Natural scene statistics and the development of the primary visual cortex

By

Jonathan J Hunt

Declaration by author

This thesis is composed of my original work, and contains no material previously published or written by another person except where due reference has been made in the text. I have clearly stated the contribution by others to jointly-authored works that I have included in my thesis.

I have clearly stated the contribution of others to my thesis as a whole, including statistical assistance, survey design, data analysis, significant technical procedures, professional editorial advice, and any other original research work used or reported in my thesis. The content of my thesis is the result of work I have carried out since the commencement of my research higher degree candidature and does not include a substantial part of work that has been submitted to qualify for the award of any other degree or diploma in any university or other tertiary institution. I have clearly stated which parts of my thesis, if any, have been submitted to qualify for another award.

I acknowledge that an electronic copy of my thesis must be lodged with the University Library and, subject to the General Award Rules of The University of Queensland, immediately made available for research and study in accordance with the *Copyright Act 1968*.

I acknowledge that copyright of all material contained in my thesis resides with the copyright holder(s) of that material.

Statement of Contributions to Jointly Authored Works Contained in the Thesis

- Hunt, J.J.[1], Giacomantonio, C.E.[1], Tang, H., Mortimer, D., Jaffer, S.,Vorobyov, V., Ericksson, G., Sengpiel, F., Goodhill, G.J. (2009). *Natural scene statistics and the structure of orientation maps in the visual cortex.* Neuroimage. 47:157-172

  JJH performed analysis of the experimental data and wrote the corresponding methods and analysis. CEG performed the elastic net simulations and wrote the corresponding methods and analysis. HT did preliminary computational analysis. SJ, VV, FS

---

[1]Joint first author

performed all experiments. FS wrote the experimental methods section. DM, GE contributed to the development of key concepts in the analysis. DM, GE, FS contributed feedback on a draft version of the manuscript. JJH, CEG, GJG wrote the introduction and discussion and revised the manuscript as a whole.

- Hunt, J. J., Bosking, W., Goodhill, G.J. (2011). *Statistical structure of lateral connections in the primary visual cortex.* Neural Systems & Circuits. 1:3
  As stated in the published manuscript: WHB acquired the experimental data. JJH, GJG designed the analysis. JJH performed the analysis. JJH, GJG, WHB wrote the paper. All authors read and approved the final manuscript.

Statement of Contributions by Others to the Thesis as a Whole

None

Statement of Parts of the Thesis Submitted to Qualify for the Award of Another Degree

None

Published Works by the Author Incorporated into the Thesis

- Hunt, J.J.[2], Giacomantonio, C.E.[2], Tang, H., Mortimer, D., Jaffer, S.,Vorobyov, V., Ericksson, G., Sengpiel, F., Goodhill, G.J. (2009). *Natural scene statistics and the structure of orientation maps in the visual cortex.* Neuroimage. 47:157-172 - Chapter 5 is based on my contribution to this publication.

- Hunt, J. J., Bosking, W., Goodhill, G.J. (2011). *Statistical structure of lateral connections in the primary visual cortex.* Neural Systems & Circuits. 1:3 - Forms the basis of chapter 6.

---

[2]Joint first author

Additional Published Works by the Author Relevant to the Thesis but not Forming Part of it

None

Irreverence is the champion of liberty
and its only sure defense.

Mark Twain

Acknowledgements

It is difficult to capture in writing the appropriate expression of acknowledgement for all those who played a role in the creation of this thesis. One would have to be a far better writer than I to adequately acknowledge the parade of conversations, laughs, smiles, encouragements, meals, meetings and beverages that are both the seemingly insignificant events of everyday life and yet the only reason a thesis is survivable. Perhaps it is fortunate that few will read this. At least I can try and capture some of the people behind the parade, although it necessarily will be an incomplete list.

Firstly, I would like to acknowledge the support of my primary advisor Prof. Geoffrey Goodhill who provided extensive support and ridiculously prompt feedback on any work. In addition, I had the good fortune to be co-advised by Prof. Jason Mattingley, particularly on chapter 7 where he provided helpful suggestions in dealing with canny humans. I collaborated with Prof. Peter Dayan on chapter 3, including a month-long visit to his lab. Peter possesses a seemingly omniscient understanding of unsupervised learning models (not to mention the entire field of computational neuroscience). Both my own understanding and chapter 3 are dramatically improved due to his input.

Frank Sengpiel provided the cat maps analysed in chapter 5 and Bill Bosking provided the lateral connectivity data used in chapter 6. In addition, both, along with Len White, provided helpful insight regarding analysis of the data.

I am grateful to all the members of the Goodhill lab for being supportive colleagues and friends. Clare Giacomantonio and Lilach Avitan graciously provided helpful feedback on this thesis at short notice. I should also mention Clare again, for the enjoyable collaboration

during the beginning of my thesis, and Duncan Mortimer and Zac Pujic who were present during the majority of my tenure and proved excellent. In addition, there are many others in the wider Queensland Brain Institute who made it a good place to work.

In the personal arena I must start by thanking my now wife, Sarah Hunt, who has been a tremendous support throughout my thesis. It would not be an exaggeration to doubt whether this thesis would have been completed without her support and encouragement. She also proofed much of my writing and tried valiantly to teach me the fine art of positioning commas correctly. Unfortunately, all remaining errors are still my own. I must also thank my parents who taught me to think for myself from a young age and now must suffer the consequences. Gratitude is also due to my youngest siblings, Sarah and Samuel Hunt, for regularly restoring an appropriate perspective on the woes of grown-up life in that way that only children are insightful enough to be able to do.

I'd also like to thank the overlapping category of the friends and flatmates I've been associated with in Brisbane for their camaraderie throughout the years.

My thanks to the organisers of the Australian Advanced Conference on Neuroscience, particularly John Bekkers, who organised for me to be an administrative assistant during their conference. This provided a delightful location and inspiring people to be around during the later stages of writing my thesis.

Everyone will be famous for 15
minutes.

Andy Warhol

Abstract

Vision is the dominant human sensory modality. Due to the relative ease with which both visual input and visual brain areas can be studied and manipulated, vision has become an important window for enlarging our understanding of the biological sensory processing. Whether artificial or biological, visual processing systems must quickly and efficiently make sense of a large volume of noisy, high-dimensional input. To do this they construct statistical models of the input and utilise these models to efficiently encode visual scenes, detect features and construct a model of the world. In this thesis, we combine the study of natural scene statistics with mathematical models, experimental analysis and visual psychophysics to glean a deeper understanding of the development and function of the mammalian primary visual cortex.

We start by considering functional models of receptive field development. We find, in agreement with previous work, that unsupervised learning models trained on natural scenes consistently learn that oriented "edges" (Gabor-like filters) are the basic features of natural scenes. The similarity between these filters and primary visual cortex receptive fields is strong evidence that primary visual cortex receptive fields are optimal encoders of visual input. We then significantly extend this work by comparing the prediction of unsupervised learning models with the receptive fields of animals reared in unusual visual environments. We find good agreement, which is evidence that aspects of receptive fields are learned during development, rather than innate. We also show that applying such unsupervised learning models to binocular visual input is not a simple extension of monocular visual input. Interocular correlations change the optimal encoding strategy of binocular input so that it depends

on edge orientation. Such functional models intriguingly predict an over-representation of vertically oriented receptive fields.

After establishing that oriented edges are the basic feature of natural scenes and the unit of primary visual cortex receptive fields, we consider the statistics of edge arrangements in natural scenes. Sigman et al. (2001) showed that edges in natural scenes over short distances tend to be tangent to a common circle, or co-circular. Edge arrangements which contain a dependence between edge position and orientation may be said to have "reduced symmetry" as they lack a symmetry in that the edge position and orientation cannot be rotated independently without modifying the statistics of the arrangement. Co-circularity is one specific type of reduced symmetry. We extend previous work on natural scene co-circularity using a noise-resistant measure of co-circularity we develop and show that natural scenes contain significant co-circularity over extremely large angular distances ($> 14°$). We also discuss preliminary work into variations in co-circularity statistics by scene type.

After establishing that co-circularity is found pervasively in natural scenes, even over large distances, we then return to the structure of the primary visual cortex, but this time at the network level. Previous work has shown that, like edges in natural scenes, V1 orientation preferences maps also have reduced symmetry. However, the details of this dependence between orientation and position have not been examined in detail. We examine cat orientation preference maps from normal, stripe and blind-reared animals and find that, although orientation preference maps do contain reduced symmetry, it is not co-circularity. Moreover, the statistics of reduced symmetry in the maps are not affected by changes to visual input during development.

Continuing our examination of V1 network structure, we consider the statistics of lateral connectivity in tree shrew V1. Previous work demonstrated that long-range V1 lateral connections are more common between regions with similar orientation preferences (Bosking et al., 1997). We re-examine this connectivity data using our noise-resistance measure of co-circularity. We find evidence that lateral connections between cells in the primary visual cortex may use two opposite wiring strategies which simultaneously facilitate quick processing of co-circular visual input while increasing the salience of the less expected deviations from co-circularity.

Finally, we use the psychophysics of binocular rivalry to test whether co-circularity statistics can affect the functional processing of visual input in humans. We show, using binocular rivalry dominance as an objective measure of salience, that randomly arranged edges are more salient than edge arrangements which contain co-circularity. This is evidence that early visual processing may be functionally utilising edge arrangement statistics. In concurrence with our findings about lateral connections, this may indicate a general strategy of increasing the salience of unexpected visual input.

Overall, we demonstrate that early visual coding uses natural scene statistics extensively. We show that oriented edges are a key currency in early visual processing. We find that the arrangement of edges in natural scenes contain rich statistical structure which influences wiring in the primary visual cortex during development and produces measurable changes in the salience of visual stimuli.

Keywords

Primary Visual Cortex, Receptive field development, Lateral connections, Neural Coding, Natural scene statistics, Co-circularity

Australian and New Zealand Standard Research Classification (ANZSRC)

010202	Biological Mathematics	50%
080106	Image Processing	25%
110906	Neuroscence Sensory Systems	25%

Contents

List of Figures

List of Tables

Sanity is not statistical.

George Orwell

1

Introduction

Life is an inherently messy affair. Everything in life, from the simplest task of spooning food into our mouths to the slightly more involved process of wooing a mate, requires us to interpret a large volume of noisy, ambiguous and sometimes conflicting input in a sensible way. Yet, at least for the simpler tasks, our brains deal so well with messy input that we are often not even consciously aware that this is a difficult problem. However, as anyone who has ever filed a tax return or vainly cajoled an automated answering system can testify, designing machines and processes which deal adequately with the ambiguities of everyday life is far from a solved problem. It is philosophically intriguing that abstracted, logical reasoning, once thought the unique domain of humans, has proven far easier to replicate mechanically than our ability, which we share with other animals, to correctly interpret ambiguous, messy input (Christian, 2011). Not only do our brains deal with mess well, even more remarkably, they also do so efficiently using less power than most modern laptops (Kandel et al., 2000;

Mahesri and Vardhan, 2005).

Fortunately for the reader, this thesis does not attempt to uncover the secret to the general problem of dealing with a messy world. Rather more modestly, here, building on previous work, we will show that our understanding of the development and function of early mammalian visual processing can be improved by examining the statistics of the natural visual world. One part of the secret to efficiently interpreting noisy, high-dimensional input appears to be the re-encoding of that input, utilising statistical models of the world to remove redundancy and noise, revealing underlying features. We will demonstrate that the mammalian early visual system has evolved specialised mechanisms to process natural visual input efficiently (we will define "efficiency" more precisely later). Further, we hope to show that elucidating biological sensory processing systems by constructing statistical models of their input is a widely applicable approach that holds promise for the understanding of sensory processing more generally.

Even at low-resolution, visual scenes are high-dimensional, particularly when a temporal dimension is included. As introduced above, a key idea that both biological vision systems and, more recently, artificial vision systems use is the exploitation of the regularity in the visual natural world in order to process visual input rapidly and efficiently. It is straight-forward to demonstrate that natural scenes contain significant statistical structure. Figure 1.1B shows a greyscale image with each pixel color chosen independently from a uniform random distribution. In principle, all natural scenes can be found in this distribution. Of course, sampling a naturalistic image from this distribution is overwhelmingly unlikely. The set of all natural scenes clearly share some characteristics that distinguish them from the wider set of all possible images. From early work on encoding TV signals (Kretzmer, 1952), people realised that pixels in natural scenes were highly correlated: neighbouring pixels are likely to be a similar colour to one another. This realisation is the basis for modern image compression. However, figure 1.1C shows a scene sampled randomly from the distribution of scenes with spatial correlation functions like that found in natural scenes. It is clear that, although this image shares more similarity with the world we encounter than the earlier pixel-independent image, both are immediately distinguishable from a real natural scene (such as in figure 1.1A) by even the most naive viewer. This demonstrates that there are

A   B   C

FIGURE 1.1: **What defines the set of natural scenes?**
A Natural scenes all share statistical properties that make them immediately visually distinct from random images. **B** One defining characteristic of natural scenes is that neighbouring pixel intensities are not independent. But, although they look more naturalistic, images which have the same spatial correlation function as natural scenes **C** are still clearly distinguishable from natural scenes.

many layers of statistical structure in natural scenes.

Although the regularity in natural scenes is probably utilised at every level of visual processing, in this thesis we restrict ourselves to examining the primary visual cortex (abbreviated as V1). V1 is one of the most well-studied brain areas experimentally and is known to depend on visual input for normal development (Hubel and Wiesel, 1970; Fagiolini et al., 1994; Gordon and Stryker, 1996; Maurer et al., 2011). This makes V1 function well-suited to elucidation by considering the statistics of its input. V1 is the earliest cortical region in the visual pathway possessing numerous afferent connections to higher visual areas. There is significant evidence that one of V1's roles is optimally recoding of visual input (Barlow, 1989; Atick and Redlich, 1992; Olshausen, 2003). An additional practical point: modern digital cameras make the acquisition of natural scenes relatively straightforward, something that is not true of other sensory modalities.

Marr and Poggio (1976) outlined four levels of understanding of neural processes:

1. Understanding the nature of the computational problem.

2. Understanding the algorithms which solve the computational problem.

3. Understanding the mechanisms with which an algorithm is implemented.

4. Understanding the processes with which these mechanisms are implemented.

All of these levels of understanding are important. In this thesis, we aim to progress our understanding primarily at levels 1-3. We aim to build this understanding by using the statistics of visual input to guide us in postulating what functions, algorithms and mechanisms are needed for early visual processes. Within this approach, our primary focus (the reason for which is demonstrated in chapter 3) is the statistics of edges in natural scenes. The thesis progresses as follows:

- In chapter 2, we provide some background material, providing an outline of what is known about the primarily visual cortex. We examine the biological characterisation of several V1 structures: receptive fields, lateral connections and visual "maps." We outline work characterising natural scenes, including more recent work considering the arrangement of edges in natural scenes.

- In chapter 3, we enlarge on earlier work that showed that filters with remarkably similar properties to the receptive fields of V1 simple cells can be learned from natural scenes by sparse coding models. We model receptive field development with unusual visual input and find that sparse coding models postdict the changes in receptive field development found in unusually reared animals. Further, we show that far from being a straightforward extension to previous models, binocular receptive field development raises new questions about the nature of visual coding in V1. Overall, this chapter cements our belief that "edges" are the low-level feature of natural scenes. This provides the focus of the remaining chapters.

- In chapter 4, we re-examine the statistics of edge arrangements in natural scenes. We consider the correlations in pairs of edges over much larger distances than previously examined. Consistent with previous findings over shorter distances, we find that natural scene edges continue to be "co-circular" over separations a large as $> 14°$. The ubiquity of this co-circularity in natural scenes leads us to examine V1 responses and wiring for evidence that it utilises these correlations.

- In chapter 5, we examine cat V1 orientation maps for evidence of co-circularity driven by visual input. We find that orientation maps have a statistical structure distinct from that of natural scenes. Using maps from animals reared in unusual visual environments we demonstrate that this structure is not driven by visual input, but may be due to spontaneous symmetry breaking.

- In chapter 6, we consider the statistics of lateral connections by exploiting previously acquired connectivity data. We find evidence that lateral connections in tree shrew appear to use two opposite wiring strategies to process co-circular edges efficiently while raising the salience of less probable anti-cocircular edges.

- In chapter 7, we use the psychophysics of binocular rivalry to show that naturalistic edge arrangements are less salient than randomly-oriented edges. This is in agreement with our work on lateral connections and demonstrates that regularity in edge arrangement is utilised by the functioning human brain. It provides evidence that entropy reduction is a key principle in early visual salience.

Finally, we conclude in chapter 8 that biological visual processing utilises the statistics of natural scenes at many different levels. We find that elucidating sensory processing by characterising the statistics of the input is a generalisable approach that holds promise and challenge for understanding other types of sensory processing.

Have no respect whatsoever for
authority.

Richard Feynmann

2

Background

Synopsis

Mammalian visual systems evolved to interpret natural visual input in order to enhance survival (Geisler, 2008). Given the evolutionary pressures underlying their development, we expect that natural vision systems will minimize their resource use by adapting to the specifics of their environment (Simoncelli and Olshausen, 2001; Geisler and Diehl, 2002). This connection between neural systems and natural scene statistics provides us with an avenue to use top-down approaches to model the functionality of biological vision systems. In the next chapter we will consider sparse coding as a model of primary visual cortex receptive fields. We will find that edges are the basic features of natural scenes. Later chapters, therefore, focus on the statistics of natural edge arrangements and their role in visual development. In this chapter, we provide a summary of the existing work relevant to the thesis.

We will first discuss the characterisation of the primary visual cortex. In the later part of the chapter, we will examine the statistics of natural scenes. In particular, we will highlight some findings about the statistics of edges in natural scenes. Finally, we will turn to examine the use of unsupervised learning models to find optimal encoders for natural scenes. We will discuss a body of work which demonstrates that such encoders learn receptive fields with similar properties to those of V1 cells.

2.1   Primary visual cortex

Biological visual systems don't just capture visual input for slow-motion replays, they analyse what they are seeing to learn useful information about the state of the environment around them. Comprehension of visual input is a difficult task because the input is ambiguous, noisy and high-dimensional. One early step is often to transform visual input into more useful representations with less redundancy (Olshausen, 2003). Although such transformations occur even in the retina (Hosoya et al., 2005), the primary visual cortex plays an important role in recoding visual input.

The primary visual cortex (V1) is the earliest brain region with markedly different receptive field properties from the retina (Kandel et al., 2000) and the first area where binocular input converges (Hubel and Wiesel, 1977). V1 receives afferent fibres from the lateral geniculate nucleus and projects efferent fibres to both the dorsal and ventral visual pathways. Located on the occipital lobe (Kandel et al., 2000), V1 is one of the most well-characterised regions of the brain.

V1 is conventionally divided into 6 layers (Purves et al., 2007), although layer 4 is usually further subdivided into 3 sublayers A-C. Each layer in V1 has distinct properties. For instance, the majority of afferent input from the lateral geniculate nucleus connects to V1 layer 4C. However, early work examining the responses of V1 demonstrated that V1 has a columnar structure. Cells in the same column, perpendicular to the surface of V1, tend to have similar responses to visual input (Hubel and Wiesel, 1968; Hübener et al., 1997), while the responses of neighbouring columns vary. This columnar structure means that, for our purposes, we can largely treat V1 as a two-dimensional surface.

Cells in V1 fire in response to particular types of visual input. Columns are said to "prefer" visual input which causes strong firing responses in their component cells. We will first outline the key features of these column preferences, before providing an overview of the experimental methods that elucidated them. Early work demonstrated that V1 columns preferred edges such as bars or Gabor patches. Columns have varying preferences for edge orientation and position (Hubel and Wiesel, 1962). Hubel and Wiesel (1962) divided responses into simple, complex and hypercomplex cells; simple cell responses are phase dependent (the spatial phase of a Gabor patch), while complex cell responses are phase independent. Hypercomplex cells modulate their response when the length of the edge is adjusted. The majority of this thesis we will focus on simple cells. The responses of complex cells can be explained by presuming complex cells are integrating the responses of several simple cells (Movshon et al., 1978). Other preferences that have been observed in V1 columns are: direction preference (Hubel and Wiesel, 1959), spatial frequency (Valois et al., 1982), ocular dominance (Hubel and Wiesel, 1968) and colour modulation (Xiao et al., 2007).

These response properties lead to the "standard model" of the primary visual cortex, which we will adopt for parts of this thesis. Under this model, we view V1 as a feedforward sensory area which receives visual input and encodes it for further visual processing. This approach undoubtably captures part of the function of V1, hence our willingness to assume it in this thesis. However, it is important to remember that this model is, at best, a simplification. This issue is excellently summarised in Olshausen and Field (2004), although it was recognised early on (Hubel and Wiesel, 1968). There is almost as much efferent input from "higher" visual areas into V1 as there is output (Stettler et al., 2002). There is substantial evidence of attention (Roelfsema et al., 1998; Gandhi et al., 1999; Ito and Gilbert, 1999; Poghosyan and Ioannides, 2008) and task modulation (Crist et al., 2001; Li et al., 2004; Chen et al., 2008) of V1 responses. Additionally, V1 is known to play an important role in contour detection and grouping (Li et al., 2006; Roelfsema et al., 1998), which is mediated by the extensive lateral connectivity (see next section). The characterisation of V1 receptive fields obtained using artificial stimuli (as commonly employed) is poor (although not useless) at predicting V1 responses to naturalistic stimuli (Felsen et al., 2005; Shalev and Paradiso, 2005). As we discuss further in chapter 8, moving from a model of V1 as a purely sensory

encoding area, to a predictive coding approach, may address some of these complexities.

The layout of columns in V1 is not random, but rather patterned. Columns with similar preferences tend to lie near one another. Feature maps show the arrangement of V1 preferences for particular features of visual stimuli. For instance, there is a relatively smooth topographic mapping from the visual field to V1 (figure 2.1A). Overlaid on this topographic mapping are various other feature dimensions. The development and arrangements of these maps has been a longstanding subject of interest.

Ocular dominance maps show which eye each region of V1 prefers. Such maps have been available for some time (Hubel and Wiesel, 1972). Figure 2.1B shows an example of an ocular dominance map. In chapter 3, we will study the development of binocular receptive fields, however, in this thesis we will not focus on the overall arrangement of the ocular dominance map.

The feature dimension whose layout we examine in detail (chapter 5) is orientation preference. Figure 2.1C shows an example of an orientation preference map. As can been seen in the figure, orientation preference maps have fascinating patterns. Orientation preferences are arranged in a smoothly varying manner around singularities, known as pinwheels (figure 2.1D). Since the physical size of V1 columns can vary significantly between species, and to a lesser degree between individuals, the spacing of these iso-orientation columns provides a better inter-species metric of length in the cortex than physical distance. More precisely, the map wavelength λ is defined as the spatial frequency (averaged over all directions, see chapter 5 methods) with the greatest power. This normalised unit of length is used in comparisons between maps. The patterning of these feature maps has long been of interest to theorists.

The existence of V1 feature maps raises several questions. The first: what is the function of these maps? While topographic maps have clear computational advantages (Diamond et al., 1999), other properties such as the intriguing structure of the arrangement of orientation preference maps do not appear to have an obvious functional advantage (Horton and Adams, 2005). Indeed, some mammals, notably rats (Girman et al., 1999) and mice (Ohki et al., 2005; Sirotin and Das, 2010), do not have orderly orientation preference maps, presenting instead more stochastic arrangements.

In an exciting recent development that might shed some light on these issues, Kaschube et al. (2010) observed that pinwheel density in orientation preference maps is near π across a range of evolutionarily distant species (when normalized by map wavelength). Further, they provided a plausible explanation of this intriguing universality by demonstrating that all pattern formation processes which met a set of reasonable symmetries, such as equal representation of orientations, would have stable solutions with such π density. This result supports the view that at least some aspects of feature map layout may simply be the result of symmetry constraints rather than directly necessary for functional properties. These results are in agreement with other work demonstrating that reduced symmetry constraints may stabilise pinwheels (Lee et al., 2003; Lee and Kardar, 2006; Wolf, 2005; Schnabel et al., 2007a). In chapter 5, we provide further evidence that some map features may arise as a side-effect of constraints on map development, rather than as input-driven functional optimisation. However, it is still unclear why such maps are not observed in other species, which presumably have similar symmetry constraints.

