schematic representation of a cat simple cell RF with a central sub-region (solid line) flanked by neighbouring opposite sign (dotted lines) sub-regions. The circles indicate the relative size and position of the centre region of ON (solid circles) or OFF (dotted circles), cells that are monosynaptically connected to the simple cell. RFs were mapped using white noise stimuli, and connectivity was determined by cross-correlation analysis. Note that although there is some spread in size of the LGN cells, the overwhelming majority of them have a diameter very similar to the simple cell sub-region width. Reprinted from Reid and Alonso 1996.

gation may be present in cat. This segregation may represent a laminar organisation for SF preference in both species, although this is not well established. Whether the segregation between LGN afferents exists or not, there is a wide range in SF preference amongst LGN cells at a given eccentricity, so much so that SF preferences may overlap across cell types (particularly apparent in primate). The actual ordering in layer 4C (layer 4 in cat) may be based on RF size rather than cell type, although this hypothesis has not yet been tested directly. As mentioned previously, topographic maps for SF preference have been found in both species in the upper layers of V1; these are discussed later in the chapter (section 2.4).

2.3 Measuring SF organisation experimentally

What the form of cortical organisation for cells with SF preference might be, and if indeed there is any, has been widely debated over the years, with a number of conflicting studies and ideas contributing to the confusion (see recent review, Issa et al., 2008). Only very recently, with a number of new optical imaging studies, has a consensus begun to emerge as to the nature of SF organisation. In section 2.4 I will review published work relating to SF organisation in V1. First, however, I summarise the experimental techniques that have formed the basis of this research, as some of the disagreement regarding SF organisation stems from specific experimental procedures.

2.3.1 Electrode recording

Extracellular microelectrodes can have a sufficiently small diameter to allow the observation and recording of action potentials of a single neuron, i.e., single-unit recording. In studies relating to vision, single unit recording is used to determine the RF or feature preferences of single cells, by correlating neuron spiking activity with the stimuli presented (see for example Jones and Palmer, 1987b). The microelectrode can be moved to different locations to determine the properties of neighbouring neurons, or multiple microelectrodes can used to determine differences in response to identical stimuli simultaneously.

Multiunit recording is performed with a slightly larger diameter electrode, which allows the aggregate spiking activity of several neighbouring neurons to be recorded. Spike sorting procedures are often used to extract individual neurons' contributions (for review see Lewicki, 1998). These procedures are often based upon assuming that action potentials from different neurons have different characteristics that remain stable across time.

Although electrode recordings offer faithful measurement of electrical activity of small numbers of cells, it is difficult to gain a good picture of the global features of cortical organisation using electrode studies, due to the extremely small sample of recordings possible compared to the number of neurons in V1. Furthermore, a number of serious sampling biases can arise. For instance, because the technique is time consuming, acquiring data from many cells takes place over a long time period, during which responses of neurons may change due to adaptation, or changes in background chemicals. It is also likely that smaller neurons are sampled far more infrequently than larger neurons, which can lead to a false impression of the constituents of a region. The technique is also invasive, with every penetration causing damage to the tissue. The location of the electrode tip is also determined only *post mortem* , if at all, so it can be difficult to know if the electrode is positioned in the correct place for the intention of the study. Despite all of these limitations, extracellular recordings have proved extremely valuable in determining single-unit function in vivo.

2.3.2 2-Deoxyglucose marker

Increased cortical activity increases energy consumption, delivered through blood flow in the form of glucose. 2-Deoxyglucose (2DG) is a glucose analog that can be radiolabeled (Sokoloff et al., 1977). 2DG is taken up by neurons with similar kinetics to glucose, but is not fully metabolised, and instead builds up intracellularly. After 2DG is injected into the cortex, neurons that are more metabolically active will take up more of the label. The brain can then be sectioned, and autoradiographs can be produced by overlaying photographic X-ray film onto the slices. Regions that took up more 2DG will then show up as dark regions in the image. Thus this technique allows imaging of the relative local patterns of activity in a brain region over approximately a 20 minute period. In a vision experiment, if the animal is stimulated with a single type of in-

put pattern during the 2DG uptake period, then neurons that are more responsive to the stimuli will show greater labels of 2DG uptake than neighbouring neurons that respond less.

The main limitations of this technique are that comparisons of the response of the same region of cortex to different stimuli cannot be made in the same animal, chronic longitudinal studies are not possible as the maps are only available *post mortem*, and the spatial resolution is low. However, 2DG has proved useful in a number of studies relating to organisation, because it was the first technique to show activity over a large range of cortex, and it is currently the only imaging technique able to show activity across all cortical laminae.

2.3.3 Optical imaging of intrinsic signals

Optical imaging of intrinsic signals (for review see Zepeda et al., 2004; for more detail see Grinvald et al., 1999) is a relatively non-invasive method for observing changes in blood oxygenation levels through the associated change in the reflectivity of blood. Changes in blood oxygenation levels are proportional to shifts in local neuronal metabolic activity, and as such can be used to measure and map changes in brain activity at a high spatial resolution and across relatively large regions of brain surface. Optical imaging of intrinsic signals has proven an excellent tool in the study of cortical mapping, for a large variety of species. As it is relatively non-invasive, it also permits chronic studies, offering insight into how cortical feature preferences develop with time.

The experimental procedure is as follows. Visual stimuli are presented upon a monitor, which the anaesthetised animal is made to fixate upon (head and eye movement is typically eliminated by use of a stereotaxic apparatus). Craniotomy is often performed to allow removal of optically dense dura, to allow good visual data acquisition. The cortex is illuminated with light of appropriately chosen wavelength for the source of intrinsic signal required, and the cortex is filmed with a digital video camera positioned on a vibration free rostrum. Initially the camera is focused upon the cortical surface to record the blood vessel pattern, and then to record activity maps the camera is focused below the cortical surface in the region 300-700μm. Typically the change in

reflected light intensity as a consequence of neuronal activity is less than 0.1%, so statistical analysis methods are required to enhance this weak signal. There are a number of sources of noise that must be controlled for, such as respiratory and cardiovascular events producing low frequency oscillations in cerebral blood flow. Large trial numbers and synchronising picture taking with heartbeat and respiration can be used to reduce these artifacts. Stimuli are presented, and the reflectance level is recorded during and after presentation. To reduce noise, the recordings are averaged over multiple trials (typically 8-32 times; Grinvald et al., 1999).

To generate the reflectance change map, this "absolute" reflectance map must be subtracted from (or divided by) a baseline reflectance map. There are a number of approaches to provide this baseline map. The use of "blank normalisation", the cortical response to a blank screen, is the most obvious approach, but it can cause systematic errors, as regions that are not selective for a presented stimulus (and thus inactive) can still show an increase in baseline activity, due to an overall elevated level of blood flow in response to the stimulus. Another approach, the "cocktail blank", is generated by summing the cortical responses to the complete stimuli set and using this as the baseline map. This can have its own difficulties, as unless the stimuli evenly cover the stimulus space, and the sum of the responses to all stimuli are homogeneous across the cortex, feature maps normalised with this "cocktail blank" can be compromised. "Difference maps" are an elegant solution to this normalisation problem, and are suitable for orthogonal stimuli that produce non spatially coincident activity. These difference maps are obtained by subtracting the maps for orthogonal stimuli, resulting in maps for preference of one stimulus relative to the other. Careful attention must also be made to stimulus-independent artifacts such as vasculature. Often these artifacts have specific spatial and temporal signatures, and various analytical techniques can be used to remove them from the data.

Of course, optical imaging of intrinsic signals is not without its limitations. The spatial resolution is limited to about 20 times the size of a neuron, and the temporal resolution is also poor. Imaging is also restricted to layers near the cortical surface (the superficial layers), i.e. 1-3. The strength of optical imaging of intrinsic signals is that it allows us to view responses of large regions of visual cortex to specific stimuli.

The ability to do this has contributed greatly to our understanding of cortical map organisation and development, providing us with a wealth of new information and opportunities.

2.4 Chronological overview of studies of SF preference organisation.

Using the above techniques, many different researchers have investigated the organisation of SF preference in cat and primate V1. The following subsections examine each of these studies. Due to the controversy surrounding this aspect of cortical processing, this review is presented chronologically to provide context and to aid the understanding of the motivations behind these investigations.

2.4.1 Maffei and Fiorentini (1977)

Maffei and Fiorentini (1977) published the first in-depth study concerning SF organisation in cat V1. The activity of single neurons in response to presentation of sinusoidal gratings at optimal orientation was recorded in penetrations perpendicular or parallel to the surface of the cortex (see figures 2.17 and 2.18). It was observed that cells in the same lamina had a matching SF preference, whilst the SF preference differed for cells in a columnar penetration. This led to suggestion of a laminar arrangement for SF preference. However, an alternative reading of this study is that while the vertical penetrations may show a lack of columnar organisation over all of the V1 layers, they do not rule out a possible within-lamina topographic organisation. Indeed, of the 12 tangential penetrations performed, 3 are reported as representative: two are in layer 4 and have track lengths of about 1mm and 1.75mm, and the other may be in layer 2/3 and has a track length of about 3mm, but within that a homogeneous SF region of about 1.75mm. The authors report that “The longest rows of constant spatial frequency... were of the order of 2mm and therefore exceeded the estimated width of a hypercolumn”. However, as we will see in later work, it may well be the case that iso-SF regions may on occasion span as much as 2mm, and it may also be the case that the longer track lengths were primarily found outside layer 2/3 (where there is more evidence to support a laminar organisation), but this is not reported. Overall, I believe

that this work indicates that there are zones of iso-SF preference, but does not clearly demonstrate a laminar organisation.

2.4.2 Tootell et al. (1981)

Tootell et al. (1981) presented drifting fixed SF gratings to cats prepared using 2DG. They observed increased uptake of 2DG in a columnar pattern, extending through all cortical layers. When exposed to gratings at all SFs, no such columnar pattern was produced. In this experiment, nine animals were each presented with a single different SF in the range 0.25 to 2 cycles/deg, and at all orientations. For each animal, densely labelled columnar patterns were produced in V1 (see figure 2.19A). The columns extend perpendicularly to and across all laminae. Layer 4 labelling was found to be considerably darker than other layers, which the authors attribute to LGN afferents and first-order cortical processing. No further laminar 2DG density differences were found to be correlated with the SF of the stimulus. The distance between the SF columns was found to be about 0.8 to 1.0mm, about the same as is found for OD and OR columns. In a control experiment four animals were tested in the same fashion, but with exposure to the entire SF range; here no columnar organisation was observed (see figure 2.19B). This suggests that all cortical units were equally excited, so no overall differences in 2DG uptake occurred.

Cats presented with only high SFs had SF columns only within the central 5° of the area centralis projection region, despite the stimulus extending 18° peripherally on all sides. In the low SF case, the columnar patterns spanned the full 36°. This is consistent with findings that high SF-selective neurons are principally found in the area centralis region, and low SF-preference cells are found throughout, which helps affirm that the observed activity patterns are primarily a consequence of SF-preference-dependent activity.

The conclusion drawn from this study is that there is a columnar organisation for SF. Even though this evidence seems compelling, there are several important uncertainties. Due to the nature of this experiment, one animal can only yield one response map, and thus whether SF is continuously mapped across cortex cannot be addressed.

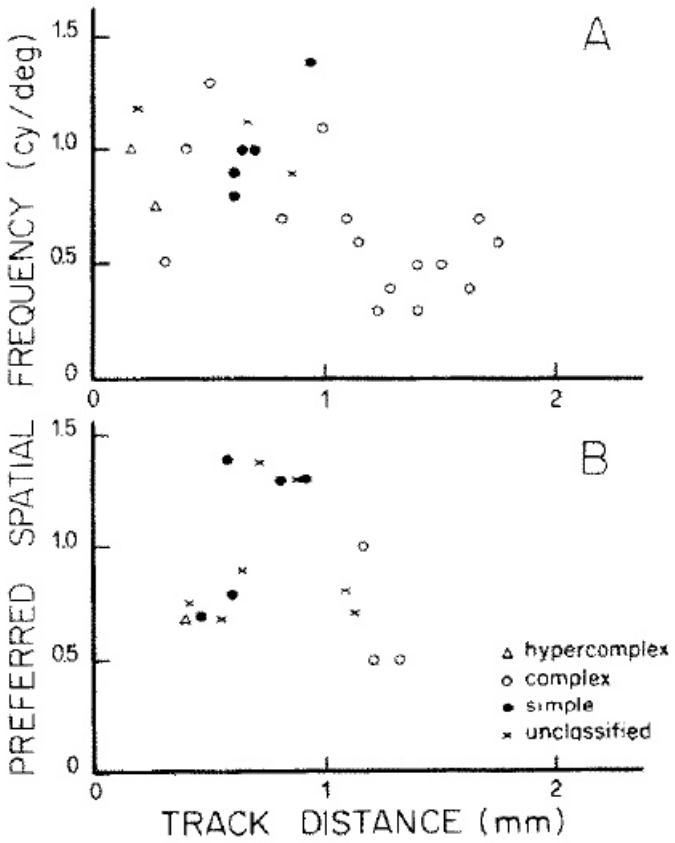

Figure 2.17: Maffei and Fiorentini (1977): Preferred SF of cells in cat V1, recorded in two penetrations perpendicular to cortical surface.

Penetrations were made within vertical and horizontal orientation preference columns (A and B respectively), and the cells' RFs were all within the *area centralis*. Over the length of each of the penetrations, the range of SF preference is high, suggesting that there is not a columnar organisation for SF. Reprinted from Maffei and Fiorentini (1977).

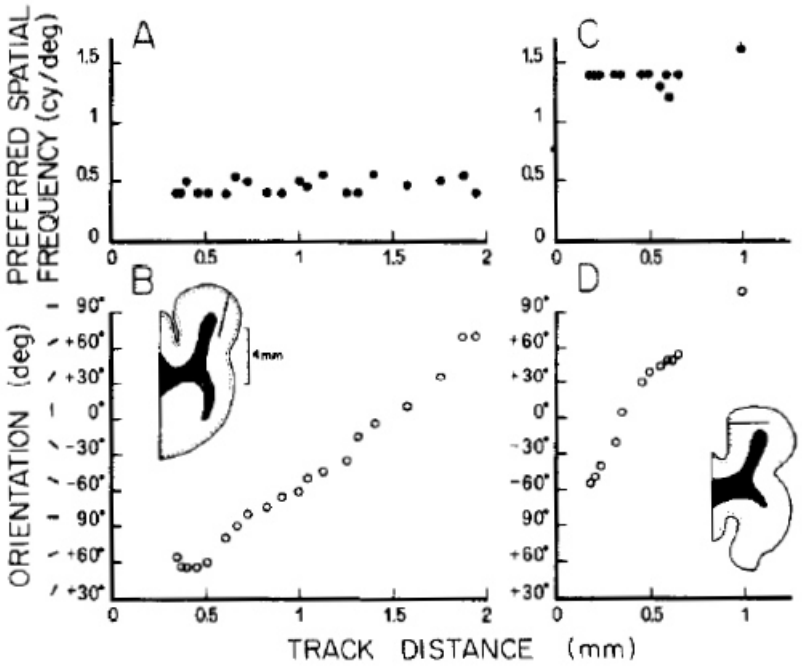

Figure 2.18: Maffei and Fiorentini (1977): Recordings of preferred SF and OR for cat V1 cells recorded in two separate penetrations tangential to the cortical surface.

(A) Penetration made in the coronal plane, near the medial surface; note the tight grouping of SF preference along this penetration. (B) Same penetration as A, showing gradual continuous shift in OR preference with penetration track length. (C) Penetration perpendicular to the medial plane; again a tight grouping in SF preference, but at a higher SF than in A. (D) Same penetration as in C; again there is a gradual continuous shift of in OR preference with penetration track length. These electrode studies provide some support for a laminar organisation for SF preference. Reprinted from Maffei and Fiorentini (1977).

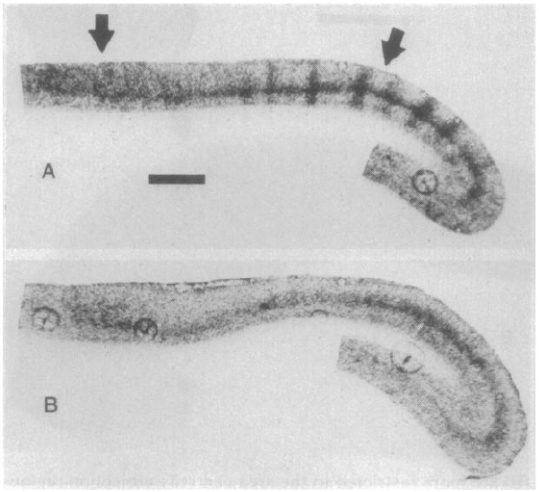

Figure 2.19: Tootell et al. (1981): Autoradiographs from cat V1.
(A) Cat A viewed a drifting single frequency sine grating at 2.0 cpd. (B) Cat B viewed a range of drifting gratings of spatial frequencies ranging from 0.25 to 2cpd. All gratings were displayed binocularly and over all orientations. Dark columns are clearly observed in the single spatial frequency case (A), but are not present in the multiple frequency case (B). The authors present this as evidence of a columnar organisation for spatial frequency in cat V1. Also note the high 2DG uptake in layer 4 in (A), but the absence of columns, which suggests a laminar organisation for SF in layer 4. Scale bar is 2mm. Reprinted from Tootell et al. (1981).

For instance, the observed maps could be the result of a binary (low and high only) mapping of SF preference. Additionally, the observed autoradiographs represent the average response of cortex to a specific SF averaged over all orientations, which can be different from an SF map constructed using optimal responses from preferred orientations (see section 2.4.10). Furthermore, it has also been suggested (Shoham et al., 1997) that the observed SF maps may in fact be related to the OD map and that these results are an artifact caused by a difference in acuity between the two eyes.

Even with these possible limitations, I consider this to be the first serious evidence for the existence of columnar organisation for SF preference. This columnar organisation would appear to extend through all layers, with the exception of layer 4. It is interesting to consider the increased 2DG uptake found in layer 4 in more detail, which I feel the authors dismiss too quickly. As cells with SF preference are known to exist

in layer 4, the uniform labelling suggests some laminar SF preference, although with this technique it is not possible to label sublaminae differentially to test this hypothesis directly. In any case, if the resolution of this experiment were greater, it could have allowed insight into whether there is a laminar organisation of SF preference in layer 4 of cat.

2.4.3 Berardi et al. (1982)

Berardi et al. (1982) conducted electrode penetration experiments in cat area 17 and 18. The investigation in area 17 focuses upon relationships between neurons separated by 200-300μm in layer 6. They find that cell pairs with large differences in OR preference tend to have small differences in SF preference, and that the converse is also the case. When the track path is taken into consideration, they find that large changes in SF preference occur more frequently in penetrations perpendicular to the cortical surface (35.3% of neurons have SF preference differing by more than one octave, from a sample of 105 neuron pairs), and less frequently for penetrations tangential to the cortical surface (3.8% of neurons have SF preference differing by more than one octave, from a sample of 335 neuron pairs). These findings are consistent with a laminar organisation for SF preference in layer 6. In examining the SF preference of neurons along track paths, they observe clustering of SF preference, but for different SFs at different depths in layer 6. They speculate that this change in SF preference may match with the electrode entering a different sublayer. It is reasonable to consider a laminar organisation for layer 6, because it receives afferent input from the LGN, and so performs first-order processing as in layer 4. An alternative interpretation is that these sharp transitions between SF preference regions may correspond to changes between SF preference columns or "blobs", as the electrode is moving obliquely through the layer.

2.4.4 Tolhurst and Thompson (1982)

A similar electrode penetration study of cat area 17 by Tolhurst and Thompson (1982) also investigated SF preference for neighbouring neurons. They observed that neurons

within about 100μm of one another had a greater likelihood of similar SF preference than neurons further apart, but even proximal neurons may have greatly differing SF preference. They could not reproduce long runs of neurons with the same SF preference along tangential penetrations (observed by Maffei and Fiorentini, 1977), possibly because they had difficulty keeping the penetration within a single lamina. Adopting a different approach, they made multiple penetrations perpendicularly to cortex and compared SF preference with laminar position across the separate penetrations, finding no clear pattern. They did observe iso-SF preference clusters of neurons with different SF preference within the same lamina, in different penetrations, which would not be expected with an overall laminar organisation. They reported little difference in the form of SF preference clustering observed in penetrations both tangential and perpendicular to the cortex, which is contrary to the findings of Berardi et al. (1982) in layer 6, although the spatial scale is different. Overall, the authors conclude that some form of local clustering (on a scale <200μm) for SF preference does occur, but they are unable to support either a laminar or columnar organisation. The SF preference clustering may be weaker than the organisation for OR preference; they frequently found neurons with very different SF preferences within a sample of neighbouring neurons.

2.4.5 Tootell et al. (1988)

Tootell et al. (1988) followed their earlier 2DG SF organisation study in cat with a 2DG study of SF in macaque. Using essentially the same methodology, they demonstrated a topographic organisation for high and low SF preference. Two new and striking features come from this work: The first is that regions of low SF preference appear to occupy the same positions as cytochrome oxidase blobs, whilst high SF preference regions surround these blobs. And the second is that a clear laminar variation in SF preference is reported between the 4Cα and 4Cβ sublaminae.

Presentation of low-SF gratings across all orientations (in 45° steps) produced high 2DG uptake in the cytochrome oxidase blob regions. The same pattern was found to extend through layers 2, 3, 4A, 4B, 5, 6 and possibly layer 1. This pattern appears to be a columnar arrangement for low SF preference, aligned with the cytochrome oxidase blobs. In the complementary experiment, presentation of high SF gratings produced

high 2DG uptake in the interblob regions, resulting in annular topography. This annular topography was also columnar, extending through layers 2, 3, 4B, 5, 6 and possibly 1 and 4Cα. These two patterns appear complementary to each other, although this observation could not be verified directly because the two maps were obtained from different animals. Repeating the experiments produced the same results, reducing the likelihood of inter animal variation. To further test this idea, they performed a split field experiment, whereby the visual field is split in two, presenting to one half a sine grating at low SF and at all orientations, and in the other half a grating at high SF and at all orientations. Figure 2.20 shows the preference for high SFs in the interblob regions, whilst low SFs are preferred in the blob regions, demonstrating complementary organisations for high and low SF preference in the same animal. In the control for these experiments, sine gratings at all ORs and SFs presented to an animal produced an even pattern of 2DG uptake, so the inference that these patterns are a result of SF preference organisation seems robust.

Investigating differences in laminar SF preference, Tootell et al. find that gratings of high SF produce greater uptake of 2DG in layer 4Cβ than 4Cα, while gratings of low SF produce the converse. High 2DG uptake was also observed in layer 4A on presentation of high SFs. These were the most notable differences, although other lamina differences were also observed on a lesser scale. These results are consistent with studies that suggest that the M pathway projects primarily to 4Cα and that the P pathway mainly projects to 4Cβ (but also to 4A); other labs also find much larger receptive field sizes in 4Cα than 4Cβ (see section 2.2.2).

Furthermore, 2DG uptake in cortex is not just determined by local position within a lamina, but by global position. The range of SF preference changes with eccentricity, with the foveal representation in cortex having neurons with very high SF preference, which are not found in the regions representing the periphery. The authors demonstrate that 2DG uptake produced by exposure to high SF gratings is greater in the foveal region, matching the previous eccentricity study in cat (Tootell et al., 1981).

These results argue for some form of dual organisation, whereby there is an overall columnar organisation for SF preference, possibly emerging from an initial laminar

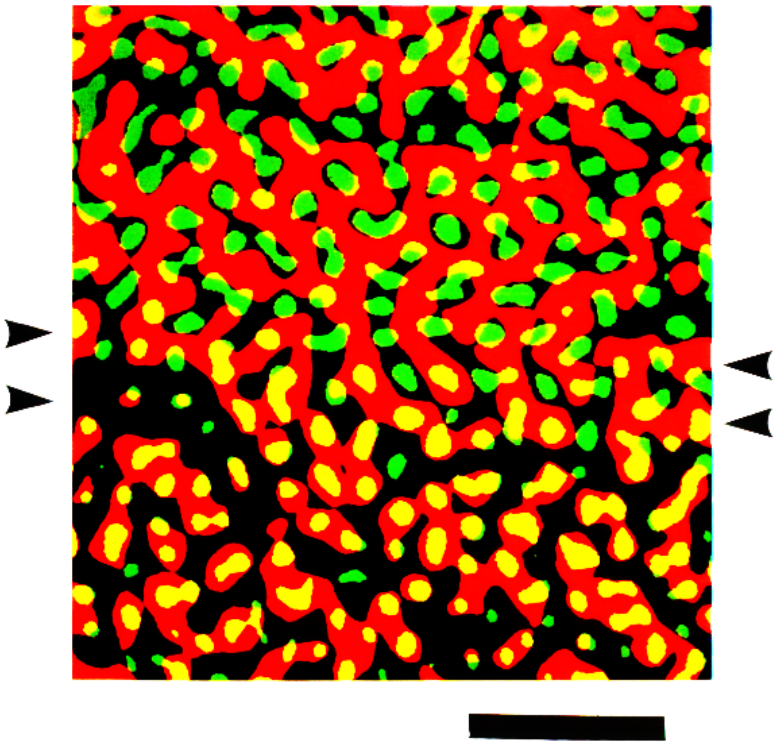

Figure 2.20: Tootell et al. (1988): Variation in 2DG uptake due to SF preference compared with cytochrome oxidase blob positions in V1 layer 3 of macaque.

High 2DG uptake regions are coloured red, high cytochrome oxidase (CO) stained regions are coloured green, and intersecting areas of high 2DG and high cytochrome oxidase are coloured yellow. The animal was presented split-field stimuli of sine gratings of 4.4cpd (medium-high SF) in one half of the visual field and 0.9cpd (low SF) in the other. The gratings were presented concurrently, binocularly and at varied orientations and drift speeds. The medium-high SF gratings resulted in the 2DG uptake in the region above the double arrows, while the low SF grating stimuli led to the 2DG uptake pattern below the double arrows. The region between the double arrows, due to a minor eye misalignment, received input from both high and low SF gratings. Note how in the upper portion of the figure areas of high 2DG uptake avoid the cytochrome oxidase blobs, while in the lower portion of the figure regions of high 2DG uptake coincide with the cytochrome oxidase blobs. Scale bar is 2.5mm. Reprinted from Tootell et al. (1988).

segregation in SF preference found between layers 4Cα and 4Cβ. The exact nature of the SF preference of the columns is not clear from these experiments, as it is not possible to distinguish between SF preference for the same region of cortex in the same animal. Due to this limitation, it cannot be determined whether within a high or low SF region the neurons are principally of a fixed SF, or whether they are clustered around an SF in a spatially random or organised fashion.

2.4.6 Silverman et al. (1989)

Utilising electrode penetrations in macaque, Silverman et al. (1989) found a strong correlation between SF preference of cells and the density of the cytochrome oxidase staining (figure 2.21). Electrode penetration recordings were made parallel to the cortical surface, near the V1 - V2 boundary, and in layers 2, 3, 5 and 6. Densitometry performed on a sectioned region of cortex along the electrode track path permitted the degree of cytochrome oxidase staining to be quantified, allowing direct comparison to SF preference of neurons measured along the same track. A remarkably high correlation is observed between these two measures – low and high SF-preference regions correspond with blobs and interblobs respectively. The periodicity of the organisation of SF was found to be in the region of 600 to 725μm. Importantly, this work also addresses the continuity of SF preference, showing that the SF preference varies continuously along the penetrations. The regions of most extreme SF preference (the highest highs and the lowest lows) are found in the centre of the interblobs and blobs. Together, these primate studies (Silverman et al., 1989; Tootell et al., 1988) provide evidence for a continuous, topographic, and columnar mapping for the organisation of SF preference in V1.

2.4.7 Born and Tootell (1991)

Born and Tootell (1991) perform a very similar electrode penetration study to Silverman et al. (1989), comparing SF preference with cytochrome oxidase uptake density in macaque. While they find a clear division in SF preference for neurons in blob (low SF) and interblob (high SF) regions, in contrast to Silverman et al. (1989) they observe

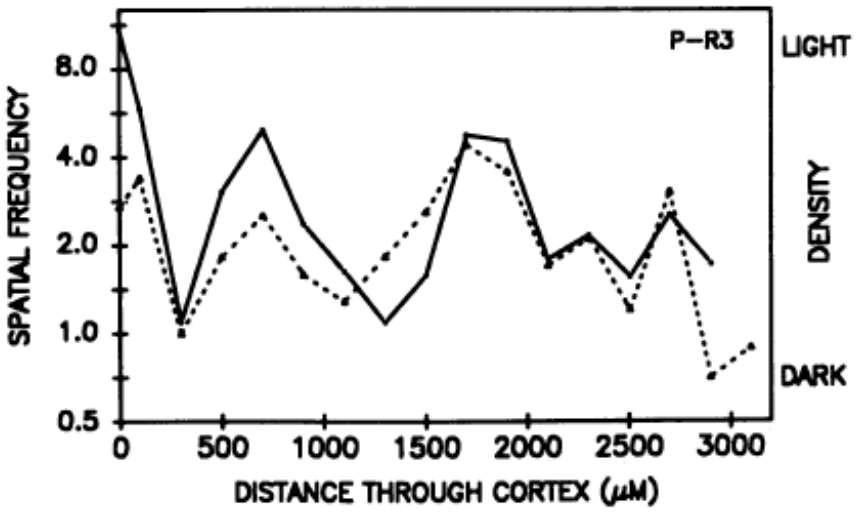

Figure 2.21: Silverman et al. (1989): SF tuning of cells and cytochrome oxidase density values with cortical position in macaque V1.

Peak SF tuning (dotted line); densitometric values for the cytochrome oxidase pattern (solid line). Correlation between the two measures indicates that cells with high SF tuning are found in interblobs, and cells with lower SF tuning are found in blobs. Reprinted from Silverman et al. (1989).

no systematic or continuous mapping. They speculate that this difference between their results and Silverman et al. (1989) may be attributable to differences in methodology. There are two main differences between these studies. Born and Tootell use single unit recording, while Silverman et al. used multiunit recordings. Born and Tootell suggest that interpreting multiunit data can be difficult for SF tuning, due to the overlap in SF tuning of different units and the difficulty in distinguishing individual cell responses. The second difference is that Born and Tootell used chromatic gratings to drive colour selective CO-blob cells (which made up 55% of all blob cells in their recording), while Silverman et al. used achromatic stimuli. Born and Tootell suggest that this factor may be sufficient to explain the differences in results. As achromatic stimuli are typically used to determine SF tuning, the use of chromatic stimuli in Born and Tootell makes it difficult to compare their results with other studies. In any case, both of these studies are in agreement upon the existence of high and low SF domains, and their alignment with cytochrome oxidase interblob and blob regions. Because of the differences in methodology and results, it is not possible to determine whether SF preference is con-

tinuous based on these electrophysiological results.

2.4.8 Bonhoeffer et al. (1995)

Bonhoeffer et al. (1995) studied the layout of functional domains in cat V1. They presented the first SF map acquired using optical imaging of intrinsic signals (see figure 2.22). Optical imaging of intrinsic signals has a significant advantage over 2DG studies, as it allows the global preference for multiple stimuli to be assessed in the same animal, although it cannot address questions of laminar organisation in layer 4. To determine the SF maps for the superficial layers, the authors present drifting gratings at different orientations and SFs and record the cortical maps generated. The SF maps were subsequently computed by summing all the OR maps of one SF frequency, and dividing by the sum of all OR maps of a second SF (i.e. a difference map, section 2.3.3). This method was later called into question by Issa et al. (2000), as it assumes SF preference and OR preference to be independent parameters, which is not generally the case. Issa et al. (2000) demonstrated that this analysis procedure can produce misleading results, not in line with the actual SF preference of cells (determined through elecrophysiology) in the map. Furthermore, use of the cocktail blank normalisation procedure polarises the map towards the extremes of SF preference.

2.4.9 Shoham et al. (1997); Hübener et al. (1997); Everson et al. (1998)

Using optical imaging of intrinsic signals, Shoham et al. (1997) found a binary mapping for SF preference in cat, within which the low SF preference regions also have a high temporal frequency (TF) preference, and the high SF preference regions have a low TF preference. They also find a correspondence between cytochrome oxidase blobs and regions of low SF preference (see figure 2.23). They suggest this is functional evidence supporting the continued segregation of the X and Y channels in cortex.

Further optical imaging studies investigating SF mapping (Hübener et al., 1997; Everson et al., 1998) use similar methods, averaging across all orientations to deter-

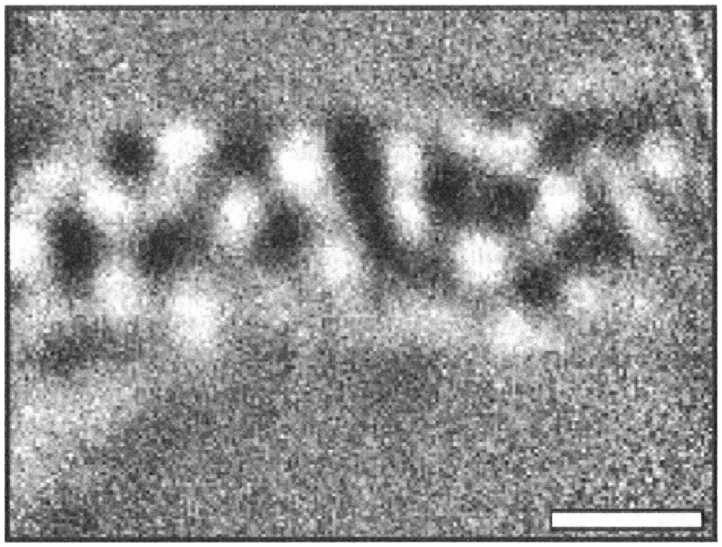

Figure 2.22: Bonhoeffer et al., 1995: SF map in cat V1.
SF map produced by dividing the sum of responses to high SF gratings at all orientations, by the sum of responses to low SF gratings at all orientations. Temporal frequency was fixed at 2 cycles/ second. Light areas indicate high SF preference, dark areas indicate low SF preference. Scale bar is 1mm. Reprinted from Bonhoeffer et al., 1995.

mine the SF map. Hübener et al. (1997) find a weak binary organisation (see figure 2.24) consisting of blobs of low SF preference in a sea of high SF preference, and with OR pinwheel centres falling in the centre of low or high SF preference regions. Everson et al. (1998) use a principal component analysis procedure on the intrinsic signal data to improve the signal-to-noise ratio. They consider the organisation to be continuous and topographic with SF pinwheels (see figure 2.25).

Overall, these studies (Shoham et al., 1997; Hübener et al., 1997; Everson et al., 1998) support a topographic organisation for SF preference. However, in light of the problems associated with the particular methods they use to calculate the maps, it is only possible to take these studies as support for there being an organisation for SF, rather than allowing meaningful conclusions to be drawn about the form of the SF map or its relations to other maps.

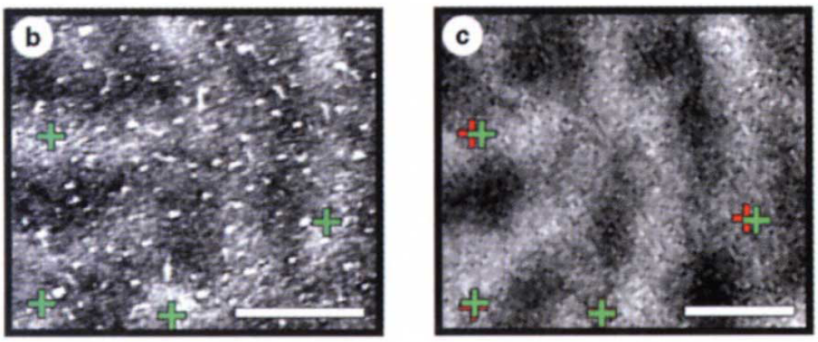

Figure 2.23: Shoham et al. (1997): Cytochrome oxidase blob positioning compared with SF map in cat V1.

(Left) Cytochrome oxidase stained region of V1. Cytochrome oxidase blobs are dark regions. (Right) SF map for corresponding region of cortex. Note correspondence between cytochrome oxidase blobs and regions of low SF preference. Low/0.2cpd SF preferences regions are dark; light regions are high/0.6cpd SF preference. Green crosses mark dye injection sites from the histological section, and red crosses mark the estimated position of the dye injection sites from the *in vivo* blood vessel pattern. The dye injections are used for aligning the histological sections with the functional maps. SF maps were produced by summing all iso-orientation maps of one SF, then dividing by the sum of all iso-orientation maps of the second SF. Scale bars are 1mm. Reprinted from Shoham et al. (1997).

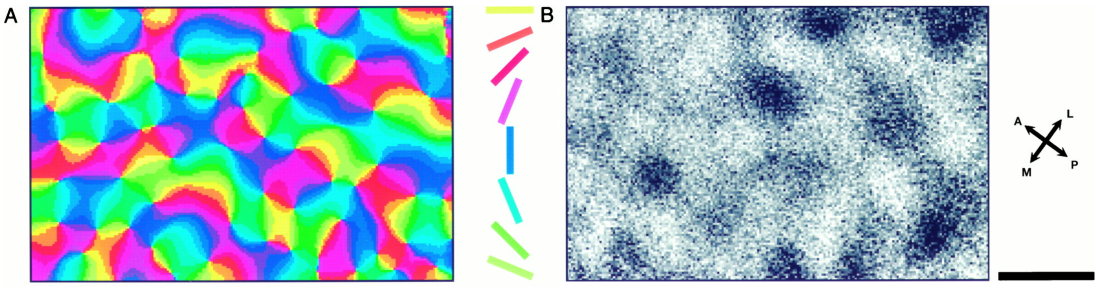

Figure 2.24: Hübener et al. (1997): Corresponding OR and SF maps in cat V1.

(A) OR preference map. (B) SF preference map for the same region of cortex. The SF map is calculated by summing iso-OR maps for a single SF, then dividing by the sum of iso-OR maps from a second SF. This method generates a SF preference map, which states how much more a pixel prefers high (light regions) to low (dark regions) SF stimuli. This approach is now known to polarise the data, producing a binary map. Scale bar is 1mm. Reprinted from Hübener et al. (1997).

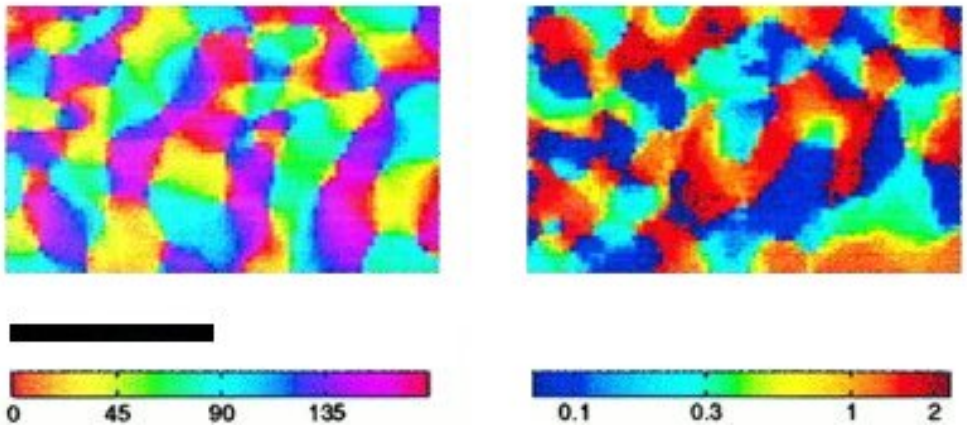

Figure 2.25: Everson et al. (1998): Corresponding OR and SF maps in cat V1. Left hand map shows orientation preference. Right hand map shows SF preference for corresponding region of cortex. SF preference is in cpd. Scale bar is 2mm. Reprinted from Everson et al. (1998).

2.4.10 Issa et al. (2000)

Issa et al. (2000) perform the first study using optical imaging of intrinsic signals that, instead of averaging SF preference across all orientations, uses the optimal SF response at the preferred OR for that particular pixel to generate the SF map. This is essentially the analogue of the method used to determine SF tuning in electrophysiological studies. Using this approach, a patchy continuous organisation for SF preference is observed in cat V1, with correlation between OR pinwheel centres and low and high SF regions (see figure 2.26). The range of SF preference of the domains is high, but domains with very different SF preference are found sufficiently close to one another that their arrangement conforms to the requirements of hypercolumn organisation and therefore good coverage. The authors verified the SF maps by using targeted microelectrode recordings.