A related area of inquiry has been the questions of the developmental time-period in which feature maps arise, and the extent to which they are driven by visual input. These questions are complicated by the fact that visual input may have a permissive effect, rather than an instructional one. Additionally, spontaneous retinal waves may play a role prior to eye-opening (Butts and Rokhsar, 2001; Cang et al., 2005b; Albert et al., 2008). Measuring V1 responses prior to eye-opening is experimentally difficult, however, significant evidence points to topographic, ocular dominance and orientation preference maps existing, in at least prototype form, prior to eye opening (Hubel and Wiesel, 1959; Chapman et al., 1996; Horton and Hocking, 1996; Gödecke et al., 1997). However, the feature maps remain at least partially plastic during development (Gordon and Stryker, 1996; Sengpiel et al., 1998, 1999; Karmarkar and Dan, 2006; Sato and Stryker, 2008; Tanaka et al., 2009; Mitchell et al., 2009). Other properties, such as direction selectivity, may develop rapidly in response to visual input (Li et al., 2006, 2008).

The final question about map structure that we will discuss is that of their local layout. Are these feature maps orderly down to the single-cell level or are there local distortions? The question is an experimentally difficult one. Das and Gilbert (1997) demonstrated distortions

in cat topography, while evidence in other species points to a precise topography (Bosking et al., 2002; Yu et al., 2005; McLoughlin et al., 2005), as has some more recent work in cats (Buzas et al., 2003; Hetherington and Swindale, 1999). Nauhaus et al. (2008) used calcium imaging to demonstrate orientation preference maps may be acutely accurate.

2.1.1   Lateral connections

An important anatomical feature that is not mentioned in our feedforward "standard model" is the lateral connectivity in V1. There are a large number of both excitatory and inhibitory connections between neurons in V1 (Gilbert and Wiesel, 1979; Kisvárday et al., 1997), some of which project for distances of over 0.5 mm (Bosking et al., 1997). These connections project between a number of different layers of V1 (Purves et al., 2007). Figure 2.2 shows an example of lateral connections in tree shrew.

The function and development of these connections has been of significant interest since they were first described (Lund et al., 1975; Rockland et al., 1982; Livingstone and Hubel, 1984; Gilbert and Wiesel, 1989; McGuire et al., 1991). They may sharpen orientation tuning of V1 cells (Crook et al., 1998; Ringach et al., 2003, 2002). There is substantial evidence that lateral connections play a key role in contour integration (Malach et al., 1993; Polat and Sagi, 1993; Stettler et al., 2002; Kisvárday et al., 1997; Cass and Spehar, 2005; Cass and Alais, 2006; Field and Hayes, 2004; Li et al., 2008), probably by mediating orientation-specific contextual interactions (Das and Gilbert, 1999; Kinoshita et al., 2009). Lateral connections mature after birth (Ruthazer and Stryker, 1996) and their development coincides, in humans, with marked improvement in contour detection (Kovács et al., 1999; Kovács, 2000).

This connection between lateral connections and contour integration, that is, grouping edge segments together, is the reason they are of interest in this thesis. A number of previous groups have studied the functional specificity of V1 lateral connections (Gilbert and Wiesel, 1979; Ts'o et al., 1986; Gilbert and Wiesel, 1989; Frostig et al., 1990; Schmidt et al., 1997; Bosking et al., 1997). This work has demonstrated that there is functional specificity in the lateral connections, particularly excitatory connectivity. Columns preferentially connect to others with similar orientation preference and along the axis of orientation. In chapter 5, we

A

B

C

D
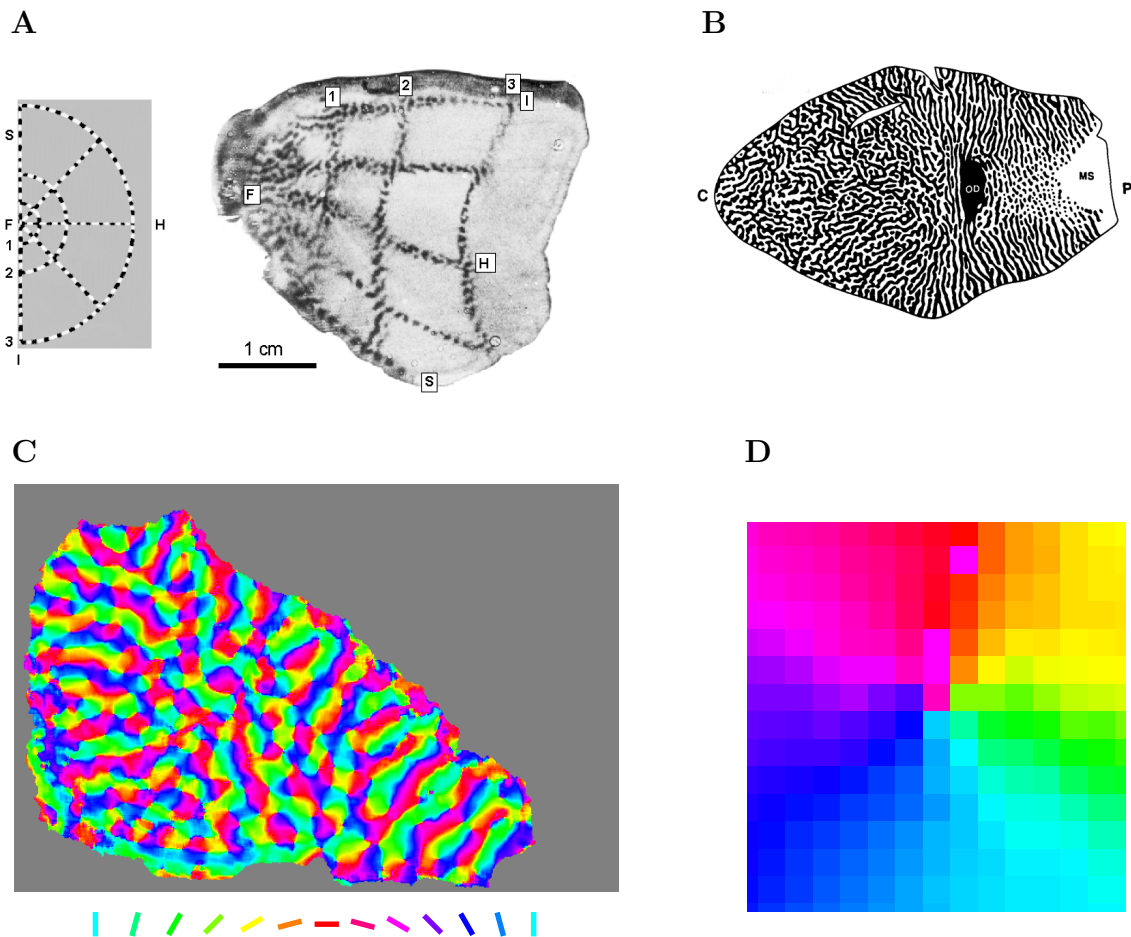
FIGURE 2.1: **Example feature maps**
A A topographic map from a macaque monkey showing the smooth mapping from the visual field (left) onto the primary visual cortex. Image taken from Swindale (1996) which is based on results from Tootell et al. (1988).
B An example ocular dominance map from a macaque monkey. The alternating black and white patterns indicate regions that prefer input from each eye. Figure from Florence and Kaas (1992).
C An orientation preference map from a tree shrew obtained using optical imaging. The colours correspond to orientation preference. Note the clustering of similar orientation responses into iso-orientation patches. Data provided by Bill Bosking from work performed with the Fitzpatrick lab (Bosking et al., 1997).
D An example of a pinwheel, a characteristic feature of orientation preference maps. Pinwheels are singularities in the orientation preference map around which all orientations are represented.

Figure 2.2: **Lateral connections in a tree shrew**
An example of the lateral connectivity in V1. Tracers were injected into the site centred at the white bars (the orientation of the bars shows the underlying orientation preference). The black dots denote the sites to which the injection site connected. Scale bar is 1 mm. Data provided by Bill Bosking from work performed with the Fitzpatrick lab (Bosking et al., 1997).

re-examine this functional specificity, guided by more detailed knowledge of the statistics of edge arrangements in natural scenes.

2.1.2 Experimental tools

The primary visual cortex is one of the most well-studied areas of the brain (Swindale, 1996). A myriad of experimental approaches have been used to elucidate its properties and function.

One advantage the visual cortex has over some other brain areas is that it is relatively straightforward to present defined stimuli to animals. Most work on V1 has used this advantage by recording from V1 while presenting visual stimuli. Animals are usually anesthetized, although some groups have developed awake, behaving setups (Vnek et al., 1999; Martin et al., 2002; Chen et al., 2002; Slovin et al., 2002; Chen et al., 2005).

Early work characterising V1 primarily used electrophysiological recordings (Hubel, 1981) and histology. Using these approaches, the layers of V1, their connectivity and the columnar structure and response preferences were accurately described.

However, when recording from only a few cells at a time it is difficult to elucidate the large-scale functionality of a cortical region. Optical imaging is a technique which provides the ability to image over significant regions of cortex simultaneously (≈ 20 mm, Ratzlaff and Grinvald (1991)). Optical imaging has been the primary method for describing the layout of the feature maps in V1 (Bonhoeffer and Grinvald, 1991; Tsodyks et al., 1999). There are two main approaches to optical imaging (Grinvald et al., 2001).

One approach is imaging with voltage-sensitive dyes (Grinvald and Hildesheim, 2004;

Yang et al., 2007). These dyes are inserted into the neural region of interest prior to imaging. Voltage-sensitive dyes can provide sub-millisecond temporal resolution, however, the dyes may be photo-toxic or impose pharmacological side-affects on the region being imaged. The spatial resolution is generally limited by the resolution of the microscope or the light-scattering properties of the tissue (Grinvald et al., 2001; Polimeni et al., 2005).

The second approach, intrinsic optical imaging, utilizes a similar imaging technique but the signal is provided by local changes in blood oxygenation and vasculature induced by neural activity (Grinvald et al., 1991; Malonek and Grinvald, 1996; Grinvald et al., 2001). This intrinsic signal is noisier and has lower spatial and temporal resolution compared with voltage-sensitive dyes. However, because it does not require the addition of a dye it ensures that neural function is not affected. It also allows for long recording times, or even chronic recordings (Grinvald et al., 2001). Intrinsic optical imaging was used to obtain the orientation preference maps we analysed from cats (chapter 5) and tree-shrew (chapter 6).

Two relatively new techniques are showing promise. Non-invasive fMRI (Vanzetta and Grinvald, 1999) has improved so that discerning the response of individual cortical columns is becoming feasible (Kim et al., 2000; Grinvald et al., 2000; Yacoub et al., 2008). Such a non-invasive technique is exciting, however, it is unlikely that it will match more invasive techniques in either spatial or temporal resolution in the near future (Logothetis, 2008).

The other technique, that has already been demonstrated to reach single-cell resolution, is calcium imaging. By inserting calcium chelating dyes, it is now feasible to image the calcium response of individual neurons (Stosiek et al., 2003). This approach allows high spatial and temporal accuracy, while also allowing one to measure a large field of view (Li et al., 2008). This technique has already been used to study the accuracy of topographic maps (Ohki et al., 2005), the structure of orientation preference maps at pinwheels (Ohki et al., 2006) and the development of direction selectivity in individual cells (Li et al., 2008). Genetic techniques permit targeting calcium dyes to cells with specific expression patterns (Miesenböck, 2009), allowing one to study a subset of cells containing specific genetic markers.

Another important approach for the study of V1 development has been the use of animals reared in unusual visual conditions (see chapter 3 for citations of some of the many examples). By rearing animals in the dark, with monocular deprivation or, even more drastically, with

blockers to remove spontaneous retinal activity (Chalupa, 2007; Elstrott et al., 2008; Feller, 2009), it has been possible to elucidate the interaction between visual input and various aspects of visual development. In chapter 5 we use maps from dark-reared animals to demonstrate that some map statistics are not driven by visual input.

2.2 Natural scene statistics

As discussed in the Introduction, the central approach of this thesis is the use of the statistics of natural scenes to guide elucidation of the development and function of the primary visual cortex. This "naturalistic" approach is not original; it has already yielded significant successes in modelling various aspects of biological vision (Geisler, 2008). Here we will provide a summary of some of the known statistics of natural scenes. The next section demonstrates one successful use of a naturalistic approach: models of receptive field development as unsupervised learning algorithms trained on natural scenes.

First we must answer the question: what is a natural scene? For our purposes, we consider the set of natural scenes to consist of all the scenes our visual system is likely to encounter in the natural world. This is a broad definition; notably, we have made little effort to consider the nuances such as the different visual experiences between species or locations around the globe. There are, undoubtably, differences in the statistics of visual input induced by such changes (Betsch et al., 2004; Webster, 2007), however, the statistics considered here appear to be universal and robust to changes in perspective or location. This allows us to dispense with such nuances, although future refinement of our understanding of natural scenes may cause such considerations to become important.

Some of the first detailed analysis of natural scene statistics was driven by the need to efficiently encode TV signals (Kretzmer, 1952). Most work has concentrated on the characteristic spatial auto-correlations function (Field, 1987; Burton and Moorhead, 1987; Ruderman and Bialek, 1994; Dong and Atick, 1995; van Hateren and Ruderman, 1998; Long et al., 2006). Spatial frequencies in natural scenes have an amplitude of $A(f) \propto 1/f$ (Field, 1987; Dong and Atick, 1995), at least when measured on large enough scales (Langer, 2000).

The focus of the thesis is on the properties of static, greyscale stimuli. This is for simplicity and because the majority of experiments elucidating V1 responses have not characterised their colour responses. However, other groups have also examined other components of natural scene statistics such as temporal correlations (Dong and Atick, 1995), seasonal variations (Webster, 2007) and colour (Burton and Moorhead, 1987). Another important area of future work is active vision. Eyes are not fixed, but rather constantly roaming, undergoing regular saccades which have a significant effect on the statistics of visual input (Rucci et al., 2007; Rucci, 2008).

Better statistical models of natural scenes can yield both an improved understanding of biological visual systems and improved artificial vision systems and image compression algorithms. Guaranteeing that one has found all the correlations in a complex, high-dimensional source such as natural scenes is an intractable problem (Tzou, 1992). It is likely that other structures in natural scenes are still waiting to be discovered.

2.2.1 Edges in natural scenes: co-circularity

Gestalt psychologists have long known that edges are an important unit of the visual world (Han et al., 1999; Elder and Goldberg, 2002). Edges are the building-blocks of contours, the units of receptive fields and a key component of textures. Fast, robust edge detection and grouping is an important step in early visual processing, both in biological and artificial visual systems (Marr and Hildreth, 1980; Fischler and Firschein, 1987; Li, 1998). Detection and grouping of edges can drastically lower the amount of information needed to describe a scene and make object recognition easier (Malik, 2001). Exploiting the statistical structure in the edge arrangements is likely to be necessary to perform this task efficiently.

As introduced above, contradistinct from earlier visual responses, V1 receptive fields are notable for their edge-detecting responses. The topic of edge arrangements in natural scenes and their role in the development and function of the visual cortex is the focus of the majority of this thesis, from chapter 4 onwards. Here we will outline some important background on the arrangement of edges in natural scenes.

Edges in natural scenes are not arranged at random. The relative position and orientation

Figure 2.3: **Co-circular edges**
A pair of edges are co-circular when they are tangent to
a common circle, as in this figure. Equivalently, using
the angles defined on the figure, edges are co-circular
when $\alpha = 2\beta$ or $\alpha = 2\beta \pm 180°$. Sigman et al. (2001)
found that edges in natural scenes had a weak tendency
to be co-circular.

of edge pairs are an important indicator of whether those edges belong to the same contour
or object (Geisler et al., 2001). For instance, edges in natural scenes are often arranged
co-linearly (that is, in lines). Edges that are co-linear are more likely to be grouped together
(Beck et al., 1989) and co-linear edges are easier to detect (Polat and Sagi, 1993).

However, more recently Sigman et al. (2001) demonstrated that there was a more subtle,
general correlation in the position and orientation of edge pairs in natural scenes. Co-linearity
only specifies the arrangement of edges of the same orientation. Sigman et al. found that
a simple geometric extension to this observation explained the correlations they observed
in natural scene edge orientations and positions. Edges are co-circular when they are both
tangents of a common circle (figure 2.3). Co-linearity is just a special case of co-circularity
(allowing an infinite circle radius). Sigman et al. observed that in natural scenes edges
tended to be co-circular. Geisler et al. (2001) further demonstrated that edges belonging to
the same contour in an image are also co-circular.

We can define co-circularity more formally. We assign each edge in a natural scene an
orientation θ and position $\mathbf{x}$ and for every ordered pair of edges $((\theta_i, \mathbf{x}_i), (\theta_j, \mathbf{x}_j))$ we define
the following:

$$\alpha = \theta_i - \theta_j \qquad \text{Orientation difference} \qquad (2.1)$$

$$\phi = \angle(\mathbf{x}_i - \mathbf{x}_j) \qquad \text{Angular separation} \qquad (2.2)$$

$$\beta = \phi - \theta_i \qquad \text{Relative angular separation} \qquad (2.3)$$

$$r = |\mathbf{x}_i - \mathbf{x}_j| \qquad \text{Distance between points} \qquad (2.4)$$

All angle differences are calculated   mod π (π rather than 2π because edges are not directional so have π symmetry). Sigman et al. found peaks in the conditional probability $p(\beta|\alpha, r)$ at $\alpha = 2\beta$ and $\alpha = 2\beta \pm \pi$.

In chapter 4, we expand on this work by examining the statistics of naturalistic edge arrangements over a much wider range of length scales r. We also introduce a measure, the degree of co-circularity, for quantitatively determining the amount of co-circularity present in a scene.

An important terminology issue to be aware of: sometimes "reduced symmetry" is used to refer to the symmetry that edges containing a pairwise correlation between orientation and position lack (as in chapter 5). Co-circularity is a specific type of reduced symmetry. This is important as several groups have studied whether orientation preference maps possess reduced symmetry (Lee et al., 2003; Lee and Kardar, 2006; Schnabel et al., 2007a; Hunt et al., 2009). Chapter 5 examines the co-circularity statistics of orientation preference maps and expands on this issue.

2.2.2 The optimal coding hypothesis and V1 receptive field development

In the introduction we outlined Marr's hierarchy of understanding in neural processes. The top two levels of Marr's hierarchy are both concerned with function rather than implementation (Marr and Poggio, 1976). This top-down understanding is complementary to elucidation of the mechanisms behind neural processes. Here we describe a class of top-down approaches which attempt to model the function of V1 simple cells as optimal encoders.

The basic postulate is that the function of V1 columns is to encode visual input in an optimal manner, removing redundancy that exists in visual input due to the correlations in natural scenes discussed above. Redundancy reduction serves two purposes: it lowers the information needed to describe the input and it leads to representations which are more useful for scene comprehension. Barlow (1961) was one of the first to enunciate this concept of optimal encoding using the language of information theory (Shannon, 1948), although he acknowledged that ideas about sensory coding were not entirely novel. Attneave (1954) and

even writings by Mach and Pearson in the late 1800s contain some of the concepts behind optimal coding.

More recently, advancements in computation and the development of efficient mathematical methods for learning optimal coding models have allowed detailed quantitative comparisons between sparse coding models and neural responses. Here we will concentrate on the use of coding models for mammalian primary visual cortex, however these approaches have also been applied to other sensory modalities and species such as human auditory (Lewicki, 2002; Smith and Lewicki, 2006) and bird auditory coding (Greene et al., 2009).

There are multiple definitions of what an optimal coding algorithm should optimize. Some groups have defined as optimal codes those which maximise the information content of the encoding (Barlow, 1989; Bell and Sejnowski, 1997; Okajima, 2004) (thus reducing redundancy). Others have focused on sparse coding models which find codes with few non-zero entries for each input (Olshausen and Field, 1996, 1997; Hyvärinen and Hoyer, 2001; Teh et al., 2003; Osindero et al., 2006; Greene et al., 2009). In practice, for natural scene statistics such approaches are similar and sometimes mathematically identical (Hyvärinen et al., 2009).

Olshausen and Field (1996) performed one of the first quantitative investigations of sparse coding for the visual cortex. They used a generative approach to model natural images, viewing each image $\mathbf{x}$ as being composed of a set of basis functions $\mathbf{w}_i$ (not necessarily orthogonal).

$$\mathbf{x} = \sum_i a_i \mathbf{w}_i \tag{2.5}$$

Here, as in later in the thesis, we are treating the two-dimensional image as a column vector. The basis functions $\mathbf{w}_i$ are the same for all images, but the mixture coefficients a_i change for each image. This generative approach is a very powerful one, used widely in many areas of statistical modelling. Olshausen and Field then searched for a sparse representation in a fairly straightforward way. They defined a cost function

$$E(\mathbf{x}) = - \left| \mathbf{x} - \sum_i a_i \mathbf{w}_i \right|^2 - \lambda \sum_i S(a_i/\sigma) \tag{2.6}$$

where the first term favours the preservation of information (by minimising the l_2 error), while the second term weights sparse representations (σ is a scaling factor). The relative importance of the two terms is controlled by λ. For $S(z)$ they used a function which favours states with few non-zero coefficients such as $S(z) = \log(1 + z)$. They used a two-step gradient descent method to minimize $< E(\mathbf{x}) >$ with respect to both a_i and $\mathbf{w}_i$, where $<>$ represents the ensemble average over the training images. They demonstrated that the sparse representation of natural scenes that this algorithm learned were Gabor-like structures similar to the responses of V1 simple cell receptive fields.

Other work has used independent component analysis to learn optimal encodings from natural scenes and found these algorithms learned receptive fields with similar properties (Bell and Sejnowski, 1997; van Hateren and van der Schaaf, 1998; Hyvärinen et al., 2001). Groups have also extended coding models with additional feature dimensions such as colour (Hoyer and Hyvärinen, 2000), temporal (van Hateren and Ruderman, 1998) or more accurate spatial layouts (Doi et al., 2003).

The majority of work using optimal coding models of V1 receptive field development has considered only monocular receptive fields and square transformations. In the next chapter we use a more recent unsupervised learning model, the product of experts (Teh et al., 2003; Osindero et al., 2006), to model binocular receptive field development. We demonstrate that inter-ocular correlations introduce some new issues into receptive coding models.

One has been endowed with just
enough intelligence to be able to see
clearly how utterly inadequate that
intelligence is when confronted with
what exists. If such humility could be
conveyed to everybody, the world of
human activities would be more
appealing.

Albert Einstein

3

Functional models of receptive field development

Synopsis

What language does the brain speak? One answer might be to say that it speaks the language of neuronal spikes and neurotransmitters. That answer would be true, but at the same level as describing the English language by listing the 50-odd phonemes in use. Obviously, a more complete description explains how these simple components, of spikes or phonemes, are assembled to describe complex phenomena. Linguists encountering a new language may start by describing the phonemes, but ultimately wish to know how these are arranged to construct words, sentences, paragraphs and ultimately convey the underlying semantic content (Akmajian et al., 2010).

Much as simple, imperative sentences might be the easiest way to start comprehending

a language, early sensory brain areas potentially provide a good place for improving our understanding of neuronal information transmission. In linguistics, plain text examples, where one knows both the input and the translation, have proved key in understanding new languages (Parkinson, 1999). Early visual areas may provide a neuronal plain text. We have a fairly good approximation of visual input and the way such brain areas represent visual input are, at least superficially, fairly well understood.

The top-down approach that we use here is to first ask what is an optimal way of representing visual input and then consider whether V1 is using such a representation. There are several definitions of optimal representations, but here we focus on sparse representations. A neuronal representation is sparse when most input is described by only a few neurons firing. Sparse representations are useful if the underlying causes of the input are sparse, which appears to be the case for natural scenes. Additionally, sparsity may be a metabolically optimal way of representing information.

This approach is not novel, several groups have compared sparse representations of natural scenes with the V1 receptive fields. These groups have demonstrated that the receptive fields learned from natural scenes are remarkably similar to the receptive fields of primary visual cortex simple cells. This is evidence that the function of the simple cells may be to transform visual input into a sparse representation.

Here we extend this work in two directions. Firstly, we examine sparse representations of binocular receptive fields. Most prior work has only considered monocular input. Secondly, we use such binocular receptive field models to consider animals reared with modified visual input, such as stripe-rearing and monocular deprivation. Modelling visual development under such unusual rearing conditions can provide new points of comparison with experiment and allow genetic influences to be teased apart from environmental factors.

We find that sparse representation models postdict the qualitative changes in receptive fields found in unusually reared animals for all six modified rearing conditions we consider. This reinforces the validity of sparsity as a principle of early sensory cortical areas and provides further evidence that V1 receptive fields are, at least in part, learned during development. However, we also find that our model predicts a significant over-representation

of vertically oriented neurons. This is due to asymmetries in the input of inter-ocular correlations at different orientations. We discuss why such vertical representation have not been observed in real animals and the implications for sparsity models of biological sensory processing.

3.1 Introduction

One of the central questions of neuroscience is: how do neurons represent, or encode, information? Code here is meant not in the sense of cryptic[1], but as an agreed way of representing information. It is clear that neurons transmit spike trains, but there are many ways of carrying information in spike trains. For instance, in principle, neurons could use Morse code, developed for transmission on telegraphs, as one way of representing alphanumerical information. It is exceptionally unlikely that Morse code is used, but its development gives some clues as to what codes neurons might use. Telegraph lines were expensive to construct, so Morse devised his code to allow the transmission of telegraphs as quickly as possible. Despite lacking the formalism to express the need clearly, Morse developed a code which optimised transmission rates by assigning short codes to commonly occurring letters (Gleick, 2011). Although absurd to propose that neurons might use Morse code, it is far more reasonable to consider whether evolution has endowed neuronal codes with optimality of some type, much as Morse attempted in his code. The hypothesis that neurons may perform optimal coding was first clearly stated by Barlow (1961).

The mathematical formalism to describe optimal coding (Shannon, 1948) was developed in an effort to efficiently encode voice and video signals (Guizzo, 2003; Gleick, 2011). The optimal coding hypothesis suggests that the primary task of early sensory processing is to recode, or transform, the input into a more useful representation. This hypothesis leads this chapter to focus on the first two levels of Marr's hierarchy (see chapter 2, Marr and Poggio (1976)) which encapsulate function. Rather than asking mechanistic questions, in this chapter we use optimality as a top-down tool to making sense of V1 responses.

What exactly is meant by an optimal code? One choice, exemplified by Morse code, is

---

[1]Though scientists may be forgiven for wondering if nature is sometimes reclusive about her workings.

a code which minimizes redundancy to allow compact representations[2]. Another approach
to optimality, that may be relevant for visual input, is sparse coding. A sparse code is one
in which the input, represented numerically, can usually be described with few non-zero
coefficients. Sparse codes are often highly redundant (Olshausen, 2003). It is possible to
draw a distinction between population sparsity and lifetime sparsity (Willmore and Tolhurst,
2001; Willmore et al., 2011). A population sparse code represents each item of input with
few non-zero coefficients. A lifetime sparse code is one where, over all visual experience,
each neuron is significantly active only a small amount of time. Since these conditions
are equivalent given appropriate normalisation conditions, which are approximately true for
natural scenes (Hyvärinen et al., 2009), we don't make a distinction between them here. A
sparse code is specific to a particular class of input, there is no universal sparse code.

What makes sparse codes interesting? There are multiple lines of argument that point
to sparse codes as a potential choice for V1 encoding (Olshausen and Field, 1997). Sparse
codes may represent a metabolic trade-off between the informational capacity of dense codes
(many neurons firing) and local codes (a small number of localised neurons firing) (Foldiak
and Young, 1995; Foldiak and Endres, 2008). However, the justification more widely used is
based on a casual generative approach to unsupervised learning.

Unsupervised learning is the problem of trying to find hidden structure in unlabelled
data. One approach, density estimation, attempts to find an underlying probability density
function from a set of examples. Such a function can reveal hidden structure in the data. One
widely used approach to density estimation is a causative generative approach. If we assume
that our data $< \mathbf{x} >$ ($<>$ denotes the average of a value over all the training examples) is
created by some underlying causes $\mathbf{s}$, then the probability density function of the data can
be written as:

$$p(\mathbf{x}) = \int p(\mathbf{x}|\mathbf{s}, \theta) p(\mathbf{s}|\theta) \tag{3.1}$$

θ encapsulates the parameters of the model. Here, $p(\mathbf{s}|\theta)$ models the distribution of causes
and $p(\mathbf{x}|\mathbf{s}, \theta)$ the effect these causes have on the observed examples. The parameters of

---

[2]Although by using a ternary symbol set Morse code, when transmitted as a wide-band amplitude modu-
lated signal, contains significant redundancy, which is the reason it was practical for transmission over noisy
telegraph systems.

the model can be fit to training data by finding the set of parameters which maximise the likelihood of the observed examples:

$$\theta^* = \max_{\theta} < p(\mathbf{x}) > \qquad (3.2)$$

Once trained, such a generative model can allow us to obtain the underlying causes of the data.

As Olshausen and Field (1996) argued, we have an intuitive idea that visual input should be composed of sparse causes, for instance, edges of a particular orientation are rarely present in an image patch, but when present can be easily observed. This intuition is behind the assumption of sparse coding. It can tested by comparing how well our generative model captures the statistics of the input when we choose a sparse distribution for the causes $p(\mathbf{s}|\theta)$. If natural scenes are composed from underlying sparse causes, then it is also natural that V1 may encode visual input sparsely. Such an encoding, which separates the underlying causes, is likely to be useful for visual comprehension (Olshausen and Field, 2004). Sparse coding has been a unifying principle for modelling the function of early sensory areas (Barlow, 1961; Olshausen and Field, 2004; Geisler, 2008; Simoncelli and Olshausen, 2001).

How do we test if sparse coding provides a good model of V1 function? There is likely to be multiple sparse codes for natural scenes. However, all sparse codes are likely to share similar properties. Here, in solidarity with previous work (Bell and Sejnowski, 1997; van Hateren and van der Schaaf, 1998; Osindero et al., 2006), we train a sparse coding model on natural scene input. The properties of the learned sparse codes are then compared with V1 receptive fields. Where the encodings are similar to V1 responses this is taken as evidence that the function of these neurons is to represent input sparsely.

As we discussed in chapter 2, sparse coding has already had marked success in modelling the response properties of the primary visual cortex (V1). Early work showed that receptive field properties of simple cells in V1 were remarkably similar to the receptive fields of sparse coding models trained on natural scenes (Olshausen and Field, 1996; Bell and Sejnowski, 1997; van Hateren and van der Schaaf, 1998). Other work has extended sparse coding models of V1 to complex cells (Hyvärinen and Hoyer, 2001), the temporal dimension (van Hateren

and Ruderman, 1998) and colour responses (Caywood et al., 2004; Hoyer and Hyvärinen, 2000). Hyvärinen et al. (2009) provides an excellent introduction to sparse coding models of the visual cortex.

Despite the contribution of sparse coding to our understanding of V1 function, there are still many unanswered questions. Are the biological receptive fields learned during development or intrinsic? What is the function of all the other wiring in V1, such as lateral connections and efferent feedback, which is not needed for sparse coding? Are the sparse codes adaptive, as are many artificial coding schemes (Ghanbari, 2003)?

A series of animal experiments provide a rich source of data for addressing these questions. A number of studies have examined the response of V1 cells in animals reared with modified visual input. For example, animals were reared in the dark, in rooms containing stripes of only one orientation, or with artificial strabismus. Table 3.1 provides a list of the six modified visual rearing conditions considered here, with references to the original experiments. Comparison between sparse coding models and the V1 responses from such animals provides new tests for sparse coding. Additionally, these results can help differentiate neuronal properties which are learned during development from visual input from those which are less modifiable.

However, in order to model such unusual rearing conditions we needed to accurately model binocular receptive field development. The majority of sparse coding work has considered only monocular field development, although the 2nd order properties of binocular visual coding have been examined (Zhaoping, 2006). Modelling binocular receptive field development both provides a more precise comparison with normal animals and allows the modified rearing conditions to be considered. Hoyer and Hyvärinen (2000) have previously examined binocular encoding and reported that learned receptive fields had similar disparity preferences to experiment. Here we performed a detailed comparison between the receptive fields of our model and experimentally recorded receptive fields. We found good agreement in many properties of the receptive fields between model and experiment. However, as we will discuss shortly, binocular sparse coding models highlight some important issues for sparse coding models of visual cortex that have not previously been reported.

We used the product of the experts model (Teh et al., 2003) for learning sparse codes, instead of independent component analysis (Bell and Sejnowski, 1997; van Hateren and van der Schaaf, 1998) which has been the model more commonly employed. The product of experts model allowed us to learn codes with significant overcompleteness, meaning the number of underlying causes (or biologically, receptive fields) is greater than the input dimensions. Such codes can be more sparse but are computationally more difficult to learn (Teh et al., 2003). There is evidence that V1 employs overcomplete coding (Olshausen and Field, 1997; Hsu and Dayan, 2007).

Hsu and Dayan (2007) compared the receptive field changes in stripe-reared animals with a monocular sparse coding model. As far as we are aware, this is the only work applying sparse coding to model receptive fields in unusually reared animals. Hsu and Dayan showed that their model matched the over-representation of the exposed orientation that occurs in stripe-reared animals. This provides evidence that V1 sparse codes are learned during development. In this chapter, by using a binocular sparse coding model, we are able to consider a wider range of unusual rearing conditions.

We compared our sparse coding model with several different types of modified rearing conditions. We found that the changes in receptive fields in real animals were qualitatively similar to those changes predicted by the model in all the modified rearing conditions. Table 3.1 shows the rearing conditions that we examined, along with a summary of the salient changes of V1 receptive fields that were found in the animals. We show here that in each of the cases, the salient changes are qualitatively reproduced by the sparse coding model. This is further evidence that V1 simple cells are learning a sparse code during development.

Additionally, we found that training sparse coding models of binocular input reveals a potential disagreement with experiment. Due to the difference in inter-ocular correlations between horizontal and vertical edges, sparse coding models learn a code with a significant over-representation of vertically-oriented edges. Such an asymmetry has not been observed in real animals. We examined this issue in detail by constructing a simplified, low-dimensional model and found that it is a robust consequence of the assumptions of sparse coding. We discuss why such an asymmetry does not occur in real animals and its implication for sparse coding as model of visual cortex later.

Of importance for the remainder of this thesis, we found, in agreement with previous results (Olshausen and Field, 1996; Bell and Sejnowski, 1997; van Hateren and van der Schaaf, 1998), that receptive fields are mostly edge-detectors. This leads us to the topic that will occupy us for the remainder of the thesis: the statistics of edges in natural scenes and the effect such statistics have on V1 structure.

Rearing type	Description of visual input	Salient changes in receptive fields	References
Normal	Normal visual input.	n/a	
Stripe	Animals were goggle-reared (or in a cylinder) exposed to a single orientation in both eyes.	Over-representation of the reared orientation. A reduced number of neurons are strongly orientation selective. Increased binocularity.	Stryker et al. (1978); Sengpiel et al. (1999); Tanaka et al. (2006, 2009)
Orthogonal	Animals were goggle-reared and exposed to horizontal orientations in one eye and vertical orientations in the other eye.	Increased monocularlity. Reduced number of neurons with well-defined orientation preferences. Over-representation of the reared orientation in each eye.	Hirsch and Spinelli (1970); Stryker et al. (1978); Tani and Tanaka (2008)
Monocular	Animals were reared with one eye occluded.	Majority of neurons are responsive to the non-occluded eye. A small minority are responsive to the occluded eye, almost all neurons are extremely monocular.	Hubel and Wiesel (1970); Wiesel and Hubel (1965a); Shatz and Stryker (1978); Stryker (1978)
Alternate-blind	Animals were reared with only one eye open at any time, but the occluded eye was regularly alternated.	Strongly monocular receptive fields but with equal representation of both eyes and all orientations.	Hubel and Wiesel (1965); Shatz et al. (1977)
Partial-monocular rearing	Animals were reared with one eye occluded but were given a small amount of binocular experience (in one experiment 1/7 of their visual experience).	Recovery of near-equal representation of both eyes, but with few binocular responses.	Wiesel and Hubel (1965b); Kind et al. (2002); Schwarzkopf et al. (2007); Vorobyov et al. (2007); Mitchell et al. (2009)
Strabismic	Nonparallel visual axis (achieved artificially by severing extraocular muscles).	Normal orientation coverage but with few binocular responses	Hubel and Wiesel (1965)

TABLE 3.1: **Rearing conditions**
In this chapter we modelled receptive field development of normal and six modified rearing conditions. The modified rearing conditions were modelled by modifying the training images to be consistent with the modifications to the visual experience of the animals. This table provides a list of the rearing conditions, a summary of the salient changes in the receptive field properties found experimentally and references to the experimental papers. For clarity and consistency, we have restricted our citations to results from cat experiments, similar results have been shown in macaque for many conditions.

3.2   Methods

Modelling receptive field development using an optimal coding model consists of two, separable steps. First, acquiring a training stimulus set and second, learning a sparse coding from the training set.

Creating stereo visual input

Most coding models of V1 receptive fields have considered only monocular receptive field development trained under normal rearing conditions. Our approach required a binocular training set and modifying the training stimuli to simulate modified rearing conditions.

A binocular training set of naturalistic image patches was developed using a similar method to Hoyer and Hyvärinen (2000). Eighteen stereo images were acquired from a commercial library of high-quality binocular images of natural scenes[3]. Each image was taken using a binocular camera with lenses spaced approximately a human-eye distance apart at varying focal distances. For each image, 5 focus points were chosen randomly, the two stereo images were aligned at the focal point and then stereo image patches were chosen randomly in a 300x300 pixel square around the focal point. This has the effect of approximating the view of an observer focussing at 5 different points in the scene. Each patch was 2x25x25 pixels. These patches naturally contain varying degrees of stereoposis. All images were first converted from color to greyscale images. A total of 100,000 training patches were created for each rearing condition.

To model modified rearing conditions the training input was modified to match the visual experience of the animals. For stripe and orthogonal rearing, we used a similar approach to Hsu and Dayan (2007). The off-axis spatial frequencies were attenuated by using an oriented Gaussian filter. For monocular rearing and alternate blind rearing, the occluded eye's images were convolved with a square kernel with a length of 150 pixels, which removed all but extremely low spatial frequencies. To simulate the strabismus, patches were chosen with an additional $\approx 10°$ of binocular disparity interjected by offsetting the binocular images (cats with artificial strabismus have disparity $\approx 20°$ but we were limited by the size of the

---

images). Figure 3.1 shows examples of the training images for all rearing conditions. In keeping with previous work, for the stripe and orthogonal reared conditions 10% of the training set was not filtered. Hsu and Dayan (2007) found that retaining 10% normal input gave a better match with experiment because it reduced the "collapse" of receptive field structure that occurred in the absence of any normal input.

Learning algorithm

We used the product of experts model (Teh et al., 2003; Osindero et al., 2006) to learn receptive fields from our training input. This algorithm has the advantage over independent component analysis that it can learn overcomplete representations of input data (Teh et al., 2003).

The product of experts models the input distribution as a product of Student-t distributions. Each training image is represented as a column vector $\mathbf{x}$. $< . >_{\text{data}}$ represents the ensemble average over all the training examples. The probability of input $\mathbf{x}$ is modelled (with n neurons) as:

$$p(\mathbf{x}) \;\; = \;\; \frac{1}{Z(\theta)} \prod_{i=1}^{n} \left[1 + y_i^2/2 \right]^{-\alpha_i} \tag{3.3}$$

$$y_i \;\; = \;\; \mathbf{w}_i^T \mathbf{x} \tag{3.4}$$

In this model, each neuron i has a receptive field $\mathbf{w}_i$ and a sparseness parameter α_i. Equation 3.4 shows that the response of each neuron is the inner product between the neuron's receptive field $\mathbf{w}_i$ and the input $\mathbf{x}$. The normalisation constant Z is dependent on all the parameters of the model α_i and $\mathbf{w}_i$. The model is trained by maximising the log-likelihood of $< \log p(\mathbf{x}) >_{\text{data}}$ with respect to the model parameters α_i and $\mathbf{w}_i$. When the number of neurons is equal to the dimensionality of the input there is a closed-form solution for $\log Z$ and the model performs independent component analysis (Teh et al., 2003).

However, when overcompleteness is introduced, meaning there are more neurons than dimensions in the input training, the model becomes much more difficult to train. In general, there is no analytical solution for the normalisation constant $\log Z$. With gradient ascent,

each training step the parameters of the model are modified along the gradient of θ to increase the probability of the training examples and hopefully converge on a local maximum.

$$\Delta\theta = \epsilon \left[< \frac{\partial \log p(\mathbf{x})}{\partial \theta} >_{\text{data}} - < \frac{\partial \log p(\mathbf{x})}{\partial \theta} >_{\text{model}} \right] \tag{3.5}$$

Here θ encapsulates all the parameters of the model, $<>_{\text{model}}$ is the ensemble average over the distribution of the model and ϵ is the learning rate. The first term increases the probability of training examples and the second term decreases the probability of the model samples. When the model distributions matches the training examples $\Delta\theta = 0$ and the gradient ascent converges.

However, since the model distribution is not normalised, sampling accurately from the model is challenging. Markov chain Monte Carlo (MCMC) can be used for this problem (Andrieu et al., 2003). Each training example can be used as a starting point for a Markov chain with transition probabilities specified by the model distribution. If this chain is followed for long enough then the final state is independent of the starting point and thus provides a good sample of the model distribution. However, following the chain for long enough to ensure independence is computationally expensive. This can be avoided using contrastive divergence (Hinton, 2002). In contrastive divergence the model distribution is approximated by following only a few steps along the Markov chain of each training example. Although this does not provide a good approximation to the model distribution at each step, Hinton (2002) showed that it approximates gradient ascent with only a small error term while providing a huge reduction in computational complexity. The reason is, briefly, that in equation 3.5 only the difference between the model distribution and the training distribution is important. If the parameters of the model are correct then there is no pressure for the MCMC to move away from the data distribution. Contrastive divergence exposes any difference after only a few steps of MCMC sampling.

We used contrastive divergence with a learning rate of 10^{-3} for all the work described in this report. As is common, training data was whitened and dimension reduced using principal component analysis before the product of experts model was applied. We retained the first 150 principal components. We used an overcompleteness factor of 2 as used by Hsu

and Dayan (2007).

3.2.1   Characterisation of receptive field properties

Once the receptive fields were learned we needed to characterise them for comparison with experiment. We primarily concentrated on two receptive field properties: binocularity (or ocular dominance), orientation preference and disparity. This is because these were the dimensions along which significant changes were measured in the animal experiments.

Binocularity is a measure of the preference a neuron has for responding to input in one eye or the other. To measure binocularity we used the definition described in Hoyer and Hyvärinen (2000):

$$b = \frac{||\mathbf{w}_{\text{left}}|| - ||\mathbf{w}_{\text{right}}||}{||\mathbf{w}_{\text{left}}|| + ||\mathbf{w}_{\text{right}}||} \tag{3.6}$$

For comparison with experiment, we binned the values of b into 7 bins with boundaries at $[-0.85, -0.5, -0.15, 0.15, 0.5, 0.85]$ (as in Shouval et al., 1996). Bins 1 or 7 correspond to highly monocular responses while values in the middle correspond to binocular responses.

We calculated the orientation preference for each receptive field as in Hyvärinen et al. (2009, ch. 6). Each eye was treated separately. We compared the response of the filter to sinusoidal gratings over a range of spatial frequencies and orientations. We used quadrature-phase gratings to find the phase invariant response. We examined spatial frequencies between 0 and 0.16 cycles/pixels with spacing of 0.5/pixels and all orientations with a spacing of 1°. In addition, we calculated the circular variance (Ringach et al., 2002) of each receptive field, which is a measure of the orientation selectivity. Circular variance is defined as $V = 1 - |R|$ where:

$$R = \frac{\sum_k r_k e^{i2\theta_k}}{\sum_k r_k} \tag{3.7}$$

r_k is the response of the neuron to stimuli of orientation θ_k. We used the responses calculated from the sinusoidal gratings to find the circular variance of each neuron (treating each eye separately). When plotting distributions of orientation preference we only included receptive fields which had a circular variance > 0.6, as receptive fields with low circular variance cannot be reliably assigned an orientation.

Disparity distribution

Our model predicted a significant over-representation of vertical orientations. We investigated whether there was significant differences in the disparity distribution across different orientations. This required characterising the disparity responses of the receptive fields. We used an approach similar to Hoyer and Hyvärinen (2000). The receptive field response was calculated over a range of disparities. This required a choice of stimuli. We used each eye's receptive field as a known optimal stimuli for presentation. The disparity tuning curve was taken as the mean response to the left-eye receptive field being presented to both eyes at different disparities and the right-eye receptive field presented to both eyes at different disparities. A baseline response was calculated as the response when presented with monocular stimuli.

Once the disparity tuning curve was calculated, neurons were classified into tuned inhibitory and tuned excitatory based on the sign of the maximum deviation from the baseline response. The preferred disparity of a neuron was taken as the position where the maximum deviation occurred. Cells with monocular receptive fields, poor spatial localisation (very low preferred spatial frequency), lacking a well-defined orientation preference in both eyes (circular variance < 0.6) or flat disparity response curves were excluded.

Statistics

In order to test whether the changes observed in the receptive fields under different model conditions were significant, each simulation was repeated 25 times with a different random seed. The seed affects both the model fit and the sampling of the training images. Error bars were calculated as the standard error of the mean. All statistical tests were performed with the two-sided Kolmogorov-Smirnov test (Marsaglia et al., 2003). This test avoids making assumptions about the distributions being tested and tests for differences in both distribution mean and shape.

3.3   Results

We examined 7 different rearing conditions (normal case and six modified rearing conditions, see table 3.1) using a sparse coding model. We show here that in each rearing condition the salient features of receptive field changes found experimentally are reproduced by our model.

Figure 3.1 shows an example of the training input for each rearing condition, along with a sample of the resulting receptive fields. As expected from previous work, the receptive fields in the normal case are localised edge detectors. Some receptive fields are Gabor like, while others are less localised. The receptive fields in each of the modified rearing conditions maintain this basic structure, although marked differences in the binocularity of receptive fields in different rearing conditions can be seen.

The binocularity in each rearing condition was quantified for a more detailed comparison. Figure 3.2 shows the result. In the normal rearing condition there is a range of binocular responses, although a deficit of strongly monocular responses. As in experiment, monocular rearing led to strongly monocular responses for the open eye and little response for the closed eye. Stripe-rearing, which increased inter-ocular correlations, resulted in increased binocular responses. Rearing conditions which decreased inter-ocular correlation, alternate blind-rearing and strabismus, had the opposite effect, as has been observed experimentally.

We also examined the binocularity of the receptive fields in the partial blind rearing case in more detail. Experimental studies have shown that only a small amount of binocular experience interspersed with monocular deprivation is needed to maintain normal or near-normal visual acuity (Kind et al., 2002; Mitchell et al., 2009; Vorobyov et al., 2007). Figure 3.5 shows the binocularity for partial blind rearing with varying degrees of binocular visual experience. As seen in experiment, only a small amount of binocular experience (10% in our model, 1/7 in experiment) is needed to facilitate good recovery of responses for both eyes. In agreement with experiment, the recovery leads to more monocular responses, but nonetheless near-normal representation of both eyes.