Issa et al. (2000) observe that OR preference is independent of the SF used during testing, whilst SF preference is dependent upon which OR is used during testing. This relationship seems quite reasonable when viewed in the light of the typical shape of simple cell receptive fields in V1. These receptive fields are often described as Gabor patches, which have a clear symmetry and orientation. Any SF grating can be used to

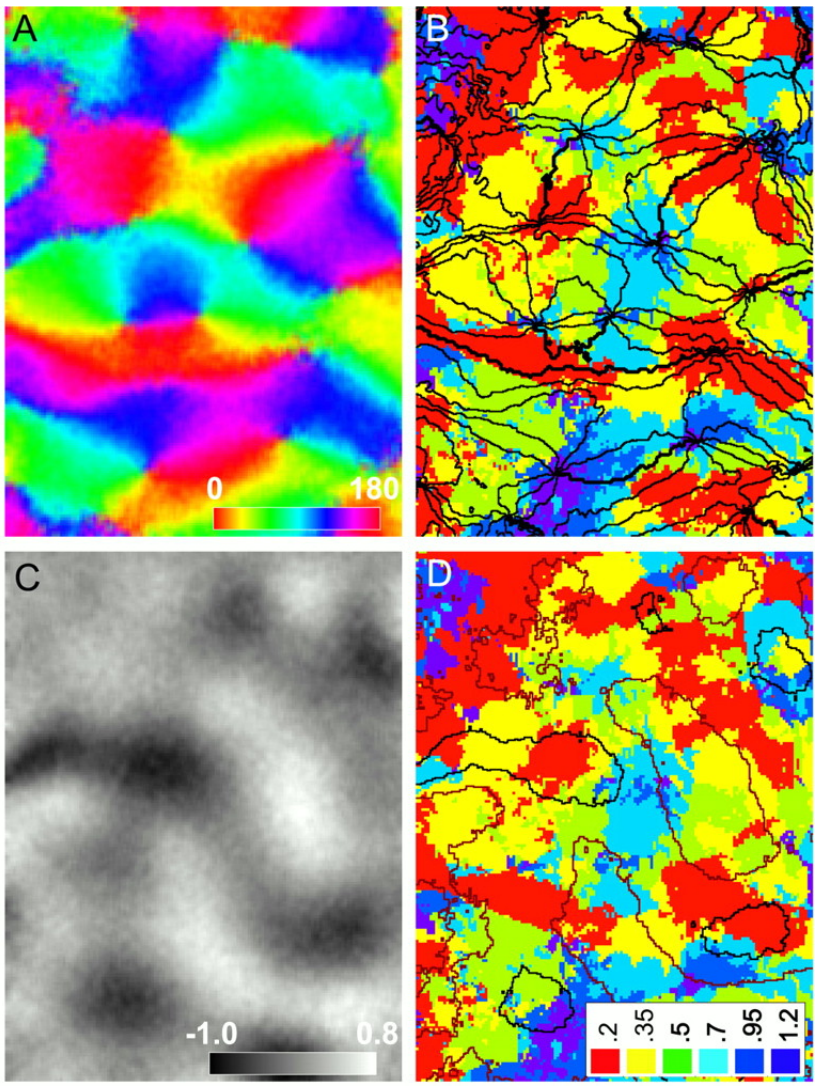

Figure 2.26: Issa et al. (2000): Corresponding OR, SF and ocular dominance maps, in cat V1. (A) OR preference map. (B) SF map with superposed iso-orientation lines. Iso-orientation lines converge at pinwheel centres. Note the tendency for pinwheels to fall on regions of extreme SF preference (e.g., high/purple or low/red). (C) Ocular dominance map; light areas are contralaterally dominated, whilst dark areas are ipsilaterally dominated. Scale bar shows the deviation of pixel intensities ($\times 10^4$) from 1.0. (D) SF map shown in B, with ocular dominance contours superposed. SF preference was defined for a given pixel as the SF of the stimulus that produces the greatest activity, out of all presented stimulus combinations. Colour bar length in A is 1mm. Reprinted from Issa et al. (2000).

determine OR preference from such a receptive field, if tested at all positions (phases). However, unless the preferred OR is used, erroneous values for SF preference will be found, even if the result is averaged over all ORs. The authors demonstrate this effect, and conclude that their definition for SF preference of a pixel being the stimulus that best activates the pixel from all ORs and SFs presented is the most appropriate. To back up this claim, the authors perform the analysis methods used in previous optical imaging studies of SF organisation (Shoham et al., 1997; Hübener et al., 1997; Everson et al., 1998) using their own data, and succeed in producing SF maps with the same features as described in these previous studies (binary, or continuous with SF pinwheels). The authors also criticise previous studies for their use of cocktail blank normalisation, which can bias SF measurements. Cocktail blank normalisation was originally introduced to avoid artifacts in OR maps measured using vector averaging. Normalising by the sum of all responses helps to remove spurious signals resulting from the increased overall level of light absorption that occurs for any visual stimulus, even in regions not actually responsive to stimuli (and in which there is no increase in neural activity). However, for SF measured using a discrete set of frequencies, behaviour of the light signal in unresponsive regions does not affect the calculation at all — only maximally activated pixels are retained from all the different single condition maps, and no averaging is performed. Issa et al. (2000) argue that the previously reported maps were misrepresentative of the actual organisation of SF preference, due to artifacts created by the analysis procedure.

Issa et al. succeed in unifying previous optical imaging studies by clarifying the process for determining SF preference. They also characterise the organisation for SF preference in the superficial cortical layers of cat as a patchy but locally continuous topographic mapping, which is generally supported by subsequent experimental work.

2.4.11 Sirovich and Uglesich (2004)

One exception to the emerging consensus is Sirovich and Uglesich (2004), who perform optical imaging of intrinsic signals in cat, and find no organisation for SF preference. They suggest that SF preferences in cortex are arranged as a random mixture of high and low SF populations, and that no systematic topographic organisation is appar-

ent (see figure 2.27). The authors believe preceding optical imaging results may have been compromised not only due to the averaging of SF preference across different orientations raised by Issa et al. (2000), but also by the failure to remove a blood-related artifact arising from the vascular bed, termed the "nonspecific response", from the data. The authors outline a procedure to remove this nonspecific response, and then process their data using singular value decomposition, which seeks to represent the data in a more concise form, thus reducing noise. This analysis results in a two SF channel representation of the information, whose spatial arrangement (see figure 2.27) is regarded by the authors as not organised. Furthermore, the authors re-analyse previously published results with their methodology, and find them to be consistent with their interpretation of SF mapping being random.

Mallik et al. (2008) later identified problems with the analysis procedure used by Sirovich and Uglesich, primarily the imposition of two basis functions upon the data, which they consider to be unjustified and inappropriately fitted to continuous data. They also point out that the SF map constructed by Sirovich and Uglesich (2004) (see figure 2.27) does appear to have some clustering for SF preference, and is not dissimilar to maps observed by other investigators.

Sirovich and Uglesich's suggestion that previous optical imaging of intrinsic signals studies are flawed due to blood flow artifacts is also countered in a later study by Mallik et al. (2008) that demonstrates topographic SF maps by using an alternative imaging technique (flavoprotein autofluorescence) that is less susceptible to vasculature artifacts. Furthermore, Mallik et al. (2008) also use optical imaging of intrinsic signals on the same cortical region and observe matching and statistically similar SF maps, thus demonstrating that the results due to optical imaging of intrinsic signals are genuine and not an artifact of vasculature.

Sirovich and Uglesich's study stands alone among recent studies in its claims for a random organisation for SF. In light of the broad consensus found amongst subsequent research, as well as Issa et al.'s re-analyses of the earlier optical imaging studies concerning SF organisation, and the criticism directed at Sirovich and Uglesich's analysis, it is reasonable to conclude that the actual organisation is not random.

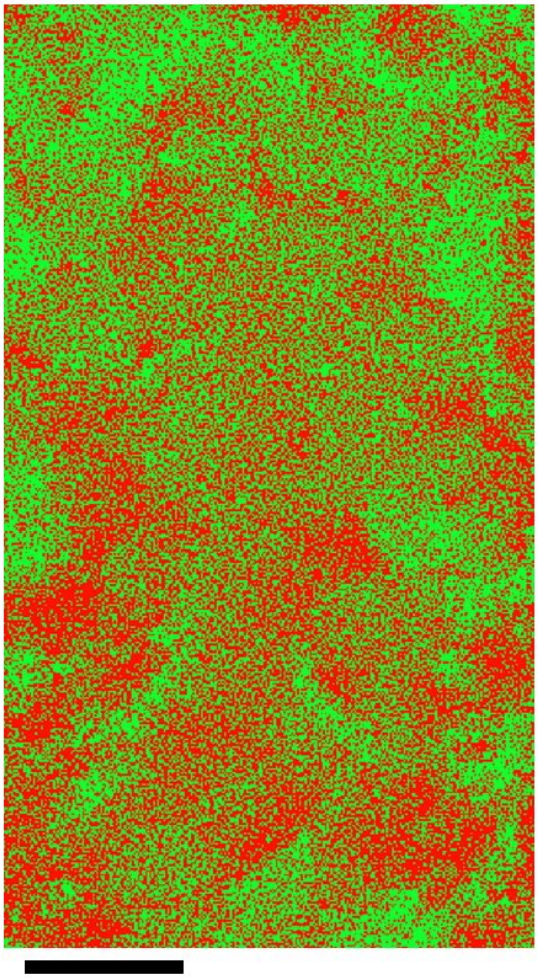

Figure 2.27: Sirovich and Uglesich (2004): SF map in cat V1.

To construct the above plot, each pixel measured using optical imaging has been subdivided into 25 subpixels, which were generated with a random colour (red for low SF preference, and green for high SF preference) with the appropriate probability, to reflect the possible underlying contributions from neurons in the high and low SF channels that collectively make up the optical imaging response at that pixel. This procedure results in what the authors argue is a random organisation for SF preference, although other authors consider this map to have evident clusters of low and high SF preference even after this artificial imposition of randomness. Scale bar is 1mm. Reprinted from Sirovich and Uglesich (2004).

2.4.12 Yu et al. (2005)

Yu et al. (2005) investigate the relationship between different feature maps using optical imaging of intrinsic signals in ferret, and what underlying principles may govern this relationship. They construct their SF map (see figure 2.28) in the same fashion as Hübener et al. (1997), by averaging over all ORs presented and subtracting the response to high SFs from the response to low SFs. The result is an SF map with high and low preference SF domains. While this analytical method has been discredited in other species due to the dependence of SF upon OR, the authors argue that in ferret SF preference is largely independent of OR, and that in this instance there is little difference in map form.

Yu et al. also construct maps for OR, ocular dominance and retinotopy. They find that their results are consistent with dimension reduction models (see section 3.2.2), and that a region of high gradient in one feature map tends to coincide with low gradient regions in the other feature maps. Where these high gradient regions do overlap, they have a tendency to do so orthogonally. The retinotopic map in ferret is strongly anisotropic, and this anisotropy is also found in the feature maps, where their isodomains are elongated such that the high gradient axis is orthogonal to the high gradient axis of the retinotopic map. This work is the first to show some form of SF organisation in ferret, and suggests that continuity and coverage are key aspects of the organisation of V1.

2.4.13 Molotchnikoff et al. (2007)

Molotchnikoff et al. (2007) test some of the claims from Sirovich and Uglesich (2004), performing an electrode penetration study on sets of neighbouring cells in the superficial layers at the area 17/18 border of cat. The authors find that only about half of nearby neurons share similar SF preference, and argue that there is thus no overall organisation for SF preference. Extracellular recordings were made at 26 sites, and spike sorting was utilised to distinguish responses from different adjacent neurons. Assum-

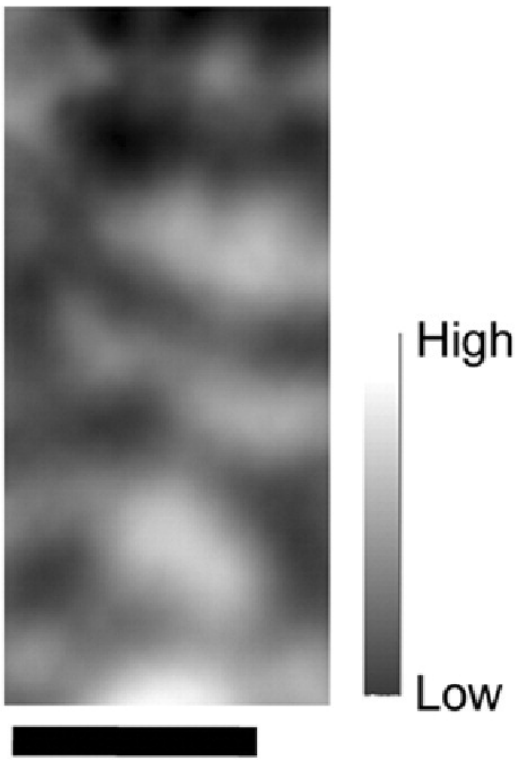

Figure 2.28: Yu et al. (2005): SF map in ferret V1.

The SF map was generated by subtracting the averaged response (over all orientations) to high frequency stimuli, from the averaged response (across all orientations) to low frequency stimuli (similar to Hübener et al., 1997). Scale bar is 1mm. Reprinted from Yu et al. (2005).

ing that the neurons in a single penetration share the same OR preference, the authors determine the compound receptive field (the sum of classical receptive fields of the individual neurons), and then calculate the preferred OR of the group of cells. They then present sine gratings at this fixed OR and a range of SFs, using sorted responses to these stimuli to determine SF preference for each of the neurons. However, is likely that the preferred OR of the compound receptive field is not shared by all of the other cells, and if so this procedure will incorrectly estimate the SF preference of those neurons. A previous similar study conducted by DeAngelis et al. (1999) found a strong clustering of SF preference for neighbouring neurons, but determined the neurons' SF preferences in the more conventional manner, by optimising OR and SF for each cell. In any case, SF preference is widely considered to be distributed, such that even within SF domains a wide range of SF preference can be observed. Thus these findings do not undermine the remaining body of research.

2.4.14 Zhang et al. (2007)

Using optical imaging of intrinsic signals, Zhang et al. (2007) studied how complex images are represented in SF domains of cat V1. They produce a simple binary SF preference map (see figure 2.29), by solely presenting two SFs. Although the work is not specifically focused on measuring SF maps, they observe a preference of the low SF domains for higher temporal frequencies and the high SF domains for lower temporal frequency preference. The investigators test the hypothesis that cortical response to complex stimuli (summed sine gratings of low and high SF and square waves), can be predicted by a linear sum of responses to the individual components of the complex stimuli. In other words, whether knowledge of the OR, SF and TF components that make up the complex stimuli is sufficient to predict cortical response, provided the corresponding feature maps are known. When the complex stimuli are presented they find that both high and low SF domains are activated, while if only a single sine grating is presented, only the appropriate SF domain is activated. The authors find that predictions generated using this linear filtering method are in good agreement with their experimental observations, and conclude that linear filtering in conjunction with feature maps is sufficient to predict responses to complex stimuli under the conditions they studied.

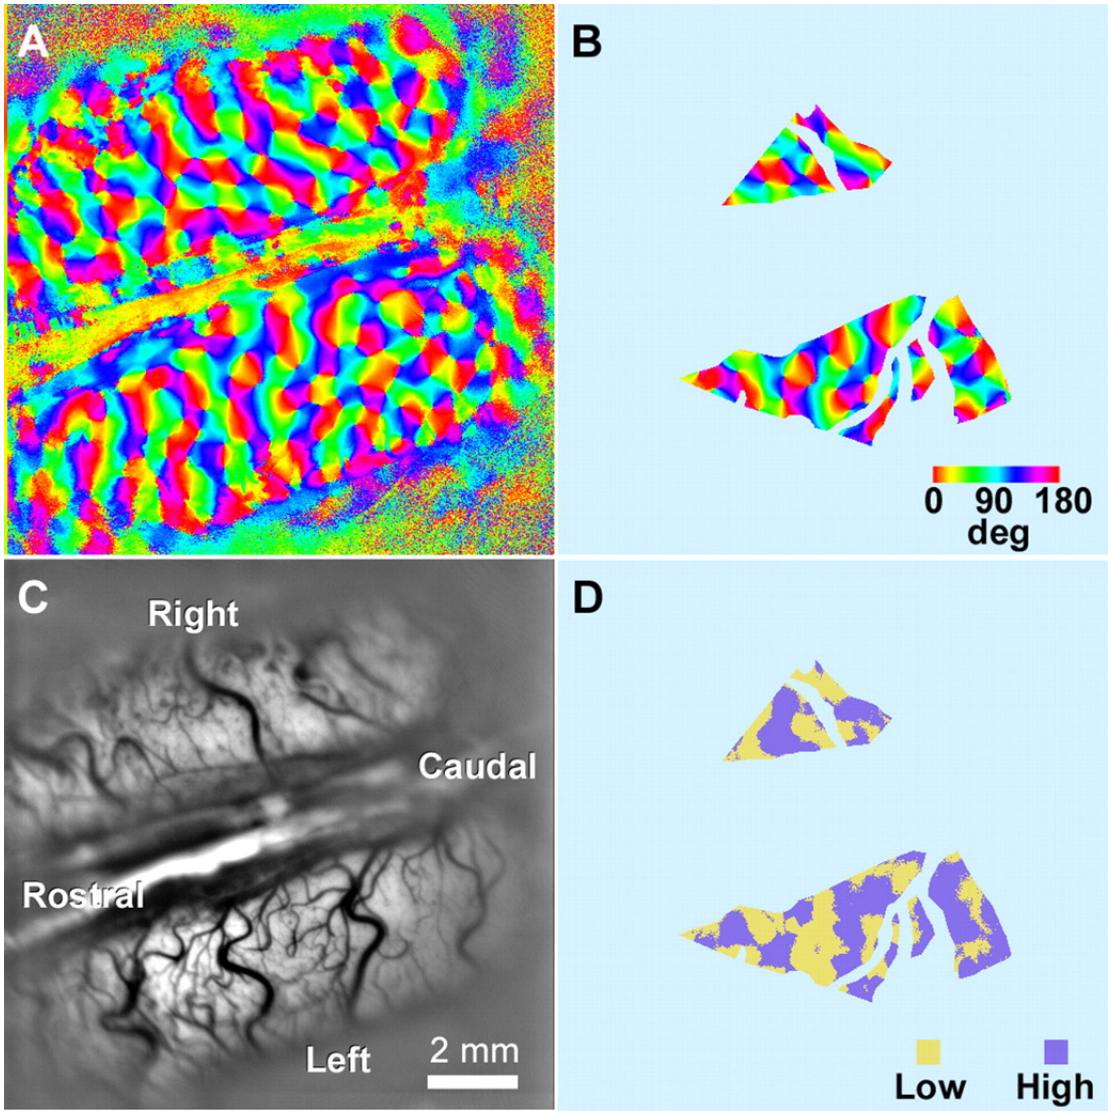

Figure 2.29: Zhang et al. (2007): OR and SF maps, in cat V1.

(A) OR preference map of both hemispheres. (C) Vascular pattern. (B) OR map with template (in light blue) overlaid. The template is used to exclude regions that are non-responsive or are likely to be compromised by vascular artifacts. (D) SF map with template overlaid. SF map was measured using two SFs: 0.3 cpd (low), and 0.9 cpd (high). Reprinted from Zhang et al. (2007).

2.4.15 Farley et al. (2007)

Farley et al. (2007) focus on how feature map organisation and interrelationships may alter with the removal of one feature map during development. They find that removal of the ocular dominance map through monocular enucleation in newborn ferrets results in a different form of organisation for the remaining feature maps. The measuring procedures are the same as used by Yu et al. (2005), and the SF map took a similar form (see figure 2.30). Removal of a feature map is found to result in an increased wavelength of the iso-value regions in the remaining feature maps. Predictions generated by a dimension reduction model (see section 3.2.2) were found to be in agreement with these experimental observations. The authors conclude that the number of dimensions being represented in cortex affects the nature of the intermap relations, and that this is predictable with the dimension reduction model. Note that SF is just one of several dimensions examined, and it is mentioned only briefly.

2.4.16 Xu et al. (2007)

In the first optical imaging study in primate to consider SF mapping, Xu et al. (2007) consider a number of organisational relationships pertaining to eccentricity. SF maps were constructed in accordance with the method of Issa et al. (2000), and reveal a continuous topographic organisation for SF preference across V1 in bush baby (see figure 2.31) and in owl monkey (see figure 2.32). In both species, more cortical area was found to be devoted to mid-range SF preference than to high and low. The overall arrangement of SF domains found is irregular in comparison to the form of OR organisation. They observed that OR iso-domains tended to span several SF domains, and that OR pinwheels tended to occur in iso-SF domains across the whole range (not only maxima, as previously reported in cat). This arrangement allows for full coverage of the feature space, in that most combinations of SF and OR can exist within a local area of the cortex. With increase in eccentricity, they found that the proportion of cortex devoted to low SF preference increases, as is expected from the known decrease in acuity further away from the foveal region. Cytochrome oxidase blobs were also found to be in good alignment with regions of low SF preference, confirming earlier findings (Tootell et al., 1988). These observations are broadly consistent with cat studies, and

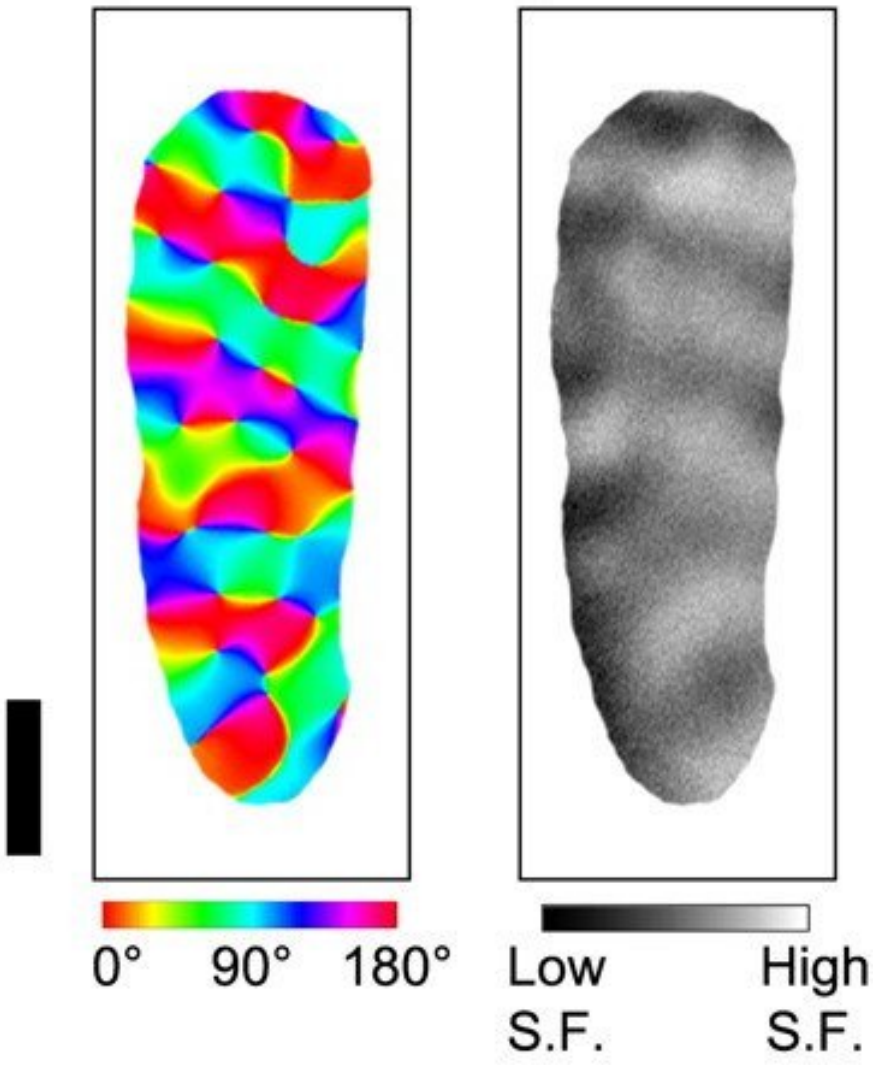

Figure 2.30: Farley et al. (2007): Corresponding OR and SF maps in ferret V1. Left map shows orientation preference. Right map shows SF preference for the same cortical region. The SF map was generated by subtracting the averaged response over all directions to high frequency stimuli from the averaged response to low frequency stimuli (similar to Hübener et al., 1997). SFs used to generate the map were 0.08 and 0.225cpd. Scale bar is 1mm. Reprinted from Farley et al. (2007).

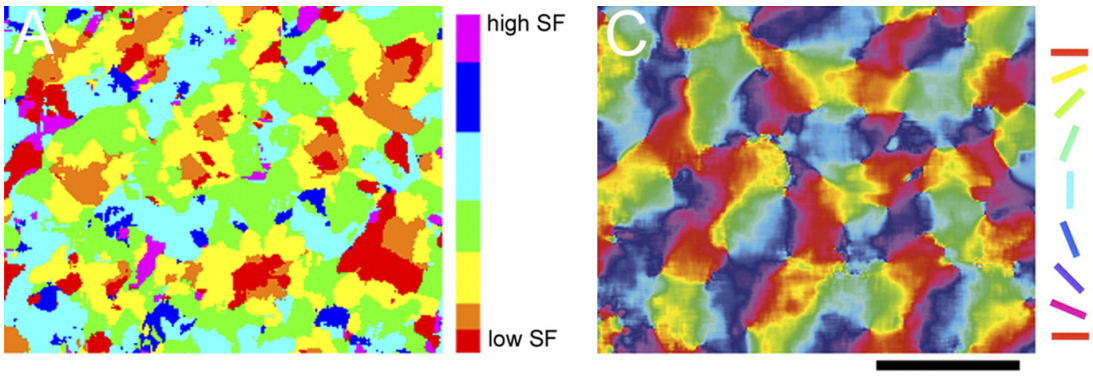

Figure 2.31: Xu et al. (2007): Corresponding SF and OR maps in bush baby V1. Left map shows cortical SF preference. SF preferences are colour coded as shown in key, where red is 0.1 cpd, orange is 0.2cpd, yellow is 0.5 cpd, green is 0.8cpd, cyan is 1.2 cpd, blue is 1.6cpd, and magenta is 2.0 cpd. Right map shows OR preference for the same region of cortex. Scale bar is 1mm. Reprinted from Xu et al. (2007).

add to the bulk of evidence supporting the existence of SF maps in a number of species.

2.4.17 Mallik et al. (2008)

As mentioned previously, Mallik et al. (2008) further investigate the validity of optical imaging of intrinsic signals as a means to obtain SF maps in cat. They compare optical imaging of intrinsic signals to an alternative and entirely independent technique, autoflourescence imaging, which they demonstrate to produce quantitatively similar results (see figure 2.33). Small differences do exist between maps produced using the different techniques, primarily due to differences of resolution. Autoflourescence imaging has a higher spatial resolution, so the map appears noisier (given that SF preference organisation is weaker than OR organisation). The artifacts also differ, as optical imaging of intrinsic signals suffers from blood flow artifacts (the nonspecific response), while autoflourescence imaging does not). See section 2.4.11 for further information.

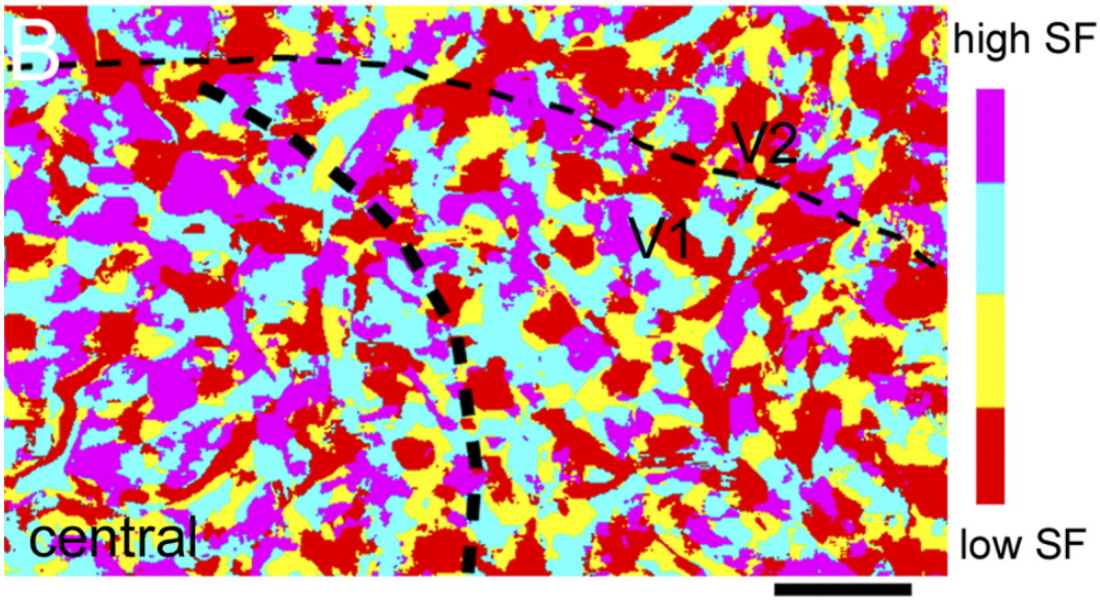

Figure 2.32: Xu et al. (2007): SF map in owl monkey V1 and V2.

Narrow dashed line indicates V1/ V2 border, and the thick dashed line separates the central and paracentral regions of V1. Note the greater representation of high SF preferences in the central region. Colour key indicates SF preference, where red is 0.2 cpd, yellow is 0.6 cpd, cyan is 1.2 cpd, and magenta is 2.4 cpd. Scale bar is 1mm. Reprinted from Xu et al. (2007).

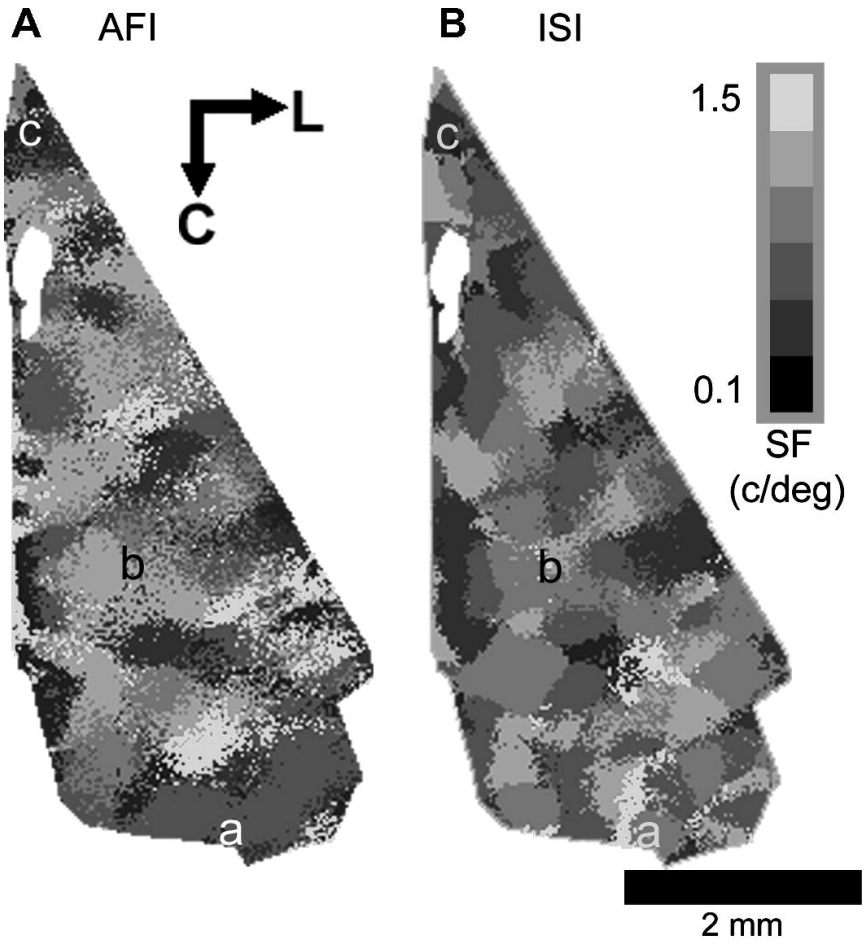

Figure 2.33: Mallik et al. (2008): SF maps in cat V1.

(A) SF map measured using autoflourescence imaging. (B) SF map for the same region of cortex measured using optical imaging of intrinsic signals. Note the broad correlation between the two maps. Reprinted from Mallik et al. (2008).

2.4.18 Purushothaman et al. (2009)

Purushothaman et al. (2009) demonstrate a method for interpreting optical imaging data called receiver operating characteristic (ROC) analysis. This analysis allows for the production of differential probability maps, whereby the probability of discriminating between two presented stimuli, given the reflectance values at that pixel, is mapped for all image pixels. They compare this method to standard difference maps and show that it has a higher signal to noise ratio, feature selectivity variance is accounted for (unlike in difference maps), and that resultant feature domains are more discernible. The authors conclude that it is a more sensitive analysis technique. As part of their investigation of this technique they measure SF maps in bush baby. They compare SF maps created using this ROC approach and the approach of Issa et al. (2000) (where the stimulus that produces the greatest activity in a pixel is regarded as the preferred stimulus). They find that the ROC approach yields maps that are less noisy and with less fractured domains. Optical imaging response to high frequency stimuli is very low, and its signal is easily lost when the single condition map is blank screen normalised. Due to the improved signal to noise ratio, the ROC procedure allows the high-frequency values to be discerned. With this increased sensitivity, they find larger high SF domains than using the Issa et al. (2000) approach (see figure 2.34). Indeed, on visual inspection there appears to be a more even distribution between SF preferences in the maps produced using ROC analysis.

In agreement with this lab's previous studies, they find a continuous organisation for SF preference across V1, with patchy iso-SF domains loosely tiling this cortical region. The authors suggest that the quality of these results is much higher than in previous studies, due to the sensitivity and power of the ROC analysis method, and as such, the most recent results are strong evidence for a continuous topographic mapping of SF in the superficial layers of V1.

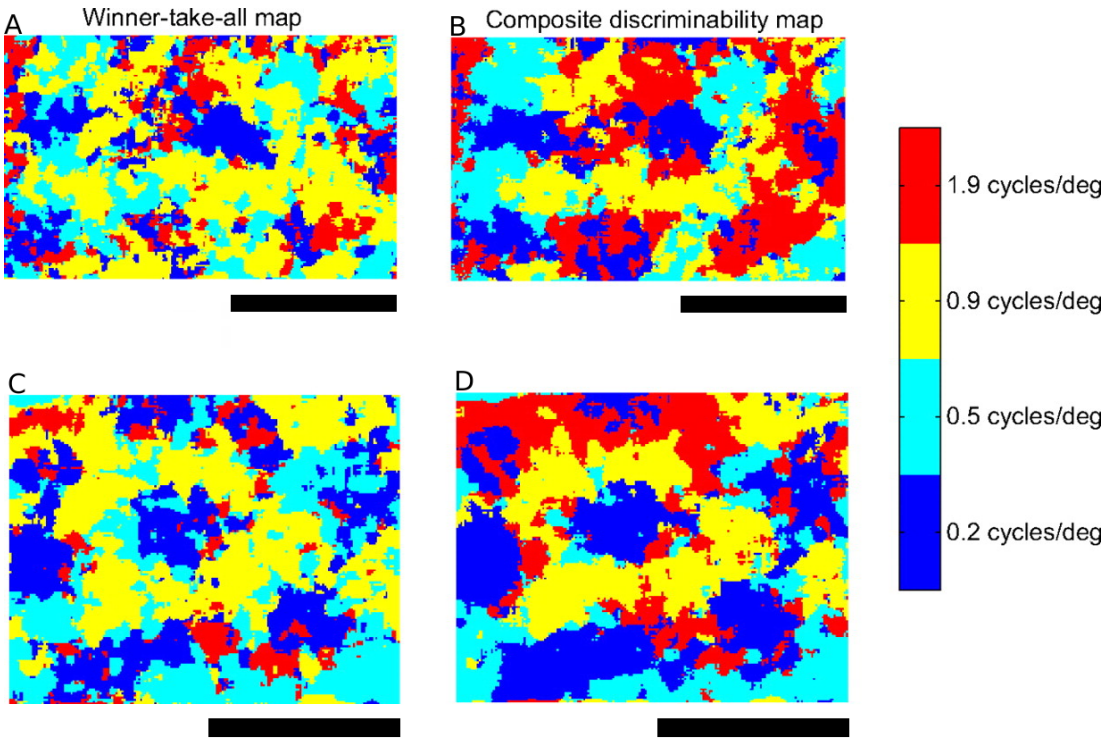

Figure 2.34: Purushothaman et al. (2009): SF maps for two regions of V1 in bush baby, determined by different methods.

(A and C) SF maps determined using winner-take-all method from Issa et al. (2000). (B and D) SF maps the from same data, determined using discrimination probabilities. Note that the composite discriminability map has more high-SF domains, and the structure of the map is somewhat less fragmented. Scale bars are 1mm. Reprinted from Purushothaman et al. (2009).

2.4.19 Summary of SF mapping experiments

Overall, the experimental evidence is heavily in favour of the existence of topographic SF maps in the superficial layers of V1. The electrode and 2DG studies broadly support this view, but due to the limitations of these techniques, it is difficult to draw clear conclusions from them alone. This problem is further exacerbated by the organisation for SF being somewhat weaker than for OR, which has contributed to early difficulties in measuring mapping for SF. More recently, all but one of the optical imaging studies find a topographic organisation of some form, and the variation in organisational structure found in the earlier studies (as well as the conflicting study of Sirovich and Uglesich, 2004) are now understood to be due to methodological differences. It appears that SF has continuous, patchy, topographic mapping in all the species investigated. Some authors report coincidence between OR pinwheels and iso-SF domains, but studies continue to disagree about this point, and the precise relationship remains unclear.

The mapping in deeper layers (not accessible to current imaging techniques) remains unresolved. The little experimental evidence that exists concerning the organisation for SF in layers 4 to 6, suggests a laminar organisation, but further investigation is required.

2.5 Development of feature preference organisation in V1

The development of cortical organisation for feature preference has been widely studied in a number of species. Three key elements have been found to drive the development of topographic maps: molecular guidance cues, spontaneous neural activity, and visually driven activity (for reviews see White and Fitzpatrick, 2007, and Huberman et al., 2008). These processes work together to create the intricate circuitry that supports cortical feature maps. While they are not entirely independent of one another, certain mechanisms are known to play a more significant role in the development of particular feature preferences than others. Retinotopic mapping is finished before eye opening,

and is dependent upon molecular cues and retinal waves (Cang et al., 2005). OD columns are formed before eye opening and are dependent upon spontaneous retinal activity; blocking of this activity disrupts the organisation of OD columns dramatically (Huberman et al., 2006). Visual experience is not required for OD and OR maps in cat; these maps can be found in kittens at 3 weeks whether their eyes were open or not (Crair et al., 1998).

The ferret, which has a visual system very similar to the cat, has been used extensively to examine the development of cortical OR maps. It has been observed that OR maps are present before eye opening (Chapman and Stryker, 1993), which takes place at around PND 30 (post natal day 30). While the neurons selectivity matures with development until it attains adult levels, the maps organisation itself is stable. Dark rearing experiments demonstrate that OR maps can develop in the absence of any light-mediated visual experience (White et al., 2001). However, without visual experience normal levels of adult OR selectivity are not reached.

Even though initial development of OR maps can occur without visual experience, spontaneous neuronal activity is required. Silencing the spontaneous activity of ON RGCs during PND 21-48 prevents the formation of OR maps (Chapman and Gödecke, 2000). Furthermore, abolition of V1 neural activity with TTX prevents the selectivity of neurons from reaching maturity (Chapman and Stryker, 1993). These experiments demonstrate that spontaneous neural activity is required for the formation of OR maps, and that visual experience is needed for normal development of neural preference, such that adult levels of selectivity are attained.

No work has been published about the development of SF maps, and the time course of their development is not known. This is likely to be due to the difficulty in measuring SF maps, which will be made significantly more difficult before maturation of SF selectivity. In a dark rearing study of SF selectivity in cat (Derrington, 1984), individual neurons were found to develop SF tuning spanning the full octave range (albeit with a lower mean SF preference and larger bandwidths, than in normally reared cats) at 4 weeks post natal, despite having no visual experience. This is not entirely surprising, given that OR preference develops in the absence of visual experience due

to spontaneous activity. A neuron that is highly selective for OR will also be highly selective for SF, particularly if it predominantly receives input from LGN cells with similar SF preference (see section 2.2.7). It seems reasonable that these two selectivities may arise together, and if that is the case then their feature maps may also.

It is likely that the distribution of SF preferences found in an SF map reflects the range of SF information present in its stimulus set. It is not known at what point in development SF maps first arise, or whether SF map formation is driven by spontaneious activity or natural images. Furthermore, it is not known how different types of retinal ganglion cells (and thus SF channels) are involved in retinal waves. As there is no data relating to the development of SF maps, the training stage within these simulations is somewhat artificial, consisting of only natural images. As visual experience is necessary for the maturation of both selectivity and OR maps it seems the appropriate stimulus to use. This thesis will consider how these preferences can develop, using input-driven, activity-based, self-organising cortical models, providing predictions to be tested in later experiments.

2.6 Conclusion

This chapter has identified and discussed attributes of the visual system that are essential for processing and representing SF information. The first stage of SF processing occurs in the retina, whereby the retinal circuitry effectively filters the image for different SF components. The retinal signal is conveyed to V1 via the LGN, which does little to the spatial difference-of-Gaussians representation of the information. In V1, simple cells have Gabor patch RFs with SF preference that broadly matches the SF preference of the afferent LGN cells that provide thalamic input. These simple cell RFs are thought to arise by superposition of aligned thalamic RFs. They are spatially organised by SF preference in cortex, possibly laminarly in layer 4C as a consequence of segregated thalamic input, and certainly topographically in the more superficial layers. This organisation appears to provide good coverage across the visual field, allowing for full representation of SF information in an image. Subsequent chapters will examine how this organisation can arise during the process of development.

Chapter 3

Theoretical approaches to the development of SF preference

3.1 Introduction

There has been a great deal of success in reproducing topographic feature maps using computational models (for reviews see Swindale, 1996, and Goodhill, 2007). Retinotopy, orientation preference, and ocular dominance maps have frequently been investigated. However, cortical spatial frequency preference representation and organisation has not previously been the main focus of computer modelling work. Where SF preference has featured, it has tended to arise from non-biologically based models, usually formulated at quite a high level of computational abstraction. For example, the elastic net algorithm (Durbin and Willshaw, 1987), which forms the basis of one of the dimension reduction models to include SF preference amongst its dimensions (Carreira-Perpiñán et al., 2005), consists of three equations with no physiological foundation. Despite this, these models have often succeeded in producing SF map results qualitatively similar to those observed experimentally.

Any model must be abstracted to some degree. A successful model is founded upon the minimum essential features of the system under consideration, and within that framework succeeds in capturing the behaviours of the system that the investigator is concerned with. From the experimental background to cortical SF processing in chapter 2, there are three clear experimental observations that a model incorporating

SF processing needs to be able to reproduce, if we are to attempt to understand the requirements for SF preference to develop in V1:

- It must develop a topographic map for SF preference organisation like those found in V1 layer 2/3 (section 2.4).
- There must be a large range of SF preference, spanning the input space (section 2.2.7).
- Its neurons must have realistic RFs (section 2.2.7).

These three key criteria (subsequently referred to as the SF modelling criteria) must be fulfilled by a model in order for it to account for the experimental data. I aim to develop a model that incorporates the important aspects of the visual system architecture, and results in these criteria being fulfilled, in order to drive future experiments to help determine the essential features of SF processing in animals.

In the following sections these previous computational investigations into the development of cellular SF tuning and organisation are introduced and described. How well they fulfil the SF modelling criteria, and how they relate to the modelling work in this thesis, is then discussed.

3.2 Previous models for SF preference

3.2.1 Sparse coding models

Sparse coding models arose from the desire to explain the form of the RFs of simple cells in V1. It is thought that V1 evolved to represent natural images in an efficient fashion, and that the shape of RFs has evolved to achieve this. One way to encode information efficiently is via a sparse code, wherein neurons in a spatial region fire sparsely and independently (Barlow, 1989). With sparse activity, features in the input space can be efficiently coded by a small set of neurons from a large assembly of neurons with differing feature preferences. Independently active neurons are less likely to be redundant, in that they do not have highly correlated activity to other neurons. Based upon the idea that images have a “sparse structure”, and that any image can be constructed

from a small number of basis functions from a large set (Field, 1994), the SparseNet model (Olshausen and Field, 1996) learns basis functions for natural images subject to the conditions of sparseness, independence and information preservation. The set of basis functions is determined by using an iterative approach to minimise the error in the image representation, while maximising the sparseness of activity. This is done by starting with randomly initialised basis functions, and by incrementally altering the basis functions, by performing gradient descent upon a cost function until the function is minimised. Figure 3.1 shows the results of this model for a natural image data set.