Next we quantified the distribution of orientation preferences. Not all receptive fields were orientation selective. Figure 3.3 shows the orientation selectivity of the neurons in each rearing condition. Modified rearing conditions resulted in less orientation-selective neurons

than in the normal case. This is in agreement with experiment for the monocular and partially monocular case. However, such a decrease in orientation-selectivity has not been reported in the stripe or orthogonal rearing cases. This may indicate additional constraints or normative effects in the animals that stop receptive field structure from collapsing in the presence of such simplified visual input.

The salient change in receptive field properties in the orthogonal and stripe-reared cases that has been reported in experiment is the over-representation of the training orientation. Figure 3.4 shows that this over-representation occurs in our model. Additionally, in the orthogonal rearing case there is a strong correlation between eye preference and orientation preference, which has been observed in experiment.

The orientation preference distributions also reveal an important mismatch with experiment that we deal with in the next section, significant anisotropies in the distribution of orientation preferences. In most rearing conditions there is a small over-representation of horizontal orientations. We discuss this issue further later, however, it may reflect an over-representation of cardinal axis in real animals (Coppola et al., 1998b). A much more dominant effect is the over-representation of vertical edges.

A normal **B stripe** **C orthogonal** **D monocular**

E alternate-blind **F partial-blind** **G strabismic**
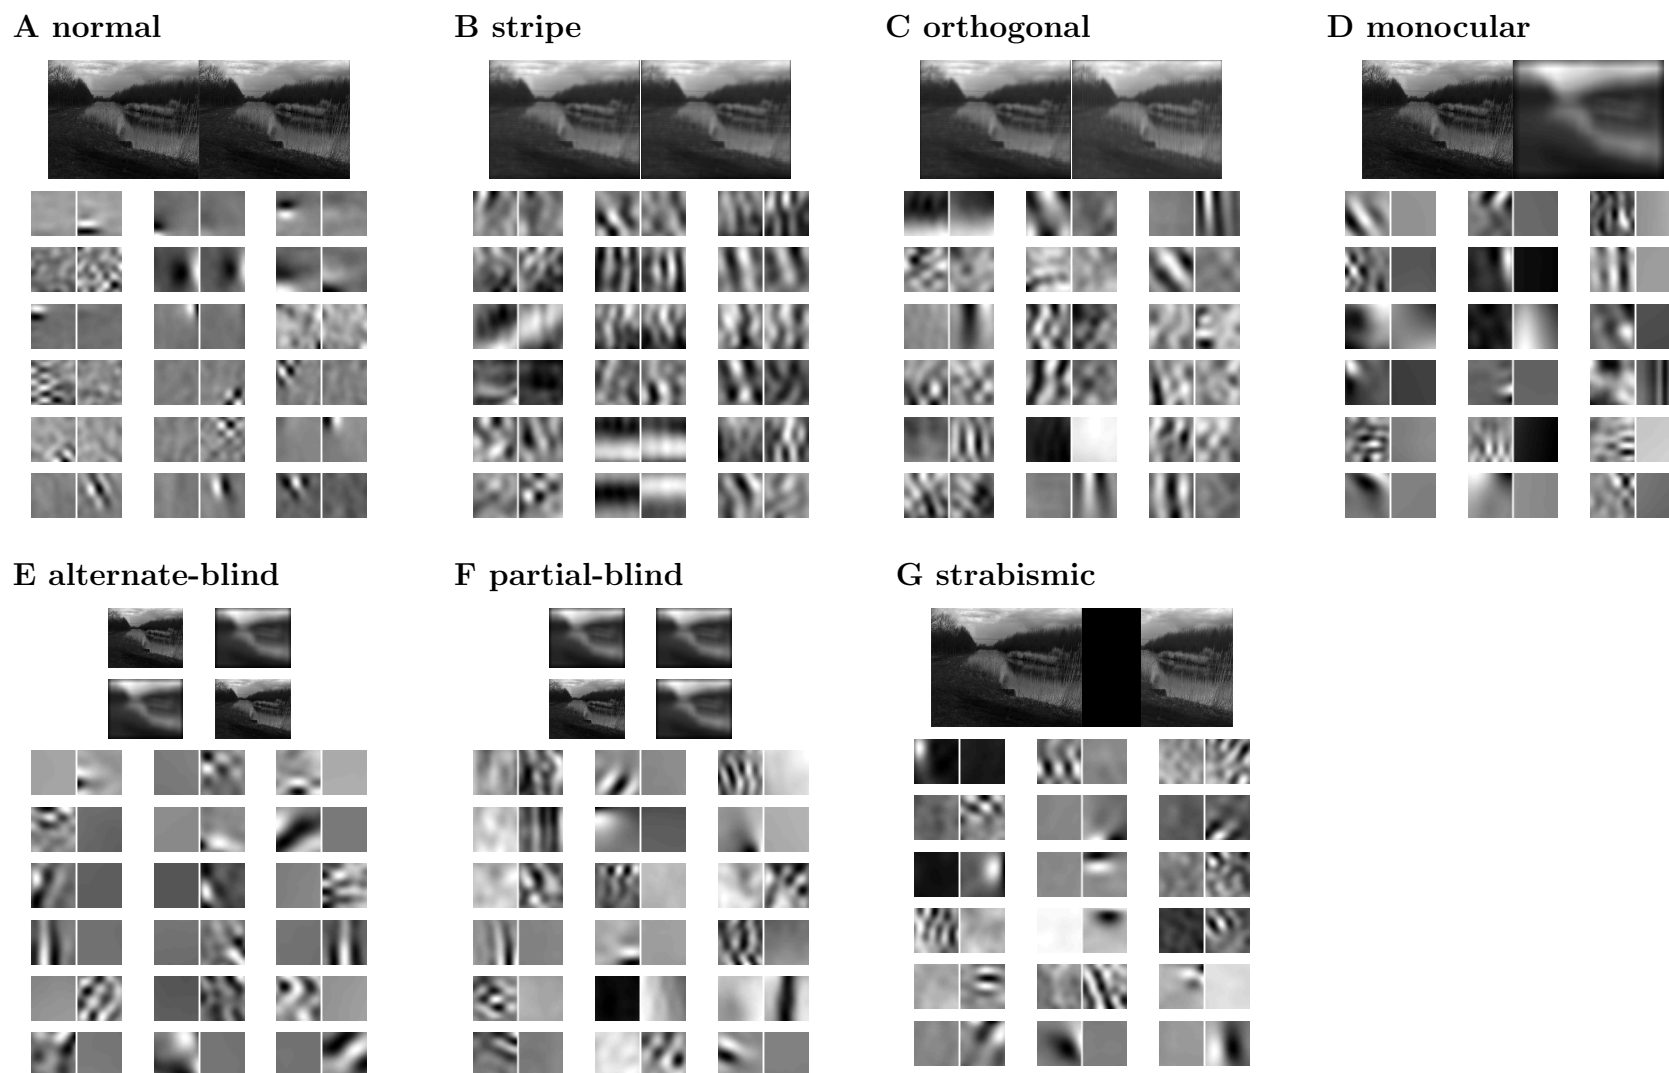
FIGURE 3.1: **Example training images and the resulting receptive fields**
Example binocular training images and the resulting receptive fields generated using the sparse coding model for all the rearing conditions considered. Observe the Gabor-like receptive field structure, along with significant changes in binocular response across rearing condition.

3.3.1 The case of the missing vertical over-representation: did disparity do it?

One, initially unexpected, result we found was that our model consistently predicted a significant over-representation of vertically oriented receptive fields. This vertical over-representation is present across the rearing conditions, including in the normal case.

Such significant vertical over-representation has not been reported in experiment. One possibility for this mismatch would be if the over-representation in the model occurred at large disparities. Most experiments of V1 orientation representation are performed at 0 disparity.

To test this, we examined the distribution of disparity preferences across orientations. Figure 3.6 shows the disparity preferences in the model neurons. Both excitatory and inhibitory disparity tuning was observed, as in experiment (Poggio and Fischer, 1977; Prince et al., 2002). One asymmetry was an increased range of disparity preferences for inhibitory neurons with orientation preferences near vertical. However, overall the disparity preferences are similar at all orientations. In particular, we also plotted orientation preferences of neurons with near 0 disparities. We continued to observe a significant vertical over-representation even in this case. This eliminates disparity preference as an explanation for the absence of vertical over-representation in experiment. Disparity didn't do it[4].

3.4 A low-dimensional model provides insight into vertical over-representation

The significant over-representation of vertical orientations we found in the model required explanation. One possibility was that the asymmetry in orientation representation might be due to the disparity asymmetry in the training input. This asymmetry arises because the eyes are spaced apart horizontally, so inter-ocular horizontal disparities are extremely common, while vertical disparities are rare. In order to investigate we constructed a simpler, low dimensional product of experts model.

---

[4] But we've still got our eye on him.
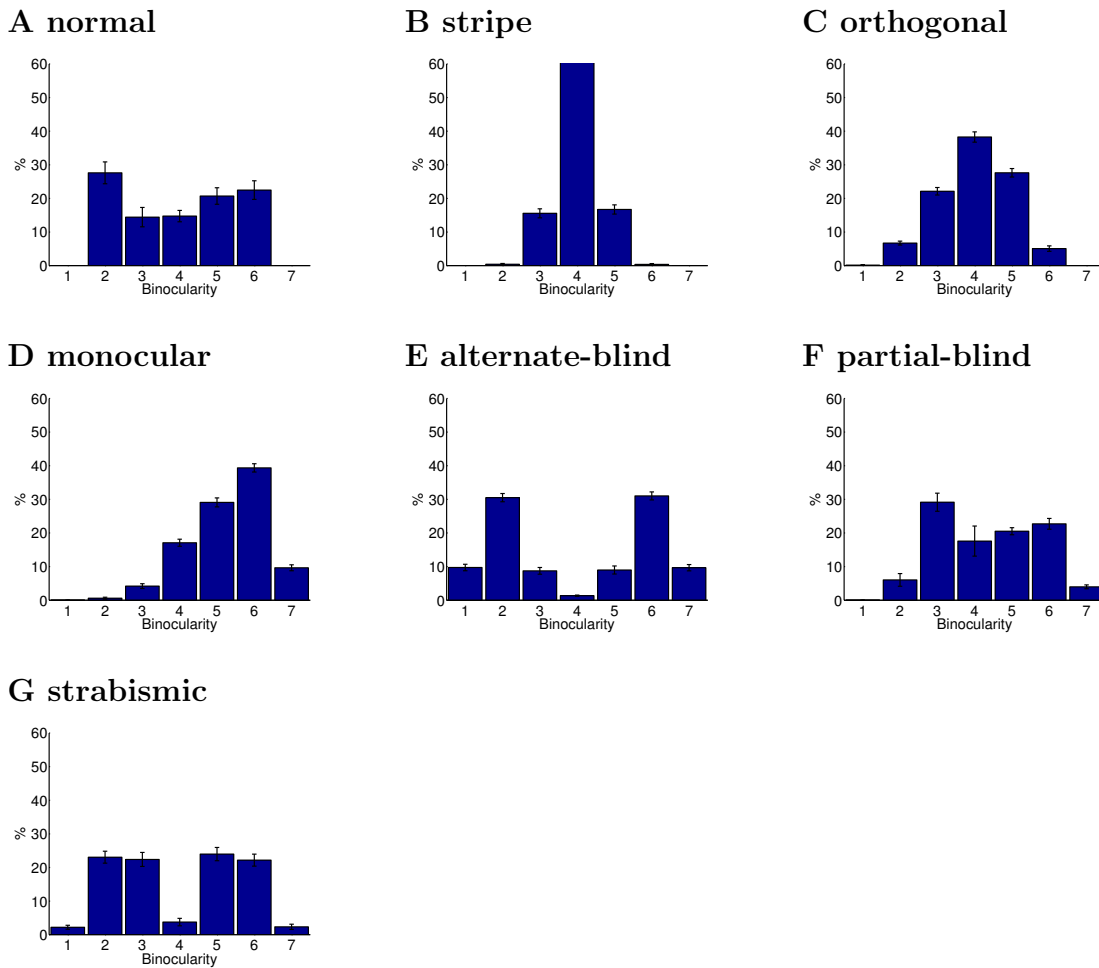
FIGURE 3.2: **Binocularity for different rearing conditions.**
The binocularity of the receptive fields in each rearing condition is shown. Binocularity was measured on a 7 point scale as in Shouval et al. (1996). Values 1 and 7 represent completely monocular responses while values in the middle correspond to binocular responses.
A In the normal rearing condition neurons had a range of binocular responses, although there are few completely monocular neurons.
B In the stripe-reared condition binocularity increases due to higher inter-ocular correlation caused by the reduction in off-axis Fourier components.
C In the orthogonal-reared condition binocularity also increases.
D In the monocular-reared condition neurons develop responses primarily only for the unoccluded eye. This leads to strongly monocular responses for the unoccluded eye.
E Alternate blind rearing removes inter-ocular correlation as each eye is presented with stimuli only when the other eye is occluded. This leads to the strongly monocular responses with neurons distributed equally between both eyes.
F Partial monocular rearing recovers receptive fields for both eyes, however, there are fewer binocular neurons.
G Strabismus decreases inter-ocular correlation and thus leads to increased monocularly.
Errorbars show the SEM. The binocularity distribution of all the modified rearing conditions were significantly different from the normal rearing condition ($p < 10^{-7}$, Kolmogorov-Smirnov).

FIGURE 3.3: **Orientation selectivity across rearing conditions**
The fraction of neurons that are orientation selective for different rearing conditions (each eye shown separately). Neurons are considered selective when their circular variance (Ringach, 2002) is < 0.6. As in experiment, blurring of the visual input significantly decreases the fraction of orientation selectivity neurons for both stripe and orthogonal rearing.
Errorbars show the SEM. All modified rearing conditions had a significantly different fraction of orientation selective neurons ($p < 10^{-8}$, Kolmogorov-Smirnov, each eye tested separately)

Firstly, we reduced our model to a "world" consisting of only binocular Gabor patches. That is, both the training input and receptive fields consisted solely of Gabor patch pairs, one patch per eye. This simplification dramatically lowered the dimensionality of the problem. This reduction was reasonable since the high-dimensional model demonstrates that receptive fields are Gabor-like. We further reduced the dimensionality by restricting orientation to a binary choice: horizontal or vertical, and by allowing receptive fields to vary in location only in one eye (i.e. allowing variation in disparity but not receptive field position). This means that receptive fields are fully described in this low-dimensional space by an orientation $\theta \in 0, \pi/2$ and disparity $d \in R$ preference. Formally, we defined the response of each neuron as:

$$y_i(\theta, d|\theta_i, d_i) = \frac{1}{\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} du \, dv \, [G(u,v|\theta_i,0)G(u,v|\theta,0) + G(u,v|\theta_i,d_i)G(u,v|\theta,d)] \quad (3.8)$$

G is the Gabor function:

$$G(u,v|\theta,d) = \exp\left[-\frac{(u-d)^2 + v^2}{2\sigma^2}\right] \cos \omega u' \quad (3.9)$$

$$u' = u\cos\theta + v\sin\theta \quad (3.10)$$

$$\omega = \pi, \sigma = 1 \quad (3.11)$$

Equation 3.8 replaces equation 3.4 in the high-dimensional model. As before, the response y_i
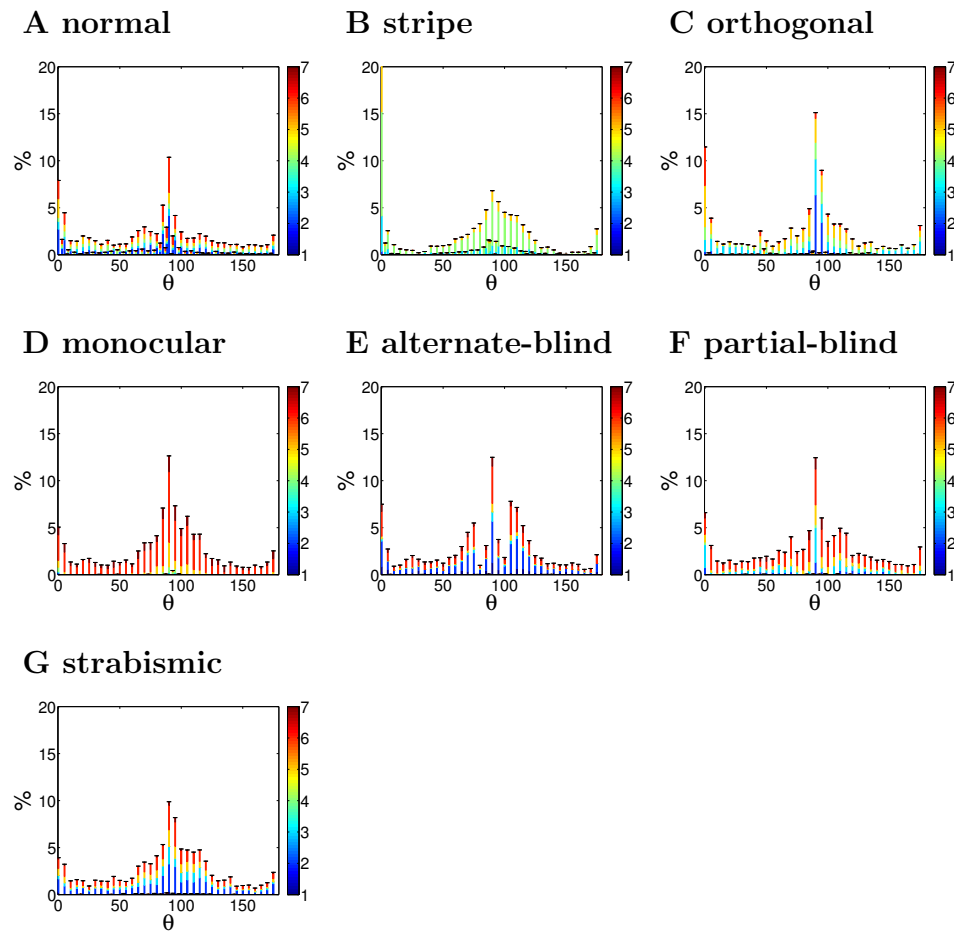
FIGURE 3.4: **Orientation preference distributions for different rearing conditions**
Only neurons that had a well defined orientation preference (circular variance < 0.6) in at least one
eye were included. The colour of each neuron indicates the binocularity of neural response. For
neurons that had a well-defined orientation preference in both eyes, we took the average orientation
preference across both eyes.
A In the normal rearing condition there is a large over-representation of vertical orientations, and
to a lesser degree horizontal orientations. Nonetheless, a range of orientation preferences are ob-
served.
B In the stripe-reared condition there is a significant over-representation of neurons responding
to horizontal lines ($0°$). These horizontal neurons are also strongly binocular. The vertical over-
representation also persists.
C In the orthogonally-reared condition there is over-representation of horizontal neurons in the
left eye and vertical neurons in the right eye. As found experimentally, these neurons are strongly
monocular for the eye that was over-exposed to their orientation.
D In the monocular-reared condition there is good representation of orientation preferences but
only for the unoccluded eye.
E In the alternate-blind-rearing case there is an even distribution of orientation selectivity and
strong monocularity.
F In the partial monocular-reared condition there is recovery of response for both eyes across the
full range of orientations.
G In the strabismic case there is an increase in monocularity.
Errorbars show the SEM. All modified rearing cases, except strabismus ($p = 0.32$) had orienta-
tion preference distributions significantly different from the normal case ($p < 0.01$, Kolmogorov-
Smirnov).
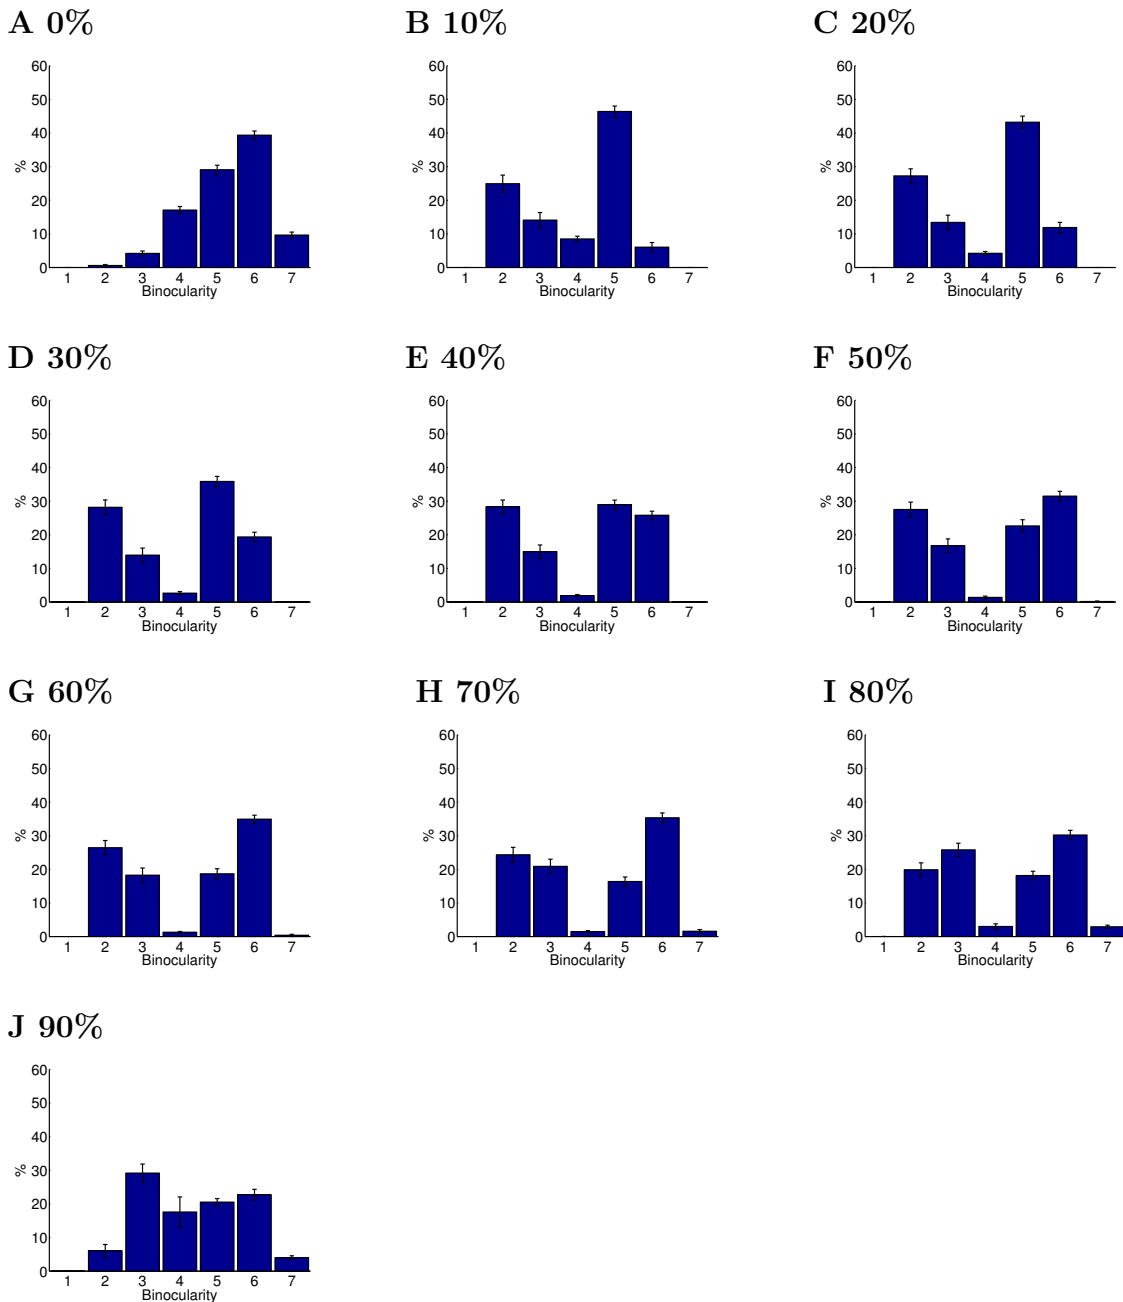
FIGURE 3.5: **Binocular recovery in partially blind reared animals**
This shows the binocularity of the receptive fields learned from input with varying fractions of binocular experience. When the model is trained on completely monocular input **A**, the resulting receptive fields are almost completely unresponsive to the occluded eye. However, even 10% binocular experience **B** leads to a substantial recovery of response to the occluded eye. However, the recovered receptive fields are more monocular than in the normal case. Further increases in the fraction of binocular experience causes a slow recovery of binocular receptive fields, however, even with 90% normal visual experience, neurons are still significantly more monocular. **C-J**.
Errorbars show the SEM. All the partial blind rearing conditions had binocularity distributions significantly different from the monocular case ($p < 10^{-12}$, Kolmogorov-Smirnov). Additionally, all cases with partial binocular experience had significant different binocularity distributions from the normal case ($p < 10^{-12}$, Kolmogorov-Smirnov).
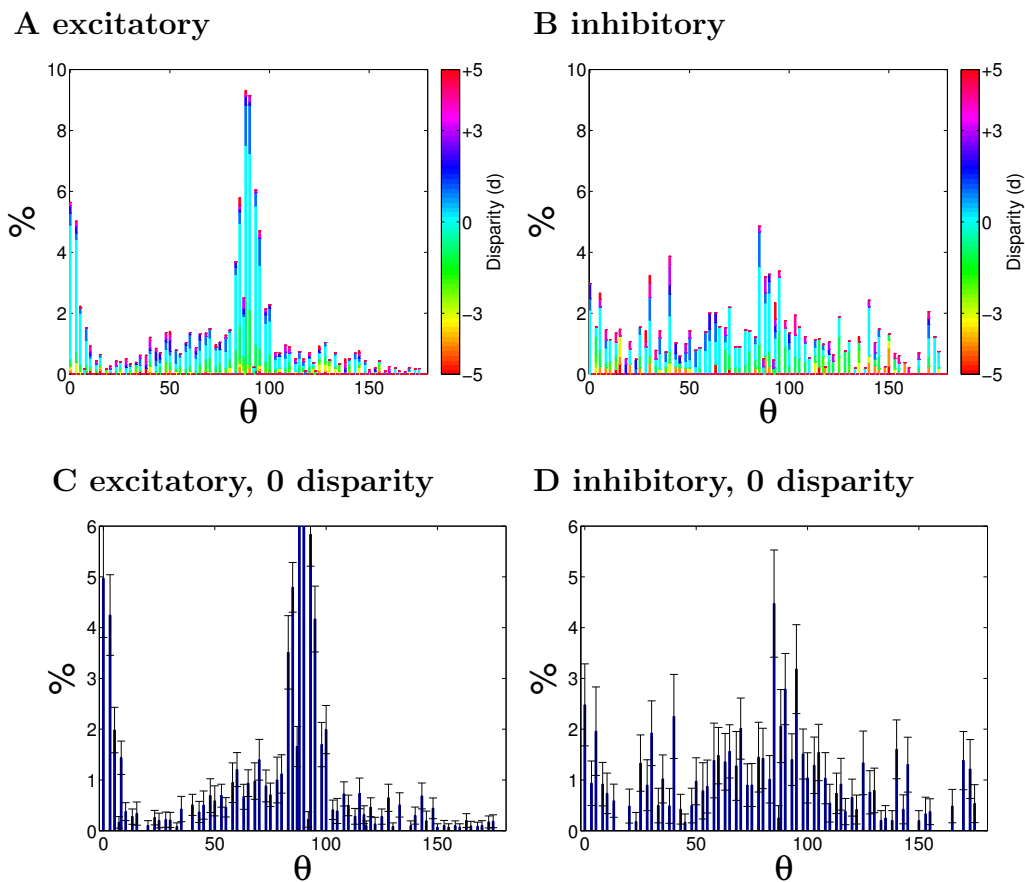
FIGURE 3.6: **Disparity distribution of neurons in high-dimensional model**
The preferred disparities of the binocular receptive fields learned on the normal binocular input
are shown here grouped by preferred orientation. Receptive fields were classified as tuned excita-
tory or tuned inhibitory. **A** shows the preferred orientations and disparities for excitatory neurons
($n = 41 \pm 1.6$). **B** shows the preferred orientation and disparities for the inhibitatory neurons
($n = 16.2 \pm 0.15$) The disparities are reported in pixels with positive values corresponding to far
disparity and negative values to near disparity preferences.
Most experiments have examined V1 response near 0 disparity. The distribution of orientation pref-
erence of neurons with disparity preferences near 0 ($< 1°$) is shown for excitatory **C** and inhibitory
receptive fields **D**. The significant over-representation of vertical orientations persists. In both
the excitatory ($p = 0.001$, Kolmogorov-Smirnov) and the inhibitory case ($p = 0.01$, Kolmogorov-
Smirnov) we found that the fraction of vertically oriented neurons was significantly higher than the
fraction of horizontal oriented neurons.
Receptive fields without binocular responses, ill-defined orientation preferences or a non-localised
(very low preferred spatial frequency) were excluded. These results demonstrate that the over-
representation of vertical orientations occurs even at near 0 disparities. A wider range of inhibitory
disparity preferences for vertical orientations can also be seen.
Errorbars show the SEM.

of each neuron is the inner product of the receptive field with the input, however, now both the receptive fields and input are continuous Gabor functions. The two terms in equation 3.8 correspond to the responses of each eye. The disparity parameter translates the Gabor in the right eye horizontally and may be positive or negative. Both input and receptive fields are fully described in the model by an orientation and disparity. Note that we are assuming the disparity is encoded by position shifts in the receptive fields between eyes rather than phase shifts, a distinction raised by Fleet and Wagner (1996). However, since inter-ocular correlations decay similarly under a phase-shift model, modifying this assumption will have little affect on the model.

found that disparity information can be encoded in binocular cell responses by either a phase-shifted or position-shifted encoding. We have not attempted to differentiate between these encodings here as both choices of encoding would be difference would not be likely affect the over-represented we have observed.

We calculated neuronal responses to input at different disparities using equation 3.8. Figure 3.7 shows the disparity tuning curve for neurons with horizontal and vertical orientation preferences. This shows a significant asymmetry between the disparity tuning of horizontally and vertically oriented neurons: the responses of vertically oriented neurons decrease rapidly as the input disparity d moves away from the neuron's preferred disparity d_i, while, for horizontal receptive fields, the decay is much slower. This result is relatively intuitive. Imagine two edges overlaid and being dragged apart horizontally. If the edges are oriented vertically then their correlation rapidly decreases, while a pair of horizontally oriented edges remains correlated much longer.

The full symbolic expression for the tuning curves were complex, so for the numerical work below we fitted a simple expression to each curve:

$$c_1 \cos c_2(\theta - \theta_i) \exp\left[-(d - d_i)^2/c_3\right] + c_4 \tag{3.12}$$

The fit was essentially indistinguishable from the full solution over the relevant range.

We used the low-dimensional model to test whether the asymmetry in disparity could result in the asymmetry in receptive field representation we found in the high-dimensional

model. Rather than training the low-dimensional model on simplified input we used a
direct sampling approach which allowed us to focus on the key question: how should such
asymmetries in disparity tuning affect the distribution of receptive field orientations. In this
direct approach neuron preferences, both disparity d_i and orientation θ_i were modelled as
random variables.

The distribution of disparity preferences was assumed to be the same as the distribution
of disparities observed in human natural scene viewing (Liu et al., 2008), which can be ap-
proximated as a Laplace distribution. The distribution of viewing disparities is a distribution
over angular separation, while disparity in our model is parametrized in spatial wavelengths.
This means that a parameter for receptive field spatial frequency ν is needed. We assumed
that receptive fields for different spatial frequencies do not interact, which allowed us to
calculate separate, low-dimensional models at varying spatial frequencies.

Orientation preference, which was limited to the binary choice of vertical or horizontal,
was assumed to be sampled from a Bernoulli distribution. A Bernoulli distribution is de-
scribed by one parameter, we termed p_{vert}, which is the probability that a neuron has a
vertical orientation preference. These assumptions allowed us to create a simplified gener-
ative model with only two free parameters: the fraction of vertically oriented neurons p_{vert}
and the model spatial frequency μ.

With this simplified model we tested whether the asymmetry in disparity tuning leads
to an over-representation of vertically oriented neurons. Monte Carlo sampling was used to
test the prediction our model made about the relative likelihood of horizontal and vertical
input. Each round of Monte Carlo 150 receptive fields preferences (d_i, θ_i) were sampled from
their random variables and then we calculated the ratio:

$$f_v = \frac{p(\theta = 0, d = 0)}{p(\theta = \pi/2, d = 0)} \tag{3.13}$$

where $p()$ is the product of experts model distribution function using the low-dimensional
receptive field response defined in equation 3.8. This ratio is the model's predicted likelihood
of vertical input relative to horizontal input. Since we only ever needed to calculate a model
ratio, we avoided all the issues surrounding the PoE normalisation. We used 50000 rounds

of Monte Carlos sampling to estimate f_v for each choice of parameters.

If neuron orientation preferences were distributed equally ($p_{vert} = 0.5$), then we consistently found that the model predicted vertical input to be more likely than horizontal input. A simple root finding method, Brent's method (Brent, 2002), was used to find the value of p_{vert} necessary for the product of experts model to predict an equal probability of input orientations ($f_v = 0.5$). This was performed separately for different spatial frequencies to an accuracy of 0.01. Figure 3.8 shows the value of p_{vert} required at different spatial frequencies. There is a significant over-representation of vertical neurons required across the physiologically relevant range of spatial frequencies in order for the model to predict an equal likelihood of vertical and horizontal input. This result is due to the difference in the decay of inter-ocular correlations between vertical and horizontal receptive fields. As expected, the over-representation increases with increasing spatial frequencies.

Although the simplifying assumptions in our low-dimensional model would lead us not to expect quantitative agreement with the high-dimensional model, there is rough agreement between the two. The over-representation of vertically oriented neurons, relative to horizontal, predicted in the low-dimensional model, is ≈ 2 over the range of spatial frequencies considered, in fair agreement with the high-dimensional model. This indicates that the over-representation in the high-dimensional model is probably due to the asymmetries in inter-ocular correlations.

3.5   Discussion

We have shown here that a sparse coding model can qualitatively explain the changes observed in V1 simple cell receptive fields for six modified rearing conditions. This provides evidence that a key function of V1 is to re-encode visual input sparsely. Additionally, it indicates that the encoding is learned, at least in part, from visual input during development. We have demonstrated that binocular sparse coding provides a unifying framework for modelling the changes in receptive field properties seen across a wide variety of experiments rearing animals with modified visual rearing conditions.

One exciting area of recent experimental research has been that of partial blind rearing

Figure 3.7: **Disparity tuning curves in the low-dimensional model**
This shows the disparity tuning curves of the receptive fields in the simplified model (equation 3.8) for the 3 possible combinations of the input orientation θ and receptive field orientation θ_i (the equation is symmetric with respect to θ and θ_i). The asymmetry in the disparity tuning curves of vertically and horizontally oriented receptive fields is evident. As expected, the receptive fields have little response to their non-preferred orientation. Vertically oriented receptive fields have a narrow response to their preferred disparity d_i, while horizontally oriented receptive fields have a much wider tuning curve for disparity.

(Wiesel and Hubel, 1965b; Kind et al., 2002; Schwarzkopf et al., 2007; Vorobyov et al., 2007; Mitchell et al., 2009). Monocular reared animals do not develop significant V1 responses to the occluded eye (Hubel and Wiesel, 1970; Wiesel and Hubel, 1965a; Shatz and Stryker, 1978; Stryker, 1978). However, recent experiments have demonstrated the recovery of near-normal visual acuity in animals allowed only a small fraction (1/7) of binocular experience. Better understanding of the mechanisms involved in such recovery could lead to important applications in the treatment of visual impairments. The results here indicate that the recovery observed in partial blind rearing experiments can be understood as a natural consequence of developmental optimisation of V1 coding, rather than due to any specialised recovery mechanisms.

3.5.1   Incomplete models: vertical over-representation

Given their simplicity, the success of sparse coding models at capturing many aspects of receptive field development is remarkable. However, despite the binocular coding model

FIGURE 3.8: **Over-representation of vertically oriented neurons in the low-dimensional model**
The asymmetry introduced by disparity being only horizontal causes an asymmetry in the representation of vertical and horizontal input. This plot shows the over-representation of vertically oriented neurons required for the low-dimensional model to maintain an equal probability of horizontal and vertical input. As would be expected, the required over-representation varies as a function of spatial frequency.
p_{vert} is the probability that a neuron's orientation preference is vertical rather than horizontal. All values are accurate to within 0.01. Values above 0.5 indicate an over-representation of neurons with vertical orientation preferences.

successfully postdicting the changes observed in unusually reared animals, there were some important disagreements with experiment. In particular, vertically oriented neurons were significantly over-represented in every rearing condition.We will briefly discuss some minor discrepancies before focussing this major issue of vertical over-representation.

Biological retinas are not arranged in square pixels, unlike most digital representations of images. Such differences in representation can have a small effect on the results of sparse coding models (Doi et al., 2003). In pixel form, at low frequencies, non-cardinal spatial frequencies are not as well represented, while at high frequencies off-cardinal frequencies are better represented. Such differences, perhaps emphasised in the binocular model due to interocular correlations, may be behind the small over-representation of cardinal axis observed in our results. However, there is evidence that cardinal axis are slightly over-represented in animals (Coppola et al., 1998b; Furmanski and Engel, 2000; Li et al., 2003) (although not in humans, see Mannion et al. (2010b,a)), possibly due to a prevalence of cardinal

edges in natural scenes (Coppola et al., 1998a; Dragoi et al., 2001; Hansen and Essock, 2004). Therefore, some degree of cardinal over-representation in the model is in agreement with experiment. Additionally, like Hsu and Dayan (2007), we found the distribution of orientation preferences contained narrow troughs of under-representation bordering over-represented orientations. Such discrepancies are probably explained by the minor differences between biological visual input and our training examples. However, they may also be hints that additional constraints play a role in biological receptive field development.

The most important feature of our results was the strong over-representation of vertical orientations. Unlike some of the other minor discrepancies, which may be due to the pixel representation of the images, the vertical over-representation cannot be dismissed easily.

The low-dimensional model we introduced allowed us to consider vertical over-representation in a simpler model with well-defined training input. This model demonstrated that vertical over-representation could arise from the asymmetry in inter-ocular correlations between horizontal and vertical edges. This asymmetry is due to the eyes being spaced apart horizontally rather than vertically. We found good agreement between the degree of over-representation predicted in the low-dimensional model and the more biologically realistic model trained on natural scenes.

As far as we are aware, this over-representation of vertical orientations in sparse coding models has not previously been reported. Our low dimensional model provides evidence that the asymmetry arises due to inter-ocular correlations. Most groups have only considered monocular receptive field development, so such correlations will not exist in their training input. Hoyer and Hyvärinen (2000) studied binocular receptive field development using a similar, independent component analysis model. They did not show the orientation distribution of their receptive fields, so it is not possible to tell if they found such an vertical over-representation in their results. Okajima (2004) used a supervised model of binocular receptive field development to examine disparity coding. Again, the distribution of orientations was not examined, however, in this case the use of supervised model imposes different constraints on the encoding so that an over-representation may not have been present.

Such an over-representation of vertical orientations has not been reported in experiment. The orientation preferences in V1 simple cells have been studied for some time (Hubel and

Wiesel, 1959). Although there is some evidence that cardinal representations are slightly over-represented (Coppola et al., 1998b), an over-representation of vertical orientations has not been reported. Our model predicts an over-representation of vertical orientations of nearly two-fold compared to horizontal orientations. Such a significant over-representation would be difficult to miss. We wondered if this mismatch could be due to most animal experiments only examining orientation preference at 0 disparity, where the asymmetry in inter-ocular correlations do not exist. However, even when we considered only receptive fields with near 0 disparities, we still found a significant over-representation of vertical orientations. It seems that such an over-representation of vertical orientations does not occur in real animals.

Such a mismatch is unlikely to be due to the specifics of our sparse coding approach. Vertical over-representation occurred even in the partial blind rearing condition, where the model is trained with only a small amount of binocular input. This indicates that the over-representation is robust to perturbations in the input. Further, the results of a low-dimensional model support the view that the vertical over-representation is a general consequence of inter-ocular correlations in sparse coding, rather than specific to our particular sparse coding model or training input. The overcompleteness of the model may impact on the amount of vertical over-representation, however Hsu and Dayan (2007) did not find dramatic shifts in orientation over-representation with varying degrees of overcomplete representations.

Spontaneous retinal waves are one possible reason for the mismatch between model and experiment. Spontaneous retinal waves occur before eye-opening and contain some important statistical similarities with natural scenes (Albert et al., 2008). Notably, spontaneous retinal waves do not contain inter-ocular correlations that generate the vertical over-representation we find in our model. It is possible that receptive fields are partially learned from the statistics of spontaneous retinal waves. By eye-opening certain aspects of the receptive fields, such as their orientation distribution may be fixed. This hypothesis could be examined by training a sparse coding model on a mixture of simulated spontaneous retinal wave input and natural scene input, although it is not obvious how to incorporate changes in plasticity into a sparse coding model.

Alternatively, this mismatch between theory and experiment may indicate additional constraints on receptive field development. It is possible that other mechanisms, such as top-down feedback, act to restrict V1 receptive field development. For instance, a recent theory of V1 orientation preference map development postulates orientation rotation invariance (isotropic orientation representation) as an important symmetry law needed for map stability (Kaschube et al., 2010). The agreement of receptive field changes between model and experiment over a wide range of rearing conditions appear to demonstrate that sparse coding is an important function of receptive field development. However, such codings may be learned with additional constraints we have not included in our model that may explain the discrepancy. Hsu and Dayan (2007) found some small discrepancies between their sparse coding model response and those reported in experiment, which they also attributed to additional constraints not included in their model.

Several groups are pursuing multi-layer, unsupervised learning models of cortical visual areas (Hoyer and Hyvärinen, 2002; Hyvärinen and Hoyer, 2001; Perrinet et al., 2004; Lee et al., 2007; Karklin and Lewicki, 2008; Köster and Hyvärinen, 2010). The observation highlighted here, particularly the over-representation of vertical orientations as a consequence of binocular correlations in unsupervised learning models, raises important questions about the possibility of additional constraints or the influence of spontaneous retinal waves on visual system learning. Resolving this issue may be important for improving such multi-layer models. At a minimum, it indicates that coding models need to be tested with binocular input for richer comparison with experiment.

3.5.2 Unsupervised learning models demonstrate the building-blocks of natural scenes are edges

When dealing with high-dimensional data, such as visual input, an apparently trivial question turns out to be difficult to answer: what are the basic features of the data? Knowing the basic features of a dataset is a prerequisite to comprehending the data. Unsupervised learning models can find the underlying features in datasets while imposing few assumptions. We have seen in this chapter that a sparse learning model captures many salient features of V1

receptive field development, including development of receptive fields in modified rearing conditions. In addition to modelling V1 responses, such learning algorithms also, implicitly, answer the question about the casual building-blocks of natural visual scenes. The consistent answer is that natural scenes are built from edges of different orientations, spatial frequencies and positions. Edges are both the sparse components of natural scenes and, not unrelatedly, the way they are coded in V1. Decomposing natural scenes into underlying components helps us reduce the complexity of visual processing.

V1 consists of much more than the feature detectors we have modelled it with in this chapter. Real V1 neurons incorporate temporal integration (Ringach, 2002), attentional modulation (Ito and Gilbert, 1999; Poghosyan and Ioannides, 2008), lateral connections between columns (Das and Gilbert, 1999; Bosking et al., 1997; Stettler et al., 2002), feature maps (Swindale, 2008), and significant feedback from other cortical areas (Lund et al., 1975). One way of modelling some of these aspects of V1 is to continue to be guided by naturalistic input statistics, but consider naturalistic statistics at different feature levels. As we will discuss in detail in the next chapter, edges in natural scenes contain important statistical relationships. In particular, edge pair relative positions and orientations are correlated. The remainder of the thesis will be examining the role these higher-order statistics play in V1 development and function.

3.5.3 Future work

It is clear that unsupervised learning will only take us so far in understanding brain function. At some point, brains have goals, seek rewards and predict the future. However, just as has proved important in articificial machine learning, it appears the brain often transforms high-dimensional input into more useful representations (Hinton and Salakhutdinov, 2006). Coupling unsupervised learning models with other types of learning is an active area of machine learning research (Hinton, 2007b,a). Such learning models may be important for extending functional models of cortical behaviour from early sensory areas. This will require advances in both machine learning and neuroscience. Early sensory coding provides a window into cortical function that is amenable to experiment and theoretical description, it will likely

play a role is guiding and refining such ambitious models.

The approach to testing sensory encoding models described here may be generally applicable to further work on both vision and other sensory modalities. Training coding models using modified visual input allows both strong tests of coding theories and a tidy way of separating intrinsic developmental processes from properties learned from external input. Audio encoding is one area with rich input properties (Attias and Schreiner, 1997) which may be amenable to such analysis. Functional models of cochlea encoding during normal development already exist (Lewicki, 2002; Smith and Lewicki, 2006).

When I am in a painting, I'm not
aware of what I'm doing. It is only
after a sort of 'get acquainted' period
that I see what I have been about.

Jackson Pollock

4

Co-circularity in natural scenes

Synopsis

Natural scenes contain statistical relationships at many different feature levels. For instance, exploiting pixel-wise temporal and spatial correlations in natural scenes is an essential component of image compression schemes. More recently, Sigman et al. (2001) examined correlations in edge orientations. Sigman et al. generalised previous work, showing that edges in natural scenes have a tendency to lie tangent to a common circle. Here we expand their foray by examining co-circularity at a wider range of length scales than previously considered, using a robust statistic for measuring co-circularity strength. We also used co-circularity to cluster natural scenes and present some preliminary results about the variation in edge statistics across the distribution of natural scenes.

We find that co-circularity is present in natural scene edges to distances of $> 14°$, significantly further than previously reported. We find no significant linear relationship between

57

co-circularity and separation distance. Additionally, we find that the amount of co-circularity in natural scenes is a sparse distribution, with the majority of scenes having a small amount of co-circularity, but with many outliers far from the mean. We find, using k-means clustering, that these outliers belong to a different class of scenes from the majority.

Co-circularity, in natural scene edges, is an important statistical property because it is present over a wide range of distance separations. Further work is needed to examine why some natural scenes have significantly different co-circularity statistics from the majority.

4.1   Introduction

As elaborated in chapter 2, the distribution of natural scenes contain correlations at many different feature levels. Some of the first modern efforts on natural scene statistics, driven by a need for better TV signal encodings, concentrated on the strong spatial and temporal correlations in pixel intensity (Kretzmer, 1952). More recently, other authors have considered these pixel-wise correlations in more detail (Dong and Atick, 1995; van Hateren and Ruderman, 1998; Long et al., 2006), as well as characterising other properties such as colour (Párraga et al., 1998; Webster, 2007; Hansen and Gegenfurtner, 2009). However, since a perfect compression algorithm would provide a solution to the halting problem (Wilson and Keil, 2001), it is not possible to prove that one has found a complete description of the statistics of a set. Noticing new statistical relationships is potentially always possible.

The predominance of co-linearity, edges aligned along their long axis to form lines, is an obvious feature of the world. This makes it all the more surprising that it was only recently Sigman et al. (2001) generalised this observation by showing that edges in natural scenes tended to lie on a common circle - co-circularity. This means there are small correlations between edge position and orientation for all pairs of edges, not just edges of the same orientation, as in co-linearity. Since Sigman et al., others have shown that co-circularity is important in human contour grouping (Geisler et al., 2001; Fulvio and Singh, 2009; Geisler and Perry, 2009) and that co-circularity can arise due to the presence of closed, smooth contours (Chow et al., 2002).

Here we extend the previous studies of co-circularity in natural scenes in two ways.