The form of the RFs is determined by the statistics of the input data set. Natural images result in Gabor-patch-type RFs, while sparse gratings (superposed sine gratings) yield grating-type RFs, and sparse pixels result in RFs where only one pixel is significant. The statistics of natural images coupled with the requirements of sparseness, independence, and image preservation are sufficient to produce localised, orientated, and bandpass RFs. The SF preference range of the receptive fields in Olshausen and Field (1996) spans about three octaves, and this range of RFs is sufficient to represent the image input data set through linear superposition, to an error in image reproduction of about 10%.

However, more detailed consideration of the RF properties in cat and macaque (Ringach, 2002) reveals a more diverse range of RFs than is produced with the SparseNet model. The RFs produced under SparseNet take the form purely of edge detectors, being somewhat more elongated than animal RFs, and without any low pass, "blob-like" RFs. This problem is thought to arise by the choice of the type of sparseness imposed upon the system. Rehn and Sommer (2007) replace the "soft" sparseness used by Olshausen and Field (1996), which limits the average neuronal activity, with a "harder" form of sparseness, whereby the total number of neurons that can be active is limited. This approach yields RFs even more closely matched to those found in animals (see figure 3.1).

These sparse coding approaches are formulated at an abstract level without any specific implementation of biological processes that could lead to sparse RFs. However, from a functional point of view, the outcome in terms of receptive field form

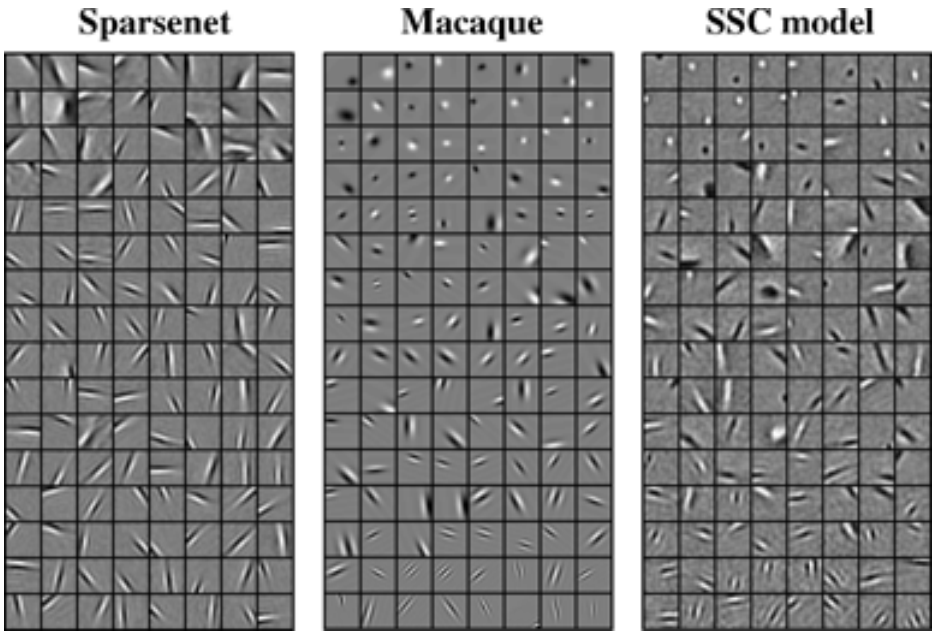

Figure 3.1: RFs from sparse-coding models and macaque V1.

SparseNet RFs are over-elongated Gabor type patches, making excellent edge detectors. Actual macaque RFs, are in general more squat and include spot-like unorientated RFs (unpublished data from Ringach (2007), obtained by reverse correlation and then fitted to Gabor patches, and adjusted for eccentricity by adjusting scaling). The SCC model (using hard sparse coding criteria) has a better match to the animal data, featuring both non-OR and OR selective units. Order is with respect to degree of OR selectivity, left to right, and top to bottom. Reprinted from Rehn and Sommer (2007).

and SF range matches that found in animals quite well and is sufficient to span by linear superposition the image space, which may be a requirement of the visual system. These results support the view that not only are Gabor type RFs an efficient coding strategy for representing natural images, but that they could be an emergent property of a system constrained by sparseness and the need to accurately represent the input space. These models are able to satisfy two of the SF modelling criteria, producing realistic RFs with a large and representative range of SFs. However, these models do not explain the organisation of SF preference, and do not provide a mechanistic account of development. The model's structure and developmental process do not parallel a living system, and the mechanisms within it do not typically have biological analogues.

Hybrid models that have incorporated biologically motivated mechanistic constraints into initially idealised sparse coding models have been designed to try to match biological results. Hyvärinen and Hoyer (2001) extend the SparseNet model (Olshausen and Field, 1996) by the addition of a second layer of "complex cells" that receive topographically pooled afferent input from the SparseNet sheet, the "simple cells". The sparseness of this entire "Topographic ICA" system is maximised during learning: both the sparseness of simple cell activity, and the sparseness of complex cell activity. This linked requirement for sparseness, coupled with the topographic pooling of simple cells to complex cells, results in the simple cells organising themselves topographically as they develop their feature preferences. The simple cells' RFs are typically Gabor type in form, as per the standard SparseNet model, and again the RFs are somewhat elongated, and low SF preference is somewhat under represented. The organisation for OR preference does have the common features found experimentally (iso-OR zones, linear zones, pinwheels, fractures, etc.). Due to the small area modelled (25 by 25 units), which is little more than one "hypercolumn" (see figures 3.2 and 3.3), it is impossible to extrapolate the organisation to a greater scale. Thus one cannot determine how well the Topographic ICA OR map truly matches experimental OR maps. SF organisation also suffers from this problem; while there is a clear organisation within this region, there is only one concentrated zone of low-SF preference. The authors suggest that larger patch sizes are likely to result in further blob formation, but this is not known for certain. Phase preference is randomly distributed, as found experimentally (DeAngelis et al., 1999). These results are consistent with biological data. However, even though the retinotopy map does have a local continuous organisation, its global organisation is distorted with discontinuities. This distortion is likely to be partly a result of using a torus (wrap around) basis for connecting the simple and complex sheets, which makes it difficult to determine how the retinotopic, OR, and SF maps would interact in a large model.

The Topographic ICA complex cells display key behaviour observed in animals, possessing phase invariance, while retaining OR and SF preferences. However, no discussion is given to the organisation of these complex cells. It would be interesting to know if they have a topographic organisation that matches the simple cell layer, which would be necessary for reproducing the key experimental observation of columnar or-

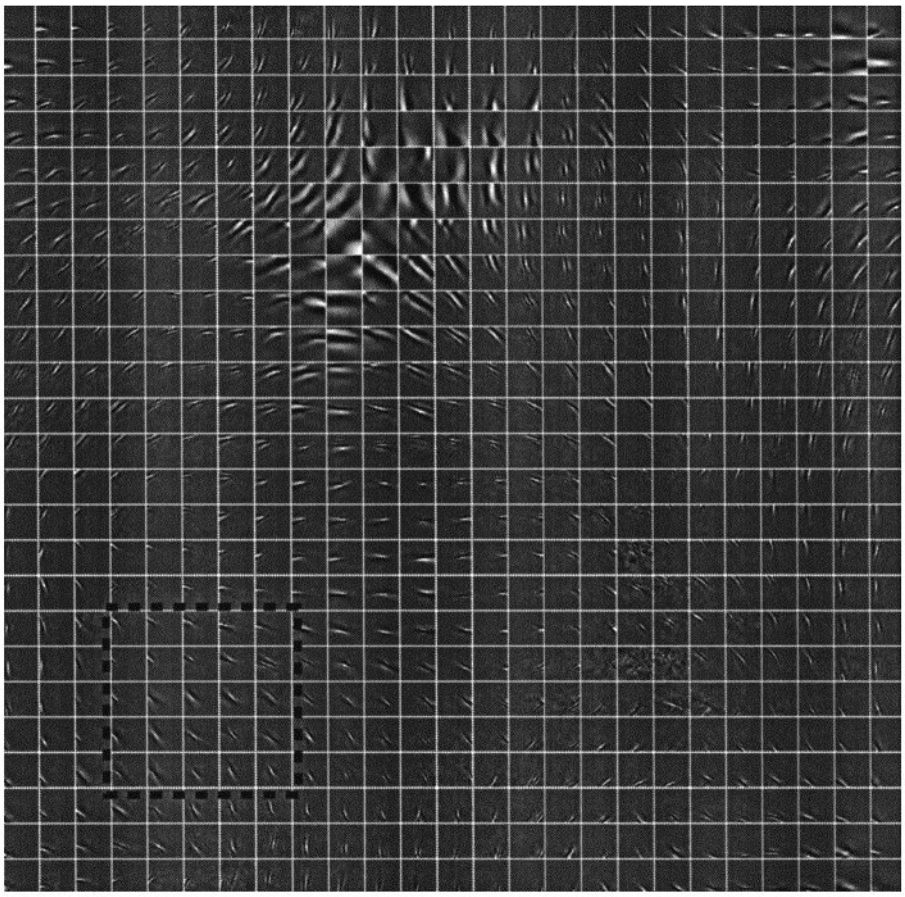

Figure 3.2: Topographically organised basis vectors from the Topographic ICA model. The basis vectors have Gabor shapes and are organised continuously, with one blob region of low SF preference. The dotted box region show an area that forms the connection field to a single cell in the complex cell layer. Note that this sheet of simple cells is "wrap-around" from top to bottom and from left to right. Reprinted from Hyvärinen and Hoyer (2001).

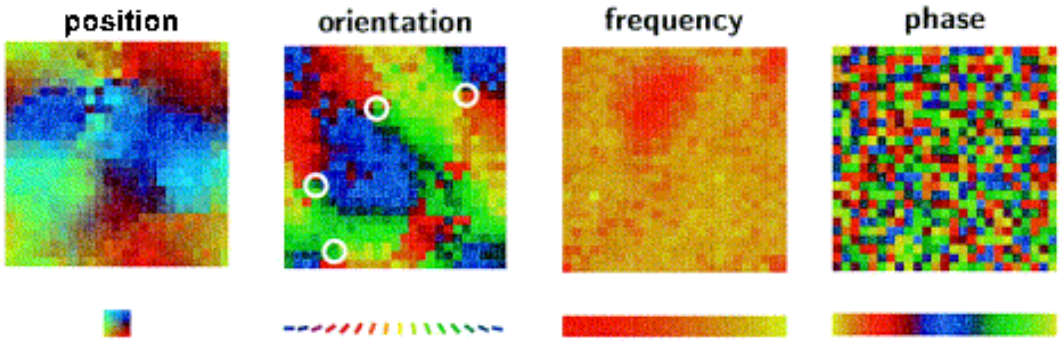

Figure 3.3: Topographic ICA model: retinotopy, OR, SF and phase maps. Retinotopy is locally continuous, but globally poorly organised, which is partly due to the "wrap-around" nature of the cortical sheets. OR preference has the key OR map features; pinwheels are marked by white circles. SF preference has iso-SF regions, but the cortical sheet size is too small to infer any global organisation. Phase preference is randomly organised. Reprinted from Hyvärinen and Hoyer (2001).

ganisation in OR preference. Overall, the significant outcomes from the Topographic ICA model are the production of small scale topographic maps for OR and SF, composed of units with Gabor type RFs, and the generation of complex cells with realistic properties. Unfortunately, the limitations of the model's scale and its toroidal cortex representation make it difficult to generate predictions about the possible global organisation for SF preference and how this may be linked with OR preference. Moreover, just as for SparseNet and SCC, this model has little physiological grounding. Instead, the results reflect the idea that the organisation of the visual system is heavily affected by the input image statistics, and the emergence of Gabor type RFs may be a direct consequence of these statistics. SparseNet results suggest that it is likely that the visual system has evolved to optimally represent natural image data, but cannot illuminate how this might be achieved in the actual animal visual system.

In a similar approach, Weber (2001) uses a combined model of sparse coding bottom-up feature detectors and an attractor with horizontal weights. Using filtered grey-scale natural images, the system succeeds in developing topographic maps (figure 3.4) of OR and SF, while phase is not ordered (as demonstrated experimentally, DeAngelis et al., 1999, and thought to be important for complex cell formation). To

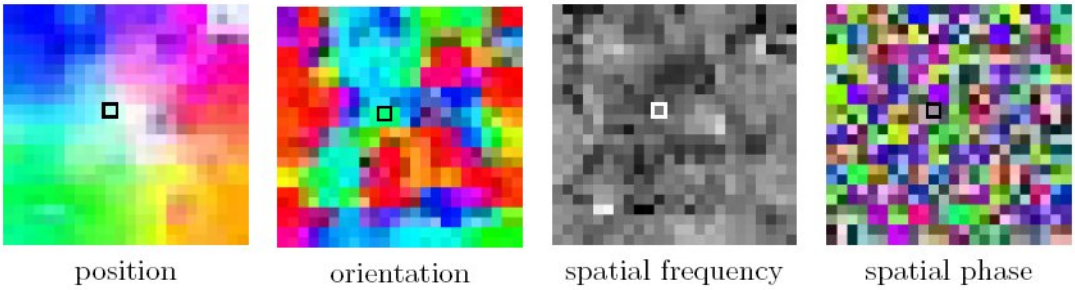

Figure 3.4: Adapted SparseNet model: retinotopy, OR, SF and phase maps. Retinotopy is continuously organised. The OR preference map is quite rough, but has the key features found experimentally. The SF map does not appear to have much structure to it, or particularly clear iso-SF zones, but it is predominantly continuous in SF preference. Spatial phase preference appears nearly randomly organised. Reprinted from Weber (2001).

develop phase invariance, the output of the feature cells is treated in a similar manner to the energy model (Adelson and Bergen, 1985), by squaring the activity of two cells with similar OR preference that are 90° out of phase (a quadrature pair), and then summing the squared activity. This approach is also used by Hyvärinen and Hoyer (2001), but it remains difficult to see how this could be achieved biologically. As for SparseNet, the resulting model is quite abstracted from the underlying physiology of the visual system. The main outcome of the work is the importance of lateral connectivity in forming attractor sites, thereby leading to edge detection, and similarity in other local feature preferences.

In addition to fulfilling the SF modelling criteria of producing realistic RFs spanning the SF input space, these hybrid sparse coding models (Weber, 2001; Hyvärinen and Hoyer, 2001) succeed in produce topographic SF maps. However, as we have seen, their SF maps are poor, and the biological foundation of the models is limited.

3.2.2 Dimension reduction topographic models

Dimension reduction models try to characterise feature maps and the interrelationships between them. Although, these models are abstract, they have had a great deal of success at reproducing features found in animal feature maps. The elastic net algorithm (Durbin and Willshaw, 1987) was first used for investigations concerning ocular dom-

inance columns (Goodhill and Willshaw, 1990), and the relationship between retinotopy and orientation organisation (Durbin and Mitchison, 1990). Carreira-Perpiñán et al. (2005) succeeded in developing multiple feature maps (see figure 3.5), including SF maps, using the elastic net model. Their results compare favourably both with observed maps and experiments relating to monocular deprivation and single orientation rearing, both in terms of qualitative form and rough equivalent time line. The results suggest that this high level abstraction may well summarise what the brain is trying to achieve in terms of dimension reduction. Their SF map is binary, with islands of high SF preference in a sea of low SF, which is in line with some experimental studies, though not with most recent SF map findings (as described in section 2.4 of chapter 2).

Just as with the previously discussed models, the elastic net algorithm has no biological analogue. Moreover, the input data used is extremely artificial, consisting of a discrete input space, with orthogonal dimensions representing each of the features (position, orientation, direction, ocular dominance, and SF) as one floating-point number apiece. The SF dimension has two possible states: "low" or "high". The algorithm covers this feature space, attempting to compress these pre-categorised feature dimensions into a singular sheet representation. This model does not help to explain how the brain extracts this information from the pixellated array in the retina representing the visual scene, due to its lack of physiological basis and its artificial (non-image-based) input, and thus much remains to explain.

In two experimental and computational modelling papers, Yu et al. (2005) and Farley et al. (2007) make direct comparisons between organisational and re-organisational relations between feature maps both in ferret and in a dimension-reduction paradigm. The model considered is a modified Kohonen self-organising feature map algorithm (Kohonen, 1982; Obermayer et al., 1992). Predictions borne from the dimension reduction model are found to be supported by their experimental observations:

- High-gradient regions (regions where the feature preference changes dramatically over a small distance, such as fractures in an OR map) of different feature maps do not spatially overlap.

- Maps for different visual features maintain orthogonal relationships in regions

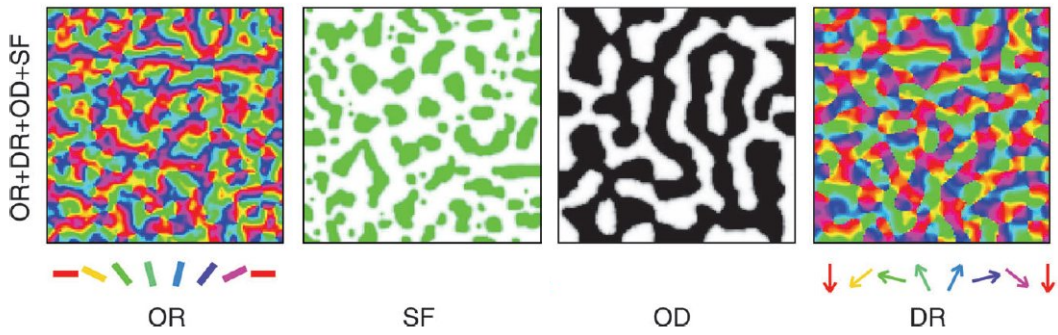

Figure 3.5: Elastic net model: OR, SF, OD and DR maps.
The stimulus space is constructed from discrete feature points spread along orthogonal feature dimensions. Spatial frequency stimuli have only two states, "high" or "low". The resultant maps, and their inter-relationships, are a result of the elastic net algorithm trying to represent this multi-dimensional space in a two dimensional sheet. Reprinted from Carreira-Perpiñán et al. (2005).

where they have coincident high gradients (for example, a coincident OR fracture and SF fracture would typically cross each other perpendicularly).

- The retinotopic map is not significantly distorted by the OR map (although see Das and Gilbert, 1997, for report of retinotopic map distortions matching OR fracture locations).

- An anisotropic retinotopic map results in an anisotropy of other feature maps, and in the iso-feature regions being elongated along the high gradient axis of the retinotopic map (see figure 3.6).

- Removal of a feature in stimulus space results in a reorganisation of the feature maps, altering spatial relationships between the maps. The actual number of feature maps in a cortical region affects the form of organisation amongst them.

Dimension reduction models have characterised a number of inter-map relationships, for which there is growing experimental support. This work strengthens the idea that coverage and continuity are organisational principles underlying feature map organisation. Coverage uniformity refers to the uniform representation across the cortex of all feature parameter combinations, while continuity is the slow change in neurons'

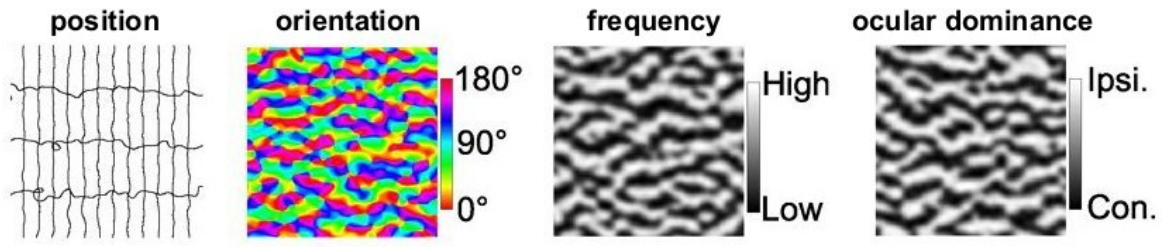

Figure 3.6: Kohonen SOM: retinotopy, OR, SF and OD maps.
(A) Anisotropic retinotopic map; lines indicate iso-elevation and iso-azimuth contour lines. Contour intervals are identical in each direction, thus demonstrating an anisotropy in the magnification factor. (B) OR map. (C) Spatial Frequency map. (D) Ocular dominance map. Note the elongation of iso-feature regions along the maximum gradient of retinotopy change, as predicted by this model, and also observed in ferret. Reprinted from Farley et al. (2007).

feature preferences across the cortex. Without coverage, parts of the cortex would be blind to particular feature combinations, while continuity is considered to be a result of minimising cortical wiring. Note that coverage uniformity (which is essential for functional vision) can be achieved without continuity, resulting in "salt and pepper" organisation, as has been observed in rat (Girman et al., 1999; Ohki et al., 2005). However, together these constraints allow for the production of realistic feature maps with realistic properties matching maps from higher mammals. The processes that drive these mapping relationships remain unknown, and because the inputs and the V1 implementation are so highly abstracted this type of model is unable to offer an insight into what these might be.

There is a difference of approach and intention between the dimension reduction models and more biologically grounded models. Dimension reduction models investigate distinct feature representation and mapping from a top-down perspective, whereas more biologically grounded work takes the reverse approach in trying to create maps using a more bottom up methodology. The top down approach is useful in trying to gauge what the overarching organisational structure might be for map development, but it is unable to provide anything close to a mechanistic understanding accounting for organising principles. The architecture of dimension reduction models does not parallel the actual visual system, so while these models can succeed in showing us

what the cortex may be trying to do, they are unable to tell us how it is accomplishing this. Thus previous models of SF maps have typically been too abstract to offer much insight into SF organisation and the mechanisms that may lead to it.

Of the three SF modelling criteria, the dimension reduction models satisfy one of them, by producing realistic SF maps, and offering a number of predictions relating to intermap relations. The range of SF preference is arbitrary due to the abstract nature of the models, and for the same reason there are no receptive fields in this context.

3.2.3 Physiological models

Physiological models are grounded in observations and inferences drawn from observing the visual system. Though mechanistic implementation of these observed processes, this model tries to capture essential features of biological systems, to account for biological results, and to generate specific predictions relating to the biological system. Determining what these "essential features" of a biological system are can be extremely difficult, especially when a system is poorly understood, or the data concerning it is conflicting or absent.

Mathur and Bhaumik (2004) present a feed-forward model (for model description see Bhaumik and Mathur, 2003) for the formation of simple cell receptive fields and subsequent SF mapping. The model architecture is composed of three regions, retina, LGN and cortex, each region having a 2D array of units. The units are spiking, and the LGN and retina consist of ON and OFF units. The model uses algorithms based on diffusive cooperation and resource limited competition for neurotrophic factors to generate weight changes between the LGN and cortical cells. This process successfully produces OR- and SF-tuned simple cells that organise into OR and SF maps (figure 3.7).

The diffusion of signal between neighbouring units leads to the organising of feature preferences. In most models, some form of lateral communication/connectivity is required for topographic map formation (although see Ringach, 2004a). In other models, direct neural connectivity is used, but the Mathur and Bhaumik (2004) model

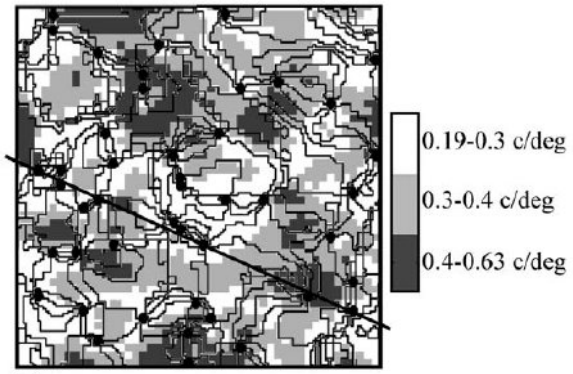

Figure 3.7: Feed-forward neurotrophic model: OR and SF maps.
Spatial frequency map, with superposed orientation contour map. OR pinwheels are shown as black circles. Reprinted from Mathur and Bhaumik (2004) .

shows that diffusion-based systems are sufficient. The authors report a close parallel between their results and those of Issa et al. (2000), in that the co-localisation of pinwheel centres and extreme iso-SF preference zones occurs to about the same extent. The form of their SF maps is also similar to the Issa et al. (2000) map, being continuous and having occasional fractures. This model meets two of the SF processing criteria well, i.e., producing topographic SF maps, and having cortical units with realistic RFs. However, the range of SF preference is limited to a range of less than 2 octaves. This may be due to the model LGN having DoG cells of only one size, unlike in animals.

In addition, the Mathur and Bhaumik model's cortical units develop their receptive fields and feature preference organisation in the absence of any visual stimuli. As a result, the model is unable to account for any visual experience dependent effects, which are well established in OR (Sengpiel et al., 1999; Tanaka et al., 2006), and ocular dominance maps (Wiesel and Hubel, 1965), and may also apply to SF (for a review regarding cortical map development dependencies upon visual experience, see White and Fitzpatrick, 2007).

Zhu et al. (2009) describe a spiking neuronal network model used to investigate SF selectivity. The model includes an LGN with single size DoG units, connected to a V1 sheet, which is made up of a grid of laterally connected integrate-and-fire excitatory and inhibitory neurons. The model has OR and SF organisation built into it by hard

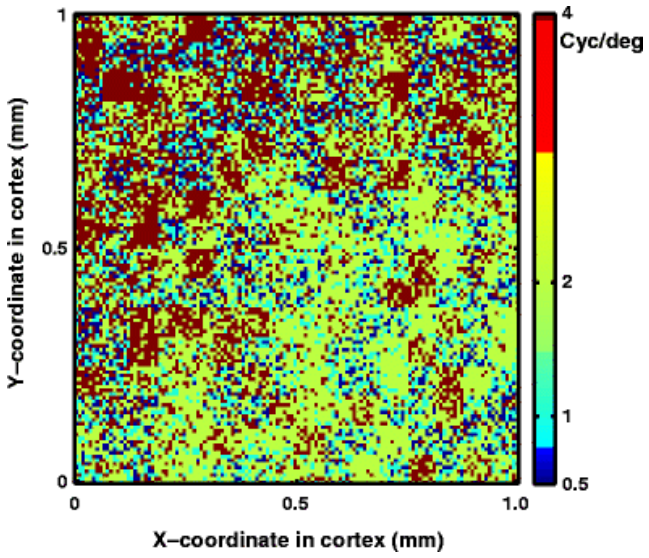

Figure 3.8: Zhu et al.'s neuronal network model: SF map.
SF preference distribution was predetermined by reproducing the statistical distribution observed in Sirovich and Uglesich, 2004. Reprinted from Zhu et al. (2009).

wiring the LGN connections to V1 units so as to form orientated and SF selective RFs. The OR organisation consists of four pinwheels, and the organisation for SF is based upon the Sirovich and Uglesich (2004) study. RFs with lower SF preference are built, by constructing the subregions from many LGN units. The model is not developmental, as all of the connections and weights are fixed, and thus it cannot explain how SF preference occurs in animals.

Zhu et al. find that LGN inputs do not fully account for the SF selectivity of V1 neurons, and that cortical inhibition is responsible for sharpening of the SF tuning. Without cortical inhibition they find the response of neurons to be much higher in the low frequency range, the cortical inhibition suppresses this, resulting in improved tuning. This research shows that a cortex without a topographic organisation for SF preference may still be able to acquire cortical neurons with a high degree of SF selectivity. As a lot of recent work (see section 2.4) suggests that an organisation for SF does exist in V1, the success of this model with a non-organised cortex may mean that the functional ability of V1 need not depend upon its organisation.

Sirosh and Miikkulainen (1996) use a simple LISSOM (laterally interconnected

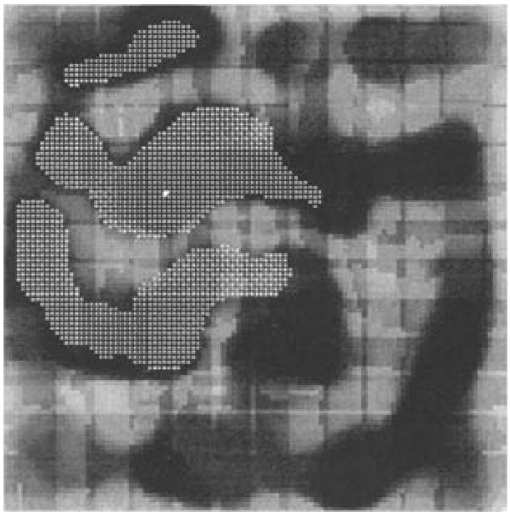

Figure 3.9: LISSOM: size selectivity map.
Size preference of neurons is indicated by colour scale, black (small) to white (large). The small white dots show the lateral input connections to the neuron labelled centrally with the large white dot. Neurons are organised for size preference, forming patches of areas with similar size preference selectivity. Lateral connections are made predominantly to units with the same size selectivity. Reprinted from Sirosh and Miikkulainen (1996).

synergetically self-organising map; LISSOM will be described in detail in section 3.3.1) model to simulate size preferences. The simulation consists of a single cortical sheet with short range lateral excitatory connections and long range inhibitory connections. The model receives afferent input from a retinal sheet, which is exposed to visual stimuli composed solely of isotropic Gaussian spots of different sizes. The model is able to develop receptive fields that have preference for different sizes of stimuli. However, as there is no OR information in the stimulus, no OR preference develops in the model. There are also no ON and OFF channels, so the RFs are single-region blobs of different sizes. The production of continuous "spot-size" preference topographic maps (see figure 3.9) is something for which there is no known correlate in animals, but it does show the potential for this form of model to be used to investigate SF preference in a more realistic architectural and experimental framework. The model also provides a credible mechanistic framework, established in a context of successful visual system modelling results, which can be used to investigate SF organisation.

As for the non-mechanistic models, none of these mechanistic models meet the

three modelling criteria. Mathur and Bhaumiks model has a significant biological basis, and generates a good SF map, but as discussed it is restrictive in its LGN representation, resulting in a narrow SF range. It also does not develop due to any visual stimulus. Zhu et al.'s model SF map is not topographic, unlike current imaging results, and is also not developmental. Finally, the LISSOM model does not model SF preference directly. No mechanistic model exists that contains the key ingredients to generate SF maps. In the next section, LISSOM is described in more detail, and its suitability for creating such a model is considered.

3.2.4 Summary of overview

As argued in the above sections, previous modelling studies have a number of limitations with regard to SF processing and map formation. Figure 3.10 allows easy visual comparison of the mapping contributions from each of these previous studies. Only the dimension reduction and feedforward neurotrophic models produce SF maps that are qualitatively similar to those observed experimentally (see section 2.4). But, as discussed, these models also have significant limitations, such as not yielding receptive fields for different SFs, and not incorporating natural image input.

3.3 Adapting the LISSOM model for SF

The lack of a biologically grounded model that is able to fulfil the SF modelling criteria drives the desire to generate a new model to investigate SF organisation. Furthermore, previous models have not included multiple SF pathways, and I consider this to be of fundamental importance in cortical SF processing. The LISSOM model (Miikkulainen et al., 2005) is an appropriate starting point for SF mapping investigation for a number of reasons:

- It is based upon a solid biological foundation: its architecture is capable of reproducing the broad underlying anatomy/physiology and with that the basic neural processing that occurs from retina through to V1 (Miikkulainen et al., 2005).
- It includes explicit RGC/LGN processing that can be extended to include separate SF pathways.

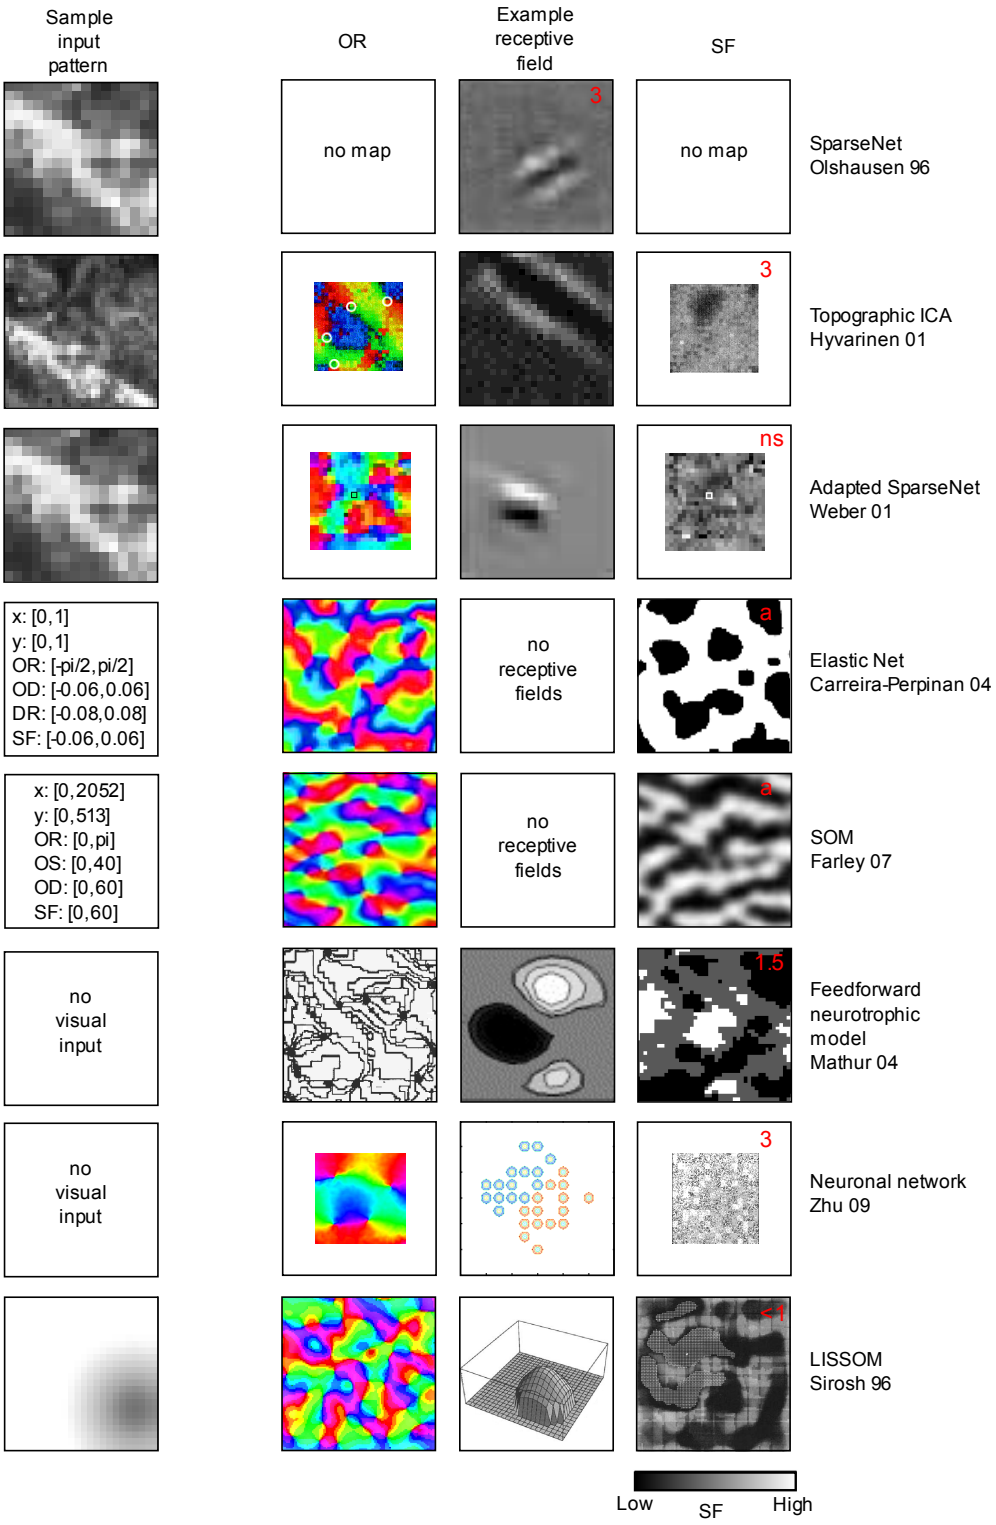
Sample input pattern
OR
Example receptive field
SF
no map
3
no map
SparseNet
Olshausen 96
3
Topographic ICA
Hyvarinen 01
ns
Adapted SparseNet
Weber 01
x: [0,1]
y: [0,1]
OR: [-pi/2,pi/2]
OD: [-0.06,0.06]
DR: [-0.08,0.08]
SF: [-0.06,0.06]
no receptive fields
a
Elastic Net
Carreira-Perpinan 04
x: [0,2052]
y: [0,513]
OR: [0,pi]
OS: [0,40]
OD: [0,60]
SF: [0,60]
no receptive fields
a
SOM
Farley 07
no visual input
1.5
Feedforward neurotrophic model
Mathur 04
no visual input
3
Neuronal network
Zhu 09
<1
LISSOM
Sirosh 96
Low
SF
High

Figure 3.10: (Previous page) Comparison between published computational SF and OR maps. Computational SF maps and matching OR maps, from published computational work on SF mapping (all original SF maps may be found elsewhere in this chapter). (First column) Example visual stimuli used to drive development in each model. Three of the models use whitened low resolution natural images as stimuli. The elastic net and SOM use low-dimensional vector input; the numbers indicate the range from which vectors are drawn. Note that for the Carreira-Perpiñán et al., 2005 model, OD, direction (DR) and SF take only binary values. (Second column) OR maps for each simulation. (Third column) Example RFs from each simulation. (Fourth column) SF maps for each simulation. Numbers in top corner of each SF map indicate the octave range shown in each map, "a" indicates arbitrary, and "ns" indicates that the range for this map was not specified and is indeterminable. For the Sirosh and Miikkulainen, and Zhu et al. models example OR maps not from their papers have been added to aid illustration (the LISSOM OR map is taken from Sirosh, 1995; the OR map for Zhu et al. is a detail from Issa et al., 2000, depicting 4 pinwheels, which the authors say they have hardwired into their simulation). Mathur and Bhaumik's maps have been heavily processed to separate them using image manipulation software. Maps have been cropped, scaled and recoloured to aid comparison.

- It has established success in producing realistic topographic feature preference maps that are qualitatively similar to experimentally observed maps (Miikkulainen et al., 2005).
- It develops from realistic visual stimuli, which can either be natural images or spontaneous patterns of neural activity. Although there is some debate regarding maps arising without visual stimulation, and the possibility that some maps may in effect be "hard-wired", there is much evidence for the emergence of certain feature preferences being dependent upon visual experience (for review, see White and Fitzpatrick, 2007).
- It is self-organising, using Hebbian learning, which is a simple principle with copious experimental support (for review see Tsien, 2000).
- It has self-organising lateral excitatory and inhibitory connections. These connections are thought to play an important role in feature map development, and the tuning for selectivity of individual neurons (see Weliky et al., 1995; Bosking et al., 1997).

A new model, formulated upon these foundations, overcomes many of the limitations of previous SF preference organisation modelling studies. This biologically motivated approach will provide for a more realistic model of this aspect of the visual system than any previously published model, and I hope it can be used to address the question "What allows SF selectivity and organisation to arise in V1?".

3.3.1 Basic LISSOM model

In the LISSOM model (Miikkulainen et al., 2005), the early visual pathway is represented by interconnected sheets of activity-driven firing-rate units. Separate sheets represent the retinal photo-receptors (the input sheet), the RGC/LGN ON and OFF cells, and the V1 simple cells (see figure 3.11). The retinal photoreceptor sheet provides the input space to the model, allowing greyscale images to be presented as visual stimuli. The retinal input pattern that maximally drives some unit in the model is referred to as that unit's receptive field (RF). Units in the RGC/LGN ON and OFF sheets have identical centre/surround DoG RFs that are a weighted sum of activity values of the photoreceptors projecting to that unit. The units that make up the projection from

one sheet to a unit in a following sheet are referred to as the connection field (CF). V1 units have CFs from both ON and OFF RGC/LGN sheets. The weight of each connection to V1 is initially random, but the weights are adjusted with time, through Hebbian learning. The V1 also units make excitatory and inhibitory lateral connections with each other.

When an input pattern is presented on the photoreceptor sheet, the responses of the photoreceptors are pooled and weighted according to the RF of the RGC/LGN units. The RF of an RGC/LGN cell is given by the difference of two normalised Gaussians. The contribution $L_{xy,ef}$ of a pixel (photoreceptor) at (x,y) in the RF of an ON centre cell (e,f) with centre (x_c,y_c) is given by:

$$L_{xy,ef} = \frac{\exp\left(-\frac{(x-x_c)^2+(y-y_c)^2}{\sigma_c^2}\right)}{\sum_{mn}\exp\left(-\frac{(m-x_c)^2+(n-y_c)^2}{\sigma_c^2}\right)} - \frac{\exp\left(-\frac{(x-x_c)^2+(y-y_c)^2}{\sigma_s^2}\right)}{\sum_{mn}\exp\left(-\frac{(m-x_c)^2+(n-y_c)^2}{\sigma_s^2}\right)} \tag{3.1}$$

where $\sigma_c$and σ_s are the standard deviations for the centre and surround Gaussians respectively, and x and y are the centre of the RF.

OFF weights are the negative of ON cell weights. As firing rates in biological systems cannot be negative, activity in either cell type is thresholded so they only have positive values. Thus ON and OFF cells in the same RGC/LGN location are complementary, as they can never be active simultaneously. The activity of a single LGN cell X_{ef} is given by the scalar product of the activity of each photoreceptor χ_{xy} in the connection field and its corresponding weight value:

$$X_{ef} = H(\gamma_L \sum_{xy} \chi_{xy} L_{xy,ef}) \tag{3.2}$$

where γ_L is a constant scaling factor, and $H(.)$ is a thresholding function:

$$H(\psi) = \begin{cases} 0 & \psi < 0 \\ \psi & \psi \geq 0 \end{cases} \tag{3.3}$$

Every iteration begins with a new image presentation to the photoreceptor sheets, followed by 9 cortical settling steps. These settling steps allow the activity in V1 to stabilise as neurons interact laterally. In each of these settling steps, the response

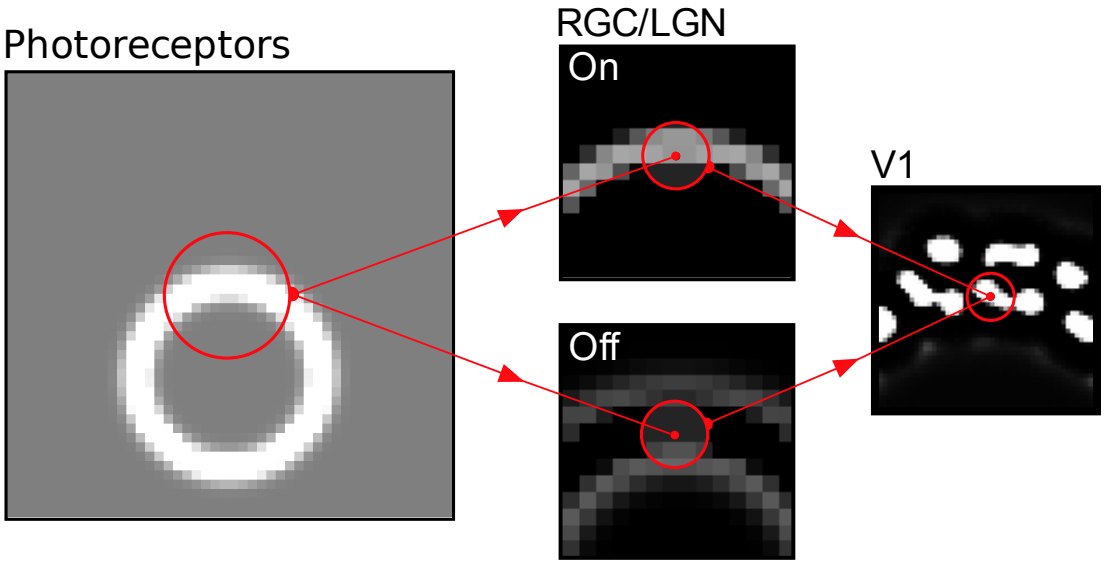

Figure 3.11: LISSOM model architecture.