Firstly, we quantify the amount of co-circularity present in natural scenes over longer distances than previously considered ($\approx 14°$ compared to $\approx 3.4°$). In order to make this computationally feasible and to allow quantification of the amount of co-circularity present in natural scenes we define a new metric: the degree of co-circularity (Hunt et al., 2009). The second extension is that we group natural scenes by their co-circularity statistics to better understand where the variance in co-circularity arises.

By utilizing both improved quantification of co-circularity and advances in computational power we demonstrate two key findings. Firstly, we demonstrate that correlations between edge orientations are extremely long-range. We consider the degree of co-circularity present between edges separated by distances varying from $< 1°$ to over $14°$ and do not find a significant decrease in the mean degree of co-circularity with distance. This indicates that co-circularity may be a particularly important natural scene statistic, as most other scene correlations are not present at these long length scales. Additionally, we demonstrate that the amount of co-circularity in natural scenes varies significantly. Since we used large photos, this variation is not due to the sample size, but rather differences in the type of scene. Supporting this, we find that the amount of co-circularity in natural scenes is an extremely sparse distribution, with the majority of scenes having a small amount of co-circularity. Most of the variance is due to a small fraction of outliers. We provide some preliminary analysis, which shows that co-circularity may be a useful statistic for classifying scene types.

4.2   Methods

There are two steps involved in examining the statistics of orientations in natural scenes. First, natural images have to be filtered to assign an orientation to each pixel. Then a statistic is calculated from the orientations. We created a metric to quantify the co-circularity of each image over a range of distances. All edge extraction and quantification was done using a custom-made MATLAB[1] script. The calculations were run on a commodity-hardware cluster, each image was run as a separate task.

---

[1] version 2008a, The Mathworks, Natick MA, USA

4.2.1 Edge extraction from natural scenes

A high-quality, artefact-free library of natural scenes is a prerequisite for accurately deter-
mining natural scenes statistics. We used the van Hateren image database (van Hateren
and van der Schaaf, 1998) which consists of 4000 images of the natural world, linearized and
with no lossy compression applied. Each image is 1536 horizontal by 1024 vertical pixels,
sampled in 12-bit greyscale.

Each image in the training set was then filtered, using 3-level steerable pyramids (Simon-
celli and Freeman, 1995), to find the energy in each orientation at each pixel. The use of
steerable filters allowed a multi-scale filter to be calculated efficiently so that the orientation
of each pixel was determined from multiple length-scales. The orientation of each pixel was
assigned to the orientation with highest energy at that pixel; the orientation selectivity of
the pixel was calculated as the fraction of total pixel energy present in the highest-energy
orientation. After this step, for each image, we had a map of orientation and orientation
selectivity at each pixel position.

4.2.2 Quantifying co-circularity

As we discussed in chapter 2, co-circularity is a specific type of symmetry-breaking. Edges are
co-circular when they are tangent to a common circle. This relationship between the relative
position and relative orientation of a pair of edges breaks a symmetry that a distribution of
randomly oriented edges possess: orientation can not be rotated independently of position
without modifying the statistics of the edges. Of course, in natural scenes edge pairs are
not perfectly co-circular, but rather they are near co-circular more often than would occur if
edge orientation was independent of position. To study the co-circularity of edges in natural
scenes we needed a robust way of measuring the amount of co-circularity in a scene.

In Hunt et al. (2009) we introduced a metric for measuring co-circularity which is straight-
forward to calculate, intuitive and robust to noise. We explain that metric, the degree of
co-circularity D_{diff}, here. Each pixel i at position r_i possesses an orientation θ_i, where
$\theta_i \in [0, \pi)$. For each pair of edges we can then define the terms listed in table 4.1 (also
shown diagrammatically in figure 4.1). The definitions are identical to those introduced in

θ_i	orientation of pixel i
θ_j	orientation of pixel j
s_i	orientation selectivity of pixel i
s_j	orientation selectivity of pixel j
r	distance between pixel's i and j
ϕ	angle of the line from pixel i to pixel j
$\alpha = \theta_j - \theta_i$	orientation difference between pixel i and j
$\beta = \phi - \theta_i$	angular position of pixel j relative to the orientation of pixel i

TABLE 4.1: **Definition of terms**

FIGURE 4.1: **Definition of terms**
Diagram showing the geometric definitions of the relative orientation and position between an ordered pair of oriented pixels. The definitions are listed in table 4.1.

chapter 2, but now we have added a weighting by orientation selectivity s.

A pair of pixels are co-circular when they satisfy:

$$\theta_j = 2\phi - \theta_i \quad (\text{mod } 180°) \tag{4.1}$$

The deviation of θ_j from co-circularity can then be described as an angle:

$$d_{\text{diff}} = |\theta_i + \theta_j - 2\phi|_{\text{ mod } 180°} \tag{4.2}$$

Note that due to the ϕ term, d_{diff} is not symmetric to permutation of i and j (since we calculate the mean d_{diff} over both permutations of pixels this lack of symmetry is not important). We use the notation $||_{\text{ mod } 180°}$ to indicate that the absolute difference is calculated modulo

180°:

$$|a|_{\text{mod } 180°} = \min(a \,(\text{mod} 180°), (-a) \,(\text{mod} 180°)) \tag{4.3}$$

Geometrically, this means that when calculating the co-circularity of a pixel pair we always consider the smallest angle between the pixel orientation and the fully co-circular orientation θ_{cc}. It follows that d_{diff} is never larger than $90°$ (for instance, if $\theta_b = 135°$ and $\theta_b^{cc} = 20°$ then $d_{\text{diff}} = 65°$).

We calculated D_{diff}, the weighted mean d_{diff} between all ordered pairs of edges binned by distance:

$$D_{\text{diff}}(r) = \frac{\sum_{i,j} s_i s_j |\theta_i + \theta_j - 2\phi|_{\text{mod } 180°}}{\sum_{i,j} s_i s_j} \tag{4.4}$$

$D_{\text{diff}}(r)$ has the intuitively appealing property that it is $45°$ when orientations are independent of position and decreases if the edges are oriented co-circularly. This calculation is performed over both permutations of i and j for each pair of pixels. This is to make the result independent of the pixel ordering, since, as previously mentioned, d_{diff} is not symmetric for i and j. The mean was calculated in bins spaced evenly $\approx \frac{5}{6}^\circ$ (50 pixels) apart positioned at $r = [50n + 25, 50n + 75)$ pixels $\approx [n\frac{5}{6}^\circ + \frac{5}{12}, (n+1)\frac{5}{6}^\circ)$ $n \in \mathbb{Z}_{17}$.

In addition to examining $D_{\text{diff}}(r)$, the mean d_{diff} at each distance bin r, we also calculated the distribution of d_{diff} at each distance. The values of $d_{\text{diff}} = |\theta_i + \theta_j - 2\phi|_{\text{mod } 180°}$ for each pair at distance r were binned to create a histogram $H(r, d)$. The r bins were the same as used for calculating $D_{\text{diff}}(r)$ and the d_{diff} bins were spaced $15°$ apart. This allowed a basic characterisation of the distribution of co-circularity values.

Kurtosis is a measure of the sparsity of a distribution. It is defined as:

$$k = \frac{E([x - E(x)]^4)}{\sigma^4} \tag{4.5}$$

For the normal distribution $k = 3$. We estimated the kurtosis of the $D_{\text{diff}}(r)$ distribution using a biased estimator:

$$k^* = \frac{\frac{1}{n} \sum_{i=1}^{n} (x_i - \bar{x})^4}{\left(\frac{1}{n} \sum_{i=1}^{n} (x_i - \bar{x})^2\right)^2} \tag{4.6}$$

Since we had a large sample, the bias in the estimator does not contribute significantly to

the result.

Statistical tests

In order to test whether the deviation from 45° of $D_{\text{diff}}(r)$ was significant, we generated a matched control for each scene by randomly permuting the positions of the edges in the scene while retaining their orientations. This meant the controls contained the same distribution of orientations, but distributed independent of position. Since the distribution was sparse, we used a two-sided, non-parametric, paired, sign-rank test Hollander and Wolfe (1999) to test whether the distribution of $D_{\text{diff}}(r)$ was significantly different from the controls. Each distance bin r was tested independently. Since the p-values were extremely low no multiple-comparison correction was explicitly applied.

Many correlations in natural scenes decay with increasing distance. A least-squares linear regression model was fit to the mean $D_{\text{diff}}(r)$ to examine if there was a linear relationship between $D_{\text{diff}}(r)$ and r.

4.2.3   Scene clustering

The distribution $D_{\text{diff}}(r)$ values across the 4000 scenes was sparse. Since each scene has a large number of pixel pairs from which the $D_{\text{diff}}(r)$ value is calculated, these deviations cannot be explained by noise present in the sampling.

We used k-means clustering (Lloyd, 1982; MacQueen, 1967; MacKay, 2003) to examine if the sparsity in co-circularity across scenes was due to a subclass of scene types. For the clustering, each scene was described by a vector of $D_{\text{diff}}(r)$ values binned as described earlier. The clustering used $k = 5$ clusters. We found that increasing the number of clusters did not reveal any new structure.

Once each scene was assigned to a cluster we used a low-dimensional visualization technique, tSNE (van der Maaten and Hinton, 2008), to display the image samples in a two-dimensional plot. As expected, clusters are separated in this low-dimensional display. The mean $D_{\text{diff}}(r)$ for each cluster was plotted separately. Additionally, we randomly chose six scenes from each cluster and displayed them to see if there were any obvious characteristics

FIGURE 4.2: **Mean $D_{\mathbf{diff}}(r)$ versus distance in natural scenes**
The circles indicate the mean value and the error bars show 1 standard deviation. A signrank test showed that the $D_{\text{diff}}(r)$ distribution was significantly different from the matched controls at all lengths $p < 10^{-74}$ (this is not corrected for multiple comparisons, but the low value indicates significance).
The black line shows a linear regression of the mean $D_{\text{diff}}(r)$. This found no significant correlation between mean $D_{\text{diff}}(r)$ and distance r ($p = 0.81$).

about the scenes belonging to each cluster.

4.3 Results

Significant co-circularity is present in natural scenes at a wide range of length scales. Figure 4.2 shows the mean co-circularity across the 4000 scenes up to $> 14°$. At every length r considered, the mean $D_{\text{diff}}(r)$ value is ≈ 43.4 and significantly different from the matched controls. As the error bars demonstrate, there is a significant variation in the co-circularity between scenes.

Figure 4.3 shows a boxplot of the $D_{\text{diff}}(r)$ distributions and their kurtosis. It is clear that the $D_{\text{diff}}(r)$ distribution at all r is extremely sparse. Interestingly, the kurtosis value peaks at $r \approx 3°$. We have, so far, not discovered the reason behind this peak.

The distribution of d_{diff} values across all the natural scenes was considered. The distribution, calculated across all the scenes we considered, is shown in figure 4.4. It can be seen that the distribution has a near-constant slope which leads to the over-representation

FIGURE 4.3: $D_{\mathbf{diff}}(r)$ **sparsity**

A boxplot and kurtosis plot of the $D_{\mathrm{diff}}(r)$ distribution as a function of separation r. The red dots indicate outliers which are defined as points $> q_3 + 1.5(q_3 - q_1)$, where q_1 and q_3 represent the lower and upper quartiles respectively, and mutatis mutandis for the low outliers. The blue boxes denote the quartiles and the red lines the mean. The whiskers extend to the furtherest point that is not an outlier. The distributions are extremely sparse, as seen by the high kurtosis and extreme number of outliers.

of co-circular edges.

4.3.1   Scene clustering

Because the $D_{\mathrm{diff}}(r)$ values were highly sparse, we used k-means clustering to see if some of the variation could be explained by subsets of scene types containing markedly different co-circularity statistics. Figure 4.5A shows a low-dimensional view of the scene clusters. The majority of scenes belong to a central dominant cluster, which, unsurprisingly, has a mean $D_{\mathrm{diff}}(r)$ near the mean value found in figure 4.2. The other, smaller, clusters of scenes have $D_{\mathrm{diff}}(r)$ distributions significantly different, including one cluster containing slightly anti-cocircular scenes. We show in figure 4.5B some random examples of the images belonging to each cluster. There does appear to be some obvious visual trends in the clusters, such

A

B
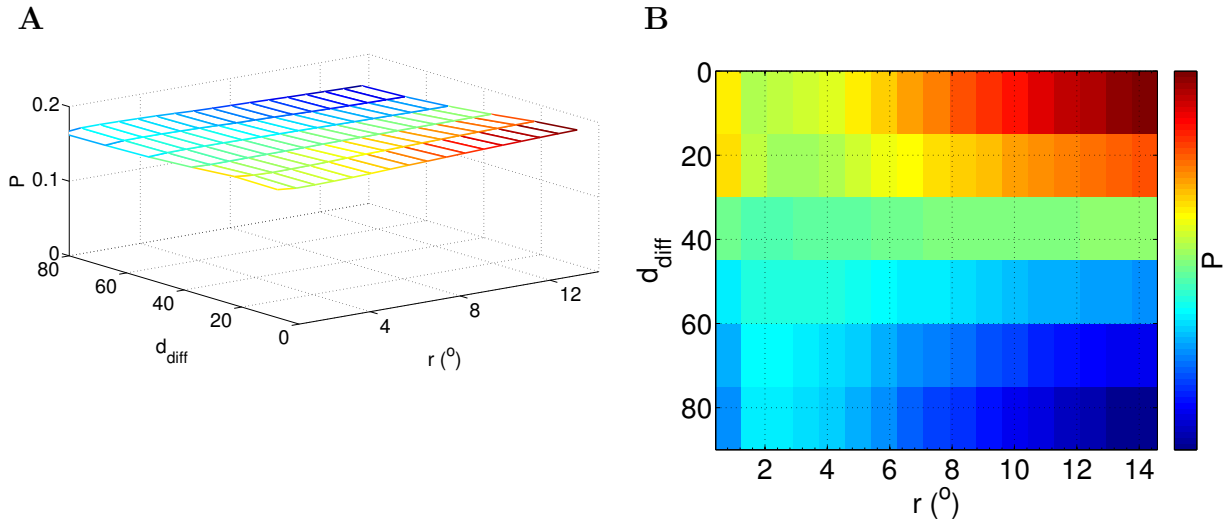
FIGURE 4.4: **Distribution of $d_{\mathbf{diff}}$ values across natural scenes**
Figures **A** and **B** both show the distribution of d_{diff} values as a function of separation r. The distribution is near flat, with a small constant decline in probability with increasing d_{diff} value.

as images dominated by a single repeating motif appearing to have low $D_{\mathrm{diff}}(r)$ values (high co-circularity). However, to draw firm conclusions would require a more in-depth objective consideration of the clusters.

4.4   Discussion

A key finding demonstrated here is that co-circularity in natural scenes is present over exceptionally long distances of at least 14°. Additionally, we found no significant linear decay in the amount of co-circularity with pair separation r. This is in contrast to correlations in pixel intensity which decay rapidly with distance (Field, 1999). There may be raised co-circularity levels at distances $\lesssim 3°$ that diminish with distance, but the $D_{\mathrm{diff}}(r)$ remains significantly below 45° at all separation distances r considered.

The initial discovery of co-circularity in natural scenes (Sigman et al., 2001) appeared to show the correlation between edge position and orientation decaying with distance. However, the plots in Sigman et al. (2001) are in Cartesian co-ordinates so that at greater separations r, the probability distribution of edge positions becomes "smeared out." Our findings here demonstrate that the correlation between edge angular position ϕ and orientation θ remains

A

B

C
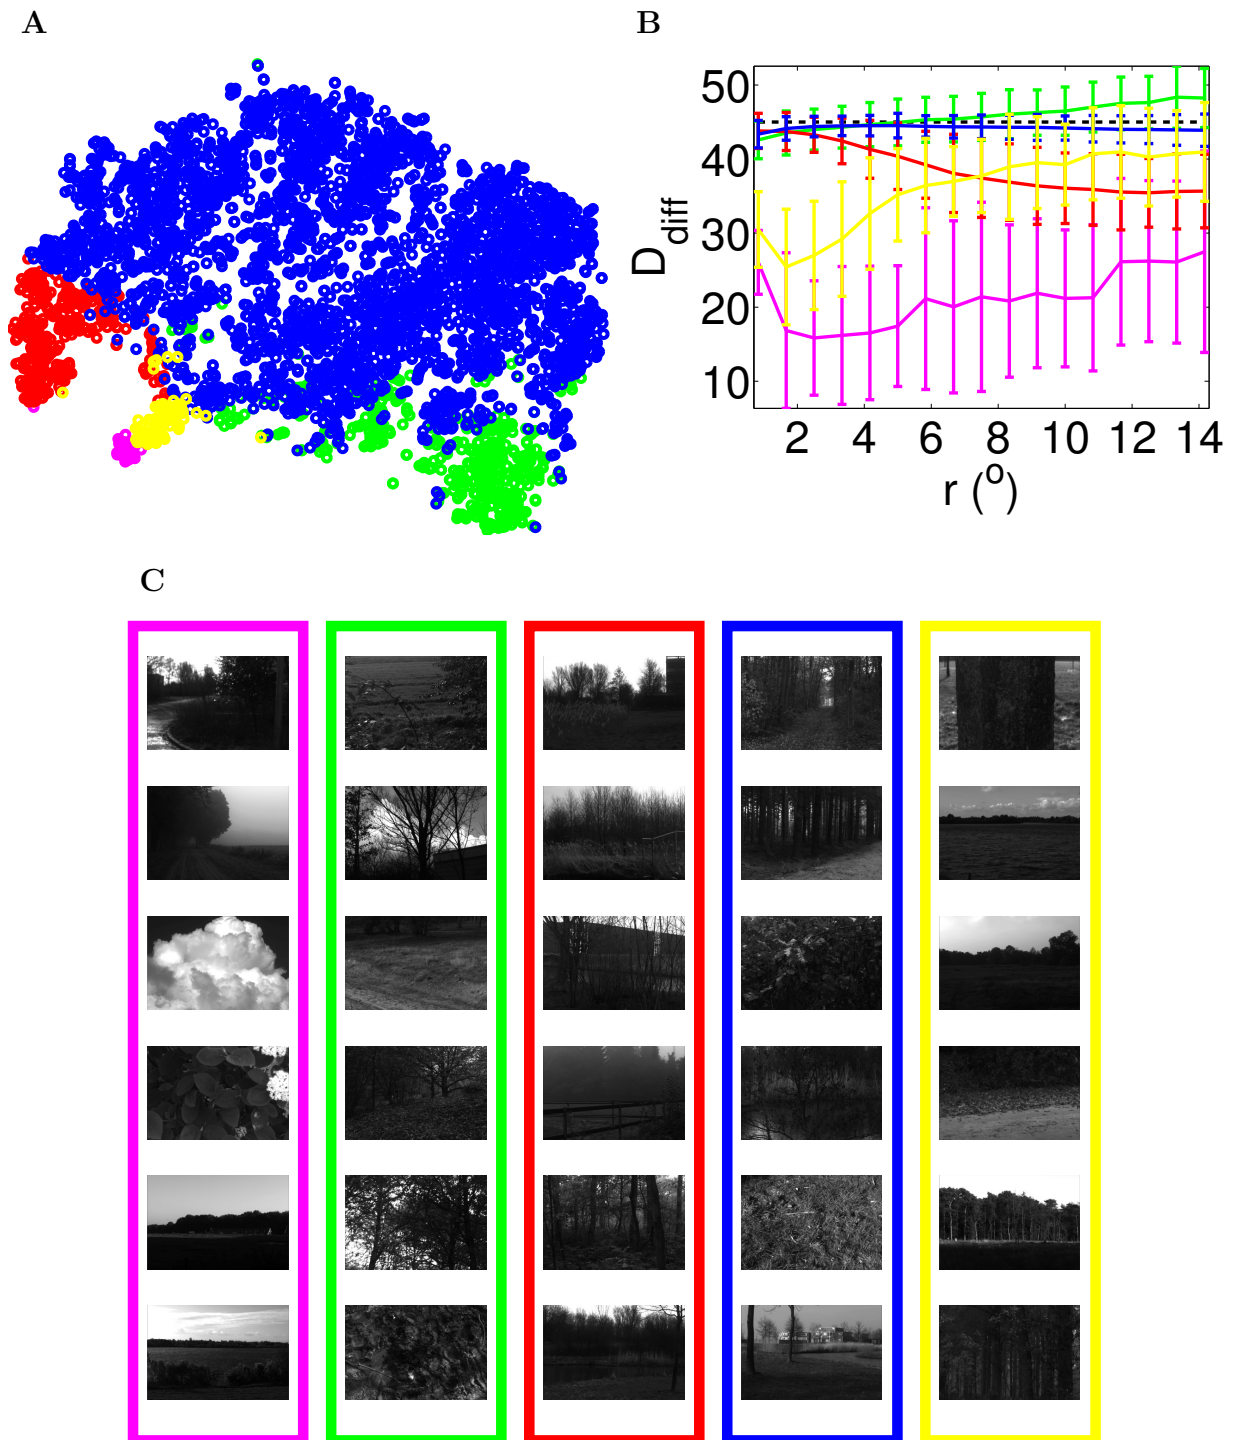
FIGURE 4.5: **Natural scenes clustered by co-circularity value**
A shows a low-dimensional plot (t-SNE, see methods) of the 5 natural scene clusters found using k-means clustering based on the scene $D_{\mathrm{diff}}(r)$ values. One can see that most natural scenes belong to a single dominant cluster.
B shows the corresponding $D_{\mathrm{diff}}(r)$ versus r plot for each cluster. It is clear the high-variance in $D_{\mathrm{diff}}(r)$ values is due to a small fraction of scenes with markedly different amounts of co-circularity. The majority of scenes have $D_{\mathrm{diff}}(r) \approx 43.4°$ (blue cluster). The error bars show $\sigma = 1$ standard deviation.
C Randomly sampled images from each cluster. The coloured frames of the image sets correspond to the cluster they were sampled from.

correlated to an approximately constant degree over all the separation distances r we considered. Presumably, the probability distributions in Sigman et al. (2001) would show a constant width, high probability region, if plotted in polar coordinates.

We also found that the amount of co-circularity varied significantly between scenes. This occurred even though each scene was large ($\approx 1.5 \times 10^6$ pixels), containing a huge number of edge pairs ($\approx 2.5 \times 10^{12}$ pairs). Our k-means clustering showed the majority of scenes contain a small co-circular preference in edge arrangement, with the variance due to a small number of outlier clusters. Co-circularity in natural scenes is present over large distances, but only sparsely, with the majority of scenes containing only a small dependence between edge orientation and position. The underlying causes behind these variations are not yet well-understood.

An obvious question that arises from this work is why edges in natural scenes are often tangent to a common circle. This question is not necessarily teleological, it is possible that co-circularity may be explained by the properties of the processes underlying the creation of the visual world. For instance, Chow et al. (2002) demonstrated that co-circularity may be due simply to the prevalence of closed, smooth contours in scenes.

In this study, as in Sigman et al. (2001), we considered the statistics of all edges in a scene. Of course, edges often belong to objects and may be grouped together to form contours or textures. Geisler et al. (2001) examined the statistics of edges that human subjects had grouped as belonging to the same contour. They found that edges belonging to the same object also tended to be co-circular. Other work has demonstrated that co-circularity may be an important property in texture recognition (Ben-Shahar and Zucker, 2004; Motoyoshi and Kingdom, 2010). Further work on the statistics of grouped edges are also needed to better understand the role co-circularity plays in contour detection and object recognition.

As discussed in chapter 3, edges are the basic independent feature of natural scenes. Here we have demonstrated that natural scene edge orientations are correlated across wide distances. This correlation is a form of symmetry breaking (Lee et al., 2003; Lee and Kardar, 2006; Tang et al., 2010), imposing a relationship between edge orientation and position such that rotation of the orientation and position independently will modify the image statistics. The remainder of this thesis builds on this work to consider how this prevailing, long-range

correlation of edge arrangements affects primary visual cortex structure and function.

The best things in life are silly.

Scott Adams

5

Circles and symmetries: examining correlations in orientation preference maps

Synopsis

Visual activity after eye-opening influences feature map structure in primary visual cortex. For instance, rearing cats in an environment of stripes of only one orientation yields an over-representation of that orientation in V1. However, whether such changes also affect the higher-order statistics of orientation maps is unknown. Previous work has demonstrated that orientation maps in normally reared animals contain a weak statistical coupling between orientation preference and position. As we have discussed earlier, natural scenes contain a particular form of this coupling known as co-circularity; however, whether this drives the development of comparable structure in V1 is unknown. Here we examine the coupling statistics of normal, stripe-reared and dark-reared cats, and find that the orientation-position

couplings in the map are not consistently related to those present in the input, or to genetic relationships. Further, we compare these results with a simple analytic model of map patterning and demonstrate that the analytic maps do not contain edge-position couplings.

5.1   Introduction

Plasticity can be observed at multiple levels of brain organisation. In chapter 3, we showed that receptive field changes in animals reared with unusual visual input can be understood by unsupervised learning models trained on similarly modified visual input. Another more macroscopic example is the maps of the primary visual cortex of animals such as cats, monkeys and ferrets which provide a paradigm example of how brain structure can be altered to reflect that statistics of sensory input (Katz and Shatz, 1996; Sur and Leamey, 2001; Sengpiel and Kind, 2002). For instance, depriving one eye of visual stimulation during the critical period leads to a reduction in the amount of cortical territory occupied by that eye's representation of visual space, with a concomitant increase in the territory occupied by the open eye (Hubel et al., 1977; Shatz and Stryker, 1978). Similarly, raising a cat in an environment consisting mostly of stripes of only one orientation causes analogous changes in the orientation map in V1; the presented orientation expands its territory at the expense of other orientations (Blakemore and Cooper, 1970; Sengpiel et al., 1999; Tanaka et al., 2006). Complete removal of the map of one feature can also lead to changes in the periodicities and interrelationships of the remaining maps (Farley et al., 2007; Swindale, 2007).

In chapter 4 we found, building on previous work (Sigman et al., 2001), that edges in natural scenes are co-circular even at large distance scales. We wondered if such statistics are seen in the patterning of visual maps. As just discussed, V1 maps are plastic to large scale changes in input statistics. Additionally, in chapter 4, we showed that receptive fields learned new encodings from perturbations in visual input. Primary visual cortex simple cells have well-defined orientation preferences overlaid on a visual topographic map. This made V1 maps an obvious candidate neural substrate that we hypothesised would be arranged to reflect to prevalence of co-circularity in visual input. Some previous work had hinted that this might be the case.

By analysing orientation maps in V1 from an individual macaque and an individual cat, Lee et al. (2003) discovered a higher-order statistical regularity of these maps. This was a weak, but systematic dependence between the preferred orientations at pairs of cortical points and the angle of the line in the cortex joining that pair, relative to a fixed but arbitrary set of axes (see also Lee and Kardar, 2006). Any such regularity means that map structure is not invariant to adding a fixed angle to the orientation preference at each position in the map, since this would change the relationship between orientation preferences and directions in the cortex. As invariance to addition of a fixed angle is a type of symmetry that orientation maps do not possess, Lee et al. (2003) referred to these maps as having "reduced symmetry" (figure 5.1A–C). Similar results have been reported for tree shrews (Schnabel, 2008). However, this does not demonstrate that such reduced symmetry is driven by natural input. The reduced symmetry found in natural scenes, co-circularity, is a specific type of reduced symmetry (figure 5.1F).

Establishing that co-circularity exists in a cortical orientation map is harder than just showing reduced symmetry. To show reduced symmetry one only needs to demonstrate some dependence between orientation difference α and relative position β (figure 5.1). To prove co-circularity is present it is necessary to also know the relationship between cortical directions and directions in visual space, since the latter is where orientation preferences are specified. Lee et al. (2003) did not have this information for the single macaque and cat map they analysed, so could only show that the maps had reduced symmetry. They postulated that the source of the reduced symmetry in the maps may be the reduced symmetry from co-circularity in visual input, but since they could not establish that the reduced symmetry in the map had the form of co-circularity, the question of whether co-circularity exists in orientation maps was not clearly answered.

In addition to our hypothesis that orientation maps contain co-circularity, we also hypothesised that the reduced symmetry seen in the maps was driven by natural scene input during development. This is a distinct, but related hypothesis; neural properties can also be selected for evolutionarily and genetically encoded. To address this second hypothesis we compared orientation map statistics from the three rearing paradigms.

Before outlining our findings, we need to introduce a slightly more general property than

co-circularity which is important for this chapter. We will refer to it as "co-circularity with offset" (figure 5.1E,F). By this we mean that the strongest relationship between preferred orientations at different cortical positions is that given by co-circularity with the addition of a specific offset angle we will call τ. The case of $\tau = 0$ corresponds to co-circularity; however, determining the actual value of τ for a specific map requires knowledge of the transformation between visual field and cortical coordinates.

In this chapter we present results suggesting that both of our hypotheses are incorrect. In particular, using our generalised, quantitative measure of co-circularity we compared reduced symmetry in orientation maps from six normal, three dark-reared and four stripe-reared cats. This analysis confirmed Lee et al.'s (2003) $n = 1$ result that reduced symmetry exists in cat orientation maps. However, we found that the specific type of reduced symmetry that our maps contained was co-circularity with offset ($\tau \neq 0$). In fact, τ appeared to be stochastically chosen with differing values between individuals and even between hemispheres on the same individual, which is inconsistent with the hypothesis that the statistics are driven by natural scenes, for which $\tau = 0$.

Addressing the second hypothesis, we found that despite the drastic differences in image statistics that the cats experienced, there were few consistent differences between the co-circularity statistics of the animals, and in particular no correlation in τ values, again suggesting no influence of input statistics on this aspect of the maps. We also tested whether genetic relationships affected the form of reduced symmetry by comparing the variation in co-circularity statistics between hemispheres of the same cat and found little significant genetic influence.

If the reduced symmetry in the maps was not driven by visual input another possibility was that it arose due to other constraints on map layout. A simple analytical model of V1 orientation patterning: "essentially-complex planar forms" (ECP) (Wolf, 2005) was used to generate some artificial maps to test this hypothesis. We found that the maps in the simple analytic model did not contain either co-circularity or co-circularity with offset, despite having otherwise similar statistical properties such as pinwheels. This provides evidence that the reduced symmetry in the maps does not arise from other known constraints on map layout. We discuss later possible reasons why reduced symmetry arises in the animal maps.
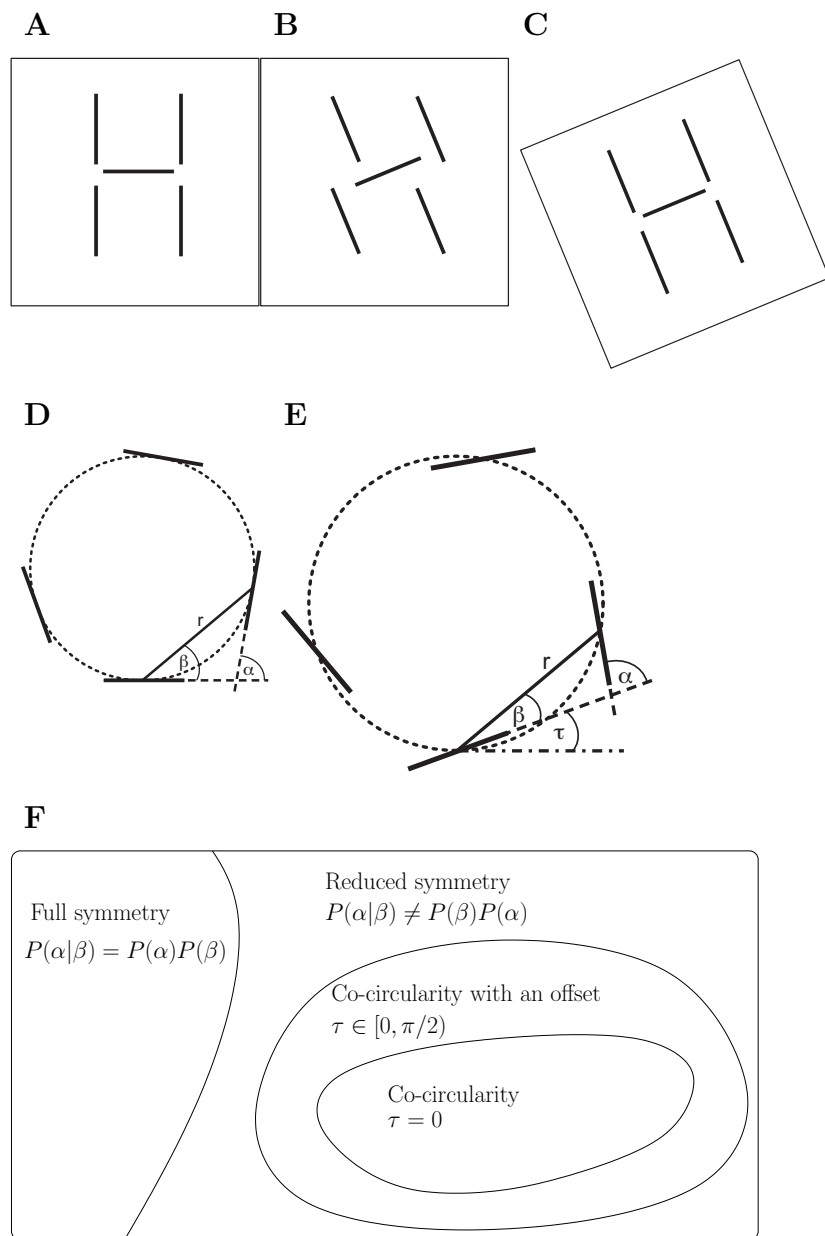
FIGURE 5.1: **A family of reduced symmetries**
A In the natural world objects can be thought of as a group of line segments whose orientations are coupled to one another and to their positions in space (the retinotopy). Following Lee et al. (2003) we refer to coupling between orientation and spatial position as reduced symmetry.
B This structure is modified if orientations are rotated independent of position.
C When the particular reduced symmetry of A is maintained, by rotating the positions of the segments together with the rotation of the individual segments, the image retains its structure (Lee et al., 2003).
D Coupling between the position and angle of oriented segments, or orientation preferences, can be described by α, the difference in orientation of the two segments, and β, the direction of one segment with respect to the orientation of the other segment. A feature of natural scenes is that segments tend to be co-circular, tangent to a common circle (dotted line) which occurs when $\beta = \alpha/2$ as shown here.
E Co-circularity with offset, when $\beta = \alpha/2 - \tau$, is a more general form of coupling between relative positions and angles of oriented segments. Here τ is $-20°$.
F A "set" diagram showing the relationship between full symmetry, reduced symmetry and co-circularity (with and without offset τ). The meaning of $P(\alpha, \beta)$ and τ is explained in the Methods section.

5.2   Materials and Methods

5.2.1   Animals and rearing

All imaging data was provided by Frank Sengpiel. All experiments were performed on cats aged 5-6 weeks. The details of the rearing and imaging methodology are provided in the associated manuscript (Hunt et al., 2009).

Orientation preference maps were created from the imaged iso-orientation maps using standard methodologies (Frostig et al., 1990; Bosking et al., 1997; Sengpiel et al., 1999). The iso-orientation maps were band-passed filtered to eliminate high-frequency noise and drifts in luminance. Polar orientation preference maps were calculated by vectorial addition of the four iso-orientation maps. Each pixel's orientation was defined by the angle of the resulting vector and the selectivity by the length of the vector. For presentation the orientation preference maps were pseudo-color coded, with the vector angle represented by the hue and the length of the vector encoded as the brightness of the colours (Bonhoeffer and Grinvald, 1993). The maps were then masked for all subsequent analysis to select only pixels clearly belonging to V1 and specifically excluding any pixels on the V1/V2 border region.

5.2.2   Retinotopic information

A complete retinotopic map for the cats in these experiments was not obtained. However, all orientation maps were from the small, exposed section of V1 on the dorsal surface of the brain, just lateral of the midline, and the retinotopic map in this area in cat is stereotypical (Tusa et al., 1978; Diao et al., 1990; Payne and Peters, 2001). In addition, the comparison of responses to two spatial frequencies ($0.5\,\text{cycles}/°$ and $0.15\,\text{cycles}/°$) allowed a functional estimate of the vertical meridian as the V1/V2 border, since the higher spatial frequency gives stronger V1 activation, and the lower spatial frequency stronger V2 activation (Bonhoeffer et al., 1995). Thus, since the approximate directions of the horizontal and vertical meridians were known, we rotated the orientation maps into alignment with the conventional visual field coordinate system. Examples of our estimated position for each map's vertical meridian are marked with white lines in Figure 5.2A–C. Rotation was performed by first separating any images

containing both hemispheres into two separate images (for some cats both hemispheres were imaged simultaneously), and then manually selecting the vertical meridian by visual inspection. Bi-linear interpolation was then used to generate the rotated map. Maps were also flipped vertically to give a right-handed coordinate system. Based on our experience with retinotopic maps obtained by optical imaging from V1 of young cats (Zepeda et al., 2003), and maps obtained with single-cell recordings (Tusa et al., 1978), we are confident that our estimates of the vertical and horizontal meridians were correct to within 10°.

While globally the retinotopic map is non-linear and anisotropic, in the small area we imaged it is on average linear and approximately isotropic (Tusa et al., 1978; Payne and Peters, 2001). Thus, we assumed that there was a linear, isotropic relationship between visual field and cortical coordinates so that, after rotation, the relative positions of cells in cortical coordinates were similar to the relative positions of their receptive fields within a constant scaling factor. This simplification has been used in previous studies (Lee et al., 2003; Lee and Kardar, 2006; Schnabel et al., 2007a).

However, in order to test whether any uncorrected anisotropy or jitter in the retinotopic map could have a major effect on our conclusions regarding co-circularity, we also examined the sensitivity of our analyses to the addition of these influences. Anisotropy was introduced by globally stretching or contracting the cortical maps by a factor ρ, which is equivalent to stretching or contracting the visual field representation by a reciprocal amount, and co-circularity statistics were then recalculated. To examine the effect of random jitter in receptive field position we added noise to the calculated position of each pixel (uniform noise $\mathcal{U}(-1, 1)$ pixels).

In general, overall map topography is determined by a combination of Eph/ephrin gradients and spontaneous waves of retinal activity (Cang et al., 2005a,b, 2008). Since there is no evidence that either of these sources of retinotopic information are disrupted in dark- or stripe-rearing, we assumed that overall map topography in these cases was the same as for the normal case. Additional post-hoc support for this assumption comes from the fact that we found no significant changes in co-circularity between the different rearing conditions.

5.2.3 Co-occurrence histograms for orientation maps

As in the previous chapter, we use the following notation to define the relative orientation and position of two pixels, except now orientation refers to the V1 orientation preference:

θ_i orientation preference of point i;

θ_j orientation preference of point j;

s_i orientation selectivity of point i;

s_j orientation selectivity of point j;

r cortical distance between points i and j;

ϕ angle of the line from point i to point j;

$\alpha = \theta_j - \theta_i$ relative orientation preference of the two points;

$\beta = \phi - \theta_i$ angular position of the second point relative to the orientation preference of the first point.

The orientation preferences θ_i and θ_j are defined in real space where the stimulus is presented, relative to zero degrees along a horizontal x-axis. The selectivity of each point is between $[0, 1]$ where 1 is a completely selective point and 0 implies lack of any orientation preference. The relative position of the cortical points ϕ is defined relative to the horizontal axis of the retinotopic map.

Reduced symmetry in natural images and orientation maps has previously been demonstrated by calculating joint histograms of the frequency at which different relative orientations α occur at different relative positions (r, β) (Sigman et al., 2001; Geisler et al., 2001; Lee et al., 2003; Lee and Kardar, 2006). We calculated a similar distribution $P(\beta, \alpha|r)$ for both experimental and simulated maps. A systematic relationship between α and β (or more precisely, that $P(\beta, \alpha|r) \neq P(\beta|r)P(\alpha|r)$) at a particular length scale r indicates reduced symmetry.

All distances on the map were normalized by the wavelength λ of the map. As there is variation in column spacing between individuals (Kaschube et al., 2002), and the imaging of each map may be at different magnifications, the use of λ as the unit of length aims to ensure that the quantification of reduced symmetry is not affected by overall changes in map scaling. To calculate the mean wavelength for each map we first calculated the Fourier

power spectrum. For the normal and stripe-reared cases, most of the power was present in a narrow wavelength band, and we therefore calculated the mean wavelength as the mean of the Fourier spectrum averaged over all directions as previously described (Carreira-Perpiñán and Goodhill, 2004). For the dark-reared cases, the power in the Fourier spectrum was much more diffusely distributed over wavelengths, making it harder to characterise these maps in terms of a single wavelength. While we used the same averaging method to determine a wavelength, we note that wavelength calculations using any method will in general be less reliable for the dark-reared case.

For every ordered pair of cortical points we measured r, ϕ, α and β. We then binned ϕ, α and β at intervals of $10°$, restricting α and β to the interval $[-90°, 90°)$ because they were periodic. We binned r at intervals centered every $\lambda/4$, where λ was the wavelength of the map. The first bin, referred to hereafter as $r = 0$, contained pairs where $0 \leq r < 1/8\lambda$, the $r = 1/4\lambda$ bin contained the pairs where $1/8\lambda \leq r < 3/8\lambda$, and so on. The pairs were weighted according to their selectivity $s_i s_j$. From this we obtained $H(r, \phi, \beta, \alpha)$, the joint probability distribution describing how often each selectivity-weighted combination of r, ϕ, α and β occurs over all pairs of points.

Next, we noted that there is a systematic non-uniformity in the spatial distribution of pairs

$$C(r, \phi) = \sum_{\beta} \sum_{\alpha} H(r, \phi, \beta, \alpha) \tag{5.1}$$

arising from the finite size of our maps, and the presence of more moderately separated than closely separated pairs. Since it has the potential to cause spurious correlations in $P(r, \beta, \alpha)$, a correction was made for this non-uniformity by normalizing H:

$$P_1(r, \phi, \beta, \alpha) = \begin{cases} \frac{H(r, \phi, \beta, \alpha)}{C(r, \phi)} & \text{for } C(r, \phi) \neq 0 \\ 0 & \text{for } C(r, \phi) = 0. \end{cases} \tag{5.2}$$

We then marginalized over the unwanted ϕ-dimension:

$$P_2(r, \beta, \alpha) = \sum_{\phi} P_1(r, \phi, \beta, \alpha) \tag{5.3}$$

and normalized to calculate the final probability distribution of interest:

$$P(\beta, \alpha | r) = \frac{P_2(r, \beta, \alpha)}{\sum_\beta \sum_\alpha P_2(r, \beta, \alpha)}. \tag{5.4}$$

The procedure outlined above is slightly different from that used previously by Lee et al. (2003) and Lee and Kardar (2006), since we normalised for non-uniform $C(r, \phi)$ and restricted α, β and ϕ to the interval $[-90°, 90°)$ instead of $[0, 90°)$.

In natural scenes, edges are co-circular more often than should occur by chance. In this notation, edges are co-circular when $\beta = \alpha/2$ and $\beta = \alpha/2 \pm 90°$ or α and β. Hence, if the reduced symmetry in V1 orientation maps is caused by reduced symmetry in natural images then we might expect ridges in the $P(\beta, \alpha | r)$ distribution along $\beta = \alpha/2$ and $\beta = \alpha/2 \pm 90°$. If the directions of the retinotopic axes in the cortex are not known, then $P(\beta, \alpha | r)$ can still be calculated, as was done in Lee et al. (2003) and Lee and Kardar (2006). However, there may be a rotation between the axis from which ϕ is measured, and the axis from which θ_i and θ_j are measured. A rotation of τ between the two coordinate systems would cause $P(\beta, \alpha | r)$ to be shifted by τ along the β axis, and thus correlations would now be expected at $\beta = \alpha/2 - \tau$. The presence of these correlations is the definition of co-circularity with offset (figure 5.1E). In a following section we describe a method for finding the rotation τ_{best} which maximises the correlations at $\beta = \alpha/2 - \tau_{\text{best}}$.

5.2.4   Measuring co-circularity in orientation maps

There are in principle many different types of reduced symmetry, i.e. $P(\beta, \alpha | r) \neq P(\beta | r) P(\alpha | r)$. We found it useful to restrict our quantitative analysis to an investigation of whether the reduced symmetry has the specific form of co-circularity (i.e. $\beta = \alpha/2$), or co-circularity with offset (i.e. $\beta = \alpha/2 - \tau$), where τ is the angle of rotation of orientation needed to produce exact co-circularity, see Figure 5.1E,F).

While $P(\beta, \alpha | r)$ for a particular map is a useful method of visualizing directly whether α depends on β for a given separation r, it is hard to compare these multidimensional distributions between conditions. We used the degree of co-circularity as a scalar measure

to allow direct comparison. As before, this was defined as:

$$D_{\text{diff}}(r) \quad = \quad \frac{\sum_{\text{pairs}} s_i s_j |\theta_i + \theta_j - 2\phi|}{\sum_{\text{pairs}} s_i s_j} \tag{5.5}$$

$D_{\text{diff}}(r)$ has the intuitively appealing property that it is 45° when no co-circularity is present (i.e. when the orientation preferences at all pairs of locations have no systematic dependence) and decreases if the relationship between the angles at the two positions becomes more co-circular. To calculate $D_{\text{diff}}(r)$ we binned r in intervals of $\lambda/4$ as for the co-occurrence histograms. The first bin, referred to hereafter as $r = 0$, contained pairs where $0 \leq r < 1/8\lambda$, the $r = 1/4\lambda$ bin contained the pairs where $1/8\lambda \leq r < 3/8\lambda$, and so on. This bin spacing balances the requirements of having a reasonable number of pairs in each bin to reduce noise, while allowing fine enough resolution to investigate trends with r.

5.2.5 Calculating the co-circularity with offset

We also wanted to calculate the co-circularity with offset τ, which is a more general type of reduced symmetry. One way this type of reduced symmetry can arise is if co-circular edges are rotated without rotating their position. We calculated a more general form of D_{diff} which incorporates an edge rotation τ.

$$D_{\text{diff}}(r, \tau) \quad = \quad \frac{\sum_{\text{pairs}} s_i s_j |\theta_i + \theta_j - 2\phi - 2\tau|}{\sum_{\text{pairs}} s_i s_j} \tag{5.6}$$

In order to determine the offset τ_{best} at which reduced symmetry is most strongly represented, we calculated $D_{\text{diff}}(r, \tau)$ for $\tau = 0$ to 85° in 5° increments. Since line elements are co-circular if they obey the relation $\beta = \alpha/2$ or $\beta = \alpha/2 \pm 90°$, searching over this range covers all possible values for τ. The rotation that gave the best match, τ_{best}, was taken to be that which gave the minimum of $\int_{r=0.75\lambda}^{1.25\lambda} D_{\text{diff}}(r, \tau) dr$ (i.e. maximum reduced symmetry at a cortical separation of one wavelength). The range of integration was chosen because we found that reduced symmetry is maximal over this range. For each map we denote $D_{\text{diff}}(r, \tau_{\text{best}})$ as $D_{\text{best}}(r)$.

5.2.6   Statistics

Since the source of reduced symmetry statistics in the maps was the object of investigation, it was not clear a priori if these were likely to be normally distributed. Therefore, all statistical tests were performed using permutation tests similar to those in Kaschube et al. (2002). The only assumption these tests make about the data is that samples are independent. No knowledge of the source distribution is assumed.