Greyscale images are presented on the photoreceptor sheet; here a simple ring-shaped pattern is shown. The red circle on the photoreceptor sheet shows the connection field for two spatially corresponding RGC/LGN ON and OFF units. Pixels in this connection field have their activity spatially filtered by point-wise multiplication with a DoG filter, centred on the connection-field centre. These weighted activities are then summed, resulting in the activity of the RGC/LGN unit. Every unit in V1 receives input from the weighted sum of all the units in each of its connection fields. Within V1 there is inhibitory and excitatory feedback, with each neuron making lateral connections to all neighbouring neurons within a defined radius. Each V1 unit thus has 4 input connection fields (2 afferent inputs from the ON and OFF channels, and 2 lateral inputs from the lateral excitatory and inhibitory connections). The diagram shows the responses of such a model to the ring stimulus.

$\eta_{ij}(t)$ of a cortical unit (i,j), at settling time step t within the iteration is given by the combined activation from all projections (from the previous settling time step, $t-1$) to that unit:

$$\eta_{ij}(t) = \sigma\left(\sum_p \gamma_p \sum_{kl \in CF_{pij}} X_{kl}(t-1)\omega_{pkl,ij}\right) \tag{3.4}$$

where p denotes the projection (ON or OFF for one of the afferent projections from the LGN, E for the excitatory feedback projection, I for the inhibitory feedback projection), CF_{pij} denotes the connection field of projection p to unit (i,j), X_{kl} is the activation of a unit (k,l) in the connection field (CF_{pij}) of the V1 unit (i,j), and $\omega_{pkl,ij}$ is the afferent weight from (k,l) to (i,j) in projection p. γ_p is a strength scaling factor appropriate to that specific projection. The projections to the unit (i,j) include any lateral connections within the model, so unit (k,l) will be on the same sheet as unit (i,j) for lateral feedback projections. The function $\sigma(.)$ is a smooth sigmoidal activation function. In standard LISSOM there is an updating schedule to adjust the parameters of the sigmoid, as well as the radii of lateral connections, over the course of development (see Miikkulainen et al., 2005). However, I have used a modified version of LISSOM extended by my colleague Veldri Kurniawan (2006), in which the sigmoid function is adjusted automatically through homeostatic plasticity of intrinsic excitability, to maintain an average firing rate (Triesch, 2005). This is done through a form of dynamic sigmoid:

$$\sigma(x) = y = \frac{1.0}{1.0 + \exp(-(ax+b))} \tag{3.5}$$

The adjustments to the sigmoidal parameters a and b at each iteration, z, are given by equations 3.6 and 3.7:

$$a_z = a_{z-1} + \Delta a \tag{3.6}$$

$$b_z = b_{z-1} + \Delta b \tag{3.7}$$

where

$$\Delta a = \beta\left(\frac{1}{a_{z-1}} + x - \left(2 + \frac{1}{\mu}\right)xy + \frac{1}{\mu}xy^2\right) \tag{3.8}$$

$$\Delta b = \beta\left(1 - \left(2 + \frac{1}{\mu}\right)y + \frac{1}{\mu}y^2\right) \tag{3.9}$$

β is a small learning rate, and μ is a unit's target mean firing rate. Without the sigmoidal function being able to adapt the mean firing rate may saturate as the model develops, this homeostatic mechanism maintains a mean firing rate, leading to sparse activation.

Initial weight values for thalamocortical connections and lateral inhibition are random within a 2D Gaussian envelope (see section 4.2), while lateral excitatory weight values are initially from a 2D Gaussian distribution. After each iteration, once the activity has settled, the connection weights of all cortical units are updated. The learning rule for both afferent and lateral weights is Hebbian with divisive post-synaptic normalisation:

$$\omega'_{pqr,ij} = \frac{\omega_{pqr,ij} + \overline{\alpha}_p X_{pqr} \eta_{ij}}{\sum_{ouv} \left(\omega_{ouv,ij} + \overline{\alpha}_o X_{ouv} \eta_{ij} \right)} \tag{3.10}$$

where $\omega'_{pqr,ij}$is the new weight from unit(q,r) in projection p to unit (i,j), p is one of the projections to this neuron (ON, OFF, E, or I) ,$\overline{\alpha}_p$ is the learning rate for projection p, o is an index varying over the projection(s) to be normalised together (either $\{I\}$ or $\{E\}$ for feedback projections, which are normalised independently, or$\{ON, OFF\}$for an afferent connection, which are normalised together), X_{pqr} is the pre-synaptic activity after settling, and η_{ij} is the activity of the neuron after settling. These equations ensure that neurons with similar activity patterns develop strong connections and that weights do not grow without bound. The feedback connections are normalised independently so that they will remain in balance, while the afferent weights are all normalised together so that neurons can learn stronger connections to a specific projection. The effect of the afferent thalamocortical normalisation is that the sum of the weights of the afferent projections to a unit (i,j) sum to 1:

$$\sum_{p \in \{ON,OFF\}} \left[\sum_{kl \in CF_{pij}} \omega_{pkl,ij} \right] = 1 \tag{3.11}$$

The afferent learning rate $\alpha_A = \alpha_{ON} = \alpha_{OFF}$ (set to the same value for all afferent projections) decays exponentially to allow settling of cortical preferences, where ϕ is the initial learning rate at $t = 0$, and λ is the time constant:

$$\alpha_A = \phi \exp(\frac{-t}{\lambda}) \tag{3.12}$$

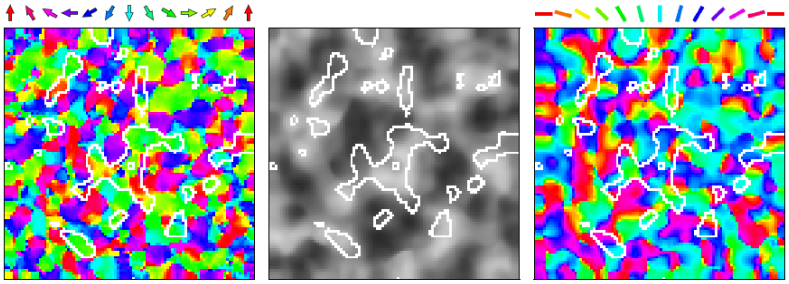

Figure 3.12: LISSOM feature maps and lateral connections.
(a) Direction of motion map. (b) Ocular dominance map, with eye preference represented on a continuous scale, with left eye preference in black, and right eye preference in white. (c) OR preference map; note the existence of key OR map features, iso-zones, linear-zones, pinwheels and fractures. For one neuron marked with a small white square, self-organised long range lateral inhibitory connections are superposed in white on each map. From Bednar and Miikkulainen (2006).

Apart from not representing SF, LISSOM provides perhaps the closest representation to an actual visual system's architecture of any map development modelling approach so far. Functionally matching the visual system architecture and developmental process leads to the generation of V1 feature maps that have a high qualitative similarity to maps measured experimentally. The architecture is also able to organise to reflect intrinsic properties of cortical organisation, such as the distribution of lateral connections (see figure 3.12).

3.3.2 Assumptions of the LISSOM model

The principal assumptions (for a more complete discussion see Miikkulainen et al., 2005) of the LISSOM model are:

- Recurent and adapting lateral connections: The excitatory connections focus activity in maximally active regions of the system, effectively generalising the "winner takes all" mechanism which is found in SOM networks. This allows efficient self-organisation. Inhibitory connections help to decorrelate activity, which increases the variation and sparsity of representation that can develop within the model. Adaptation of lateral connections aids this decorrelation. Ex-

periments in cat have observed restructuring of lateral connections over development (Callaway and Katz, 1990).

- Normalisation of connection weights: This is required to prevent unbounded growth of connection weights. In LISSOM multiplicative normalisation is utilised; in some other models (Goodhill, 1993; Miller et al., 1989) a subtractive mechanism is used. However, subtractive normalisation would not work in LISSOM, as it leads to the development of binary weight values, which loses much of the information normally represented in the weight pattern.

- Short range excitatory and long range inhibitory connections: For strong (high-contrast) stimuli, this circuitry is needed for the creation of activity bubbles which allow LISSOM to successfully self organise. Other self-organising models also use this structure (von der Malsburg, 1973; Miller et al., 1989; Miller, 1994). However biological observations suggest that the actual horizontal circuitry found in cortex is more complex than this. Work concurrent with this thesis by my colleague Judith Law has shown that self-organisation can still be achieved in a much more complicated model that includes separate inhibitory interneurons receiving long-range excitation. The results suggest that the LISSOM circuitry (with long-range inhibition) accurately models the net effect of lateral inhibition for the high-contrast areas that drive self-organisation. Thus for simplicity this thesis focuses on the basic, less complex model.

- Parameter adaptation - The standard LISSOM model features scheduled changes to the radius of the lateral excitatory connections, as well as changes to the learning rate of the neurons. These were implemented to represent the maturation of the cortex and allow for fine tuning of map development, implementing the critical periods for developmental plasticity. It is known that early in development there is a critical period during which the cortex is highly plastic, and outside of which the cortex is less plastic (Hensch, 2004). This gradual reduction in the learning rate allows for the maturation of neurons' selectivity within the cortex, and leads to greater stability later in development. The actual form of the reduction in learning rate is thought to be noncritical.

Many of the assumptions of the LISSOM model are shared with other computational models, such as SOM (Kohonen, 1982). However LISSOM incorporates many more

assumptions than less biologically realistic models, such as the dimension reduction models (described in section 3.2.2). This is neccessary to match the level of detail required to generate a biologically motivated model, partially because in developing a model that captures certain well understood aspects of a system, other elements of the system (at the same depth of complexity) may not be so well understood, so well founded assumptions about them are required.

In addition to the assumptions already described, for the standard LISSOM model, the models described in this thesis make use of one further and important assumption:

- Homeostasis - It is known that neurons' intrinsic excitability is adaptable, and it is thought that this type of plasticity may be important for keeping a neurons' firing rate around a certain average range (Turrigiano and Nelson, 2004). Otherwise, it may too frequently either be inactive or oversaturate. This type of adaptation, which adjusts the frequency of a neurons firing over time so that an average firing rate is maintained, is called homeostasis (Turrigiano, 1999). Using a homeostatic mechanism in the model is a way of ensuring that all units within the model are active; without a mechanism such as this, it is possible for units to fail to develop. The type of homeostatis mechanism used is from Triesch (2005), but the form of the mechanism is unlikely to be critical.

3.3.3 Expansion of LISSOM to allow modelling of SF preference

When SF preference is measured for units in the simple LISSOM model (described above), it is found that there is a very narrow variation in SF preference between different units, and that their preference is centred around the SF preference of the RGC/LGN sheet units. To investigate SF preference organisation in the context of the LISSOM model, a biologically motivated approach is to extend the RGC/LGN representation by introducing a variation in RGC/LGN receptive field sizes. This will allow for visual stimuli of differing SFs to drive the RGC/LGN units and in turn the cortex. For LISSOM, this is accomplished by expanding the RGC/LGN to include additional sheets with different SF tunings, thus creating a variety of separate SF pathways.

In animals, the LGN makes connections to V1 layer 4C, and there is evidence

both for and against the continued segregation of SF processing within V1 (section 2.2.7). In formulating the model architecture, we will in turn investigate two different architectures provoked by the possible nature of these connections:

1. V1 units receive convergent afferent input from all the different SF channels (chapter 4).

2. LGN afferent input to V1 is segregated by SF channel into separate sub laminae. Units in a sub lamina receive afferent input from only one SF channel (chapter 5).

In the following two chapters I use these models to investigate how SF preference and organisation can arise and might be influenced as a result of these architectures, and whether these models are able to meet the SF criteria.

Chapter 4

Single layer V1 model

4.1 Introduction

To gain insight into how SF preference may develop and become organised in V1, the models presented in this thesis aim to reproduce the functional processing of SF found in the early visual system. I assume that the breadth of SF tuning found in V1 is a consequence of the large range in size of ON and OFF LGN cells that make afferent connections to V1 (see section 2.2.7). Before extending the model to reproduce a realistic range of SF preference similar to that found in V1, it must first be decided how to include in the LISSOM simulation LGN cells of different RF sizes. There are two main practical approaches: including heterogeneous RGC/LGN sheets of units with a range of sizes of RFs, or including multiple homogeneous RGC/LGN sheets of units with fixed size RFs.

A single heterogeneous RGC/LGN sheet could consist of a large range of SF preferences. However, a specific point on retina will have restricted coverage for SF, determined by whatever size LGN cell happens to be covering this region. This would result in poor coverage of the SF input space, leading to poor image representation. This could be overcome, but would require a prohibitively large LGN sheet with units having randomised positions, in order to achieve good coverage. The coverage could be improved by adding additional heterogeneous RGC/LGN sheets, thus increasing the number of LGN cells covering a point on the photoreceptor sheet. However, this seems a clumsy approach, because the segregation into sheets has no obvious physiological

basis.

Alternatively, multiple heterogeneous sheets, each with ON/OFF cell populations drawn from a different size distribution, could produce good coverage for SF preference. Such an architecture may more closely parallel that found in cat and primate, but analysis is still difficult because two units in different sheets might differ in size only by a tiny amount.

For clarity and ease of analysis, the LISSOM model is extended with multiple homogeneous SF pathways, each with a fixed, but different, SF tuning. This approach allows the breadth of SF preference to be represented in a easily controlled manner, while capturing what appears to be the essential feature of the system, i.e., multiple SF representations for each spatial location. It also allows for straightforward architectural segregation between SF channels in V1, which is known to occur in cortical layer 4, as will be considered more fully in the following chapter.

This chapter considers the case in which thalamic afferent connections of all sizes are initially convergent upon V1 neurons; each V1 unit receives input from all the SF channels. I show that given this architecture, a Hebbian learning network can organise topographically for SF preference (as in Farley et al., 2007). However, the representation is not ideal: coverage of the SF input space is limited. By introducing a pruning mechanism, the range of SF preference represented in cortex can be increased, while also allowing the development of more OR-selective RFs. However, this procedure detrimentally affects the organisation for OR preference.

The results presented in this chapter demonstrate that this single layer V1 model with multiple SF pathways is successful in developing topographic organisations for SF and OR preference, covering a wide range of spatial frequencies. The success of this model suggests that some form of segregation of SF channel inputs may be necessary in order for V1 neurons to develop realistic RFs by Hebbian learning to achieve preferences covering the whole SF range.

4.2 Single layer V1 model

To adapt the LISSOM model to incorporate multiple SF pathways, further RGC/LGN ON/OFF sheet pairs are added. In principle, any number of SF pathways can be included. Here, to keep the computational cost feasible, the model consists of 4 distinct pathways (figure 4.1), spanning a total SF preference range of 2 octaves. Units in each RGC/LGN ON/OFF sheet pair have centre/surround RFs. The SF pathways are identical except for the scaling of the retinothalamic and thalamocortical connection fields and projections, which increase by a factor of 1.59 between adjacent pathways. This geometric increase in size across the SF pathways allows for an even SF representation of 2 octaves across the 4 pathways. The size of the LGN sheets also increases as a consequence of edge buffering, which is necessary to eliminate cropping of the afferent connection fields of the cortical sheet. The sheet sizes are rounded to a multiple of the reciprocal of the appropriate sheet density, to make the grid spacing even.

Every cortical unit receives afferent projections from each of the SF pathways. As in the standard LISSOM model, cortical neurons make lateral excitatory and inhibitory connections with neighbouring neurons. The initial weights for the thalamocortical and lateral inhibitory connection fields are drawn from a uniform random distribution within a 2D Gaussian envelope:

$$\omega_{kl,ij} = \text{rand}(0,1)\exp\left(-\frac{(x_k-x_i)^2+(y_l-y_j)^2}{2\sigma^2}\right) \tag{4.1}$$

where $\omega_{kl,ij}$ denotes the weight between unit (k,l) in the source sheet of a given projection and unit (i,j) in V1, and $\text{rand}(0,1)$ is a function that generates a uniform random number between 0 and 1. The Gaussian is centred at (x_i,y_j) which corresponds to the cortical position of unit (i,j), and (x_k,y_l) is the sheet position of unit (k,l). σ is the standard deviation, set to be equal to the radius of the appropriate connection field bounds. The lateral excitatory connection field weights are drawn from a Gaussian distribution with standard deviation 0.02. The weights of all the connection fields develop through Hebbian learning with divisive normalisation (see section 3.3.1). The afferent thalamocortical weights are jointly normalised, which means that all the afferent weights from all SF channels are grouped together when they are normalised. This

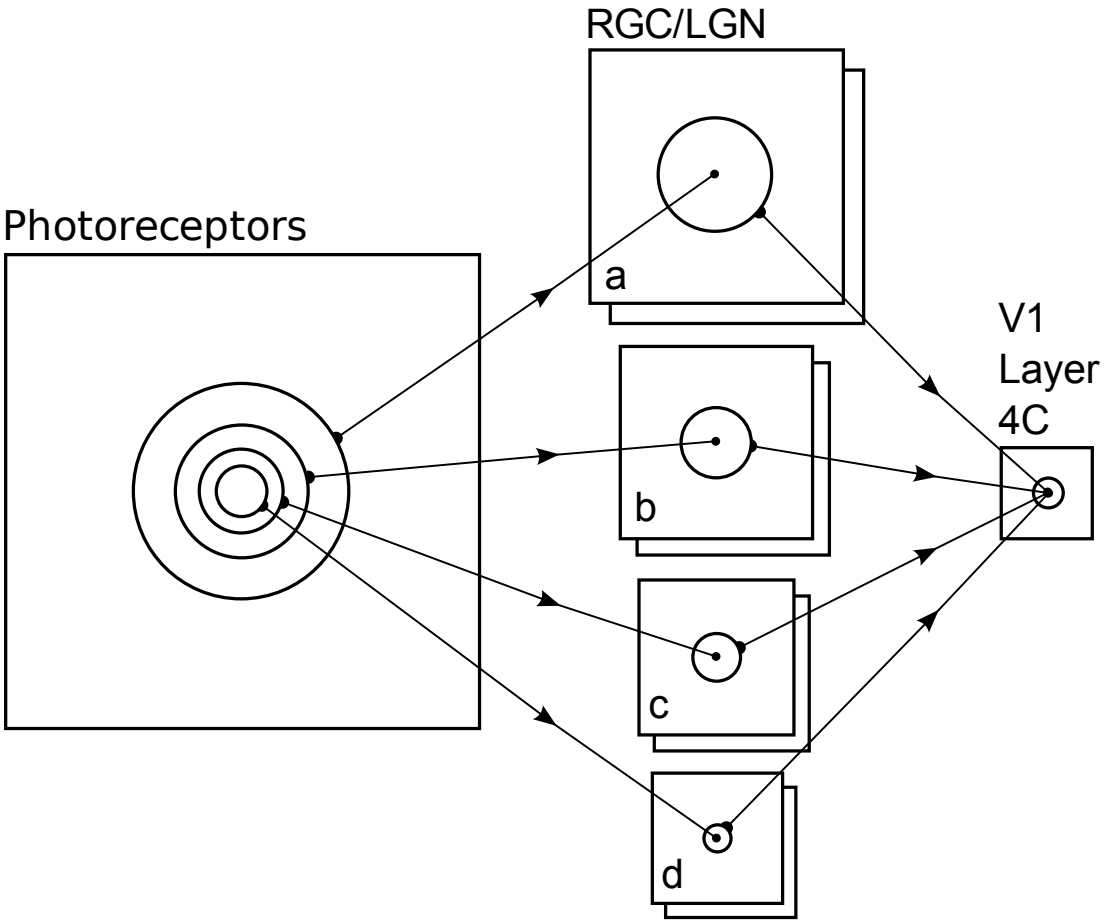

Figure 4.1: Single layer model, schematic.

Grey-scale natural images are presented on the photoreceptor sheet, which consists of 116 $\times$ 116 pixels. There are 4 distinct SF channels, each with an ON and an OFF sheet (OFF sheets are shown in background; connections to and from the OFF sheets are not shown), and each tuned to a different SF. Each RGC/LGN unit is activated from its centre/surround (DoG) receptive field on the input sheet. Receptive field sizes are uniform within a sheet, but differ between the different pathways, resulting in different SF preferences for each pathway. Every unit in V1 layer 4C receives input from units in each of the 8 RGC/LGN sheets. Additionally, in V1 layer 4C there is inhibitory and excitatory feedback. Each V1 unit has 10 input fields (8 afferent and 2 laterals).

grouping allows for the development of stronger connections to units in any specific SF channels. The cortical density is the number of neurons per unit size, parameter N_{V1}, unless stated otherwise is 96 for all models. Further parameter values for the single-layer V1 simulation can be found in tables A.1 and A.2, in appendix A.

4.2.1 Experimental procedure

Each model was presented with 5000 grey-scale, 116×116 pixel natural image patches sampled from 37 pictures, each 576×768 pixels, taken from the landscapes and foliage data sets in the McGill image bank (Olmos and Kingdom, 2004), available at http://tabby.vision.mcgill.ca/. These images were taken with Nikon Coolpix 5700 digital cameras callibrated according to Paraga (2003). Camera settings were fixed and are specified on the website. The images' original size was 1920×2560 pixels, and the only processing done on the images was to rescale them to 576×768 pixels; this was done by the image publishers. Prior to sampling a patch, each image was rotated by a random orientation drawn from a Gaussian distribution with standard deviation of 20°, to make these aesthetically framed photographs more representative of visual experience. For a discussion of the training data spatial properties, see appendix C.

After learning was completed, as described in section 3.3.1, maps were measured for SF and OR preference. These were measured by disabling plasticity and then recording all cortical units' activity levels on presentation of sine gratings at 4 orientations (0°, 45°, 90°, and 135°), each at 15 phases (analogous to drifting gratings). For SF maps each of these combinations was tested at 40 spatial frequencies (logarithmically in geometric progression, from $10^{-0.30}$cycles per unit distance (cpud), to $10^{0.675}$cpud, at $\times 10^{0.025}$cpud intervals), while for OR maps 4 SFs were used, (0.87, 1.36, 2.21, 3.55 cpud - roughly the SF preferences for each SF channel). Measuring of preference maps using larger ranges of OR and SF was also investigated, but found not to significantly alter the observed maps. The intensity of the sine grating was set to the highest value that avoided saturating the cortical activity levels, typically 0.17.

To generate OR preference maps, a unit's OR preference is determined by activity vector summation of all responses to grating stimuli, where the magnitude of the vector

is the response activity, and the orientation of the vector is given by double the stimulus orientation (this maps the periodic range of sine grating orientations on to the 2π range of vector directions). The orientation preference of that unit is then defined as half the orientation of the resulting vector. This method is termed taking the weighted average of responses for OR (see Miikkulainen et al. 2005, section G.1.3). Doubling the stimulus orientations before summation, and then halving them afterwards, eliminates the difficulty that can arise if stimulus orientations differ by over 90 degrees (which would result in standard vector addition giving the wrong (orthogonal) orientation).

As SF is not a cyclic quantity, taking a weighted average would systematically under represent both high and low SF preferences, introducing biases. Instead, SF preference was defined as the sine grating SF that generates the highest-activity response for a given unit. This is the discrete pattern method described in Miikkulainen et al. (2005) G.1.2, and identical to the approach used by Issa et al. (2000). It should be noted that each unit is responsive to a broad range of SFs, with peak responsiveness corresponding to its SF preference, but with significant responses to other SFs as well.

Unit RFs were determined by reverse correlation (Jones and Palmer, 1987b), with cortical plasticity disabled. Due to cortical units receiving afferent input from channels with different SF preferences, it was necessary to use visual stimuli that would on average equally activate all SF pathways. Pink (1/f) noise was used as the stimulus, which has a power spectrum very similar to natural images (Field, 1987); other solutions such as using sine gratings as described by Ringach and Shapley (2004) may also work. To build an individual unit's RF, all presentation patterns are multiplied by the cortical unit's response to that pattern and then summed. This calculation is equivalent to sparse-noise reverse correlation described by Jones and Palmer (1987b) (see appendix B for more detailed desciption).

4.2.2 Pruning methodology

The experimental work discussed in section 2.2.7 suggests that for V1 to develop a realistic range of SF preference, V1 neurons must predominantly receive thalamic input from neurons with similar SF preference. Here I introduce a pruning process, that

severs cortical connections, allowing V1 neurons to receive connections from only one of the SF channels. This allows us to compare results between the simulation under its normal operation (where V1 neurons retain connections to all SF channels), and the more extreme pruned case, where singule SF channel connections are enforced.

For a given unit, the pruning process consists of three steps. The first is to sum the weights for each SF channel's thalamocortical projections to that unit. Weights in the largest projection to a unit are then retained, while the other weights are all set to zero, and kept at zero for the rest of the simulation. All weight values are then re-normalised, resulting in an increase in value for the retained weights, which ensures that cortical activity level values are maintained across the simulation. The third step is to set to zero the lateral excitatory and inhibitory connections from neighbouring units that favour a different SF channel. This pruning process results in cortical units only receiving afferent input from a single SF channel (the channel favoured at the time of pruning), and only having lateral connections to other cortical units that favour the same SF channel. Essentially, the pruning process allows us to compare an architecture with parallel independent SF channels in V1, with one where the channels are mixed. The following text outlines the pruning process is more detail.

The response of a cortical unit (i, j) was given in chapter 3 by equation 3.4. For clarity, the equation is expanded below:

$$\begin{aligned} \eta_{ij}(t) &= \sigma\Big(\sum_{p\in[a,b,c,d]} \gamma_{LGNp} \sum_{kl\in CF_{pij}} X_{pkl}(t-1)\omega_{pkl,ij} \\ &\quad + \gamma_E \sum_{kl\in CF_{Eij}} X_{Ekl}(t-1)\omega_{Ekl,ij} + \gamma_I \sum_{kl\in CF_{Iij}} X_{Ikl}(t-1)\omega_{Ikl,ij} \Big) \end{aligned} \tag{4.2}$$

The first summation term gives the contribution from the 4 thalamocortical afferent projections from RGC/LGN sheets a, b, c, and d (see figure 4.1). The last two summation terms refer to the lateral excitatory feedback and the lateral inhibitory feedback.

In the first step of pruning a cortical unit (i, j), all afferent weights are set to zero except those in the afferent connection field with the largest sum of weights:

$$\omega'_{pkl,ij} = \begin{cases} 0 & p \neq m_{ij} \\ \omega_{pkl,ij} & p = m_{ij} \end{cases} \tag{4.3}$$

where $m_{ij}$is the projection containing the connection field whose sum is largest for unit (i, j). For example, if for unit (i, j) the sum of weights for the connection field from sheet b is the greatest, then all weights for connection fields for unit (i, j) from sheets a, c, and d are set to zero. The weights for the winning projection are then renormalised as in equation 3.11.

Finally the lateral connections are pruned, so that all weights within the lateral excitatory and inhibitory projections to unit (i, j) which correspond to cortical units which after the first stage of afferent pruning do not share the same afferent input are set to zero. The lateral connection fields are then renormalised.

Pruning, when used, took place at time step 4000, with training then continuing until time step 5000 using the same input patterns as for the unpruned case. The choice of step 4000 was based on the maps being reasonably settled and stable by this point, while allowing the simulation to reorganise after pruning. Longer simulations were conducted with pruning taking place at different time steps and with greater reorganisation periods, but with little difference to the qualitative nature of the results.

4.3 Results

In this section results are presented for the unpruned model, a model with pruning, and a model with two widely separated SF preferences (for comparison). The results suggest that to obtain a wide range of SF preference in V1, it is necessary for each V1 neuron to receive inputs from a single SF channel, rather than combining multiple channels. The methodology for measuring RFs is also discussed.

4.3.1 Single layer V1 model - unpruned

After presentation of 5000 grey scale natural images without pruning, afferent weights for V1 units have developed OR selectivity, and a clear global organisation for OR has developed (figure 4.2A). The measured OR map bears a strong similarity to animal OR

maps (see figures and discussion in section 2.3, or explicit comparison in figure 4.12), and includes the signature features of orientation pinwheels, iso-orientation zones, linear zones, and fractures. The inclusion of multiple SF channels increases the number of feature dimensions that drive cortical activity. This could potentially have led to a breakdown in organisation for OR preference, but instead the model is still able to develop a well organised topographic OR map.

The SF preference map, figure 4.2B, has a topographic organisation with iso-SF areas and linear zones, as seen in the recent experimental maps. However, the spectrum of SF preference represented in cortex is quite narrow (histogram in figure 4.2C). It spans less than 2 octaves, with most of the neurons favouring intermediate SFs in the central region of the SF input range, and few neurons with SF preference at the higher and lower ends of this range. In this model, RGC/LGN sheets a and d have DoG RFs with SF preferences of 0.9 cpud and 3.6 cpud respectively, and we would expect this range to be reflected in the cortical sheet.

The RFs of cortical neurons in the unpruned model do have distinct ON and OFF sub-regions, and there is some variation in RF size, and in SF preference tuning, (figure 4.3). The RFs are typically blobs, or two or three region Gabor type patches (see figure 4.4), although the RFs are not always well formed.

As discussed in section 2.4, the experimental OR and SF maps tend to have high-gradient regions that do not overlap. Figure 4.5 shows that this property also hold for the model maps. It has also been observed that when high gradient regions from different feature maps do spatially coincide there is a tendency for the gradients at these points to be orthogonal (Yu et al., 2005). Determining whether this property holds for the model would need more extensive analysis, as upon inspection there is no obvious relationship between the two gradient maps. The gradient in these maps was defined at a point (i, j) as:

$$g_{ij} = \sqrt{\Big((F(i+1,j)-F(i,j))^2+(F(i,j+1)-F(i,j))^2\Big)} \tag{4.4}$$

where $F(i, j)$ is the value of the feature preference at that point (same procedure as

Blasdel, 1992). For OR maps F is cyclic so any difference in OR preference exceeding $\frac{\pi}{2}$ is subtracted from π, ensuring that all differences are acute (and not a mixture of acute and obtuse), and thus remain within the permissible range.

Alignment of OR pinwheel centres with high and low SF preference iso-zones has been reported in some previous experimental and modelling work, but in this simulation we have not found any relationship between pinwheel position and SF features (see figure 4.5, bottom row). This is discussed further in sections 5.3.4 and 5.4 in chapter 5, and in section 6.3.4 of chapter 6.

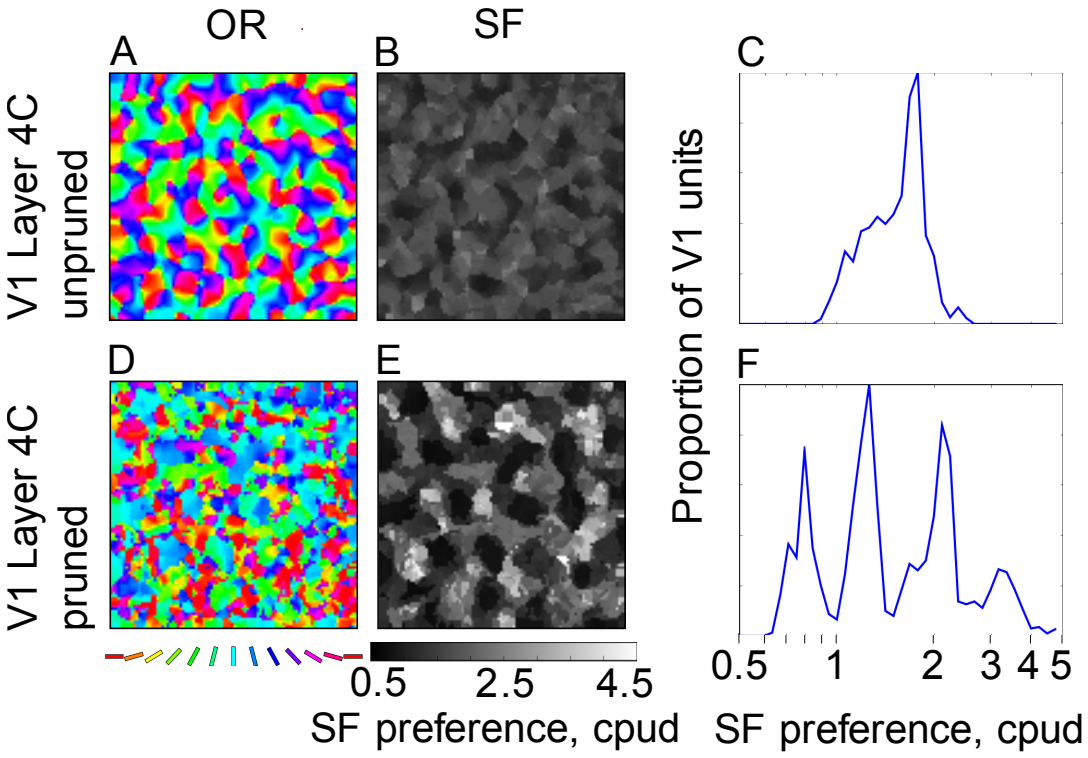

Figure 4.2: Results for single layer V1 model.

(A and D) V1 OR preference maps with and without pruning. (B and E) SF preference maps with and without pruning. (C and F) Corresponding normalised SF preference histograms for V1 with and without pruning. Pruning increases the range of SF preference, so that it spans the full range of the SF channels, at a cost to the smoothness of the OR map. The four peaks observable in the pruned SF histogram correspond to the preferred SF tunings of each of the four pathways.

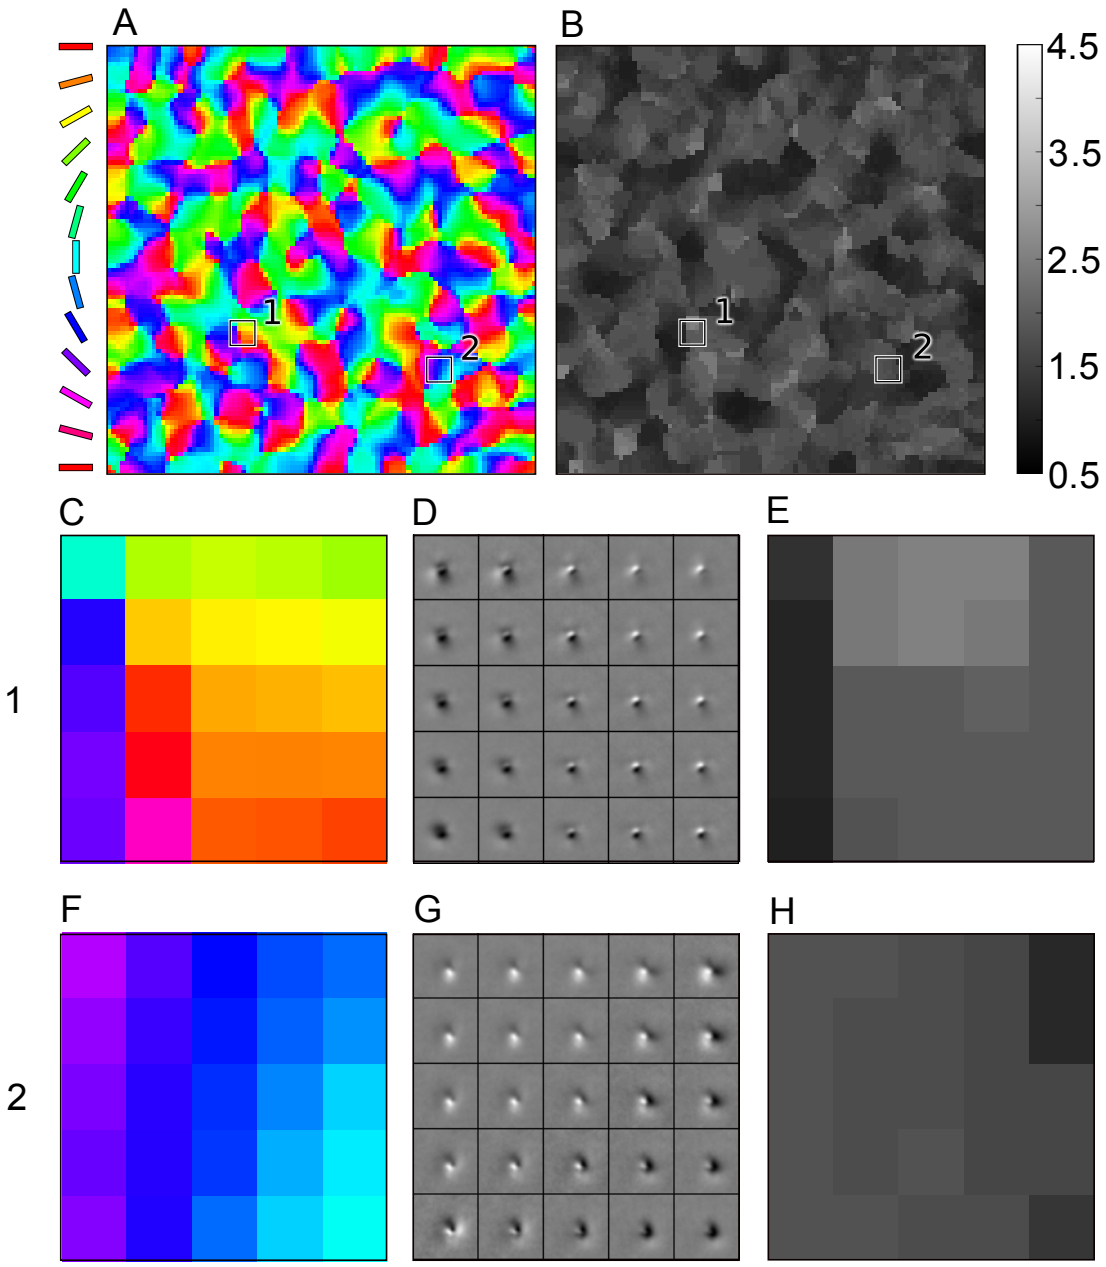

Figure 4.3: RFs for unpruned single layer V1 model.

(A and B) OR and SF maps for unpruned single layer V1 model. (C, D and E) Details of OR map, corresponding RFs, and SF map for boxed region 1 in A. (F, G and H) Details of OR map, corresponding RFs, and SF map for boxed region 2 in A. Note the small variation in size of RFs, and the small range of SF preference values in E. The RF plots have been cropped to aid viewing, and are each shown on a 75×75 pixel region of retina. Scale shows greyscale SF preference in cpud for B, E, and H.

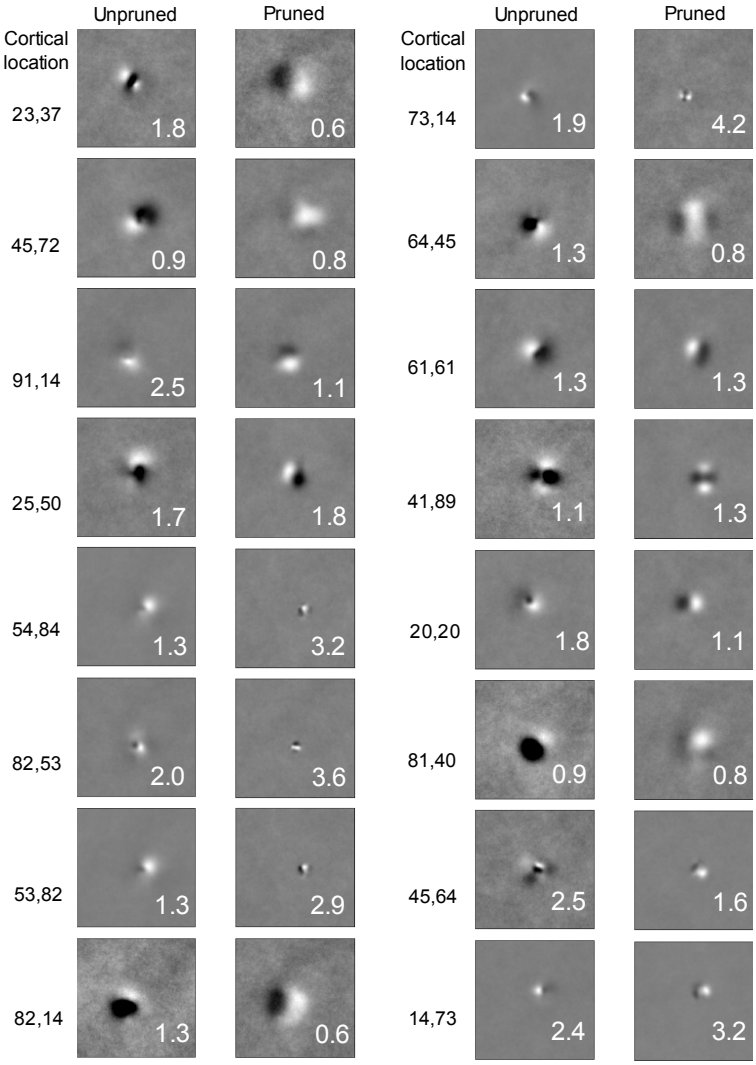

Figure 4.4: Single layer V1 model, Representative RFs.

Representative RF plots from the unpruned and pruned simulations, for neurons at the specified cortical locations. A range of RF profiles exist: blobs, 2 lobes, and 3 lobes. Numbers inset into the RF plots are the neuron's SF preference in cpud from the SF map. The increase in range of SF preference after pruning is reflected in the diversity of the RF plots.

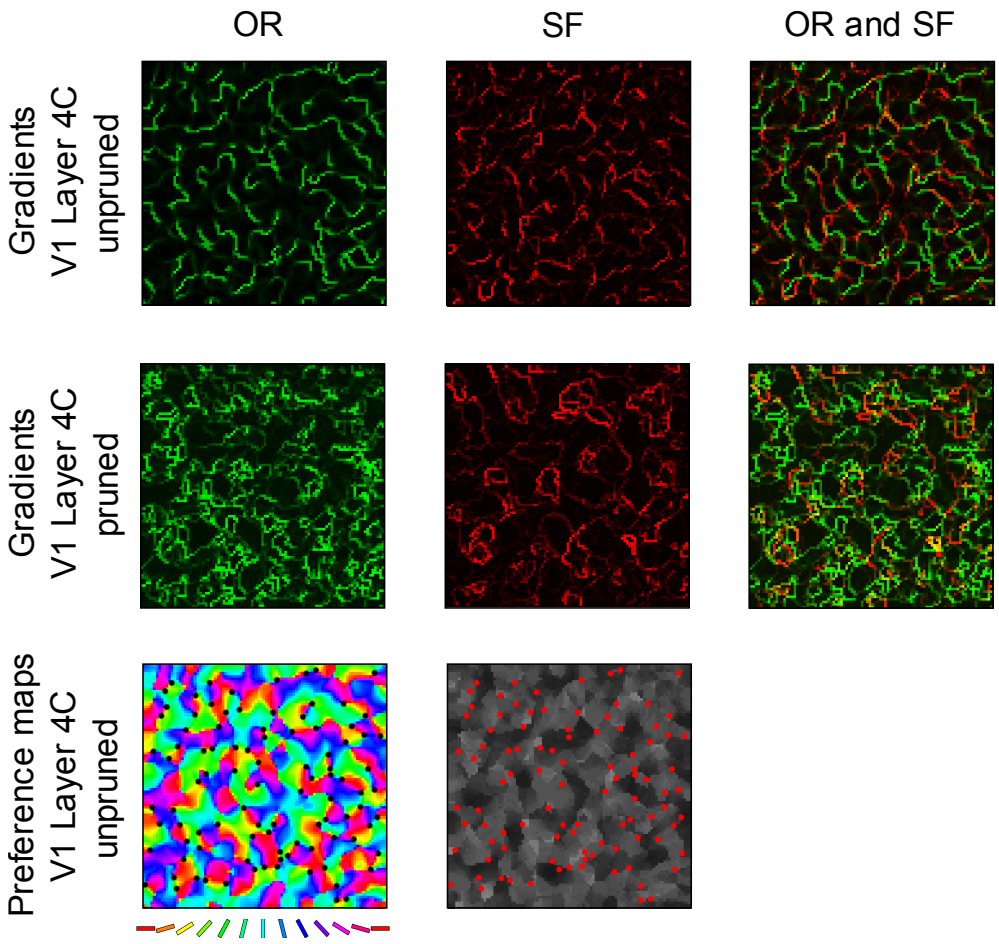

Figure 4.5: OR map and SF map intermap comparisons.

(1st row) Gradient comparisons for single layer V1 model without pruning. There are many fractures in the OR and SF maps, illustrated by the prevalence of high gradient regions. In the gradient overlay plot, OR and SF fractures are typically non-coincident; areas of overlap are rare and appear yellow. (2nd row) Gradient comparisons for single layer V1 model with pruning. High-gradient regions in the maps correspond to fractures and pinwheels. (3rd row) OR map for unpruned single layer V1 model, with pinwheels marked by black dots. Pinwheels are marked with red dots on the corresponding SF map, and do not appear to have any consistent relationship with it.

4.3.2 Single layer V1 model - pruned

To understand the role of each SF channel, the simulation was repeated, but with pruning (see section 4.2.2) at time step 4000. The maps were measured again after a total of 5000 image presentations (figures 4.2 D, E, and F). The pruning process affects the OR organisation, resulting in a more disorderly OR map. For instance, iso-orientation zones are reduced in size, there are more fractures between iso-orientation zones, OR pinwheels are less frequent, and there are more isolated islands and misplaced units. However, such features are often seen in animal maps depending on how much the animal maps are smoothed by the measurement process, and some damage to the OR organisation is to be expected due to the severity of this artificial pruning process (see discussion below).

The pruning also has a large impact on the SF map, which is now quite different - it still features clear iso-SF zones, fractures and linear zones, but it now captures the entire range of SF preference found in the RGC/LGN sheets. In contrast to the unpruned results, there are many cortical neurons with extreme SF preference, as seen in the SF preference histogram (figure 4.2F) with 4 peaks, one for each of the SF pathways. There are roughly equal numbers of cortical neurons that receive projections from each of the SF pathways after the pruning process. This shows that Hebbian learning, homeostasis and normalisation have to some extent succeeded in representing the breadth of SF available in the input space, but that the architecture leads to intermediate feature value preferences. By changing this architecture (through the pruning process), the underlying even distribution across the SF channels is revealed.

In the pruned model, the range in size of the V1 unit RFs is dramatically increased, with high and low SF preference units often having well formed two-or three-region Gabor-like patches (figure 4.6), similar to experimentally measured RFs (e.g., Jones and Palmer, 1987b). Figure 4.4 includes representative neurons from both the pruned and unpruned simulations. In the pruned case the range in SF preference and RF size is much greater, and the overall OR selectivity is also higher. Once the cortical neuron's input is restricted to a single SF pathway, it is better able to organise, leading to the development of RFs more tightly tuned for OR as well as SF.

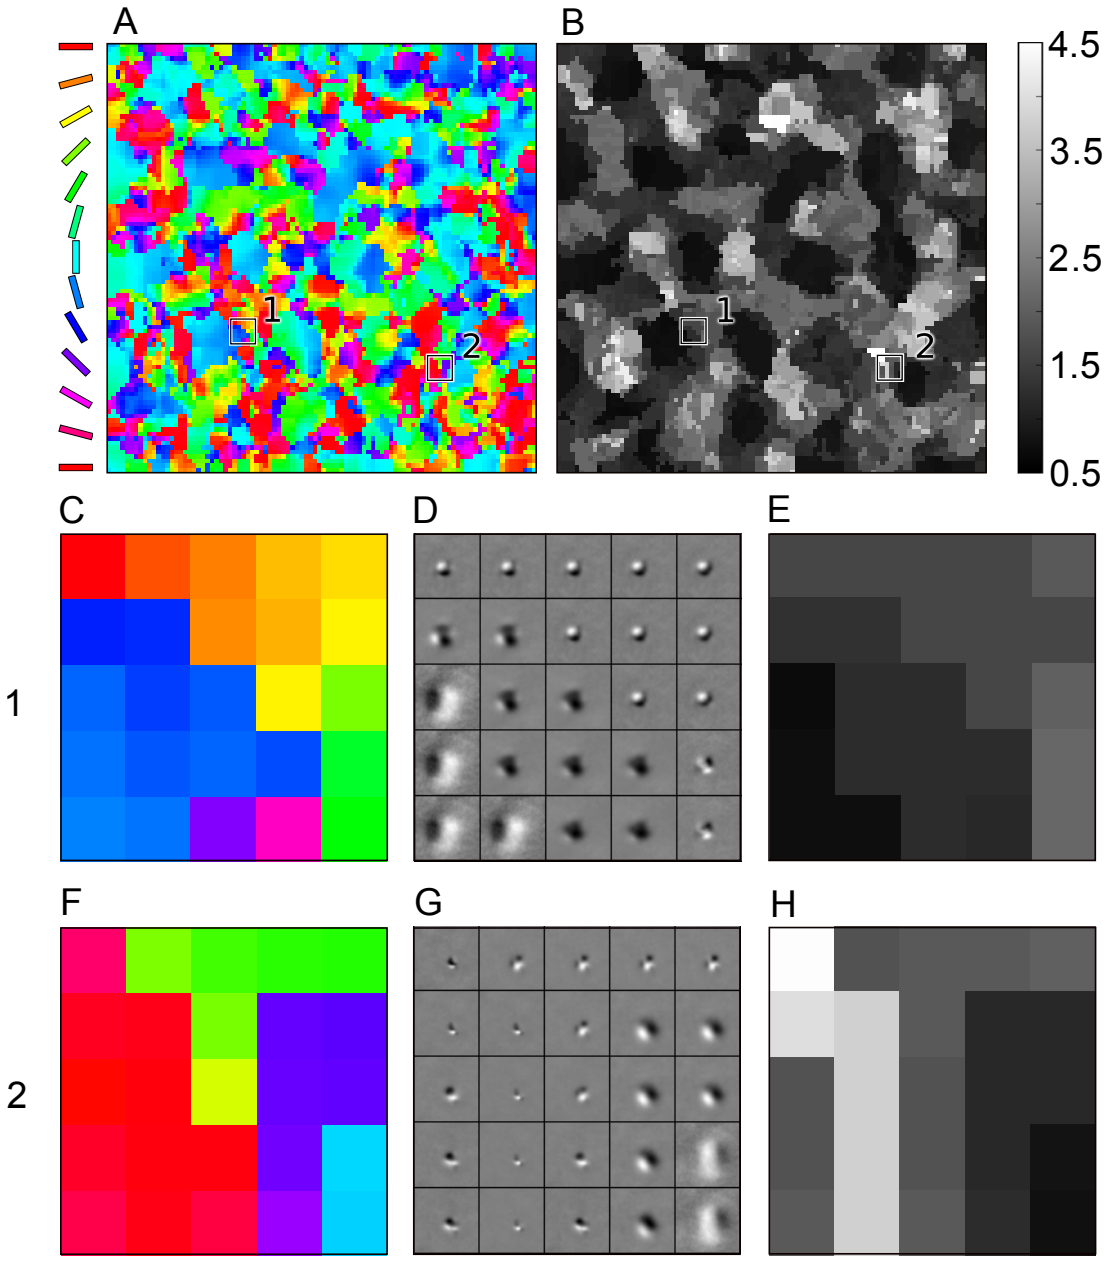

Figure 4.6: RFs for pruned single layer V1 model.

(A and B) OR and SF maps for pruned single layer V1 model. (C, D and E) Details of OR map, corresponding RFs, and SF map for boxed region 1 in A. (F, G and H) Details of OR map, corresponding RFs, and SF map for boxed region 2 in A. Note the large variation in receptive field size (compare with figure 4.3), and the good matches of RFs to their corresponding values in the OR and SF maps (as we expect). The RF plots have been cropped to aid viewing, and are each shown on a 75×75 pixel region of retina. This is the same scale as figure 4.3, which emphasises the dramatic change as a result of pruning. Scale shows greyscale SF preference in cpud for B, E, and H.

The increased disorganisation of the OR map due to the pruning process is evident in the gradient plot (middle row, figure 4.5), where there are many more high gradient regions than in the unpruned gradient plot. In the pruned network, no relationship is obvious between the OR and SF gradient maps, due to the larger degree of disorder.

4.3.3 Binary single layer V1 model

As a control to verify the necessity for multiple SF channels, the single layer V1 model architecture was modified by removing the central two SF pathways, leaving just the highest and lowest SF preference RGC/LGN sheets and their projections (we refer to this model in the text as the binary LGN model). Running this model for 5000 image presentations allowed it to organise well for OR and SF (figure 4.7). The OR map contains the key features found in observed topographic maps. The SF map is binary, featuring only high and low SF preference regions, as observed in some studies. Its appearance is similar to other binary maps, such as the ocular dominance map. This two-state organisation is a direct consequence of the reduction in the overlap of activity between channels, and thus weaker correlations, due to the increased size separation between the SF pathways.

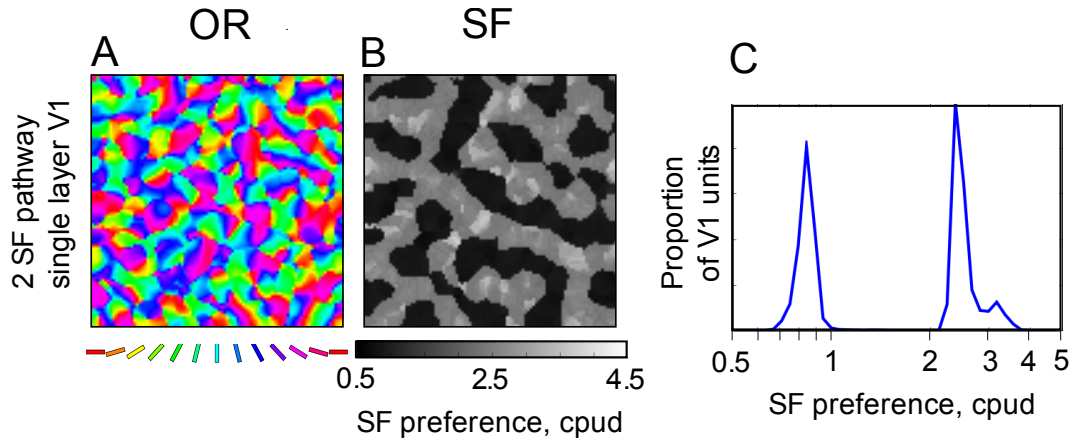

Figure 4.7: Binary SF organisation, for 2 SF channel single layer V1 model.
(A) Orientation preference map. (B) SF preference map. (C) SF preference histogram. The SF map is binary, lacking units with intermediate SF preference, but it covers a broad range of SFs.

Interestingly, the SF range is greater in the binary model than the range found in the single layer V1 model, despite both having identical high and low SF pathways. Also of importance is the failure of intermediate SF preference units to arise, given only these two SF channels. This result further supports the argument that V1 units' SF preference is dependent upon the SF preference of the LGN cells from which they receive afferent connections. Note that cortical units still have very broad tuning and have not learnt a single, exclusive SF channel preference; the preference a unit has for one channel over another is typically no more than 20%.

4.3.4 Partial pruning

The pruning process was designed to be a clear illustration of the effects of segregating pathways, but such clear segregation is unlikely to be realistic. By altering the pruning regime so that at time step 4000 only 20% of V1 neurons undergo the pruning process, it is possible to generate SF maps which have a wide range of SF preference while retaining a well organised OR map. The pruned neurons are selected randomly, and their afferent and lateral connections are adjusted in the same fashion as in the fully pruned case; the pruned units are only connected to lateral units which themselves have been pruned and receive the same afferent SF channel's input. As previously described, the simulation is then run to time step 5000 to allow it to re-organise. Figure 4.8 shows the results of this procedure. Note the increased range of SF preference within the well organised SF map, while the OR map remains well structured, but with some units dissimilar to their neighbours.

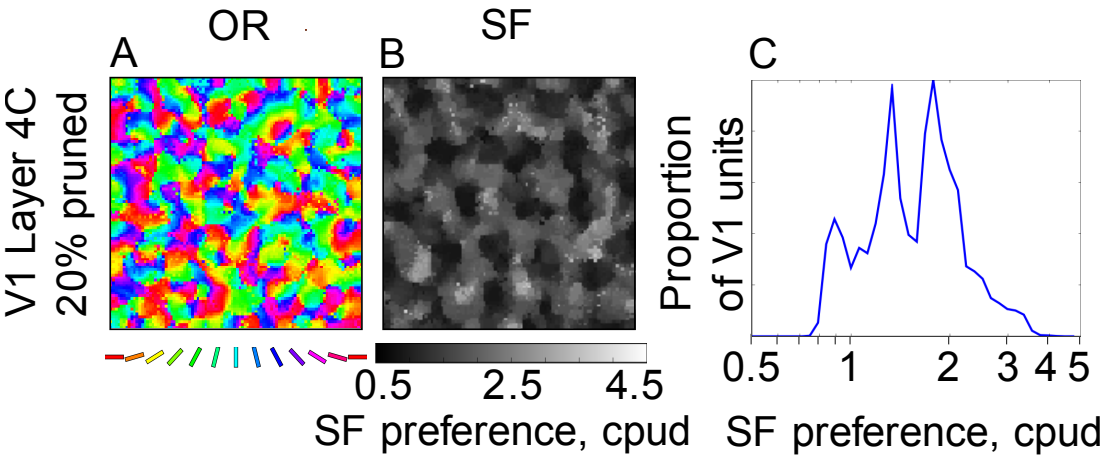

Figure 4.8: Single layer V1 model, with 20% pruning.

(A) Orientation preference map. (B) SF preference map. (C) SF preference histogram. 20% of the neurons in V1 have been pruned. The OR map is well organised, but with some mismatched pixels. The SF map is also continuous and well organised. The distribution of SFs within the map is notably increased (compared to figure 4.2C) by this mild pruning, without significantly disrupting the OR map.

4.3.5 Map development from spontaneous retinal activity

As discussed in section 2.5 in chapter 2, it is known that OR maps first develop before eye opening and that spontaneous retinal activity is necessary for this. Although it is not known whether SF maps exist prior to visual experience, it is also possible that spontaneous retinal activity may be important for the formation of SF maps, as neurons with SF selectivity do arise without visual experience (Derrington, 1984). The bulk of the modelling work contained in this thesis uses natural images as visual stimuli, as it is known that normal visual experience is required for the maturation of neuron selectivity and for maps to fully develop. To test whether the model results depend on the specific structure and other properties of natural images, a single layer V1 model (but with a reduced cortical density, $N_{V1} = 48$) was trained on patterns containing randomised elongated Gaussians (see figure 4.9A).

These stimuli are clearly very different from natural images, retaining mainly spatial smoothness, spatial locality, orientations, and a range of feature sizes. As figure 4.9 demonstrates this form of stimuli can drive the model as well as natural images, resulting in OR maps with the key features, as well as a continuous SF map. As in the natural images case, the model needs pruning to represent the full SF range in the map. These results demonstrate that only a few features of natural images are important for the development of SF and OR maps in the model, and thus that such maps are likely to develop from spontaneous activity as well, with a range of preference values reflecting the range present in the activity patterns.

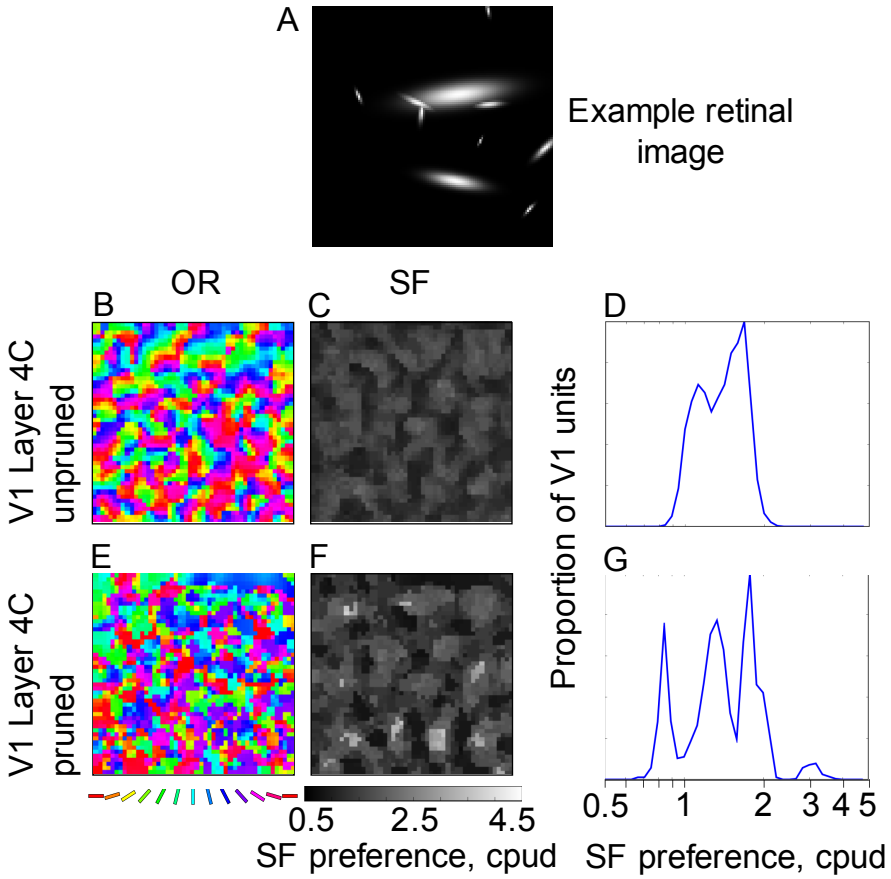

Figure 4.9: Sample stimuli, and results for single layer V1 model with randomised Gaussian presentations.

(A) Example retinal image. Each pattern presented was a randomly generated composite of 10 elongated Gaussians. Each Gaussian had an aspect ratio of 4.0, and was randomly positioned and orientated within the window. The Gaussians had size constraints, such that 4 of them had standard deviations randomly chosen between 0.02 and 0.05, 3 Gaussians had standard deviations between 0.05 and 0.1, 2 between 0.1 and 0.2, and the final one between 0.2 and 0.4. This distirbution roughly equalises the activity for each channel. The Gaussians were not added together, but were composited such that a given pixel would take the maximum value from each individual element. (B and E) V1 OR preference maps with and without pruning. (C and F) SF preference maps with and without pruning. (D and G) Corresponding normalised SF preference histograms for V1 with and without pruning. Pruning increases the range of SF preference, so that it spans the full range of the SF channels, at a cost to the smoothness of the OR map. The four peaks observable in the pruned SF histogram correspond to the preferred SF tunings of each of the four pathways. The results from these simulations show that the basic properties of map development depend only on a few features of the input images, and thus that such maps are also likely to develop from spontaneous visual system activity before eye opening.

4.3.6 Contrast invariance

For simplicity and to focus on map development, the models within this study do not incorporate contrast gain control. As a result, learning and development are dominated by the high contrast features in images. It is possible that this results in a reduction in the selectivity of neurons for SF, if the information in low-contrast regions has higher diversity in SF preference. To investigate this possibility, simulations of the single layer V1 model were run (with a reduced cortical density, $N_{V1} = 48$), but with the input images processed using the Naka-Rushton compressive nonlinearity equation:

$$L(i) = \frac{i^n}{i_{50}^n + i^n} \tag{4.5}$$

where L is the output pixel activity level after the nonlinear scaling, i is a pixel's input activity level, i_{50} is the input pixel level that will generate half-maximal response from the photoreceptors, and n is an exponent that defines the gradient of the function. The Naka-Rushton equation has been shown to describe the contrast response functions found in V1 for some animals (Albrecht and Hamilton, 1982), although within any individual animal, cells vary greatly in the nature of their contrast response. Here, the equation is used merely to pre-process the training images, as the model does not include retinal contrast adaptation, and thus is applied to pixel activations rather than visual contrast. Simulations were attempted for a range of parameters of the Naka-Rushton equation, including values matching those for cat (Zhang et al., 2007), $n = 1.6$ and c_{50}=27% (the contrast that generates half-maximal response), corresponding to $i_{50} = 0.27$ for the model (for pixel activity on a scale from 0 to 1). An example of a filtered input image is shown in figure 4.10A.

As can be seen from the results summary (figure 4.10B) the compressive nonlinearity filtering of the input images does not increase the SF preference range of the developed maps. However, as the function becomes increasingly step-like, there is a shift towards lower SF preference in the map. This occurs because the step-like nonlinearity has the effect of removing a lot of the detail found in low contrast regions of the input images, causing the images to become dominated by lower spatial frequencies, resulting in lower SF preference in the map.

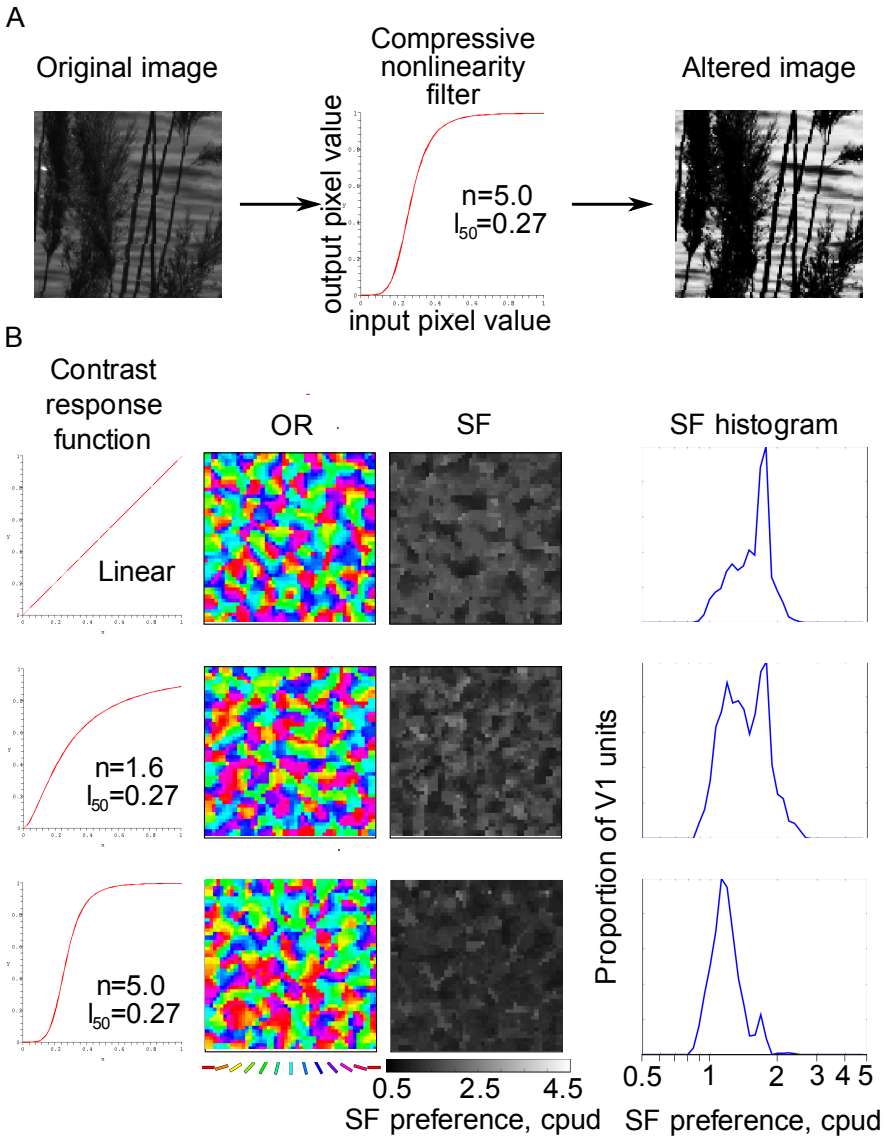

Figure 4.10: Unpruned single layer V1 model OR and SF maps, resulting from input data filtered using a compressive nonlinearity.

(A) Example training data filtered by the Naka-Rushton compressive nonlinearity. (B) Each row shows the image filter used for a particular simulation run, along with the OR and SF map results. Note the shift in SF preference towards lower SF in the simulation with more extreme contrast response functions; realistic amounts of contrast nonlinearity improve the spatial distribution but do not otherwise alter the main conclusions.

4.3.7 Bandwidth

In addition to feature preference maps, their associated bandwidth maps are needed to provide a full understanding of units selectivity and response profiles. This information is necessary to predict how units respond to stimuli (Baker and Issa, 2005; Mante and Carandini, 2005; Zhang et al., 2007). Bandwidths were calculated for OR and SF preference by presenting sine gratings at fixed contrast, over a range of SFs and ORs, in the same fashion as for calculating OR, and SF maps (described in section 4.2.1). Due to the lack of contrast invariance in the model, test stimuli were presented at fixed contrast. The bandwidth for a unit was defined as the standard deviation of its response to OR, or SF presentations. Plots of the OR and SF bandwidths for the pruned and unpruned case of the single layer V1 model (with a reduced cortical density, $N_{V1} = 48$) are found in figure 4.11. In the pruned case, the range of bandwidths increases, for both OR and SF preference. There are more units that are more tightly tuned, with smaller bandwidths, but also more units with broader tuning and larger bandwidths. The narrower range of bandwidths found in the unpruned case is likely to arise due to the correlation between the SF channels' activity.

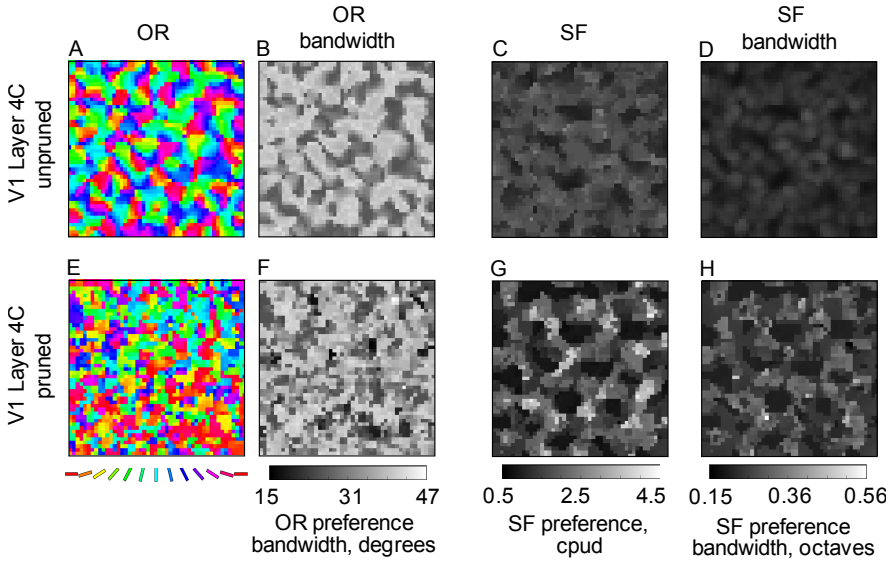

Figure 4.11: OR, SF and bandwidth maps for the single layer V1 model.

(A, B, C, and D) OR, OR bandwidth, SF, and SF bandwidth maps respectively, for the unpruned single layer V1 model. (E, F, G, and H) Corresponding plots, for the pruned single layer V1 model. Here bandwidth is defined as the standard deviation. Note the increase in range of SF bandwidth after pruning.

4.3.8 Comparison to experimental data

In comparing the results from these simulations to experimental results from this part of the visual system, we are asking whether the model succeeds in fulfilling the modelling criteria introduced in chapter 3. Although the comparison between this work and experimental results is not quantitative, the results show the model meets each of the modelling criteria, albeit with a clear caveat. This caveat is the requirement of SF pathway separation to achieve a wide range of SF preferences, through either explicit pruning or a bimodal and extreme RGC/LGN size distribution.

In figure 4.12, details from published SF maps (and where possible matching OR maps) that were discussed in detail in section 2.3 are presented on the same spatial scale, along with details from the simulation map results from figure 4.2. Despite the experimental SF maps being produced using a variety of methodologies, often in different species, and displayed using different SF keys, there are striking feature similarities amongst most of these maps. In several, the size of iso-SF zones is typically comparable to the size of iso-orientation zones, and there are SF preference fractures and SF preference linear zones. These features are all also present in the single V1 model results. The computational SF maps have qualitatively similar spatial arrangements within the range of the experimentally measured SF maps, including a loose tiling across the cortex of iso-SF zones, which is observed in all the experimental maps, and is required for coverage. The computational OR maps are not as smooth as the experimental maps, but this is to be expected because the model results are not smoothed, while experimental maps are the product of averaging over many trials and are thus implicitly heavily smoothed.

With the pruning process, the model fulfils the remaining criteria, by attaining a representative broad SF preference range in V1. The range of size of RFs and their range of SF preference spans the range of preference in the RGC/LGN layer. The RFs in both the pruned and unpruned models resemble simple cells with two or three elongated ON and OFF subregions, but in the pruned model the RFs are often more selective for OR. As discussed below, the impact the pruning process has upon the OR preference organisation is likely to be due to the form of its implementation, which may be minimised by altering this mechanism.

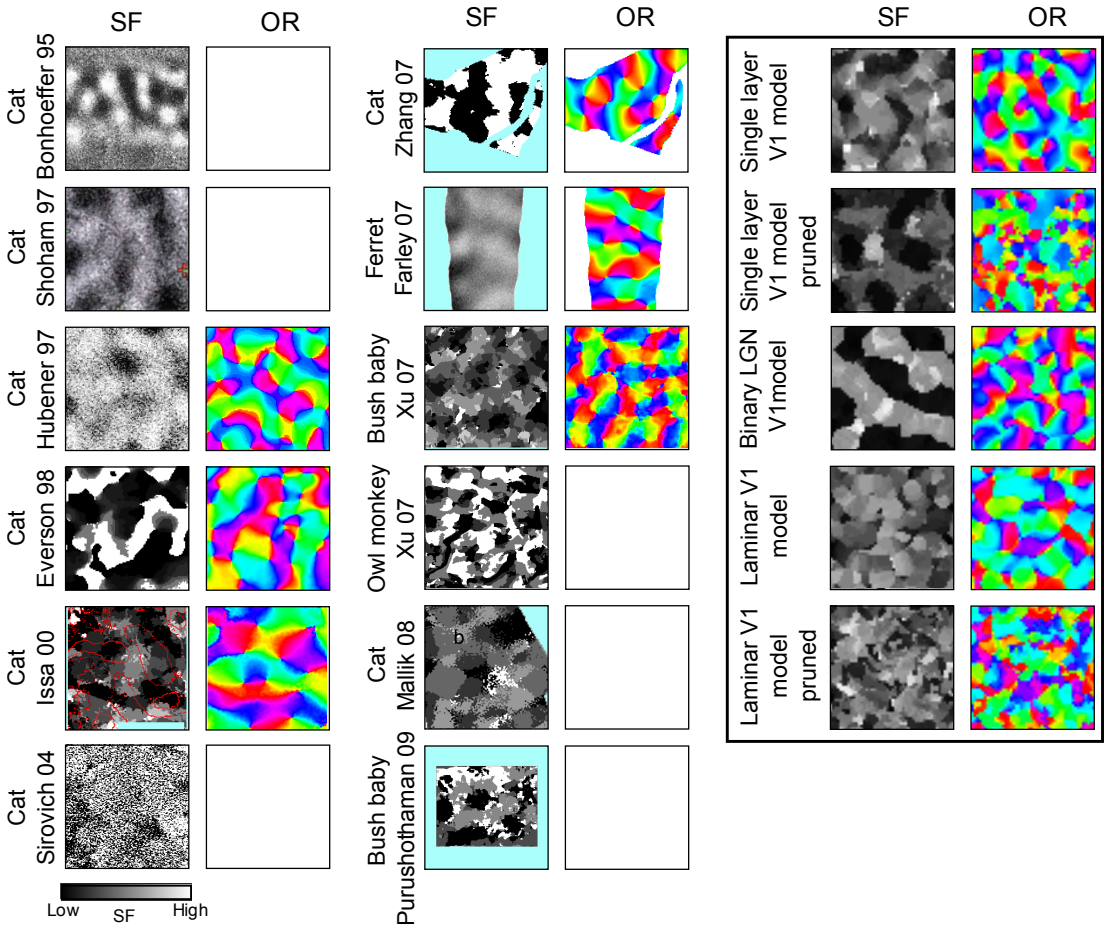

Figure 4.12: Comparison between experimental and computational SF and OR maps. Representative SF maps and (where available) matching OR maps from the same region of cortex, from published experimental work on SF mapping. Maps have been cropped to fit in 2.5mm square boxes. SF maps have been recoloured to grey-scale to aid visual comparison, although they are still on different numerical scales (note that different methodologies were used to construct each map, making them not usually directly comparable). The OR colour scales also differ between the various maps. All original scales and maps can be found in chapter 2. The box at the right shows details from the central regions of each modelling result shown in figures 4.2 and 5.8 (to be discussed in chapter 5, but included here to allow a complete comparison of results).

4.4 Discussion

Although the single layer V1 simulation is able to produce topographic OR and SF preference maps, it is only able to produce neurons that span the entire SF input space through a separation of the SF pathways, achieved in this case by the pruning of connections. If the individual contributions to a cortical units RF from each SF pathway are considered in isolation, we find that each SF pathway contributes its own RF, which has a size appropriate to its SF pathway, and that the final cortical RF is in fact a superposition of these distinct channels' RFs. There is also superposition of neighbouring cortical units' RFs, due to the lateral inter-connectivity.

Although cortical units have not developed a preference for only one channel, a difference in channel strength has developed through Hebbian learning. This difference is typically quite small. The binary simulation demonstrates that the model does have the capacity for adaptation to favour a particular channel through Hebbian learning, but between-channel correlations prevent it from developing exclusive connections to a single channel. This situation frequently arises in two ways:

- Due to the convergence of SF pathways upon a cortical unit, a unit will be maximally activated by patterns that activate all or most of the channels. Extreme SF stimuli activate mainly the extreme SF channels, which results in less excitation than for broadband stimuli and thus less Hebbian adaptation.
- As a direct consequence of using a DoG model for the LGN, the range of SFs to which an LGN cell will respond is very large. Although the different SF channels have differing SF preferences, and thus differ in the SF to which their response is highest, there is a large overlap in their SF response range. Extreme SF stimuli will thus not solely activate the more extreme SF channel, but also immediately adjacent SF channels, producing correlations in their activity.

These factors make it unlikely for stronger connections to develop to a more extreme SF channel over its neighbour, and discourage preference for a single SF channel to develop. Nonetheless, a small preference for one channel over the others does develop, and it is this that is accentuated by the pruning mechanism.

In order to try to increase the small SF channel preference range that arises, simulations using alternative parameters were run. Systematically altering parameters within the sigmoid activity equation (3.5), the learning rates, and the projection strengths, allowed an exploration of this multivariate space. This is a large parameter space to explore, and cannot feasibly be explored exhaustively. Instead, we ran simulations for different ranges of one parameter or two parameters at a time. This leads to a known range for which successful model organisation is known to occur for that particular parameter subject to the other parameter values. Parameters are often not independent of one another in their effects upon the model; for example, increasing the strength of the excitatory lateral connections can be offset to some extent by increasing the strength of lateral inhibitory connections. The inter-dependence of parameters upon the model was explored on an ad hoc basis. On the whole, the model is quite robust to changes in parameter values, but no range was found which allowed exclusive SF channel connections to develop. The failure to find a range of parameters that allow the model to develop a full range of SF preference without pruning, led to the conclusion that Hebbian plasticity alone is insufficient to do this for this model architecture, and motivated the need for some additional mechanism to achieve full range SF preference.

These results suggest that Hebbian learning may not be a suitable learning mechanism for capturing the full range of a non-cyclic feature space sampled densely, where the stimulus results in overlapping activity responses in the separate pathways. Simulations of ocular dominance mapping using LISSOM (see Miikkulainen et al. 2005) also find that no neurons occupy the extreme ends of this non-cyclic preference space when inputs are correlated, which in this instance is appropriate as V1 neurons are typically binocular. However, when the stimuli presented are not correlated, the binocular LISSOM model is able to represent the extrema, and the resulting ocular dominance map is dominated by neurons with single eye preference exclusively. The binary LGN V1 model (figure 4.7) demonstrates that SF preference can develop in a similar way to ocular preference, i.e., with no units with intermediate SF preference. In the case of SF preference, it is necessary that the entire feature space is represented, which requires additional SF channels, but as we have seen introducing more channels results in a narrower representation due to the activity overlap.

As previously acknowledged, the pruning process is highly detrimental to the organisation of the OR map. The pruning process is artificial and extreme, severing three quarters of all of the cortical connections in one time step! Also, the basis of the severing decisions is purely a comparison of weights strength between different SF channels. There are various changes that could be made to the pruning mechanism to try to retain the same broadening for cortical SF representation, while lessening the impact upon the OR map. For instance, increasing the time period over which pruning takes place may allow the cortex time to adapt to the change to its input structure so that it may re-organise. This could be done by pruning a fraction of neurons every time step (e.g., stochastically), over a long time period, say 1000 iterations, or, instead of severing a connection, its weight values could be gradually reduced over time.

Alternatively, instead of pruning all neurons, a selection process could be introduced such that only neurons with a severe imbalance in afferent weights are considered for pruning. Or only neurons where all the afferent pathways share a common OR tuning could be pruned, which would mean the OR map would be unaltered. Some combination of these approaches may also be possible. Also, it may not be necessary to prune all connections. A modest amount of pruning would still allow many units to attain preference for SF at the high and low ends of the scale, which may be sufficient to reflect the requirements of SF coverage in V1. The criterion for selecting the pathway to retain, by comparing the total sum of weights in each channel, was chosen arbitrarily and may not be the best approach. Other selection criteria, such as retaining the pathway with the most highly tuned OR preference, may also yield SF organisation with less disruption to the OR mapping.

The preceding discussion focuses on ways to implement pruning directly, each of which requires various strong and potentially implausible assumptions about the underlying biological mechanisms. Rather than a physical separation between the SF channels arising through some kind of pruning process, there may be an alternative form of separation between the channels that may serve the same end. For instance, there is evidence to suggest that similarly sized LGN cells have the same temporal firing properties (latency and duration, see Weng et al. 2005), and that larger LGN cells have shorter response latency than smaller LGN cells. This could potentially pro-

duce an inherent temporal segregation between SF channel activities. Effectively, these differences in activation pattern may allow Hebbian learning processes to distinguish between the SF channels, letting preference for an individual channel develop.

Note that the organisation of TF in V1 is controversial. There have been three published optical imaging studies relating to TF mapping, and they have reported different results. Shoham et al. (1997) observed correlations between SF, TF and cytochrome oxidase blobs, with low SF preference regions in cortex coinciding with cytochrome oxidase blobs, and also having high TF preference. Khaytin et al. (2008) found a uniform organisation for TF with no spatial domains. The same group (Purushothaman et al., 2009) reconsidered this data using what they consider to be a superior analysis approach that allows for variance in the data set, as well as conducting a further study. They conclude that there are TF domains in both data sets, and find their results to be in line with those of Shoham et al. (1997). The model simulations so far take no account of the temporal properties of neurons, and it remains important to try to address this issue in future work.

In primate and cat the range of SF preference at a given eccentricity can be as much as 4 octaves, but the model described above accommodates a range of 2 octaves. This range restriction has been imposed by the computational requirements of such a model, because early attempts at representing the full SF range turned out to be impractical on current computer hardware. However, there is no reason to suggest the inclusion of additional SF channels to extend the SF range would be unsuccessful, given sufficient computing power and memory.

4.5 Conclusion

This work shows for the first time a mechanistic developmental model of the visual system that is able to organise realistic topographic OR and SF maps by learning from natural images. The model is able to fulfil the full modelling criteria (introduced in section 3.1), but doing so requires an enforced separation of the SF pathways achieved by explicit pruning. The pruning process is necessitated by the limited representation

for SF preference that results from purely Hebbian learning. To achieve a wide range of SF specialisation, mechanisms additional to Hebbian learning of the spatial patterns may be involved. For instance, temporal differences between LGN neurons of different SF preference may be sufficient to produce segregation between the afferent activity signals, allowing Hebbian processes to distinguish different SF pathways and develop strong SF preference. Further experimental investigation into whether there is a developmental link between SF and TF preference may resolve this issue.

The potential requirement for a functional pathway separation for the processing of SF is not a new idea, but the separation at the individual neuron level has not been fully considered and deserves experimental attention. I suggest that at some point in development, V1 neurons that receive afferent LGN input may have connections to a wide range of LGN cells with different SF preferences, but as they develop SF preference, their afferent connections change so they primarily have connections with LGN cells with similar SF preference. Further experimental study into the early and developed thalamocortical connections with respect to SF preferences is required to address this point. Might this potential separation of SF pathways extend to lateral connectivity? In the model, pruning of the lateral connections is required to elliminate the superposition of adjacent neurons RFs, but only pruning of the afferent connections is required to generate a large SF preference range. Investigating the interconnectivity between iso-SF zones would aid in determining whether lateral connections may be similarly segregated, which is achieved explicitly in the current model but may arise gradually in models with other types of pruning or using temporal correlations.

Chapter 5

Laminar V1 model

5.1 Introduction

Now that it is clear what can be achieved with a single-layer V1 model, in this chapter we test whether a multi layer V1 model can achieve both laminar and topographic organisation for SF, as suggested by experiments on primates. As we saw in section 2.2.3, primate V1 is composed of many different neuron types, with a myriad of pathways between the multiple laminae, as well as receiving afferent and feedback connections from many other brain regions. The intention of this model is to extract a simple subset of this complex system that is able to reproduce the SF and OR organisation observed or suggested experimentally. If successful, this model will point to a possible architecture for SF processing, tying together layer 4C and the more superficial layers.

The laminar V1 model was designed as a simple expansion of the single layer V1 model described in chapter 4. The cortical sheet representing all of V1 in the single layer V1 model is treated as a building block in the multilayer model, and used to represent different laminae in V1. This architecture allows for laminar SF input segregation and the development of topographic maps for SF and OR. Layer 4C, which receives the main LGN afferent connections in macaque, is represented as 4 interconnected cortical sheets. An additional "layer 2/3" cortical sheet is also included in the model to allow feed-forward convergence of the SF pathways from layer 4C. As topographic SF organisation has been observed in the superficial layers, a single layer that receives this combined SF input from the deeper layers is necessary to allow this SF

map to develop. This model allows further investigation into how separate processing streams converge to allow columnar OR and topographic SF representations to emerge.

This chapter presents the first computational model to generate a laminar organisation for SF preference, while organising for OR preference both topographically and columnarly. In addition the model is the first to demonstrate that a topographic organisation for SF preference can emerge in the more superficial layers from laminar SF organisation input, while maintaining a columnar topographic representation for OR preference throughout the cortex. As in the single layer model (chapter 4), we find that segregation of SF pathways gives more realistic RFs.

5.2 Laminar V1 model

The laminar model uses the same retinothalamic model (photoreceptors and RGC/LGN) as in the single layer V1 model (chapter 4). The difference occurs in the form of the V1 representation. V1 layer 4C is represented as 4 cortical sheets, each of which receives afferent input from only one of the SF pathways. By setting up the model in this way, each of these cortical sublaminae are effectively hardwired to the SF preference of their input SF channel. A further sheet representing V1 layer 2/3 receives combined afferent input from all of these V1 layer 4C sheets (see figure 5.1). All the cortical sheets have adaptive lateral long range inhibitory and short range excitatory connections. This within-lamina feedback (WF) is the same as described for the single layer model, in section 4.2. Additionally, reciprocal feedback connections (both excitatory and inhibitory) are implemented *between* adjacent layer 4C cortical sheets. This between-lamina feedback (BF) is of the same form as the WF, such that connection fields between laminae are the same size, shape, and sign as within laminae. As discussed in section 2.2.3, the macaque cortex 4Cα and 4Cβ layers are very close to one another, so it is reasonable that the lateral connectivity in a layer will "spill" into neighbouring layers. Moreover, as we see in the results, this pattern of activity will allow the two layers to develop similar feature maps through co-organisation.

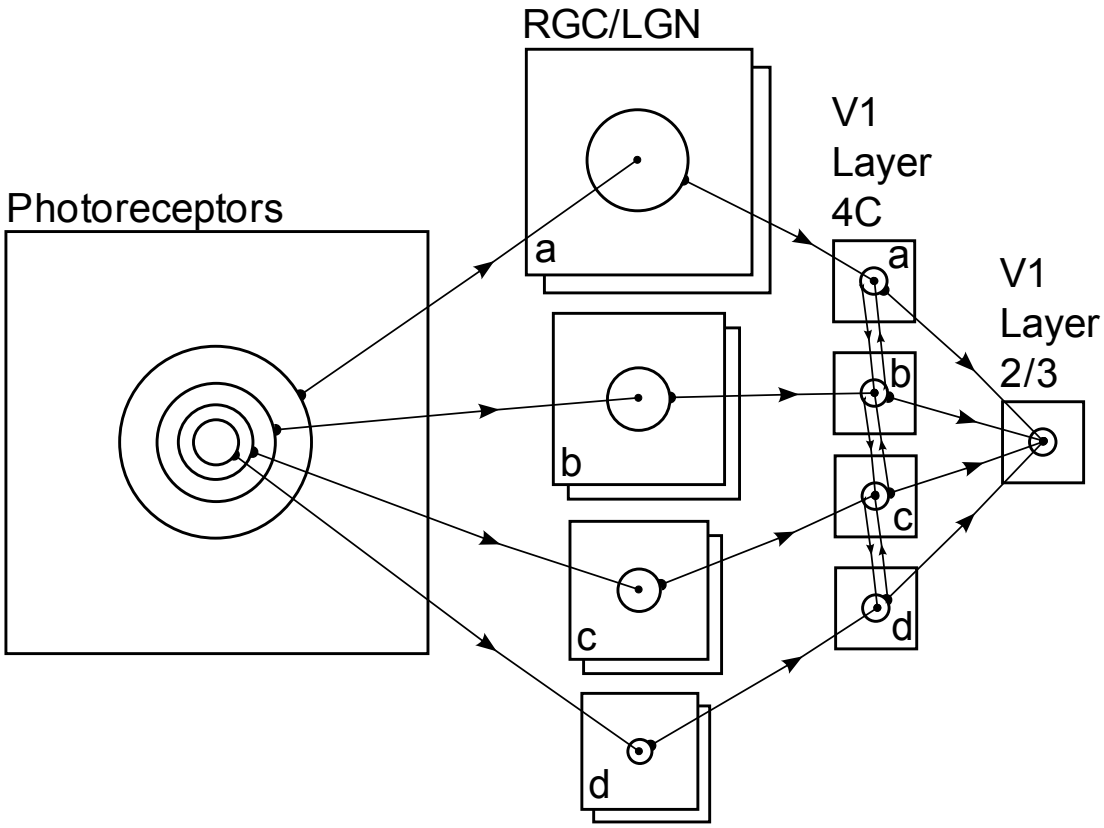

Figure 5.1: Laminar V1 model, schematic.

Greyscale natural images are presented on the photoreceptor sheet, which consists of 116 $\times$ 116 pixels. There are 4 distinct SF channels, each with an ON and an OFF sheet (OFF sheets are shown in background; connections to and from the OFF sheets are not shown), and each tuned to a different SF. Each LGN unit is activated via a centre/surround (DoG) receptive field on the input sheet. Receptive field sizes are uniform in each sheet, but differ between the different pathways, resulting in different SF preferences for each pathway. V1 layer 4C consists of 4 sheets of units, each of which receives input from its corresponding SF channel in the RGC/LGN. In addition to the lateral within-lamina feedback (WF) seen in the single V1 model shown as circles, the sheets have excitatory and inhibitory feedback connections between each other (illustrated as arrows connecting adjacent laminae). The strength of this between-lamina feedback (BF) can be adjusted so that there is a greater correlation in activity between the laminae. The final sheet in this model represents V1 layer 2/3, and receives input from each of the V1 layer 4C sheets. The layer 2/3 sheet also has within-lamina excitatory and inhibitory feedback.

The layer 2/3 sheet receives afferent connections from each of the layer 4C sheets.

The connectivity from layer 4C to layer 2/3 is such that units in layer 2/3 receive input from one unit from each of the four 4C sheets. This is the simplest connectivity between the sublaminae in layer 4 and layer 2/3 that will allow layer 2/3 neurons to become selective for any of the SF pathways. Layer 2/3 has WF of the same form as the layer 4C cortical sheets. Parameter values for the laminar V1 model can be found in the tables in appendix A.

5.2.1 Experimental procedure

The experimental and map measuring procedures are identical to those described in the previous chapter, section 4.2.1. An additional analysis procedure was required to compare similarities between feature maps. This similarity comparison, illustrated in figure 5.2, allows simple pixel by pixel comparison between any two matching size images. The resultant histogram and similarity index indicate the extent of similarity between two OR maps.

The similarity index between two pixels is determined by first cyclically subtracting their OR preference values. This OR difference is then scaled to a 0 to 1 range, such that 0 indicates the pixels preferences are 90° apart (the maximum possible), and 1 indicates they have identical preference. The similarity histogram and the similarity index (the mean similarity value) give a good indication of how well the maps match. A similarity index value of 0 corresponds to two maps that are 90° out of phase with each other, 1.0 would signify two identical maps, and 0.5 would be expected for two maps with no correlation, or that have a 45° phase relationship. (The later two cases are readily distinguished by considering the histogram, which would show an overall uniform distribution in the case of no correlation, and a central spike in the unlikely event of a 45° phase relationship).

5.2.2 Feedback within layer 4C

The strength of connectivity between neighbouring units in a layer 4C cortical sheet is important for determining the behaviour of the laminar V1 model. In this model, this connectivity is determined by the balance of strength between WF and BF. In the results section that follows, results are presented from simulations using different values

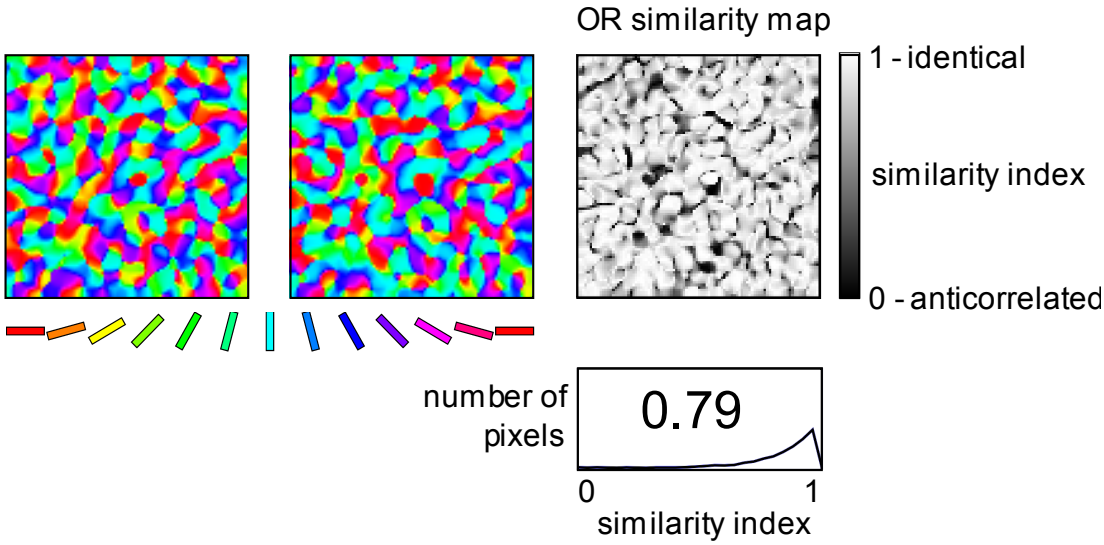

Figure 5.2: Method for comparing similar OR maps.

Two orientation maps shown next to their corresponding similarity map, produced by cyclically subtracting one map from the other. The similarity map is greyscale, with white pixels (value 1) representing corresponding pixels with identical orientation preference in the two maps, while black pixels (value 0) have orthogonal orientation preference. All the pixel data is collated in a histogram shown below the similarity map, where unrelated maps would have a uniform (flat) distribution. Here we see two highly correlated maps, which yield a histogram heavily weighted towards the "correlated" end of the curve. The mean numeric value of the histogram is shown in the middle of the histogram box, and is termed the "similarity index".

of WF and BF. These parameters are systematically adjusted to illustrate how different levels of connectivity between adjacent sub-laminae may influence cortical map development. To ensure that the total amount of excitation or inhibition a unit receives as a consequence of lateral connectivity (whether WF or BF) is roughly constant, the increases in BF were always offset by decreases in WF such that they always sum to 1.0.

5.2.3 Pruning methodology

The pruning mechanism is essentially the same to that described in section 4.2.2. However, in this laminar V1 model the SF pathways converge between layer 4C and layer 2/3, and so the pruning affects connections to layer 2/3.

5.3 Results

The results are presented in three sections. First we consider the organisation just of the layer 4C sublaminae, then we review the unpruned layer 2/3 organisation, and finally the pruned layer 2/3 organisation.

5.3.1 Layer 4C sublaminae

Under all lateral feedback conditions the sub laminae in layer 4C successfully organise for OR preference (figure 5.3), producing topographic OR maps, with the characteristic features found in animal OR maps (orientation pinwheels, iso-orientation zones, linear zones, and fractures). With zero BF, the model successfully organises for OR in each sheet, but the individual OR maps are not significantly correlated with each other. As the BF increases, the similarity between maps in adjacent sub-laminae increases, as does the similarity between the maps at the different ends of the SF spectrum. This increasing correlation can be observed by careful inspection of adjacent maps in a simulation run, and by the similarity histograms and their mean values. This is an important result as it demonstrates that topographic columnar organisation of OR preference can arise as a result of feedback between sub-laminae with different subcortical inputs.

Organisation for SF preference within each sub-lamina is topographic, with iso-SF zones, continuous zones and SF-fractures. However, the range of SF preference in each lamina is narrow. Across the four separate 4C laminae, the entire range of SF preference found in the RGC/LGN sheets is represented, each of the laminae representing primarily its own SF preference as determined from its geniculate input. As the value of the BF is increased, some overlap in SF preference representation occurs between the laminae (visible as a broadening of the peaks in the histogram in figure 5.4). Although the sheets essentially retain their SF preferences, the total range of SF represented does narrow slightly as fewer units have more extreme SF preference in the 0.50 BF case than the 0.0 BF case.

To illustrate the division in SF and OR processing in this model, figure 5.5 shows the response of the sub-lamina sheets in the 0.2 BF model to bar and edge stimuli. Line stimuli of different widths are presented to the laminar V1 model, to gauge response at different cortical depths. The line stimuli predominantly activate cortical units with the appropriate OR tuning, within the appropriate individual SF channel. Activity overlaps between the sub-laminae, because of the broad tuning of each channel, even though the peak SF preference of each sheet is quite different (as demonstrated in figure 5.4). The edge stimulus strongly activates all sheets, as a sharp edge contains all spatial frequencies.

Figure 5.6 further illustrates the simultaneous columnar organisation for OR preference and laminar organisation for SF preference that arises in layer 4C of the laminar V1 model. The RFs are shown for neurons at two locations in iso-orientation zones with matching OR preference, but different sizes and different SF tuning, across the four sub-laminae. The results in figure 5.6 are typical for this model, as can be seen by the high correlation in OR preference between laminae (figure 5.3), yet the wide distribution in SF preference between laminae (figure 5.4).

These results clearly illustrate that with segregated input the sublaminae are capable of organising topographically for OR preference and laminarly for SF preference. In addition, even minimal feedback between laminae is sufficient to allow a columnar organisation for OR preference to arise.

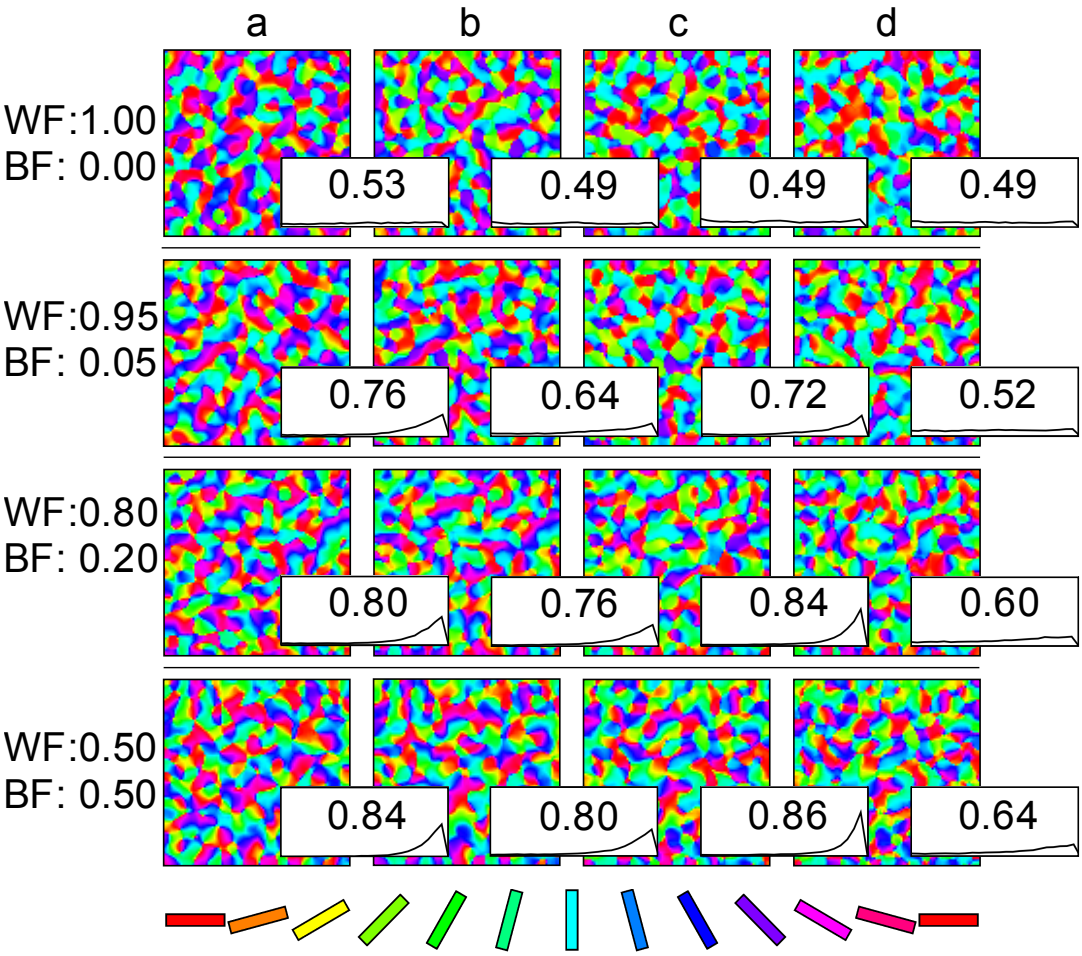

Figure 5.3: Effect of between-layer feedback on development of laminar OR maps. Each row shows OR maps for each of the V1 layer 4C sublaminae (a to d, as in figure 5.1), and the corresponding map similarity comparisons, for different relative amounts of within-lamina and between lamina feedback (as shown on left). The inset boxes are similarity histograms for corresponding pixels in neighbouring maps (from 0 (90 deg out of phase), to 1 (identical); left to right). Numbers in boxes show the mean of the distribution: 0.00 (anticorrelated maps) to 1.00 (identical maps), with 0.50 indicating uncorrelated maps. The far right-hand-side box shows the relationship between outermost sheets a and d. With zero BF (top row) the cortical sheets develop realistic but entirely uncorrelated OR maps, even though they share the same visual input. As the amount of BF is increased, the similarity between OR maps in different sheets increases. With large BF (0.5), the maps are visibly similar and are highly correlated, thus demonstrating a columnar, topographic mapping for OR.

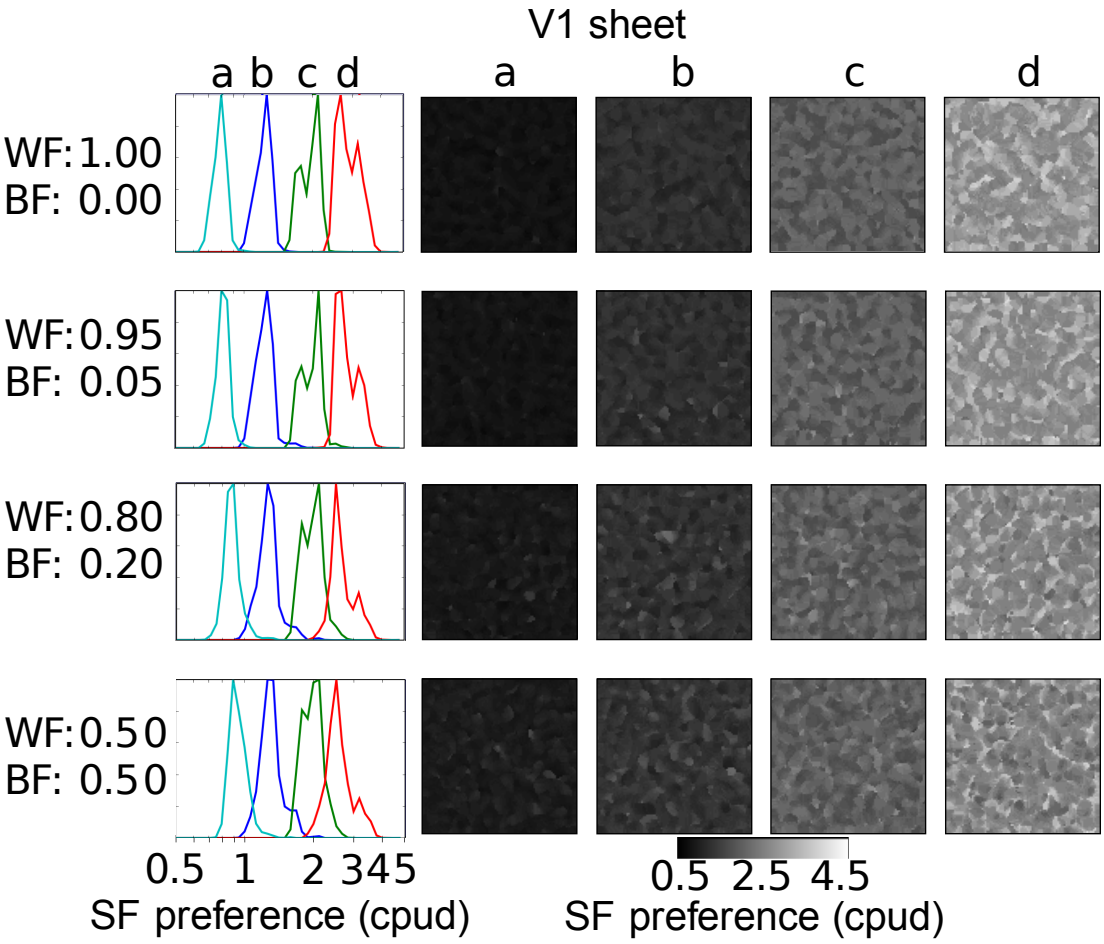

Figure 5.4: SF preference in the laminar V1 model.

Each row shows SF preference maps for different amounts of BF (as indicated on left), with corresponding overlaid normalised SF preference histograms. As feedback between sublaminae is increased, there is a small decrease in the range of SF preferred by each sublamina. Each layer has some topographic variation in SF preference, and thus an SF map, but the variation in preference between layers is much greater.

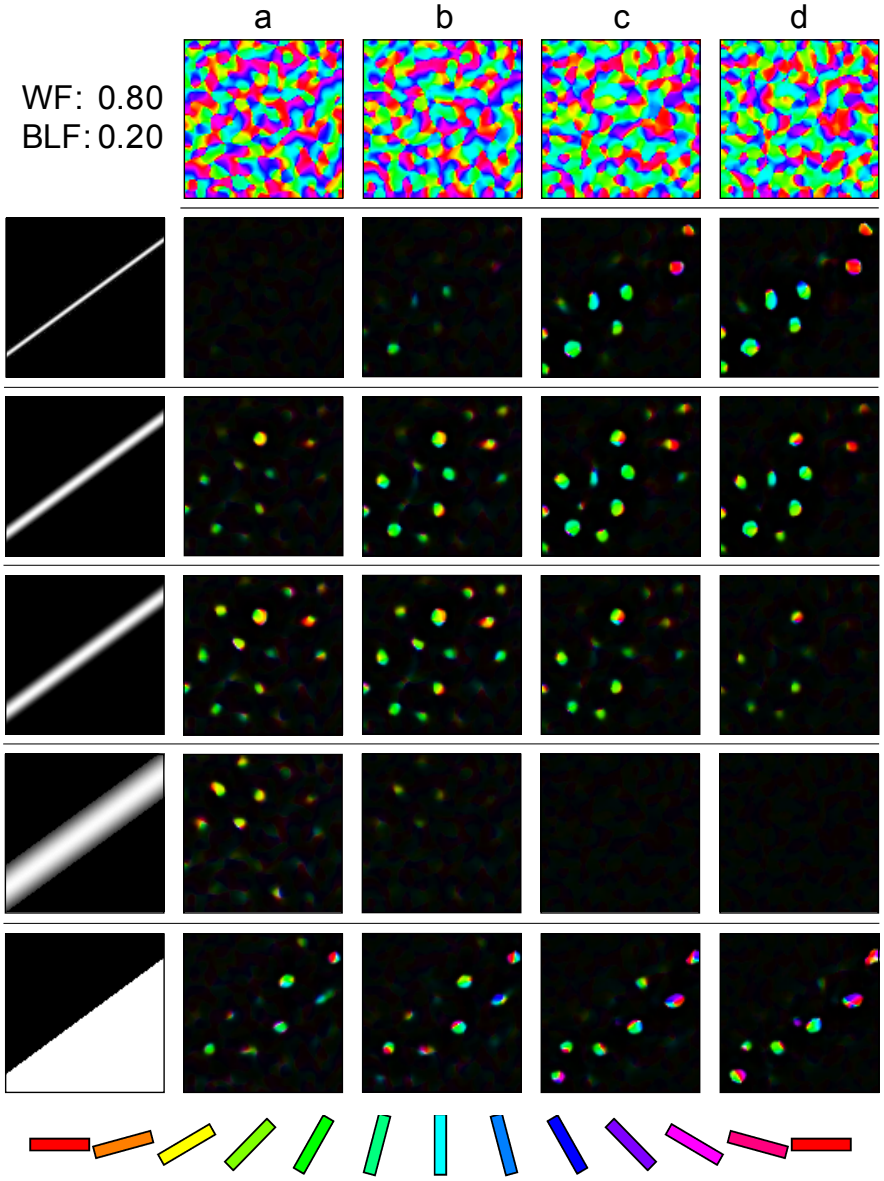

Figure 5.5: Laminar V1 response to bar stimuli.

Top row shows the OR maps for the layer 4C sheets at 0.20 BF. Subsequent rows show the corresponding cortical response (coloured according to the OR map, and more brightly coloured for more active neurons) to the presentation of bar stimuli shown at left. In each case, most of the units that respond have the corresponding OR preference (green). As the width of the line bar stimulus is increased, sheets with units that prefer lower SFs are activated. In addition, predominantly units with the appropriate OR tuning, matching the stimuli, are excited. An edge stimulus (bottom row) excites all laminae equally, because it is a good match to RFs of all sizes. These results demonstrates that cortical units tuned for both OR and SF co-exist in this lamina-based model, as consistent with experimental results.

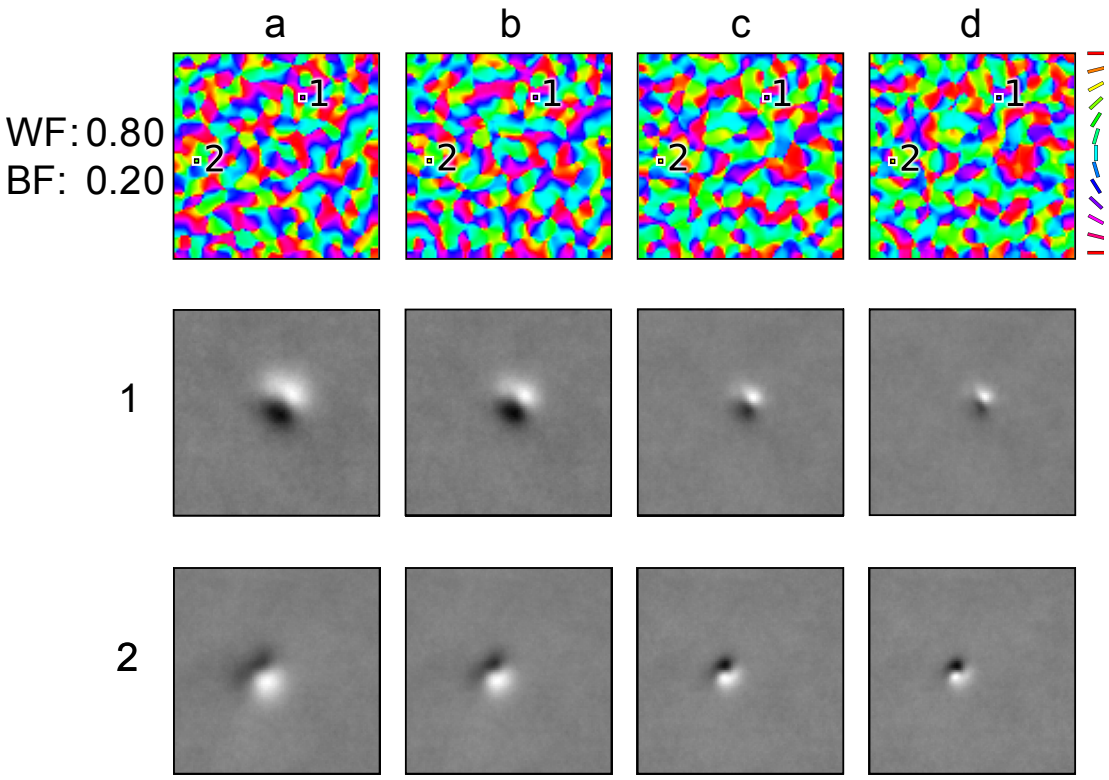

Figure 5.6: Matching units across laminae in the laminar V1 model.

RFs were measured using reverse correlation for the units indicated at positions 1 and 2. Below each OR map, the retinal RFs from the specified units are shown. These plots illustrate that the units in each sheet have similar OR preference, but different SF preference, as expected from animal experiments.

5.3.2 Layer 2/3 unpruned

Layer 2/3 organises for OR preference topographically with a high degree of similarity to the layer 4C sub-laminae (figure 5.7). The layer 2/3 map contains the typical OR map features, and in the 0.2 BF case has a mean similarity index of 0.79 to layer 4C (indicating a high degree of correlation).

The SF preference organisation of layer 2/3 is continuous and topographic (figure 5.8), with iso-SF zones, continuous zones and fractures. The range of SF preference is broad, spanning the range found in the RGC/LGN sheets. These results show that the SF segregation in layer 4 is sufficient to ensure broad coverage in a single layer 2/3, without pruning.

Comparison of the OR and SF preference gradient plots (figure 5.9) suggests that regions of high gradient in both maps tend not to coincide. Orientation pinwheels predominantly fall on the boundaries of iso-SF zones, and there are some regions where iso-SF zones and iso-OR zones spatially co-exist to some extent. Possible relationships are considered further in section 5.3.4.

The RFs of layer 2/3 neurons in the unpruned model do have distinct ON and OFF sub-regions, and there is some variation in RF size, and in SF preference tuning (figure 5.10). The RFs are typically blobs, or two or three region Gabor type patches, although the RFs are not always well formed. As described in chapter 4, SF channel superposition artifacts are observable in some RFs plotted using single pixel or white noise stimuli, which is due to cortical units not developing exclusive connectivity to a single SF channel (see section B.1).

Note that, to examine the RFs of units in layer 2/3, due to processing time and memory constraints on the machines available, it was necessary to re-run the simulation with a reduced cortical density. A cortical density of 48 was used in layer 4C and layer 2/3 to generate figures 5.10 and 5.11; all other parameters remain unchanged. While the RFs typically do have one or two sub-regions, they are not always well formed, which is a consequence of the neurons not developing exclusive connectivity to a single SF sub-lamina (see section 4.4). When mapped using white noise or single

pixel stimuli SF channel superposition artifacts are observable (see section B.1), but RFs mapped using pink noise have a wide spread of SF preference and match the OR and SF maps well (figure 5.10).

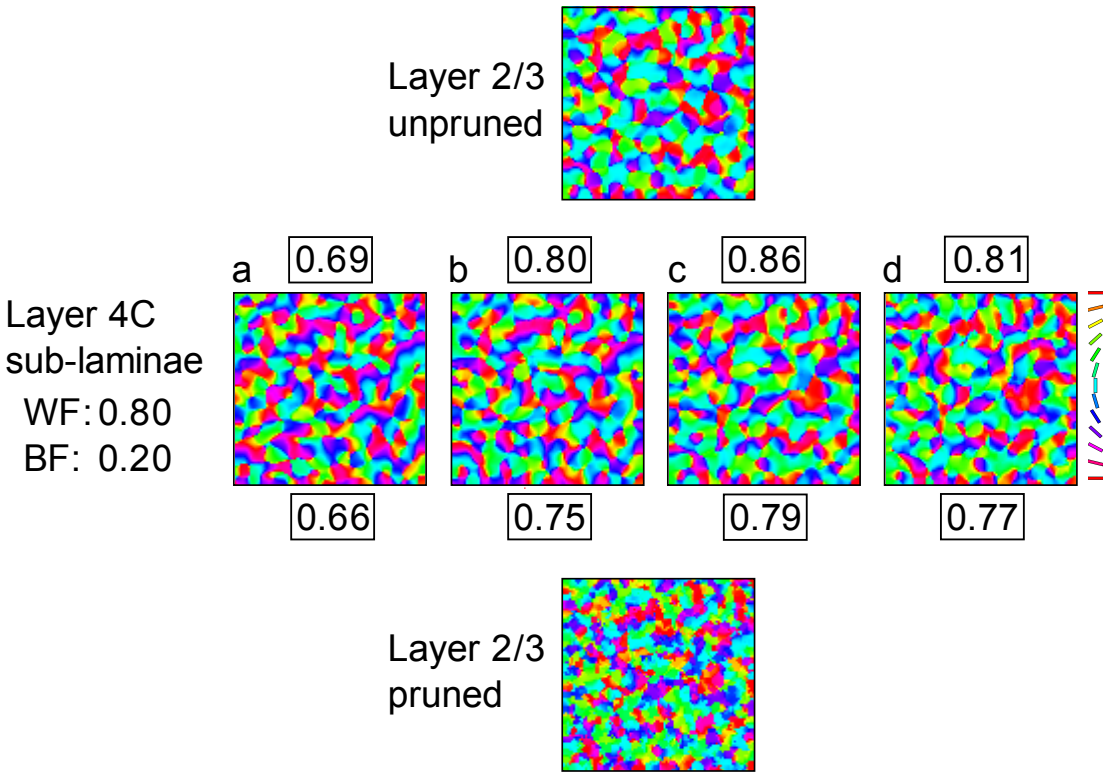

Figure 5.7: Columnar mapping of OR preference in V1.

Pruned and unpruned OR maps for V1 layer 2/3, with 0.2 BF layer 4C sub-laminae. Numbers in boxes above sub-laminae OR maps are the similarity index values between that sub-lamina and the unpruned layer 2/3 OR maps. Boxed numbers below the sub-laminae OR maps are the similarity index values between that sub-lamina and the pruned layer 2/3 OR maps. There is a high degree of visible similarity between all maps, also reflected in the similarity index values. Thus a strong columnar organisation for OR preference arises in this model architecture.

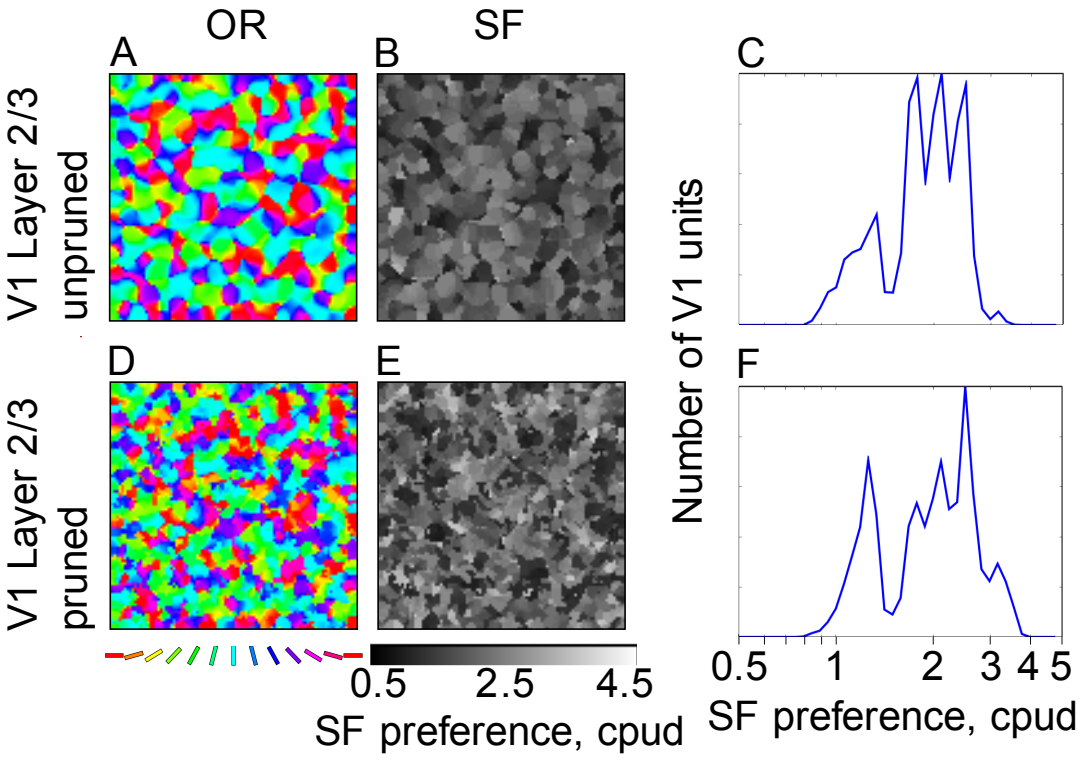

Figure 5.8: Laminar V1 model layer 2/3.

(A and D) Layer 2/3 OR preference maps with and without pruning. (B and E) SF preference maps with and without pruning. (C and F) Corresponding normalised SF preference histograms for layer 2/3 with and without pruning. The range of SF preference in the unpruned model spans the input space. Pruning very slightly increases the number of units representing the higher and lower SF preference extremes, but also results in a more disordered SF preference map. The pruning process reduces the smoothness of the OR preference map, which becomes a little rougher, but it is much less dramatic a shift than in the single V1 model pruning case.

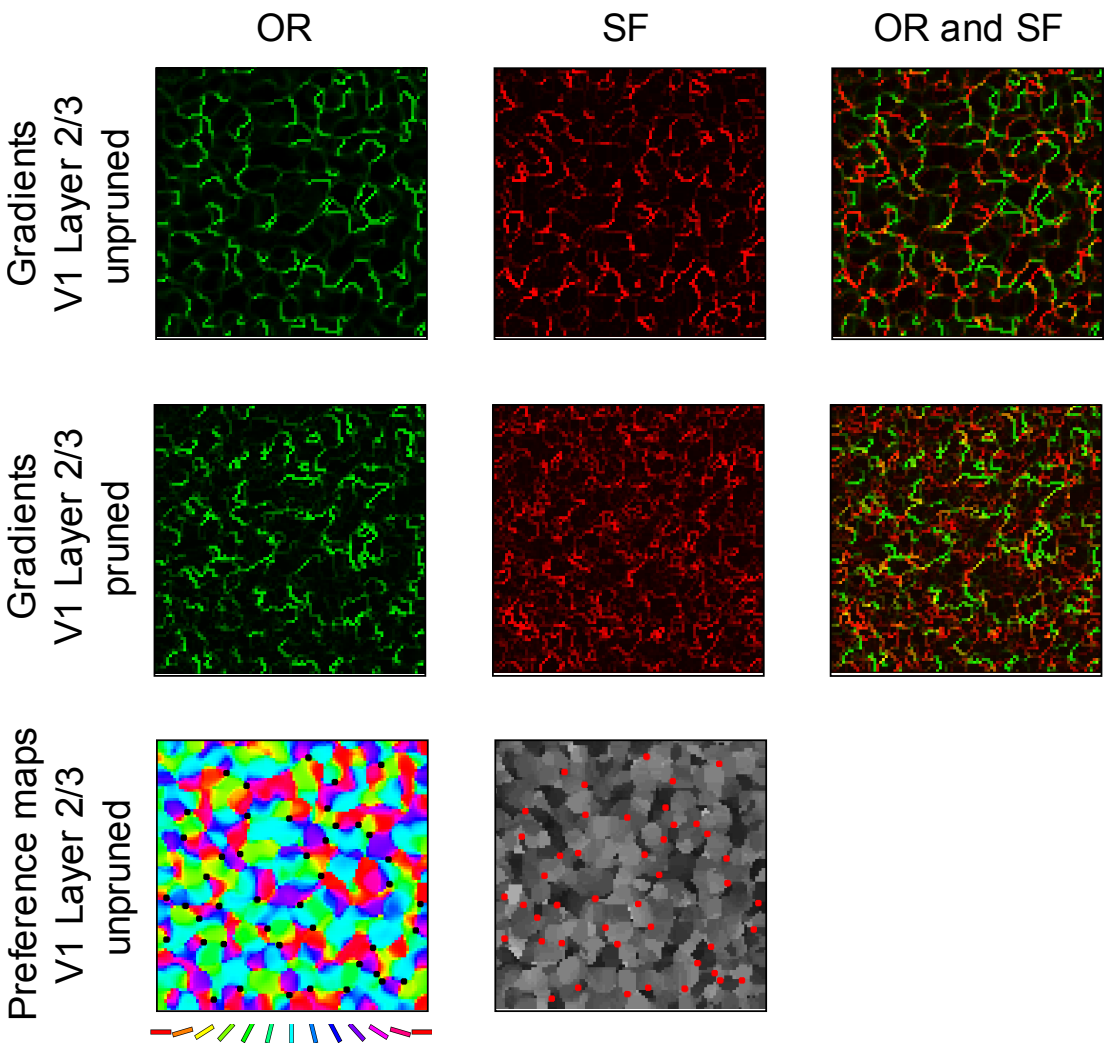

Figure 5.9: Comparison of layer 2/3 OR and SF gradients and pinwheels.

In rows 1 and 2, the OR and SF gradient maps are colour coded to aid visual comparison in the overlaid plot. (1st row) Gradients for laminar V1 model layer 2/3 without pruning. There are many fractures in the OR map, illustrated by the prevalence of high gradient regions. In the gradient overlay plot, OR and SF fractures are typically non-coincident. (2nd row) Gradient comparisons for laminar V1 model layer 2/3 with pruning. The higher level of disorder in these maps is illustrated by the increased number of fractures in these maps. High-gradient regions in the maps correspond to fractures and pinwheels. (3rd row) OR map for non-pruned laminar V1 model layer 2/3, with pinwheels marked by black dots. Pinwheels are marked with red dots on the corresponding SF map, and do not appear to have any simple relationship with it.

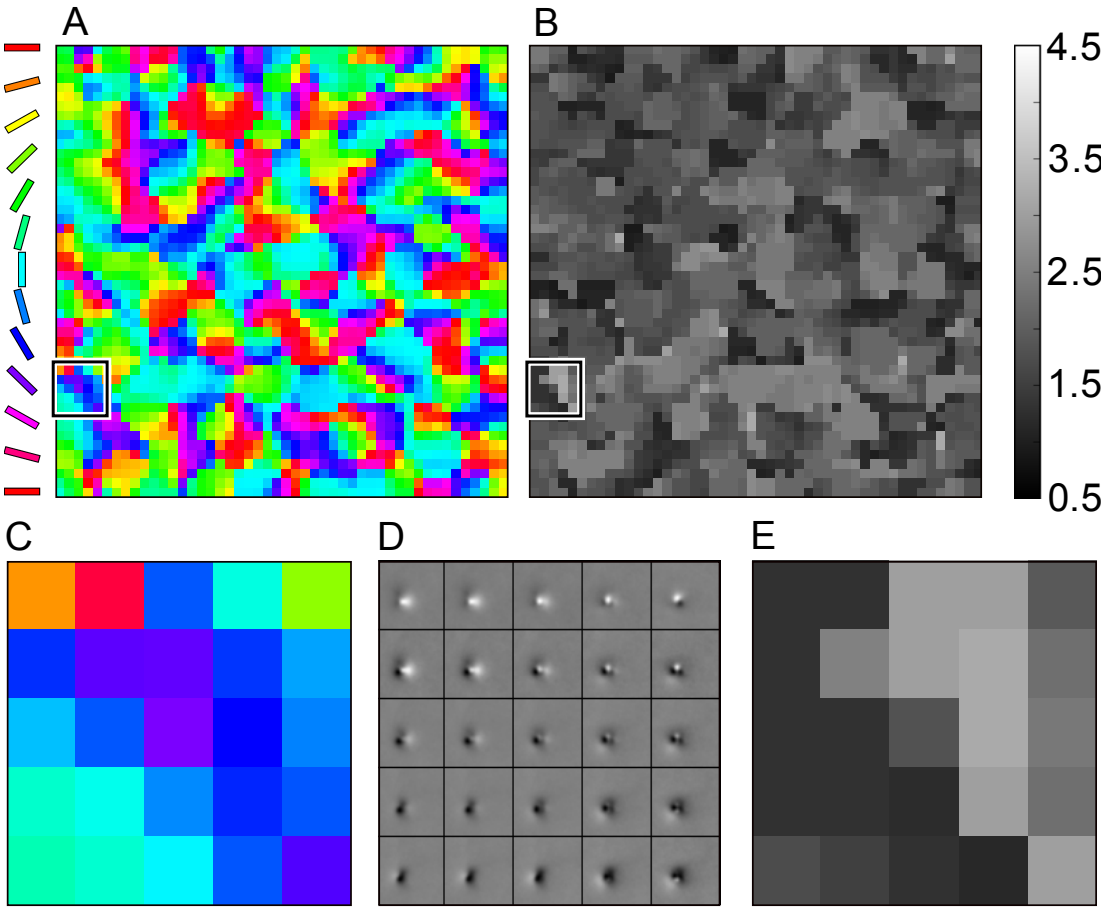

Figure 5.10: RFs from layer 2/3 in unpruned laminar V1 model.

(A and B) OR and SF maps for unpruned layer 2/3, with cortical density 48×48. (C) Detail of OR map from boxed region in A. (D) RFs for layer 2/3 cortical units in boxed region, measured using reverse correlation using pink noise. (E) Corresponding detail of SF map from boxed region in B. Note the variation in size of RFs, and the match to their corresponding values in the OR and SF maps. The RF plots have been cropped to aid viewing, and are each shown on a 75×75 pixel region of retina. Scale shows greyscale SF preference in cpud.

5.3.3 Layer 2/3 pruned

For the pruned case, all but the strongest afferent and lateral projections to each neuron in layer 2/3 were eliminated at iteration 4000 as described in section 4.2.2 in chapter 4. After pruning, layer 2/3 units received afferent input from a single unit in one of the layer 4C sheets, while their lateral connections came from only other units connected to the same layer 4C sheet. Image stimuli continued to be presented until time step 5000. As in the non-pruned case, the OR map has a high similarity to the layer 4C maps (figure 5.7), but with some disorder on the edges of iso-orientation zones, and some misplaced units. However, the damage to the organisation of the OR map is slight compared with that in the single layer model (figure 4.2). The mean similarity index to layer 4C is 0.74, indicating a strong correspondence between layer 2/3 and layer 4C for OR organisation, which is clearly visible to the eye.

The effect of the pruning process upon layer 2/3 organisation for SF preference is to increase disorder in the map: fractures are more frequent, iso-SF zones are smaller and more poorly formed, and there is less structure overall. There is also a greater representation for high and low SF preference, with more units favouring these extrema SF preferences, visible from the SF preference histogram (see figure 5.8). The gradient plots for OR and SF preference (second row in figure 5.9) indicate no obvious relationship between high gradient regions, although regions of high gradient in both maps tend not to coincide. Pruning the connections to layer 2/3 units leads to an improved range of RFs (figure 5.11). The range of sizes and SF preference is greatly increased and the overall quality of RFs is also improved.

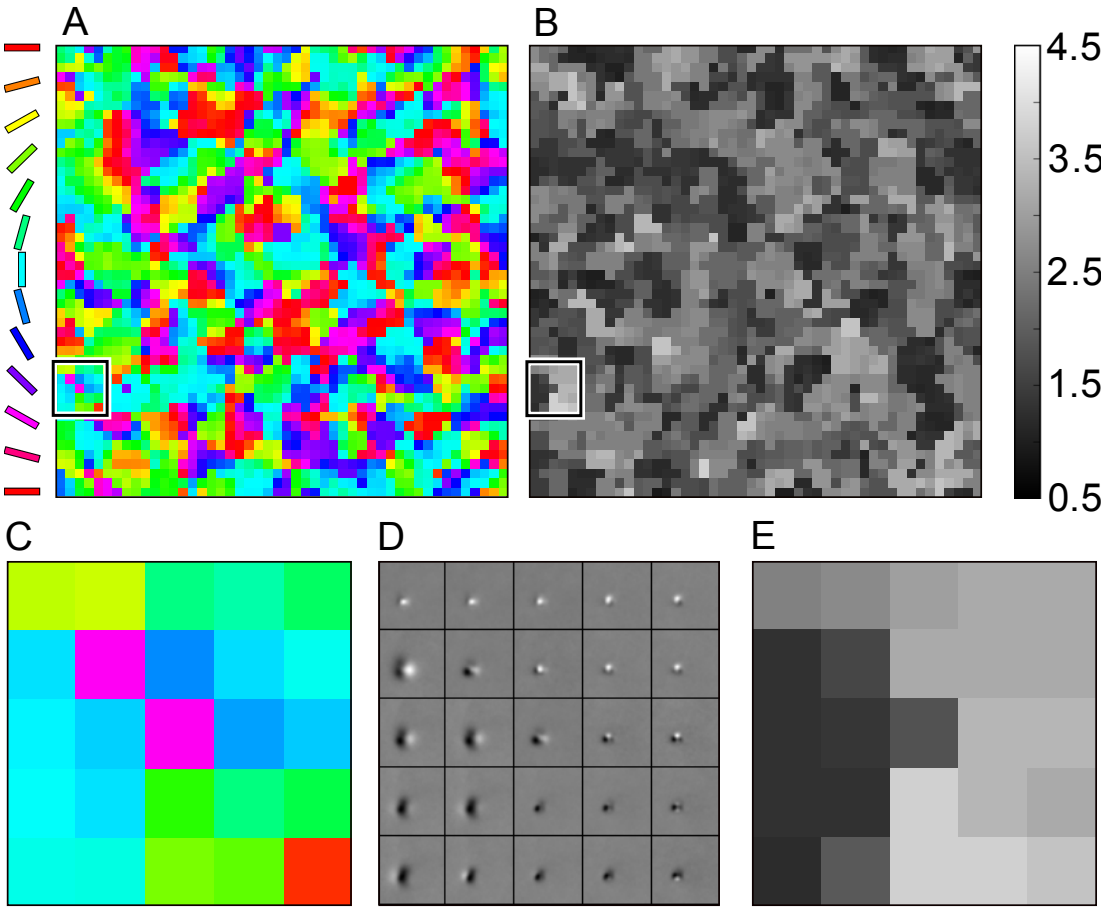

Figure 5.11: RFs from layer 2/3 in pruned laminar V1 model.

(A and B) OR and SF maps for layer 2/3. (C) Detail of OR map from boxed region in A. (D) RFs for cortical units in boxed region, measured using reverse correlation. (E) Corresponding detail of SF map from boxed region in B. Note the large variation in receptive field size, and the good matches of RFs to their corresponding values in the the OR and SF maps (as we expect). The RF plots have been cropped to aid viewing, and are each shown on a 75×75 pixel region of retina. Scale shows greyscale SF preference in cpud.

5.3.4 Comparisons to experimental data

Figure 4.12 in chapter 4 compares the laminar V1 model's layer 2/3 SF and OR maps with details from a representative selection of previously published SF maps and (where applicable) corresponding OR maps; these maps were discussed in more detail in section 2.3. The laminar V1 layer 2/3 OR and SF maps qualitatively fall within the range of diversity measured experimentally. The overall size of the iso-SF zones is within the range of experimental maps, although at the smaller end of the scale than for the single layer V1 model. The scale of the laminar V1 maps is more directly comparable with the primate maps.

In the unpruned layer 2/3 results, there are regions where iso-SF zones and iso-OR zones spatially co-exist to some extent; some iso-SF blobs that have the same position and shape as iso-OR blobs. OR pinwheels also typically occur at iso-SF boundaries or at SF pinwheels (points at which iso-SF zones from the entire range of preferences meet). Neither of these relationships have been reported experimentally, and these findings appear to conflict with experimental observations that regions of change of one feature tend to fall in an iso-value region for another feature preference, and that OR pinwheels are more likely to occur within iso-SF domains (see figure 5.12). These initial observations are not investigated further as the OR pinwheel locations are not identifiable in pruned maps as the maps are too fracturous, and in the unpruned models the SF maps do not have a range of SF preference that matches the input space. These relationships still need to be fully characterised experimentally, particularly in light of Purushothaman et al.'s recent study that suggests that previously measured SF maps may misrepresent the size of higher frequency SF domains, and that iso-SF domain sizes may be more uniform for different frequencies.

5.4 Discussion

The laminar V1 model allows us to consider how separate SF processing streams can be consistent with shared feature preferences across these different laminae, such as columnar OR representation. We see that OR maps in the model layer 4C only start to match when there is connectivity between the cortical sheets (i.e., BF is non zero).

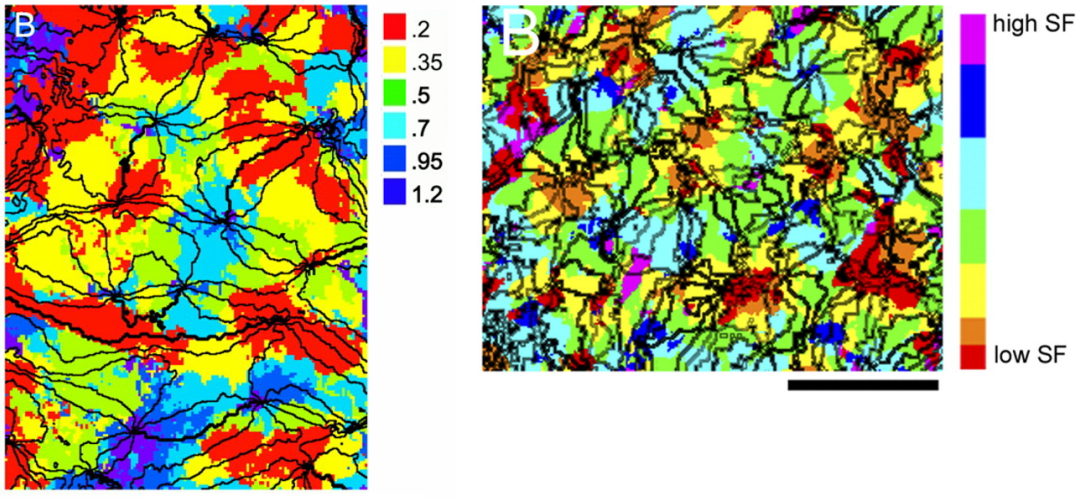

Figure 5.12: Cat and bush baby spatial frequency maps with iso-OR contours overlaid. Scale bar for both maps is 1mm. (Left) In cat, OR pinwheel centres tend to lie in regions of high or low SF preference; reproduced from Issa et al. (2000). (Right) In bush baby, the relationship between pinwheels and SF domains is less obvious; reproduced from Xu et al. (2007). The layer 2/3 laminar model results, which suggest that OR pinwheels typically occur at iso-SF boundaries or at SF pinwheels, may be in conflict with these experimental findings.

When there is no BF, the OR maps are well organised but there is no similarity between them, suggesting that matching OR maps do not arise as a consequence of similar activity patterns generated from the same image stimuli, but must be due to BF. Developing matching OR maps across the different laminae is an important validation of the model, as the columnar nature of OR preference maps is well established in many species (see section 2.2.6). This is a surprising result, as we have found in chapter 4 that the input from the different LGN SF channels is too similar for cortical units to distinguish between the channels, thus necessitating pruning. Despite this correlation, here we find that the input from different LGN channels is sufficiently different to generate different OR maps, because the overall shape is determined by the arbitrary initial patterns to activate each channel (Miikkulainen et al., 2005), rather than the long-term correlations. In any case, the simultaneous topographic and columnar organisation for OR and the laminar organisation for SF in layer 4C with BF fits with a number of experimental observations (section 2.4).

This model highlights an important idea, i.e., that feature preference organisation

and coverage in the cortex may not be restricted to two dimensions. Layer 4C in the laminar V1 model spans the entire feature space using three spatial dimensions. While OR preference is topographic, SF preference is organised by lamina, as is dictated in this model by the architecture. Each lamina's SF preference is closely tied to the SF channel it receives input from. Even with high levels of BF, the average SF preference of each lamina remains close to the preference of the afferent SF channel. In primate, the SF preference of neurons found in the 4C sub-laminae may match the SF preference of the LGN cells providing afferent input to that sub-laminae (see 2.2.7). OR preference is also thought to be topographic and columnar at this depth, although there are fewer OR-selective cells (see section 2.2.6). It may be that prior to both of these feature preferences (OR and SF) being represented in two dimensions, in a single layer 2/3 lamina, they are represented in this three dimensional fashion in layer 4C. The reduction to a two dimensional representation may be an emergent consequence of this architecture.

In order for the laminar V1 model to satisfy all three of our modelling criteria (section 3.1), the model requires pruning to generate a complete separation between the SF pathways in layer 2/3, allowing more realistic RFs to form. Topographic SF organisation with a representative range of SF preference has developed in layer 2/3, but without pruning, when measured using single pixel or white noise stimuli the RFs have the SF channel superposition artifacts outlined in section B.1.

As the layer 4C sheets all have highly correlated OR preference maps, the OR preference map for layer 2/3 is similar to these maps. In contrast to the single-V1 model (chapter 4), here the pruning process has a minimal impact upon the general organisation of the layer 2/3 OR map. The overall smoothness of the map is reduced, but there remains a clear organisation that is similar to the unpruned simulation. While in the unpruned model the layer 2/3 RFs do span a range of sizes and SF preferences, they also have the SF channel superposition artifacts observable under single pixel or white noise reverse correlation (as described in appendix B). This effect is due to the same channel superposition issue encountered with Hebbian learning of non-cyclic dimensions in chapter 4 (see section 4.4). Thus the discussion in section 4.4 of the implications of SF channel separation is also relevant for this model.

This model only contains simple cells in V1, despite observations suggesting that layer 2/3 is predominantly composed of complex cells. This limitation is for two reasons: The image set is not designed to drive complex cells, in that there is no phase relationship between successive images, so no neurons develop preferences for different phases but similar stimuli. Additionally, previous LISSOM simulations have found neurons to be organised topographically by phase similarity (Miikkulainen et al., 2005), because units with similar phase preference are activated simultaneously. Recent work by Antolik and Bednar (2008) has resulted in a developmental LISSOM-based model for complex cell formation. By merging these models, the laminar V1 model could be extended to support the formation of complex cells, without in any way compromising the results of the current model.

Further work is required to characterise the nature of the relationship between the OR and SF maps in this model. Preliminary simulations suggest that this overlap in topographic characteristics of the SF and OR maps is due to learning of a few strong stimuli early in development, leading to a correlation between the SF and OR maps at that location. This is likely to be addressable in future work through parameter adjustment, and does not impinge upon the relevance of other results from this model, such as the need for segregation of SF channels for realistic RF formation.

5.5 Conclusion

The laminar V1 model is the first developmental multi-laminate model of V1 to organise topographically and columnarly for OR preference, demonstrating the likely importance of feedback between sub-laminae to achieve this columnar structure. The model also shows that this form of OR organisation can co-exist with a laminar representation for SF preference. Furthermore, laminar SF preference representation can give rise through a feed-forward system to a topographic organisation for SF preference in higher layers. The results again suggest that a separation between SF pathways is an essential requirement for generating realistic RFs.

This developmental model is the first to be consistent with our interpretation of known experimental data relating to OR and SF mapping. Further experiments examining the laminar organisation of SF and OR in layer 4C, and the extent of segregation between SF pathways within layer 4C, are required to confirm or refute the coexistence of topographic OR and laminar SF organisation. Such experiments are difficult due to the depth of layer 4C from the cortical surface. Current high resolution imaging techniques (such as optical imaging of intrinsic signals and 2-photon microscopy) are limited in depth to about 300 - 500μm below the cortical surface, and layer 4C starts at about 1mm. However, if the upper layers of cortex could be removed, without overly disturbing layer 4, the predicted organisation of layer 4 could be tested. Techniques such as using lower frequency light (which can penetrate further into cortex), or confocal microscopes, may allow measurements at deeper layers using optical imaging (for review see Grinvald et al., 1999).

Chapter 6

Discussion and future work

6.1 Introduction

The results described in this thesis raise a number of interesting questions and point the way to further research, both experimental and computational, that may help in determining the nature of SF organisation in the visual system. The results demonstrate that continuous topographic maps for SF preference can organise in a Hebbian developmental model of the visual system. However, they suggest that a separation of SF pathways is required to allow a full range of SF preference to develop within a single V1 lamina. The results also show how a laminar organisation for SF preference can coexist with a columnar topographic organisation for OR preference, in sublaminae of layer 4. Furthermore, they show that a topographic organisation for SF can arise in the superficial layers as a result of input from these layer 4 sublaminae.

This chapter proposes experiments that may help to test these predictions and their implications upon our understanding of cortical SF organisation and development. Future extensions to the models are then described. These extensions would allow new computational models to address currently unanswered questions relating to SF processing, as well as other aspects of processing. Unfortunately, many of the suggested experiments are not yet practical, but the extensions to the computational models are quite feasible and some of them may allow us to further articulate and refine the experimental questions so that the predictions become testable. It is in this regard that computational simulations are so powerful.

6.2 Suggested experiments

In this section a number of animal experiments are suggested that may help determine and characterise the nature of SF organisation in V1.

6.2.1 Development of thalamocortical connectivity

It is a common feature of the models presented in this thesis that some form of separation of connections from units of differing SF preferences is required in order to obtain a wide range of SF preference from Hebbian learning of natural images in V1. This separation can occur by projection to different sub-laminae of layer 4C, as is widely reported to occur in macaque (Tootell et al., 1988), leading to a laminar organisation for SF, or by some form of pruning, whereby cortical units afferent connections are eliminated, leaving principally connections to a narrow range of cell sizes in LGN.

If SF inputs to V1 neurons are separated, to the extent that thalamorecipient neurons primarily receive input from LGN neurons with the same SF preference (Alonso et al., 2001), how does this develop? There are two main possibilities. One is that segregation is hard coded, such that thalamic afferents with shared SF preference collectively target the same V1 neurons from the start. The other possibility is that V1 neurons initially form connections to thalamic afferents of many different SF preferences, but over time develop stronger connections to neurons that share feature preferences such as SF preference.

By determining the RFs of LGN cells and their connectivity to V1 cells at different points in development it may be possible to determine which of these possibilities is the case. This could be done by repeating Alonso et al. (2001)'s study at different points in development. An alternative experiment that may help to determine the development of thalamocortical connectivity would be to separately stain LGN cells within certain bandwidths of SF preference. However, doing so would require the development of

more specific labelling techniques.

6.2.2 Temporal latency

If thalamic recipient neurons initially receive input from LGN neurons with a wide range of SF preferences, but then over time develop stronger connections to those neurons with similar SF preference how do these connections become more specific? The observed overlap in activity between the model SF preference channels for natual images suggests that Hebbian learning may not be sufficient to allow SF specialisation to occur (see discussion in chapter 4).

Numerous studies have established links between temporal latency and cell type (Croner and Kaplan, 1995; So and Shapley, 1979); e.g., LGN magnocells generally respond faster than parvocells. More recently, Weng et al. (2005) have shown a correlation between RF size and response latency, with larger RFs having faster response times. This difference means that neurons with similar size RFs projecting to an individual cortical neuron may be more highly correlated than their spatial similarity alone would suggest. This additional within-channel correlation may provide a means for a neuron to distinguish between SF pathways, allowing Hebbian mechanisms to strengthen a single SF pathway. If afferent thalamic connections to a cortical neuron could be determined and mapped sufficiently early in development, then this hypothesis could be tested using chronic recordings of the afferent and recipient neurons and observing how the connection strengths and RFs change over the course of development.

More directly, by isolating the animal from visual input, and electrically stimulating the LGN cells projecting to a single V1 neuron, it may be possible to manipulate the development of the V1 neuron's RF. In this way, correlating or de-correlating firing latency with LGN RF size may generate reproducible and different V1 RFs. This could potentially demonstrate the importance of the link between RF size and temporal latency, and how it may be important in development.

6.2.3 Is SF topographically or laminarly organised in layer 4?

The modelling results show for the first time that a laminar organisation for SF preference is viable, and can successfully be integrated into a topographic columnar model for SF in the upper layers. Due to the current depth limitations of two-photon microscopy and optical imaging of intrinsic signals, it is difficult to know how the SF preference of sub-laminae in layer 4 could be determined. Attempts at 3D sectioning (described in Grinvald et al. (1999)) have not succeeded in getting below layer 3. Perhaps by thinning or silencing the upper layers of V1, it may be possible to observe these deeper layers and determine whether laminar organisation for SF preference and columnar topographic OR organisation coexist, as suggested by the current experimental data and validated in the model.

Apart from non-invasive imaging, it may be possible to change the termination location of the thalamocortical afferents such that they terminate in the superficial layers. This may be possible through ablation of layers 4 to 6, using methods like those in Sur et al. (1988). Alternatively, early removal of the superficial layers (1-3) may be a means to provide access to investigate the development of this part of the visual circuitry. Of course, the results would be difficult to interpret, becuase it would not be clear whether the underlying developmental process would remain similar.

6.2.4 Is feedback important for columnar organisation of OR preference?

The results of chapter 5 demonstrate the importance of inter-laminar feedback for the generation of correlated OR maps in the laminar model. The results suggest that if the thalamorecipient layers receive sufficiently uncorrelated input, abolishing or disrupting the inter-laminar feedback in primate layer 4C may allow separate uncorrelated OR maps to develop. Uncorrelated OR maps in layer 4 might, in turn, further disrupt columnar organisation in the superficial layers. If this prediction is possible to test and verify, it would demonstrate the importance of inter-laminar feedback in the development of columnar OR organisation.

6.2.5 Determining RFs using reverse correlation with pink noise

Comparison between RFs measured using white noise and pink noise may reveal discrepancies if units have connections to multiple SF channels. As high SF stimuli would tend to activate higher SF preference afferents more strongly, differences in the receptive field plots (such as those in figure B.1) determined by white noise and pink noise should be evident, if units receive input from afferents with different SF preferences. If RFs measured using both forms of stimuli are the same, it will add further weight to the prediction that thalamorecipient neurons receive input primarily from neurons with matching SF preferences. In general, pink noise may be a more suitable stimuli to use to measure RFs, as it may reduce the number of presentations that are required to determine the RFs of neurons that have lower SF preference, because it is not scale dependent.

6.3 Future work

The models described in this work are the first to incorporate multiple SF pathways in a developmental simulation. Although these models offer insight into possible features of cortical SF development, for practical purposes they contain a number of simplifications. With advances in the speed of computing and amount of memory available, it becomes increasingly practical to create more sophisticated models, incorporating additional features of the system. In this section a number of possible extensions to the models are described. Their implementation will offer further understanding of SF processing and organisation.

6.3.1 More realistic LGN representation

The models in this thesis contain four discrete, homogeneous sheets of DoG RGC/LGN cells, with SF preference spanning two octaves. In principle this range can easily be extended to be more realistic, simply by the inclusion of more pathways with different SF preferences. This extension would allow for more direct comparison to experimental results in macaque. However, simulating the extended models would be very

computationally demanding. Increasing the laminar V1 model to span 4 octaves would require roughly 6 times as much memory, resulting in a 12GB simulation, and would be about 24 times slower to run than the current simulation (based on extrapolations from the current running times).

Once a full range of SF preferences is practical to simulate, the distribution of LGN cells could be matched to specific species, which is important because there can be large differences in the properties of cell types between species. Creating models based more closely upon specific animals may serve to explain particular features of organisation specific to that species.

Variation with eccentricity of DoG size could also be incorporated into the model, if the overall retina and cortex area were increased. This extension would allow for comparison to experimental work that has studied the variation in feature maps with eccentricity (such as Xu et al., 2007). Again, the computing requirements are currently prohibitive, but will be practical in the near future. Increasing the simulation's cortex size to roughly the full macaque V1 area (which is about 1000mm^2) will require a simulation 160 times larger, consuming 320 GB of memory for the current 2-octave SF range. Also incorporating the full SF preference range (as described above) would result in a 1920GB simulation and take roughly 3840 times longer to run than the current laminar V1 model, which will not be practical for quite some time. However, simulating a narrow strip from the fovea to the monocular region at a lower density would be a feasible short-term approach.

6.3.2 Temporal processing

Adding temporal processing to the models would allow investigation into the possible correlation between the RF size of LGN cells and their response latency. As speculated above, this correlation may allow for development of full range SF maps by temporally segregating activity between channels, thus eliminating the need for a pruning mechanism. To model this, each channel would need a different delay, which is simple to do in the Topographic simulator. However, making this extension would require incorporating models of motion direction preference (Bednar and Miikkulainen, 2003), which

will greatly increase the computational requirements and model complexity.

Once extended in this way, it will be possible to investigate organisation of temporal frequency processing in V1 as well. A recent paper (Purushothaman et al., 2009) argues that there is an organisation for temporal frequency in bush baby, which can be tested in the model. A model containing maps for OR, SF and TF would also allow investigation into links between SF and TF organisation that have been observed in some previous studies. It would also provide a test bed to investigate whether cortical response patterns are determined by the intersection of distinct feature maps for OR, SF, TF and their bandwidths (Baker and Issa, 2005; Mante and Carandini, 2005; Zhang et al., 2007). An alternative hypothesis is that cortex is best described by a single map of spatiotemporal energy (Basole et al., 2003; Adelson and Bergen, 1985), whereby each neuron has a preference for a local region of a large, multidimensional space. It has been suggested that these two models are equivalent (Baker and Issa, 2005), which depends upon the separability of the feature dimensions and could be investigated rigorously in such a model.

6.3.3 Alternative pruning mechanism

If temporal latency differences are not sufficient to separate the channels, alternative pruning mechanisms need to be investigated that have a less severe effect upon the OR organisation. A more graceful pruning procedure would allow SF maps with a large range of SF preference, without dramatically affecting the quality of the OR maps. This would permit more detailed studies of the relationship between these two feature maps. Possible alternative pruning mechanisms are listed below, and discussed in more detail in section 4.4, in chapter 4:

- Increasing the time period over which pruning takes place.
- Only pruning neurons with a severe imbalance in afferent weights.
- Only pruning neurons where all the afferent pathways share a common OR tuning.
- Only pruning a fraction of units.

- Different pruning criteria, such as retaining the pathway with the most highly tuned OR preference.

In any of these cases, the emphasis will need to be on exploring biologically plausible mechanisms for how such pruning could occur, and whether the mechanisms are specific to SF cell classes or can apply to feature processing in general.

6.3.4 Quantitative feature map and RF comparisons and characterisation

If temporal segregation or a pruning mechanism that is less disruptive to the OR organisation are successful, then a more extensive study of the relationship between OR and SF organisation in the model can be conducted. Fully characterising the relationships between these maps in the model would allow better comparison to experimental results. This may provide a powerful means to study observed intermap relationships, such as the positioning of OR pinwheels over high/low SF regions, which this study only touches on due to the fractured OR map in the pruned models.

In addition, if the LGN representation were extended to match the distribution of sizes found in animals, then a more systematic analysis of the RFs and their distribution could be performed. Fitting Gabor filters to RFs would allow for RF population characterisation, and quantitative comparison could then be made with the experimental observations (such as those of Ringach, 2002). This would allow us to understand the detailed differences in SF processing that arise from the differences in subcortical processing of each species.

6.3.5 Colour maps and CO blobs

Several studies have found a correlation between cytochrome oxidase blobs and lower SF preferences (Tootell et al., 1988; Silverman et al., 1989; Born and Tootell, 1991; Shoham et al., 1997; Xu et al., 2007). Other studies have found that blobs contain large numbers of colour-selective cells, which are often not OR selective (Livingstone and

Hubel, 1984). Cells not selective for OR will appear to have lower SF preferences, and so these cells may give confusing results for SF map measurement. In future work, it will be important to extend the model to include colour image inputs and colour map development in V1 (Bednar et al., 2005). With such an extension, it should be possible to clearly separate differences in SF preference due to colour selectivity from those due to having larger OR-selective RFs, and thus explain the observed relationships between SF maps and CO blobs.

6.3.6 More detailed V1 circuitry

The current models have a very simple cortical circuit, and do not include all the layers found in V1, or details of the form of connectivity between them. The form and strength of interconnections between different laminae are not well known or understood, even though a great deal of attention has been paid to them (for review see Sincich and Horton, 2005). Recent research (Binzegger et al., 2004) has taken a different approach to characterising and quantifying interlamina connectivity. By determining the number and distribution of synapses across laminae for different neuron types, and then extrapolating based upon published studies of the distribution of neuron types in each lamina, Binzegger et al. estimate the connectivity between neurons in different layers. By including the additional V1 layers in a model and using hardwired limits to the strength or the extent of connectivity between laminae (taken from this type of analysis), it may be possible for the model to develop cortical circuitry that more closely parallels that found in animals. This may allow for the emergence of other forms of feature preference at different points in the cortical circuit. For example, my colleagues Antolik and Bednar (2008) have demonstrated that altering the feedback relationships between different laminae can lead to phase preference being randomly distributed in the thalamorecipient layers, resulting in complex cell formation in the superficial layers. Note that if combining the laminar V1 model with the Antolik and Bednar (2008) complex cell model, the size of afferent connections from layer 4C to layer 2/3 would need to be increased. In the current laminar V1 model phase preference is topographically mapped, so adjacent layer 4C neurons have very similar RFs, and thus little would be achieved by increasing the afferent connection field size to layer 2/3 in the current SF models.

Modelling all of the laminae in V1 is important, because neurons in several specific layers and compartments of V1 (i.e., not just the superficial layers) project to higher cortical regions associated with visual processing. As well as CO blobs and inter-blobs, projections to V2 arise from the deeper layers 4A, 5 and 6. This suggests that neurons in these different laminae may have response properties that are importantly distinct from each other. Layer 6 receives thalamic input from M and P cells, while layer 4A receives input from K cells. This first-order thalamic input may lead to different response properties for these neurons, e.g. in SF preference, that are important for visual processing. Building a more comprehensive simulation of V1 would allow investigation into subtle differences between neurons' response properties in different laminae that could then be tested in experiments.

6.4 Conclusion

The suggested experiments in this chapter allow the predictions of this modelling work to be tested. The outcome of these or similar experiments will aid in determining the mechanisms for SF processing and organisation. Particularly of interest is how specialisation for SF preference can arise for individual units, and whether temporal processing is important for this. Future computational simulations may be able to examine these issues before experimental work can, due to the technological hurdles faced by experimentalists in imaging the deeper cortical layers. Such simulations may help experimentalists to refine the questions that they seek to answer.

Chapter 7

Conclusion

7.1 Conclusion

By using computer models based upon animal visual systems, the work in this thesis shows how SF maps may arise in animals from visual experience. The results suggest that a separation between SF channels is required in order for a Hebbian learning network to develop SF maps covering a large range of SFs.

Chapter 4 demonstrates that a mechanistic developmental model of the visual system can organise realistic topographic OR and SF maps through learning from natural image stimuli. V1 neurons in the model develop realistic Gabor patch type RFs. The single-layer V1 model is the first of its type to include multiple SF pathways. Due to the substantial overlap in activity between SF pathways, pruning of connections to V1 neurons is required for V1 to develop a range of SF preference that spans the input space. Alternative types of pruning that may cause less disruption to the OR organisation are suggested for future implementation. It is also possible that temporal firing differences between thalamic neurons with different SF preferences may be sufficient to allow thalamorecipient neurons to develop more specialised preferences for SF.

Chapter 5 shows how segregation of SF channel input to different cortical sublaminae can result in a laminar organisation for SF preference, and that with a small amount of feedback between these laminae columnar topographic organisation for OR can also arise. This for the first time demonstrates how columnar organisation for orientation

may be dependant upon feedback, because the afferent input to separate laminae may be different. The model also demonstrates how a topographic organisation for SF preference can emerge in the higher layers, through feed-forward input from the layer 4 sublaminae. As in chapter 4, the RFs for units in these different layers are generally Gabor-shaped.

This work makes two central predictions: segregation of SF input to cortical neurons is a requirement for the generation of RFs with SF preference spanning the input space, and laminar organisation for SF preference arises in the thalamorecipient sublaminae of layer 4, due to the known segregation in input to these sublaminae. In chapter 6, experiments are suggested that would allow investigation into these issues. Extensions to the models are also discussed, which can generate further questions and may provide further advances in our understanding of SF organisation. These simulations can be a very useful approach, as the suggested experiments may prove impractical.

Overall, this body of work provides a detailed and systematic exploration of SF organisation in V1. The results have generated testable predictions relating to the separation of SF processing, and the possible organisation of laminar processing for SF. The success of the described models in reproducing observed experimental results validates the assumptions upon which the models are based, and suggests that the models do capture important aspects of the visual system. Primarily in relation to SF processing, the need for multiple SF pathways, so that the full range of SF preference in V1 can be represented, had not previously been addressed in a computer model. I hope further work will extend this approach and help us to unravel the complexities of visual processing and organisation.

Appendix A

A.1 Model parameters

This section lists the parameter values for all simulations in this thesis. Table A.1 lists the parameters for V1 in the single-layer V1 model (chapter 4) and for layer 2/3 in the laminar V1 model (chapter 5). These parameters also apply to the V1 sublaminae in the laminar model (chapter 5), except for the lateral connection strengths, which are listed in tables A.3 and A.4. In table A.4, the edge sheets (a and d) in layer 2/3 receive double the feedback strength. They both only have one neighbouring sheet, so this ensures that each cortical sheet receives roughly the same overall amount of input activation.

All values are in sheet coordinates unless otherwise noted. For instance, with a V1 size of 1.0 and a density of 96, the V1 array will be a 96×96 array. A radius $r_E = 0.03$ represents 0.03×96=2.88, i.e., a circle of radius 2.88 neural units. Learning rates are scaled by the number of units before application, i.e., $\alpha_E = 1.8087$ for a connection field with 21 input units gives a per-connection learning rate of $\overline{\alpha}_E = 1.8/21 = 0.0861$. Units are considered part of a connection field if their centre falls within the specified boundary radius. Roughly, the number of units in a given connection field can be calculated by multiplying the area of the connection field by the appropriate sheet density. For instance, the number of units in the lateral excitatory connection field is roughly given by $\pi \times r_E \times r_E \times N_{V1} \times N_{V1}$, i.e. 26 in this case. The actual number of units can be obtained by counting each unit inside the radius, and for this example is 21. The number of units in each thalamocortical and cortical-cortical connection field is given in tables A.1 and A.2.

Gaussian clouds are used to intialise the weight patterns for lateral inhibitory connections, and afferent thalamocortical connections. A Gaussian cloud pattern is constructed by multiplying a 2D isotropic Gaussian pattern with uniform random noise in the interval [0,1.0]. For Gaussian functions, radius refers to the standard deviation. Circular boundaries, specified by appropriate boundary radii are used to provide hard limited boundaries for connection fields, outside of these boundaries all weights for the applicable connection field are zero.

Parameter	Value	Description
$s_R = s_{V1} + 2.0r_{V1} + 2.0r_{LGN}$	4.875	Retina size
s_{V1}	1.0	V1 size, in sheet coordinates
R_d	24	Retina density (neurons per unit size)
N_{V1}	96	V1 density (neurons per unit size)
r_E	0.03	Boundary radius of lateral excitatory connection field
r_I	0.1718	Boundary radius of lateral inhibitory connection field
n_E	21	Number of units in lateral excitatory connection field
n_I	861	Number of units in lateral inhibitory connection field
γ_E	1.3	Strength of lateral excitatory projection, in equation 3.4.
γ_I	-1.7	Strength of lateral inhibitory projection, in equation 3.4.
α_E	1.8087	Learning rate of lateral excitatory weights, in equation 3.10.
α_I	1.8087	Learning rate of lateral inhibitory weights, in equation 3.10.
r_{GE}	0.02	Radius of Gaussian pattern for initial lateral excitatory weights
r_{GI}=r_I	0.1718	Radius of Gaussian cloud for initial lateral inhibitory weights
mu, μ	0.05	Target average firing rate, in equations 3.8 and 3.9.
β	0.0002	Learning rate for homeostatic plasticity mechanism, in equations 3.8 and 3.9.
a	14.5	Initial $\boldsymbol{a}$ parameter value in sigmoid function, in equation 3.5.
b	-4	Initial $\boldsymbol{b}$ parameter value in sigmoid function, in equation 3.5.
λ	640	Learning rate time constant, in equation 3.12.
phi, ϕ	0.4845	Initial learning rate value, in equation 3.12.

Table A.1: Cortical sheet parameters.

Parameter	Value for a pathway	Value for b pathway	Value for c pathway	Value for d pathway	Description
c	4.0	2.52	1.59	1.0	LGN size factors
$s_{LGN} = s_{V1} + 2.0 r_{V1}$	2.625	2.024	1.646	1.406	LGN sheet size (rounded to be a multiple of the reciprocal of the density)
L_d	24	24	24	24	LGN density (neurons per unit size)
$r_{LGN} = 0.2813c$	1.125	0.7088	0.4472	0.2813	Boundary radius of LGN DoG connection field
$\sigma_c = 0.0277c$	0.1108	0.0698	0.0441	0.0277	Radius of LGN DoG centre Gaussian
$\sigma_s = 0.1108c$	0.4431	0.2792	0.1762	0.1108	Radius of LGN DoG surround Gaussian
$r_{V1} = 0.2031c$	0.8125	0.5119	0.3230	0.2031	Boundary radius of thalamocortical connection field
n_{V1}	1201	489	177	69	Number of units in thalamocortical connection field
$r_{GCF} = r_{V1}$	0.8125	0.5119	0.3230	0.2031	Radius of Gaussian cloud for initial weights in V1-LGN connection fields
γ_R	2.33	2.33	2.33	2.33	Strength of Retina - LGN connection
γ_L	1.0	1.0	1.0	1.0	Strength of LGN - V1 connection

Table A.2: SF LGN pathway parameter values.

Parameter	100WF, 0BF	95WF, 5BF	80WF, 20BF	50WF, 50BF	Description
γ_{EW}	1.3	1.235	1.04	0.65	Excitatory strength within lamina sheets
γ_{IW}	-1.7	-1.615	-1.36	-0.85	Inhibitory strength within lamina sheets
γ_{EB}	0.0	0.065	0.26	0.65	Excitatory strength between lamina sheets
γ_{IB}	0.0	-0.085	-0.34	-0.85	Inhibitory strength between lamina sheets

Table A.3: Sublaminae feedback strengths.

Parameter	Value	Description
	γ_{EW}	Strength of within sheet excitatory lateral connection
	γ_{IW}	Strength of within sheet inhibitory lateral connection
	$\gamma_{EB}/2$	Strength of excitatory connection from sheet A to B
	γ_{EB}	Strength of excitatory connection from sheet B to A
	$\gamma_{EB}/2$	Strength of excitatory connection from sheet B to C
	$\gamma_{EB}/2$	Strength of excitatory connection from sheet C to B
	$\gamma_{EB}/2$	Strength of excitatory connection from sheet C to D
	γ_{EB}	Strength of excitatory connection from sheet D to C
	$\gamma_{IB}/2$	Strength of inhibitory connection from sheet A to B
	γ_{IB}	Strength of inhibitory connection from sheet B to A
	$\gamma_{IB}/2$	Strength of inhibitory connection from sheet B to C
	$\gamma_{IB}/2$	Strength of inhibitory connection from sheet C to B
	γ_{IB}	Strength of inhibitory connection from sheet C to D
	$\gamma_{IB}/2$	Strength of inhibitory connection from sheet D to C
γ_{23}	1.0	Strength of sublaminae to layer 2/3 connection
r_{23}	0.01	Boundary radius of sublaminae sheets to layer 2/3 connection field

Table A.4: Between sheet connection parameters

A.2 Dependence of results on parameters

This type of large scale developmental model is not amenable to exhaustive parameter searches. Most parameter values are non-critical and can generally change by around 50% without significantly affecting the results. Many of the parameters within the model are inherited from the previous LISSOM model, and were not tweaked or optimised for this model. On the occasions where parameter searches were conducted, a single parameter would be changed at a time, across a specific range of interest. Otherwise identical simulations would be run, and the effect of the variation in this parameter over the organisation of the model was then examined. Where appropriate, further simulations would be conducted, using a reduced or different range of values for the parameter under consideration. This is an iterative optimisation approach. There are difficulties in this type of parameter search, because altering one parameter may influence the optimal range for other parameters. This behaviour can be very difficult to predict and to address. In the following paragraphs, the derivation of the model paramater values (those not already discussed in chapters 4, and 5) are discussed.

In the previous LISSOM model, the radius of the lateral excitatory connections, r_E, was adjusted dynamically from 0.1 to 0.01 over 5000 iterations, but in this model, this adjustment was found to be unnecessary due to the homeostatic plasticity mechanism (Kurniawan, 2006). Both the values for r_E, and r_I are fixed and non-critical.

The scaling factors for the excitatory, γ_E, and inhibitory, γ_I, projections were determined through a 2 dimensional search. The scaling factors γ_R, and γ_L are inherited from the LISSOM model. Changes of over about 30% to any of these values can have quite dramatic affects upon the models, leading to poorly organised maps, and so it is crucial that they remain in balance.

The parameters determining the learning rates of the excitatory, inhibitory and afferent connections (α_E, α_I, ϕ, and λ) are all inherited from the previous LISSOM model.

The initial weight patterns of lateral and afferent connections are not critical. Al-

tering their initial organisation may change details of the maps that develop, but will not change the qualitative nature of the results.

In the previous LISSOM incarnation that provides the foundation for these models, a piecewise linear approximation to a sigmoid was used as the output function from V1, which caused saturation of V1 at their maximum activity levels. Scheduled changes over the development period adjusted the parameters to this function. In the models within this thesis, the piecewise linear approximation has been replaced by a smooth sigmoid function, which changes dynamically to maintain homeostasis (as described in section 3.3.1). The initial parameters for this sigmoid function, a and b, were matched so that it resembles the initial parameters for the piecewise linear function. The target firing rate, μ, was matched to the average firing rate from the previous LISSOM model, and the learning rate, β, was adjusted such that the model would be responsive to developmental changes. Overall, these values produce similar mean V1 activity levels over the course of development, but without the need for scheduled changes to the parameters.

Appendix B

B.1 Reverse correlation stimuli for mapping RFs

Reverse correlation is an experimental process whereby the RF for a neuron can be determined (Jones and Palmer, 1987b). Stimuli are presented to an animal and the response of a neuron to each presentation is recorded. At the end of this process the stimuli are weighted by the neuron's response and averaged. The resulting response-weighted average of the stimuli is an estimate of the neuron's RF, provided the stimuli form an orthogonal set, and the neuron's responses are linear with respect to spatial patterns (Willmore and Smyth, 2003). Using this approach in a computer simulation has the advantage that the RFs for all of the neurons within a simulation can be calculated simultaneously.

In experiments, RFs are often measured using input patterns consisting of single pixels (sparse noise) or arrays of white noise pixels (dense noise) (Ringach, 2004b). Both types of patterns primarily contain very small features (depending on the pixel size), which most strongly activate the high-frequency channels in the model. Because all channels converge on each neuron in the unpruned model, such test patterns will result in over-representation of the high SF channels in RF plots.

In many cases, RFs measured in this way show only high frequency patterns, even though the model neuron actually responds more strongly to lower frequency patterns (e.g. neuron 23,37 in figure B.1). In others, both large and small patterns are visible in the RFs (e.g. neuron 45,72 in figure B.1). To avoid these artifacts, the model RFs elsewhere in this thesis were measured using pink noise, which has patterns at all spatial

scales.

To illustrate why the stimulus matters, each row of figure B.1 shows a unit's RF measured with single pixel stimuli, white noise, and pink noise; the inset numbers indicate the measured RF's preferred SF (determined from hand fitting a Gabor patch to the RF). The column to the right shows the preferred SF of the unit determined from the automatic SF map measurement using sine grating responses. In each case, measuring with pink noise resulted in Gabor fittable RFs that yielded a reasonable match to the preferred SF of the unit. The third and fourth rows illustrate that in the pruned case (where there is only one SF pathway to cortex), the stimuli used to generate the RF are less critical, although for low frequency units white noise may take many presentations, due to its inefficiency in stimulating such units. Note that both afferent and lateral connections must be pruned to avoid artifacts for sparse noise and white noise.

These results suggest that the type of stimulus used may be less crucial for animal neurons, which primarily receive inputs from similar sized LGN neurons (Alonso et al., 2001). The superposition effect seen in neuron (45,72) in figure B.1 is also likely to be hidden in animals, where LGN neurons cover a range of sizes rather than having one of four discrete sizes as in the model. Even so, the model results suggest that future RF mapping experiments could utilise pink noise, to avoid artifacts in some cases and to generate accurate RFs with fewer presentations for low-SF neurons.

Experimentally, white or sparse noise is often used to determine RFs, for measuring across a range of SFs, such patterns are not appropriate due to their inefficiency in driving multiple SF pathways. Unlike natural images, white noise has a flat power spectrum, which means that on average all SFs are represented equally. If reverse correlation is performed with an alternative stimuli set, then to avoid bias, the results must be corrected by their power spectrum. There are a number of methods for accomplishing this (for discussion see Willmore and Smyth, 2003). Here the power spectrum correction method is outlined. This correction involves dividing the Fourier transform of the reverse-correlation RF estimate by the mean power spectrum of the stimuli. $\hat{f}_b$, the corrected estimate of the RF, is given by:

$$\mathcal{F}[\hat{f}_b] = \frac{n_s \mathcal{F}[\hat{f}_a]}{\sum_{i=1}^{n_s} \mathcal{F}^2[S_i]} \tag{B.1}$$

where n_s is the total number of stimulus presentations, S_i is the ith stimulus, $\hat{f}_a$ is the uncorrected estimate of the RF, $\mathcal{F}[x]$ is the Fourier spectrum of x, and $\mathcal{F}^2[x]$ is the power spectrum of x. This method requires that there is no spatial dependence of the image statistics, this means that on average all parts of the image have the same spatial distribution of features.

As a means of demonstrating the correction procedure, it was conducted upon RFs from a binary model (identical to 4.3.3, but with cortical density $N_{V1} = 24$). Figure B.2 outlines the results of the correction. Far more presentations are required to produce RFs with the correction, and the resulting RFs are more noisy. In the pruned case the corrected and uncorrected RFs match well, but more presentations are required to allow the low frequency RFs to emerge. In the unpruned case, there is some correspondence between the RFs, but they are dominated by the higher SF channel that is more effectively driven by white noise patterns, and they are still very noisy and require more presentations to fully characterise them. The RF results presented in this thesis do not include the power spectrum correction, because the vastly increased processing time required required makes them impractical to calculate for larger simulations.

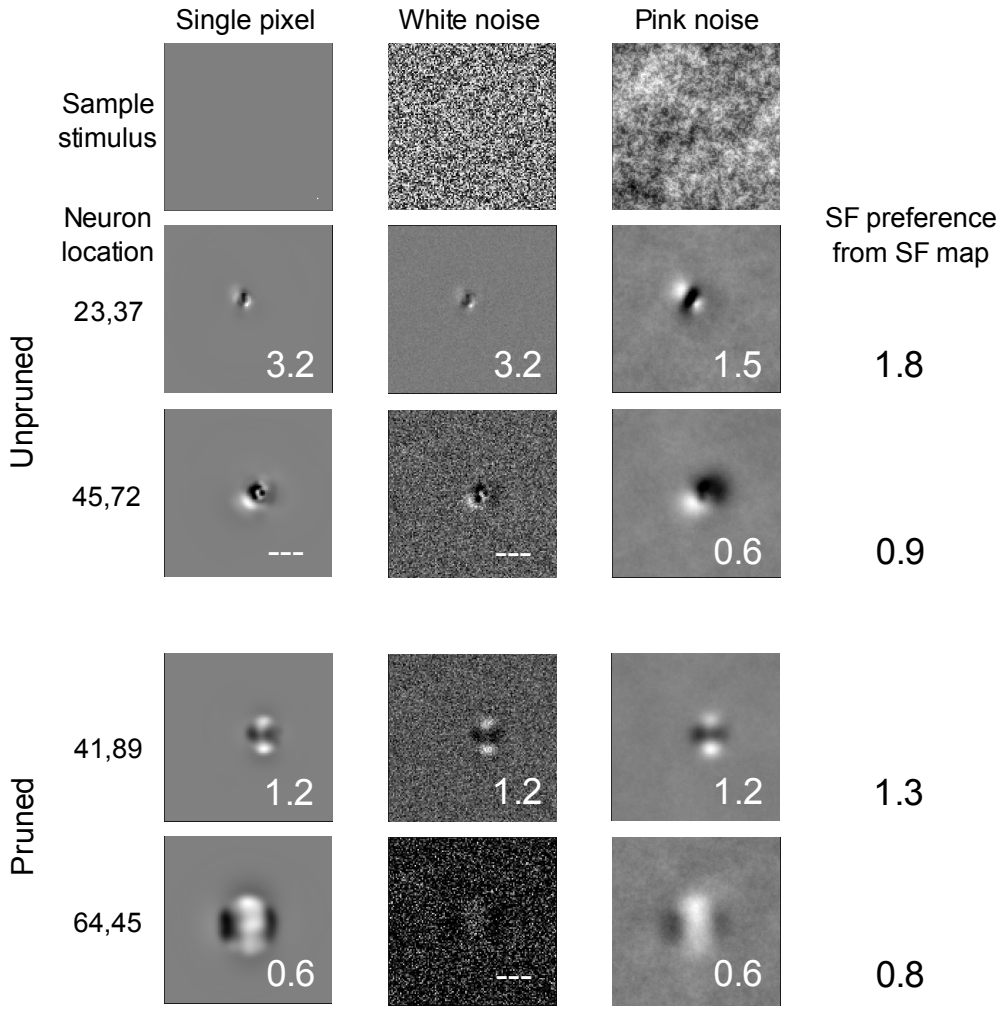

Figure B.1: Single pixel, white noise and pink noise stimuli reverse correlation comparison. Each of the bottom four rows shows RF plots for one neuron, measured using single pixel, white noise or pink noise stimuli (as shown in the examples in the top row). White numbers inset into the plots are the preferred SF (in cpud) of a hand-fitted Gabor patch to that RF plot. Where no number is shown, a Gabor patch was a poor fit to the plot. Black numbers in final column are the preferred SF (in cpud) for that unit taken from the SF map; these represent the frequency of the sine grating that maximally activates this unit. For unpruned neurons, the high frequency input to the neuron dominates the sparse and white noise RF plots, as the stimuli drive this pathway better than the lower frequency pathways. Due to pink noise on average having equal energy at all SFs, the pink noise RF plot's SF preference matches the SF map. The last two rows are for pruned neurons, where the stimulus is less critical, but white noise requires many presentations to generate RF plots for low frequency units. For the model, pink noise thus yields more accurate RF plots than sparse or dense noise.

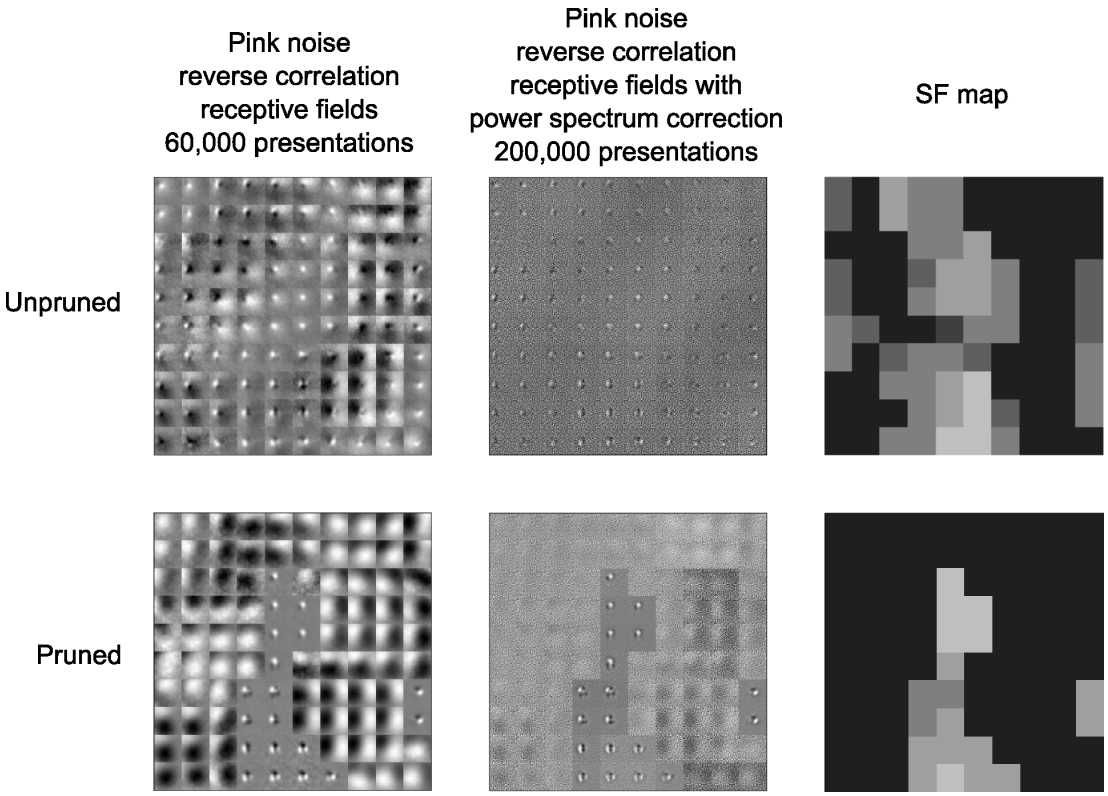

Figure B.2: RFs calculated using pink noise with and without power spectrum correction.

(Top row) Details from an unpruned binary simulation (with cortical density, $N_{V1} = 24$). The first plot show RFs generated after 60,000 presentations of pink noise. The second plot shows RFs generated after 200,000 presentations of pink noise, corrected by the power spectrum. The third plot shows the SF map for the corresponding region of cortex. (Bottom row) Corresponding plots from a pruned binary simulation. Note how in the pruned case, the RFs generated using the corrected procedure match the uncorrected RFs, particularly for the smaller RFs. The larger RFs are beginning to emerge, but require even more presentations.

Appendix C

C.1 Training stimuli

To determine the overall distribution of SFs and ORs presented to the models, Fourier analysis was performed on the image set used in the simulations. By summing the Fourier transform of every image presented, we are able to generate cumulative Fourier transforms (CFT) for the presented images. Figure C.1 displays these CFTs for the images as presented in the simulations (first column), the images without any rotation (second column), and the images randomly orientated from a uniform distribution, so as to fully eliminate any orientation bias (third column). The fourth column shows a 1/f distribution for comparison. The non-periodic nature of the image stimuli at the aperture boundaries of the images creates an edge effect, which displays itself in the CFT as a cross-like artifact (evident in the upper row). By masking the images with a soft edged disk, this edge effect can be largely eliminated (as demonstrated in the lower column), which gives a better representation of the actual prevalence of horizontal and vertical orientations. Note how all of the CFTs are broadly comparable with the 1/f plot, demonstrating the similarity of the power spectrum of natural images to a 1/f distribution (Field, 1987). The CFTs for the stimulus set with no rotation show that there is somewhat more power in the horizontal and vertical directions. Simulations using this data set produce OR maps which are dominated by preferences for these orientations, to a greater extent than those seen from monkey, so it is important to add a small amount of rotation to better model typical visual experience.

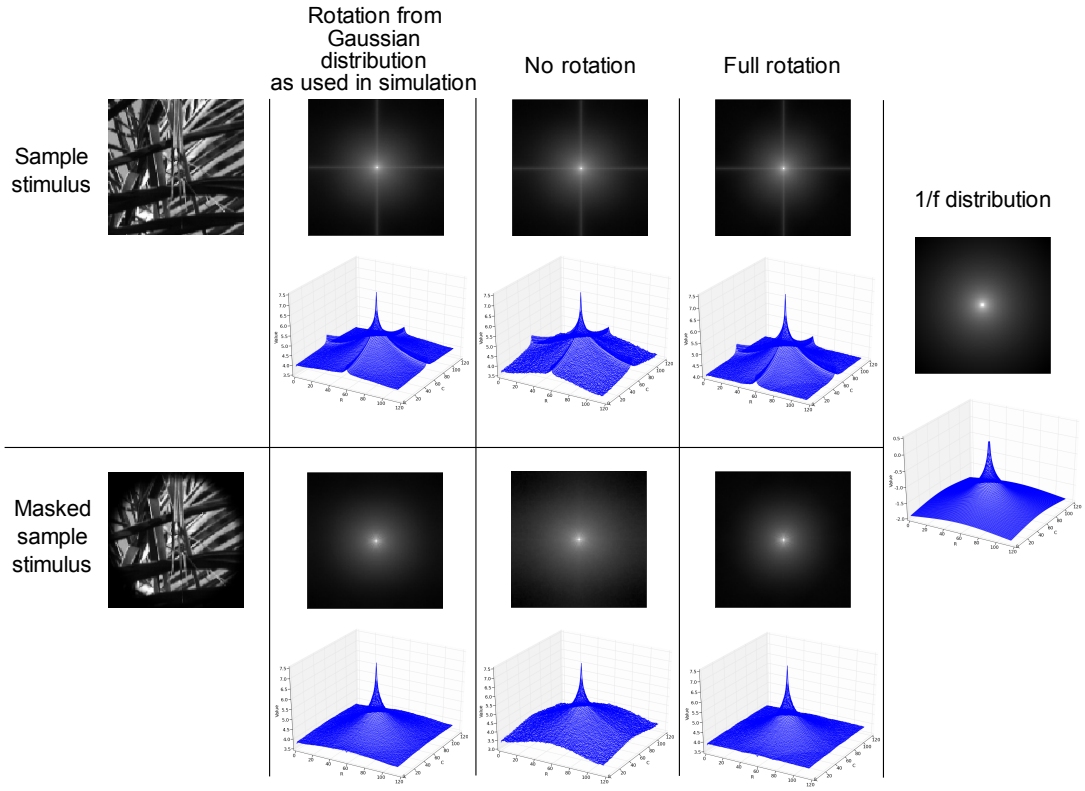

Figure C.1: Cumulative Fourier transforms of the input images.

(Upper row, left to right) Sample stimulus, CFT (Cumulative Fourier transform) of images as used in the simulations in this thesis, CFT of images without rotation, CFT of images fully rotated. The logarithm CFT plots are shown in both 2D and 3D to aid visualisation, as it is difficult to discriminate between the 2D plots. (Lower row, left to right) Sample stimulus with soft circular window to remove edge effects, CFT of windowed images as used in the simulations in this thesis, CFT of windowed images without rotation, CFT windowed of images fully rotated. (Far right) 1/f distribution. Note the similarity of all plots to the 1/f distribution.

Bibliography

Adelson EH, Bergen JR (1985) Spatiotemporal energy models for the perception of motion. *J Opt Soc Am A* 2:284–299.

Albrecht DG, Hamilton DB (1982) Striate cortex of monkey and cat: contrast response function. *J Neurophysiol* 48:217–237.

Alonso JM, Usrey WM, Reid RC (2001) Rules of connectivity between geniculate cells and simple cells in cat primary visual cortex. *J Neurosci* 21:4002–4015.

Alonso JM (2002) Neural connections and receptive field properties in the primary visual cortex. *Neuroscientist* 8:443–456.

Angelucci A, Levitt JB, Walton EJS, Hupe JM, Bullier J, Lund JS (2002) Circuits for local and global signal integration in primary visual cortex. *J Neurosci* 22:8633–8646.

Angelucci A, Sainsbury K (2006) Contribution of feedforward thalamic afferents and corticogeniculate feedback to the spatial summation area of macaque V1 and LGN. *J Comp Neurol* 498:330–351.

Antolik J, Bednar JA (2008) Developing maps of complex cells in a computational model of V1. In *Society for Neuroscience Abstracts*. Society for Neuroscience, www.sfn.org, Program No. 811.5.

Baker TI, Issa NP (2005) Cortical maps of separable tuning properties predict population responses to complex visual stimuli. *J Neurophysiol* 94:775–787.

Barlow HB (1989) Unsupervised learning. *Neural Computation* 1:295–311.

Basole A, White LE, Fitzpatrick D (2003) Mapping multiple features in the population response of visual cortex. *Nature* 423:986–990.

Bednar JA, Miikkulainen R (2006) Joint maps for orientation, eye, and direction preference in a self-organizing model of V1. *Neurocomputing* 69:1272–1276.

Bednar JA, Miikkulainen R (2003) Self-organization of spatiotemporal receptive fields and laterally connected direction and orientation maps. *Neurocomputing* 52-54:473–480.

Bednar JA, Paula JBD, Miikkulainen R (2005) Self-organization of color opponent receptive fields and laterally connected orientation maps. *Neurocomputing* 65-66:69–76.

Berardi N, Bisti S, Cattaneo A, Fiorentini A, Maffei L (1982) Correlation between the preferred orientation and spatial frequency of neurones in visual areas 17 and 18 of the cat. *J Physiol* 323:603–618.

Bhaumik B, Mathur M (2003) A cooperation and competition based simple cell receptive field model and study of feed-forward linear and nonlinear contributions to orientation selectivity. *J Comput Neurosci* 14:211–227.

Binzegger T, Douglas RJ, Martin KAC (2004) A quantitative map of the circuit of cat primary visual cortex. *J Neurosci* 24:8441–8453.

Blasdel GG (1992) Orientation selectivity, preference, and continuity in monkey striate cortex. *J Neurosci* 12:3139–3161.

Blasdel GG, Lund JS (1983) Termination of afferent axons in macaque striate cortex. *J Neurosci* 3:1389–1413.

Bonhoeffer T, Kim DS, Malonek D, Shoham D, Grinvald A (1995) Optical imaging of the layout of functional domains in area 17 and across the area 17/18 border in cat visual cortex. *Eur J Neurosci* 7:1973–1988.

Born RT, Tootell RB (1991) Spatial frequency tuning of single units in macaque supragranular striate cortex. *Proc Natl Acad Sci U S A* 88:7066–7070.

Bosking WH, Zhang Y, Schofield B, Fitzpatrick D (1997) Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *J Neurosci* 17:2112–2127.

Boycott BB, Wässle H (1974) The morphological types of ganglion cells of the domestic cat's retina. *J Physiol* 240:397–419.

Bressloff PC, Cowan JD (2003) A spherical model for orientation and spatial-frequency tuning in a cortical hypercolumn. *Philos Trans R Soc Lond B Biol Sci* 358:1643–1667.

Cai D, DeAngelis GC, Freeman RD (1997) Spatiotemporal receptive field organization in the lateral geniculate nucleus of cats and kittens. *J Neurophysiol* 78:1045–1061.

Callaway EM, Katz LC (1990) Emergence and refinement of clustered horizontal connections in cat striate cortex. *J Neurosci* 10:1134–1153.

Callaway EM (2005) Structure and function of parallel pathways in the primate early visual system. *J Physiol* 566:13–19.

Campbell FW, Cooper GF, Enroth-Cugell C (1969) The spatial selectivity of the visual cells of the cat. *J Physiol* 203:223–235.

Cang J, Rentería RC, Kaneko M, Liu X, Copenhagen DR, Stryker MP (2005) Development of precise maps in visual cortex requires patterned spontaneous activity in the retina. *Neuron* 48:797–809.

Carreira-Perpiñán MA, Lister RJ, Goodhill GJ (2005) A computational model for the development of multiple maps in primary visual cortex. *Cereb Cortex* 15:1222–1233.

Chapman B, Gödecke I (2000) Cortical cell orientation selectivity fails to develop in the absence of on-center retinal ganglion cell activity. *J Neurosci* 20:1922–1930.

Chapman B, Stryker MP (1993) Development of orientation selectivity in ferret visual cortex and effects of deprivation. *J Neurosci* 13:5251–5262.

Crair MC, Gillespie DC, Stryker MP (1998) The role of visual experience in the development of columns in cat visual cortex. *Science* 279:566–570.

Croner LJ, Kaplan E (1995) Receptive fields of P and M ganglion cells across the primate retina. *Vision Res* 35:7–24.

Das A, Gilbert CD (1997) Distortions of visuotopic map match orientation singularities in primary visual cortex. *Nature* 387:594–598.

DeAngelis GC, Ghose GM, Ohzawa I, Freeman RD (1999) Functional micro-organization of primary visual cortex: receptive field analysis of nearby neurons. *J Neurosci* 19:4046–4064.

Derrington AM (1984) Development of spatial frequency selectivity in striate cortex of vision-deprived cats. *Exp Brain Res* 55:431–437.

Douglas RJ, Koch C, Mahowald M, Martin KAC, Suarez HH (1995) Recurrent excitation in neocortical circuits. *Science* 269:981–985.

Durbin R, Mitchison G (1990) A dimension reduction framework for understanding cortical maps. *Nature* 343:644–647.

Durbin R, Willshaw D (1987) An analogue approach to the travelling salesman problem using an elastic net method. *Nature* 326:689–691.

Everson RM, Prashanth AK, Gabbay M, Knight BW, Sirovich L, Kaplan E (1998) Representation of spatial frequency and orientation in the visual cortex. *Proc Natl Acad Sci U S A* 95:8334–8338.

Farley BJ, Yu H, Jin DZ, Sur M (2007) Alteration of visual input results in a coordinated reorganization of multiple visual cortex maps. *J Neurosci* 27:10299–10310.

Ferster D, LeVay S (1978) The axonal arborizations of lateral geniculate neurons in the striate cortex of the cat. *J Comp Neurol* 182:923–944.

Field DJ (1987) Relations between the statistics of natural images and the response properties of cortical cells. *J Opt Soc Am A* 4:2379–2394.

Field DJ (1994) What is the goal of sensory coding? *Neural Computation* 6:559–601.

Field GD, Chichilnisky EJ (2007) Information processing in the primate retina: circuitry and coding. *Annu Rev Neurosci* 30:1–30.

Gilbert CD (1977) Laminar differences in receptive field properties of cells in cat primary visual cortex. *J Physiol* 268:391–421.

Gilbert CD, Wiesel TN (1989) Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *J Neurosci* 9:2432–2442.

Girman SV, Sauvé Y, Lund RD (1999) Receptive field properties of single neurons in rat primary visual cortex. *J Neurophysiol* 82:301–311.

Goodhill GJ (1993) Topography and ocular dominance: a model exploring positive correlations. *Biol Cybern* 69:109–118.

Goodhill GJ, Willshaw DJ (1990) Application of the elastic net algorithm to the formation of ocular dominance stripes. *Network: Comput. Neural Syst.* 1:41–59.

Goodhill GJ (2007) Contributions of theoretical modeling to the understanding of neural map development. *Neuron* 56:301–311.

Grinvald A, Shoham D, Shmuel A, Glaser D, Vanzetta I, Shtoyerman E, Slovin H, Wijnbergen C, Hildesheim R, Sterkin A, Arieli A (1999) *Modern Techniques in Neuroscience Research*, chapter 34: In-vivo optical imaging of cortical architecture and dynamics, pp. 893 – 969 Springer, Berlin.

Gur M, Kagan I, Snodderly DM (2005) Orientation and direction selectivity of neurons in V1 of alert monkeys: functional relationships and laminar distributions. *Cereb Cortex* 15:1207–1221.

Hendrickson AE, Wilson JR, Ogren MP (1978) The neuroanatomical organization of pathways between the dorsal lateral geniculate nucleus and visual cortex in Old World and New World primates. *J Comp Neurol* 182:123–136.

Hensch TK (2004) Critical period regulation. *Annu Rev Neurosci* 27:549–579.

Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 160:106–154.

Hubel DH, Wiesel TN (1972) Laminar and columnar distribution of geniculo-cortical fibers in the macaque monkey. *J Comp Neurol* 146:421–450.

Hubel DH, Wiesel TN (1974) Uniformity of monkey striate cortex: a parallel relationship between field size, scatter, and magnification factor. *J Comp Neurol* 158:295–305.

Hubel DH, Wiesel TN, Stryker MP (1978) Anatomical demonstration of orientation columns in macaque monkey. *J Comp Neurol* 177:361–380.

Hübener M, Shoham D, Grinvald A, Bonhoeffer T (1997) Spatial relationships among three columnar systems in cat area 17. *J Neurosci* 17:9270–9284.

Huberman AD, Feller MB, Chapman B (2008) Mechanisms underlying development of visual maps and receptive fields. *Annu Rev Neurosci* 31:479–509.

Huberman AD, Speer CM, Chapman B (2006) Spontaneous retinal activity mediates development of ocular dominance columns and binocular receptive fields in v1. *Neuron* 52:247–254.

Humphrey AL, Sur M, Uhlrich DJ, Sherman SM (1985) Projection patterns of individual X- and Y-cell axons from the lateral geniculate nucleus to cortical area 17 in the cat. *J Comp Neurol* 233:159–189.

Hupé JM, James AC, Girard P, Bullier J (2001) Response modulations by static texture surround in area V1 of the macaque monkey do not depend on feedback connections from V2. *J Neurophysiol* 85:146–163.

Hyvärinen A, Hoyer PO (2001) A two-layer sparse coding model learns simple and complex cell receptive fields and topography from natural images. *Vision Res* 41:2413–2423.

Issa NP, Trepel C, Stryker MP (2000) Spatial frequency maps in cat visual cortex. *J Neurosci* 20:8504–8514.

Issa NP, Rosenberg A, Husson TR (2008) Models and measurements of functional maps in V1. *J Neurophysiol* 99:2745–2754.

Jones JP, Palmer LA (1987a) An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *J Neurophysiol* 58:1233–1258.

Jones JP, Palmer LA (1987b) The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *J Neurophysiol* 58:1187–1211.

Kaplan E, Shapley RM (1982) X and Y cells in the lateral geniculate nucleus of macaque monkeys. *J Physiol* 330:125–143.

Khaytin I, Chen X, Royal DW, Ruiz O, Jermakowicz WJ, Siegel RM, Casagrande VA (2008) Functional organization of temporal frequency selectivity in primate visual cortex. *Cereb Cortex* 18:1828–1842.

Kohonen T (1982) Self-organizing formation of topologically correct feature maps. *Biol. Cybern.* 44:59–69.

Kremers J, Weiss S (1997) Receptive field dimensions of lateral geniculate cells in the common marmoset (*Callithrix jacchus*). *Vision Res* 37:2171–2181.

Kurniawan V (2006.) Self-organizing visual cortex model using a homeostatic plasticity mechanism. *Master's thesis, The University of Edinburgh, Scotland, UK, 2006.* .

Lachica EA, Casagrande VA (1992) Direct W-like geniculate projections to the cytochrome oxidase (CO) blobs in primate visual cortex: axon morphology. *J Comp Neurol* 319:141–158.

Land MF, Fernald RD (1992) The evolution of eyes. *Annu Rev Neurosci* 15:1–29.

LeVay S, Ferster D (1977) Relay cell classes in the lateral geniculate nucleus of the cat and the effects of visual deprivation. *J Comp Neurol* 172:563–584.

Levitt JB, Schumer RA, Sherman SM, Spear PD, Movshon JA (2001) Visual response properties of neurons in the LGN of normally reared and visually deprived macaque monkeys. *J Neurophysiol* 85:2111–2129.

Lewicki MS (1998) A review of methods for spike sorting: the detection and classification of neural action potentials. *Network* 9:R53–R78.

Livingstone M, Hubel D (1988) Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science* 240:740–749.

Livingstone MS, Hubel DH (1984) Anatomy and physiology of a color system in the primate visual cortex. *J Neurosci* 4:309–356.

Lund JS, Henry GH, MacQueen CL, Harvey AR (1979) Anatomical organization of the primary visual cortex (area 17) of the cat. A comparison with area 17 of the macaque monkey. *J Comp Neurol* 184:599–618.

Maffei L, Fiorentini A (1977) Spatial frequency rows in the striate visual cortex. *Vision Res* 17:257–264.

Mallik AK, Husson TR, Zhang JX, Rosenberg A, Issa NP (2008) The organization of spatial frequency maps measured by cortical flavoprotein autofluorescence. *Vision Res* 48:1545–1553.

Mante V, Carandini M (2005) Mapping of stimulus energy in primary visual cortex. *J Neurophysiol* 94:788–798.

Martinez LM, Alonso JM (2001) Construction of complex receptive fields in cat primary visual cortex. *Neuron* 32:515–525.

Martinez LM, Wang Q, Reid RC, Pillai C, Alonso JM, Sommer FT, Hirsch JA (2005) Receptive field structure varies with layer in the primary visual cortex. *Nat Neurosci* 8:372–379.

Mathur M, Bhaumik B (2004) Study of spatial frequency selectivity and its spatial organization in the visual cortex though a feedforward model. *Neurocomputing* 65-66:85–89.

McMahon MJ, Lankheet MJ, Lennie P, Williams DR (2000) Fine structure of parvocellular receptive fields in the primate fovea revealed by laser interferometry. *J Neurosci* 20:2043–2053.

Miikkulainen R, Bednar J, Choe Y, Sirosh J (2005) *Computational maps in the visual cortex*. Springer.

Miller KD (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 KD, Keller JB, Stryker MP (1989) Ocular dominance column development: analysis and simulation. *Science* 245:605–615.

Mishkin M, Ungerleider LG, Macko KA (1983) Object vision and spatial vision: two cortical pathways. *Trends in Neurosciences* 6:414–417.

Molotchnikoff S, Gillet PC, Shumikhina S, Bouchard M (2007) Spatial frequency characteristics of nearby neurons in cats' visual cortex. *Neurosci Lett* 418:242–247.

Obermayer, Blasdel, Schulten (1992) Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. *Phys Rev A* 45:7568–7589.

Ohki K, Chung S, Ch'ng YH, Kara P, Reid RC (2005) Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature* 433:597–603.

Olmos A, Kingdom FAA (2004) McGill Calibrated Colour Image Database. *http://tabby.vision.mcgill.ca* .

Olshausen BA, Field DJ (1996) Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381:607–609.

Paraga A (2003) Is the human visual system optimised for encoding the statistical information of natural scenes? Ph.D. diss., University of Bristol.

Peichl L, Wässle H (1979) Size, scatter and coverage of ganglion cell receptive field centres in the cat retina. *J Physiol* 291:117–141.

Purushothaman G, Khaytin I, Casagrande VA (2009) Quantification of optical images of cortical responses for inferring functional maps. *J Neurophysiol* 101:2708–2724.

Rehn M, Sommer FT (2007) A network that uses few active neurones to code visual input predicts the diverse shapes of cortical receptive fields. *J Comput Neurosci* 22:135–146.

Reid RC, Alonso JM (1995) Specificity of monosynaptic connections from thalamus to visual cortex. *Nature* 378:281–284.

Reid RC, Alonso JM (1996) The processing and encoding of information in the visual cortex. *Curr Opin Neurobiol* 6:475–480.

Ringach D, Shapley R (2004) Reverse correlation in neurophysiology. *Cognitive Science* 28:147–166.

Ringach DL (2002) Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *J Neurophysiol* 88:455–463.

Ringach DL (2004a) Haphazard wiring of simple receptive fields and orientation columns in visual cortex. *J Neurophysiol* 92:468–476.

Ringach DL (2004b) Mapping receptive fields in primary visual cortex. *J Physiol* 558:717–728.

Ringach DL, Shapley RM, Hawken MJ (2002) Orientation selectivity in macaque V1: diversity and laminar dependence. *J Neurosci* 22:5639–5651.

Rodieck RW (1965) Quantitative analysis of cat retinal ganglion cell response to visual stimuli. *Vision Res* 5:583–601.

Ruthazer ES, Stryker MP (1996) The role of activity in the development of long-range horizontal connections in area 17 of the ferret. *J Neurosci* 16:7253–7269.

Sandell JH, Schiller PH (1982) Effect of cooling area 18 on striate cortex cells in the squirrel monkey. *J Neurophysiol* 48:38–48.

Sengpiel F, Stawinski P, Bonhoeffer T (1999) Influence of experience on orientation maps in cat visual cortex. *Nat Neurosci* 2:727–732.

Sherman SM, Guillery RW (1998) On the actions that one nerve cell can have on another: distinguishing "drivers" from "modulators". *Proc Natl Acad Sci U S A* 95:7121–7126.

Sherman SM, Koch C (1986) The control of retinogeniculate transmission in the mammalian lateral geniculate nucleus. *Exp Brain Res* 63:1–20.

Shoham D, Hübener M, Schulze S, Grinvald A, Bonhoeffer T (1997) Spatio-temporal frequency domains and their relation to cytochrome oxidase staining in cat visual cortex. *Nature* 385:529–533.

Silverman MS, Grosof DH, Valois RLD, Elfar SD (1989) Spatial-frequency organization in primate striate cortex. *Proc Natl Acad Sci U S A* 86:711–715.

Sincich LC, Horton JC (2005) The circuitry of V1 and V2: integration of color, form, and motion. *Annu Rev Neurosci* 28:303–326.

Sirosh J, Miikkulainen R (1996) Self-organization and functional role of lateral connections and multisize receptive fields in the primary visual cortex. *Neural Processing Letters* 3:39–48.

Sirosh J (1995) A Self-Organizing Neural Network Model Of The Primary Visual Cortex. *PhD Dissertation, Technical Report AI-95-237* .

Sirovich L, Uglesich R (2004) The organization of orientation and spatial frequency in primary visual cortex. *Proc Natl Acad Sci U S A* 101:16941–16946.

So YT, Shapley R (1979) Spatial properties of X and Y cells in the lateral geniculate nucleus of the cat and conduction veolcities of their inputs. *Exp Brain Res* 36:533–550.

Sokoloff L, Reivich M, Kennedy C, Rosiers MHD, Patlak CS, Pettigrew KD, Sakurada O, Shinohara M (1977) The [14C]deoxyglucose method for the measurement of local cerebral glucose utilization: theory, procedure, and normal values in the conscious and anesthetized albino rat. *J Neurochem* 28:897–916.

Solomon SG, Lennie P (2007) The machinery of colour vision. *Nat Rev Neurosci* 8:276–286.

Sterling P (1983) Microcircuitry of the cat retina. *Annu Rev Neurosci* 6:149–185.

Sterling P, Demb JB (2004) *The Synaptic Organization of the Brain*, chapter 6: Retina, pp. 217–269 Oxford University Press.

Sur M, Garraghty PE, Roe AW (1988) Experimentally induced visual projections into auditory thalamus and cortex. *Science* 242:1437–1441.

Swindale NV (1991) Coverage and the design of striate cortex. *Biol Cybern* 65:415–424.

Swindale NV (1996) The development of topography in the visual cortex: a review of models. *Network* 7:161–247.

Tanaka S, Ribot J, Imamura K, Tani T (2006) Orientation-restricted continuous visual exposure induces marked reorganization of orientation maps in early life. *Neuroimage* 30:462–477.

Tolhurst DJ, Thompson ID (1982) Organization of neurones preferring similar spatial frequencies in cat striate cortex. *Exp Brain Res* 48:217–227.

Tootell RB, Silverman MS, Hamilton SL, Switkes E, Valois RLD (1988) Functional anatomy of macaque striate cortex. V. Spatial frequency. *J Neurosci* 8:1610–1624.

Tootell RB, Silverman MS, Valois RLD (1981) Spatial frequency columns in primary visual cortex. *Science* 214:813–815.

Triesch J (2005) A gradient rule for the plasticity of a neuron's intrinsic excitability In *Artificial Neural Networks: Biological Inspirations - ICANN*, pp. 65–70.

Tsien JZ (2000) Linking Hebb's coincidence-detection to memory formation. *Curr Opin Neurobiol* 10:266–273.

Turrigiano GG (1999) Homeostatic plasticity in neuronal networks: the more things change, the more they stay the same. *Trends Neurosci* 22:221–227.

Turrigiano GG, Nelson SB (2004) Homeostatic plasticity in the developing nervous system. *Nat Rev Neurosci* 5:97–107.

Usrey WM, Reppas JB, Reid RC (1999) Specificity and strength of retinogeniculate connections. *J Neurophysiol* 82:3527–3540.

Valois RLD, Albrecht DG, Thorell LG (1982) Spatial frequency selectivity of cells in macaque visual cortex. *Vision Res* 22:545–559.

Van Hooser SD, Heimel JAF, Nelson SB (2003) Receptive field properties and laminar organization of lateral geniculate nucleus in the gray squirrel (*Sciurus carolinensis*). *J Neurophysiol* 90:3398–3418.

Vanduffel W, Tootell RBH, Schoups AA, Orban GA (2002) The organization of orientation selectivity throughout macaque visual cortex. *Cereb Cortex* 12:647–662.

von der Malsburg C (1973) Self-organization of orientation sensitive cells in the striate cortex. *Kybernetik* 14:85–100.

Weber C (2001) Self-organization of orientation maps, lateral connections, and dynamic receptive fields in the primary visual cortex In *Lecture Notes in Computer Science. Artificial Neural Networks - ICANN*, pp. 1147–1152.

Weliky M, Kandler K, Fitzpatrick D, Katz LC (1995) Patterns of excitation and inhibition evoked by horizontal connections in visual cortex share a common relationship to orientation columns. *Neuron* 15:541–552.

Weng C, Yeh CI, Stoelzel CR, Alonso JM (2005) Receptive field size and response latency are correlated within the cat visual thalamus. *J Neurophysiol* 93:3537–3547.

White LE, Coppola DM, Fitzpatrick D (2001) The contribution of sensory experience to the maturation of orientation selectivity in ferret visual cortex. *Nature* 411:1049–1052.

White LE, Fitzpatrick D (2007) Vision and cortical map development. *Neuron* 56:327–338.

Wiesel TN, Hubel DH (1965) Comparison of the effects of unilateral and bilateral eye closure on cortical unit responses in kittens. *J Neurophysiol* 28:1029–1040.

Willmore B, Smyth D (2003) Methods for first-order kernel estimation: simple-cell receptive fields from responses to natural scenes. *Network* 14:553–577.

Xu X, Ichida JM, Allison JD, Boyd JD, Bonds AB, Casagrande VA (2001) A comparison of koniocellular, magnocellular and parvocellular receptive field properties in the lateral geniculate nucleus of the owl monkey (*Aotus trivirgatus*). *J Physiol* 531:203–218.

Xu X, Anderson TJ, Casagrande VA (2007) How do functional maps in primary visual cortex vary with eccentricity? *J Comp Neurol* 501:741–755.

Yabuta NH, Callaway EM (1998) Functional streams and local connections of layer 4C neurons in primary visual cortex of the macaque monkey. *J Neurosci* 18:9489–9499.

Yamada ES, Marshak DW, Silveira LC, Casagrande VA (1998) Morphology of P and M retinal ganglion cells of the bush baby. *Vision Res* 38:3345–3352.

Yu H, Farley BJ, Jin DZ, Sur M (2005) The coordinated mapping of visual space and response features in visual cortex. *Neuron* 47:267–280.

Zepeda A, Arias C, Sengpiel F (2004) Optical imaging of intrinsic signals: recent developments in the methodology and its applications. *J Neurosci Methods* 136:1–21.

Zhang JX, Rosenberg A, Mallik AK, Husson TR, Issa NP (2007) The representation of complex images in spatial frequency domains of primary visual cortex. *J Neurosci* 27:9310–9318.

Zhu W, Shelley M, Shapley R (2009) A neuronal network model of primary visual cortex explains spatial frequency selectivity. *J Comput Neurosci* 26:271–287.