In order to test for statistically significant differences in $D_{\mathrm{diff}}(r)$ or $D_{\mathrm{best}}(r)$ between sets of maps, the values at different r were treated independently. Given two sets of N measurements at a particular r, $\{D_i\}$ and $\{D_j\}$ (e.g. $D_{\mathrm{diff}}(r = \lambda)$ for normal and dark-reared cats), the difference in the mean values between the two sets was calculated:

$$\Delta Q = \frac{\sum_i D_i}{N_i} - \frac{\sum_j D_j}{N_j} \tag{5.7}$$

In order to determine if the difference in mean values between the sets was significant we calculated the likelihood of this difference in means arising if both sets are sampled from the same original source. This was done by generating pseudo-random permutations between the sets and examining how probable such a difference in means is for a randomly sampled permutation. If both sets are sampled from the same original source then randomly permuting the sets should give a similar difference in mean values. If, however, the sets are sampled from different sources then the permuted sets should tend to have a smaller difference in means. This is calculated as

$$p = < H(|\Delta Q| - |\Delta Q_t|) >_t \tag{5.8}$$

where $H(.)$ is the Heaviside function (i.e $H(x) = 0$ for $x < 0$, $H(x) = 1$ for $x \geq 0$), ΔQ_t is the difference in means between the permuted sets t, and $< . >_t$ is the ensemble average of pseudo-randomly generated permutations between the sets D_i and D_j. Permutations are created by grouping all the values of both sets together and then pseudo-randomly assigning each value to two pseudo-sets of equivalent sizes as the original sets. The use of the Heaviside function means this equation returns the fraction of permutations which

have a mean difference greater than that of the original sets. Thus, a low value indicates the original mean difference was significant. p-values were calculated using 10^5 randomly generated permutations.

If reduced symmetry in the maps arose from the co-circularity of natural scenes, one would expect $D_{\text{best}}(r = \lambda)$ for each map to occur when $\tau_{\text{best}} = 0$ (no rotation required to bring the $P(\beta, \alpha | r)$ distributions into alignment with co-circularity). A systematic error in orienting retinotopic axes in the cortex would result in τ_{best} clustering around another angle. A modified permutation test was therefore used to detect any significant clustering of τ_{best}. Given a set of τ_{best} values, the mean difference between all possible pairs i, j was calculated

$$\Delta M = \frac{\sum_{i,j} \min(\tau_{\text{best}}^i - \tau_{\text{best}}^j, \tau_{\text{best}}^j - \tau_{\text{best}}^i)}{N_{\text{pairs}}} \tag{5.9}$$

where the subtractions between τ values were calculated modulo $90°$. This arrangement corrects for the $90°$ symmetry in τ when calculating the mean difference. In order to determine if this mean difference indicates any significant grouping the probability of observing such a grouping based on sampling from a uniform random spread of probabilities across all angles was calculated as

$$p = < H(\Delta M - \Delta M_t) >_t \tag{5.10}$$

where $< . >_t$ is the ensemble average across pseudo-random samples of sets of τ values the same size as the set of interest, but randomly distributed across all angles. If the grouping of the sample is significant then the pseudo-set mean difference ΔM_t will tend to be greater than the sample set mean difference ΔM. Low p-values indicate statistically significant clustering. p-values were calculated using 10^5 permutations.

Similarly, permutation tests were used to test the significance of hemispheric grouping. To test grouping of a value between hemispheres given two ordered sets of values A and B (where these values might be D_{diff} or D_{best} at a particular value of r or τ_{best}) where corresponding entries in each set are pairs of values from the same cat, the mean difference

between pairs was calculated as

$$\Delta M = \frac{1}{N} \sum_i |A_i - B_i| \tag{5.11}$$

where N is the number of pairs. For τ_{best} the modulo 90° was used when calculating the mean difference. The significance of the grouping was then calculated as the probability of such a difference arising randomly. This is given by

$$p = < H(\Delta M - \Delta M_t) >_t \tag{5.12}$$

where $< . >_t$ is the ensemble average across pseudo-random pairing generated by permuting set B.

Essentially-complex planar form (ECP) maps

Essentially-complex planar forms are simple analytical models of orientation map layout that describe the attractors of time-evolving fields under particular symmetry considerations (Wolf, 2005) (others have also noted that orientation maps can be described simply in Fourier space (Niebur and Worgotter, 1994)). Here we use them as a simple way of generating maps without consideration of training input. As commonly done (Swindale, 1996), we represent orientation maps as complex fields $z(\mathbf{x})$, where the orientation preference at position $\mathbf{x}$ is given by $\arg(z(\mathbf{x}))/2$ and the selectivity $|z(\mathbf{x})|$. In this representation, ECP maps are defined by a sum of wavevectors as follows:

$$z(\mathbf{x}) = \sum_j^n e^{il_j \mathbf{k}_j x} \tag{5.13}$$

where the n wavevectors $\mathbf{k}_j$ are distributed equidistantly on the upper half of the circle and

$$\mathbf{k}_j = k_c(\cos(j\pi/n), \sin(j\pi/n)). \tag{5.14}$$

The binary variables $l_j = \pm 1$ determine whether the wavevector or its negative is used. We used $n = 15$ to generate the ECP maps in the paper. The values for l_j were assigned

randomly. It can be shown by symmetry considerations that for $n = 15$ there are 612 different ECP maps (Wolf, 2005), although we only generate 40 distinct maps.

For comparison with the experimental maps, the ECP map wavelengths were scaled by adjusting to the length of the wavevectors k_c to generate wavelengths close to the experimental maps. The maps were truncated, after adding a 100 pixel offset, to 128×128 pixels, a similar size to the cat maps.

5.3 Results

5.3.1 Experimental maps show reduced symmetry but not co-circularity, and reduced symmetry is not affected by rearing condition

We analysed data obtained from thirteen individual cats. For some cats both hemispheres were imaged so that a total of eleven normal, seven stripe-reared and six dark-reared hemispheres were available. Some animals used in this study were litter-mates. For this analysis we treated the hemispheres of the animals as independent (see later for justification). Figure 5.2A–C shows a sample orientation preference map from each of these three experimental conditions.

For each cat we calculated the probability that different relative orientations α occur at different relative positions (r, β) (see Figure 5.1D and Methods for angle definitions). Figure 5.2D–I shows example distributions for each experimental condition. These joint probability distributions provide a visualisation of any reduced symmetry; in particular, any systematic relationship between α and β (the presence of which is the definition of reduced symmetry) is apparent in the shape of these plots. Co-circularity with offset τ would be visible as diagonal ridges, spaced 90° apart, running along the lines $\beta = \alpha/2 - \tau$ and $\beta = \alpha/2 - \tau \pm 90°$. The smoothness of orientation maps meant that at close range ($r < \lambda/8$, where λ is the wavelength of the orientation map), α was independent of β, indicating an absence of reduced symmetry (figure 5.2D–F). This independence can be seen in the figures; every slice along α (constant β) is identical to the other slices (modulo noise). For points
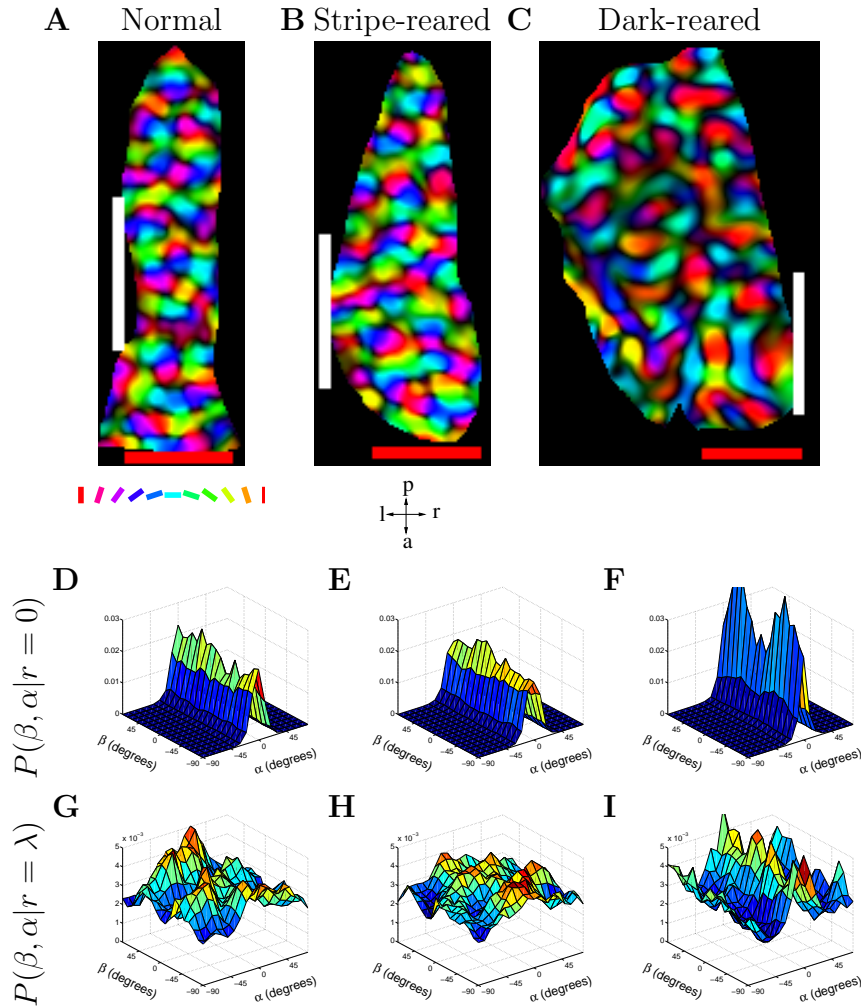
FIGURE 5.2: **Reduced symmetry in normal, stripe-reared and dark-reared cats**
Representative examples of cat orientation preference maps and joint probability distributions
$P(\beta, \alpha | r)$ at $r/\lambda = \{0, 1\}$ for the three rearing conditions. Orientation maps are polar; the colour
encodes the orientation preference according to the key under A and the brightness encodes the
orientation selectivity. Our estimate of the position of each map's vertical meridian, which was
used to orientate the maps retinotopically, is marked with a white line. The red scale bars are 1
mm. These maps have been flipped vertically after rotation to give them a right-handed coordinate
system and then masked to eliminate points outside V1. Only the unmasked points were used in
all subsequent analysis.
(**A,D,G**) Representative example from a normal cat. **A** Orientation preference map. **D** $P(\beta, \alpha | r =
0)$. At short distances there is little β-dependence in the probability distribution, with almost all
the probability near $\alpha = 0$, indicating that the orientation preferences of pixels close to one another
tend to be similar. **G** $P(\beta, \alpha | r = \lambda)$ demonstrates a β-dependence in the probability distribution,
indicating that reduced symmetry exists at separations of approximately λ.
(**B,E,H**) Representative example from a stripe-reared cat. This cat was reared in an environment
rich in 135° lines. **B** Orientation preference map. **E** $P(\beta, \alpha | r = 0)$. **H** Ridges in $P(\beta, \alpha | r = \lambda)$
indicate the presence of co-circularity with offset. While the proportions of the map preferring each
orientation changes under stripe-rearing conditions, the joint probability distributions $P(\beta, \alpha | r)$ did
not demonstrate qualitatively different properties to normally reared cats.
(**C,F,I**) Representative example from a dark-reared cat. **C** Orientation preference map. **F**
$P(\beta, \alpha | r = 0)$. **I** $P(\beta, \alpha | r = \lambda)$. Despite the maps of dark-reared cats showing visually obvi-
ous differences, such as less clearly defined pinwheels and a "blobbier" structure than those from
normally reared cats, they also demonstrated reduced symmetry at separations of approximately
λ.

separated by around one wavelength, we found reduced symmetry of varying strength in individual maps from all three experimental conditions (figure 5.2G–I). This can be seen in the figures; for each slice along α, the shape depends on the choice of β, so α and β are not independent. However, we did not find the high probability ridges at $\beta = \alpha/2$ and $\beta = \alpha/2 \pm 90°$ that are the signature of co-circularity. The absence of co-circularity was confirmed by examining the measure $D_{\text{diff}}(r)$ for each rearing condition (figure 5.3A). For each map we generated a control map by pseudo-randomly permuting the positions of the pixels. This destroyed any reduced symmetry and ensured that for these maps, α and β were independent. We used these control maps to test for significance of $D_{\text{diff}}(r)$. No rearing condition had a $D_{\text{diff}}(r)$ significantly different from control near $r = 1\lambda$. In addition, there was no significant differences between rearing conditions except at $r = 0.5$ for normal compared to dark-reared and $r = 2.5$ for stripe-reared compared to dark-reared.

To confirm that a fixed rotation of cortical axes relative to visual field axes (as would arise from, for instance, a systematic error in our estimate of the vertical meridian in the cortex) could not produce significant co-circularity, we found the offset angle for each map, τ_{best}, which maximised our measure of co-circularity, $D_{\text{diff}}(r)$, between $r = 0.75\lambda$ and $r = 1.25\lambda$. The angle of rotation required for each individual map is plotted in Figure 5.3B,D. The rotation needed to match the co-circularity of natural images was not clustered at a constant value for all maps in any one experimental condition, or overall. This indicates that, while our experimental data display co-circularity with offset, there was no consistent offset angle that might arise from, for example, a systematic error in our estimate of visual field axes in the cortex.

Figure 5.3C shows $D_{\text{best}}(r)$, the form of $D_{\text{diff}}(r)$ after rotation by τ_{best}, for the different rearing conditions. All the maps demonstrated co-circularity with offset as shown by the significant dip in each rearing condition near $r = \lambda$. The significance of this dip was established by comparing each rearing condition with the permuted control maps. Although there were significant differences in $D_{\text{best}}(r)$ between dark-reared cats and the other rearing conditions at some r values, the deviation from 45° was approximately the same; we believe these differences are due to the difficulty in calculating the average wavelength of maps for the dark-reared cats.
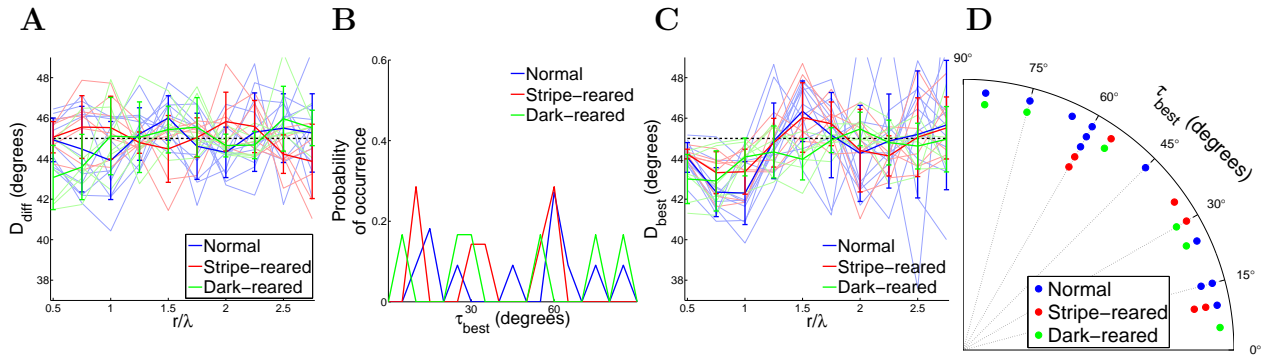
FIGURE 5.3: **Analysis of reduced symmetry in experimental maps**
A Co-circularity strength $D_{\text{diff}}(r)$ calculated for varying pair separations for the cat orientation preference maps under different rearing conditions. Thin lines are the results from individual cats; thick lines are the means for the different rearing conditions. The dashed black line at $45°$ shows the expected mean value when no co-circularity is present. There is no significant difference in the mean $D_{\text{diff}}(r)$ for any r between any rearing condition and their randomly permuted controls except at $r = 1.5\lambda$ ($p = 0.01$) for normal maps, $r = 2.5\lambda$ ($p = 0.05$) for stripe-reared and $r = 0.5, 2.5$ ($p = 0.01$) for dark-reared. This indicates that no rearing condition shows significant co-circularity at $r = \lambda$, where it would be expected if the reduced symmetry in the maps was driven by co-circularity in natural scene input. Between normal and stripe-rearing, the only positions displaying significant difference in D_{diff} is $r = 1.5, 2.0\lambda$ ($p = 0.03$). There are significant differences between normal and dark-reared cats at $r = 0.5$ ($p = 0.01$) and stripe-reared and dark-reared cats $r = 2.5$ ($p = 0.01$).
B Density plot showing the distribution of τ_{best}, the angle of rotation of each orientation map required to bring it into maximal agreement with co-circularity for $r = \lambda$. There is no significant grouping of τ_{best}, either by rearing condition or cumulatively ($p > 0.6$). This indicates that the reduced symmetry in the maps is unlikely to have arisen due to natural scene input, since this would have defined a consistent coupling between retinotopic and cortical axes. Because τ is circularly symmetric at $90°$, the ends of the x-axis can be thought of as joined and τ values near $90°$ are also close to 0.
C Maximum co-circularity strength $D_{\text{best}}(r)$ after rotation by τ_{best} for all rearing conditions. The dashed black line at $45°$ shows the expected value when no co-circularity is present. There are significant differences between normal results and control (randomly shuffled) maps at $r = 0.5, 0.75, 1, 1.5\lambda$ ($p < 0.01$), between stripe-reared and control maps at $r = 0.5, 0.75, 1, 2.25\lambda$ ($p < 0.03$), and between dark-reared and control maps at $r = 0.5, 0.75, 1, 1.5\lambda$ ($p < 0.01$). This indicates that all rearing conditions demonstrate statistically significant reduced symmetry near 1λ, but not co-circularity (as demonstrated in B). There is no significant difference in the mean $D_{\text{best}}(r)$ between normal and stripe-reared cats for any r. There are significant differences in the mean $D_{\text{best}}(r)$ between normal and dark-reared cats at $r = 0.5, 1, 1.5\lambda$ ($p < 0.04$), and between stripe- and dark-reared cats at $r = 0.5, 1.5\lambda$ ($p = 0.01, 0.02$ respectively). The differences may be due to difficulties in defining a value for λ in the dark-reared maps. Note that the minimum value for D_{best} is similar for the dark-reared cats compared to the other rearing conditions, indicating that the overall strength of reduced symmetry in the maps is not changed by rearing condition.
D The same data (τ_{best}) as in (B) but presented as a polar histogram.
All statistical tests as described in the methods section.

5.3.2 Reduced symmetry in maps from the same cat are no more similar than maps from different cats

We hypothesised that reduced symmetry in V1 orientation maps may be directly genetically encoded, as an alternative to being driven by co-circularity in natural scenes. For a total of eleven cats, maps of both hemispheres were obtained (to maximise the number of pairs for comparison, rearing conditions were pooled for this analysis). Since these paired maps are more genetically similar than unpaired maps, we tested for grouping of reduced symmetry properties between hemispheres of the same individual. The results are shown in Figure 5.4. While there was a statistically significant grouping in the strength of co-circularity between pairs at $r = 0.5, 1.75, 2.25\lambda$ (figure 5.4A), this grouping was not near a separation of one wavelength where co-circularity is strongest. The grouping may be due to the small sample size. There was no significant grouping in co-circularity with offset between pairs except at $r = 0.5\lambda$ (figure 5.4B). Similarly, there was also no statistically significant grouping in τ_{best} between hemispheres from the same cat (figure 5.4C). Overall, the results indicate that genetics have little effect on reduced symmetry at the cortical range at which it is expected to be strongest. Although some cats were littermates, there was not enough data to allow a robust comparison of these more subtle genetic relationships. However, one would expect a genetic component of co-circularity to be demonstrable by hemispheric clustering.

5.3.3 Loose ends: map distortions and constraints

Finally, we attempted to eliminate two "boring" explanations for our findings. Firstly, we know that the real topographic map is not completely smooth and uniform as we have been forced to assume. Might these topographic distortions be the reason that we do not detect co-circularity in the maps? Secondly, if the specifics of the reduced symmetry in the maps is not driven by visual input or genetically encoded, is it just a consequence of other constraints on visual maps, such as the formation of pinwheels or equal representation of orientations? We considered both of this possibilities.

Firstly, we examined the effect of jitter and anisotropy on our results. Although there is substantial experimental evidence from other studies to support our assumption of mostly
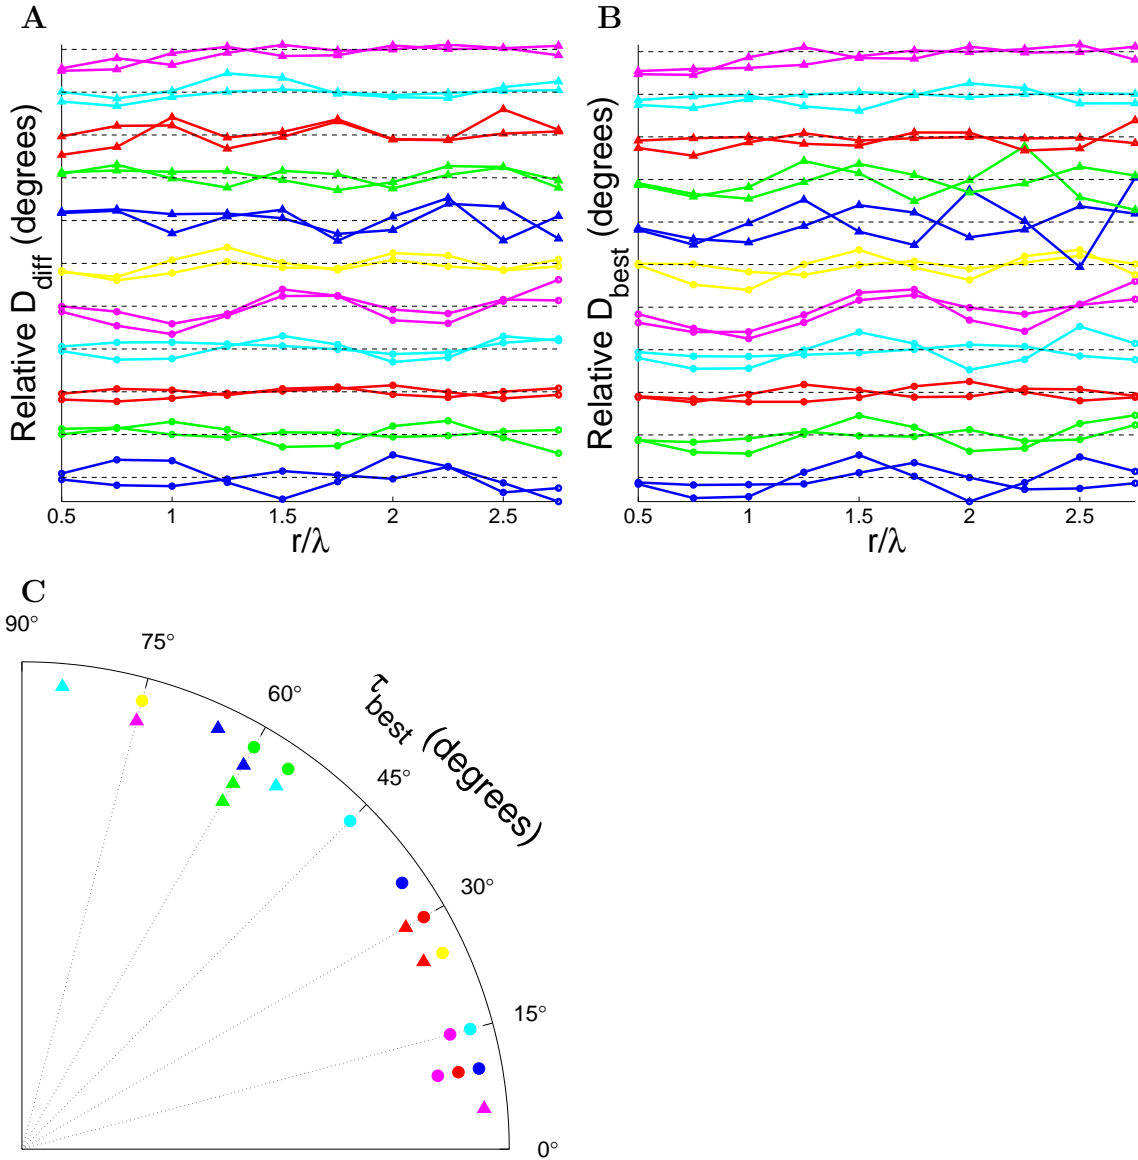
FIGURE 5.4: **Influence of genetics on reduced symmetry**
A Co-circularity strength $D_{\mathrm{diff}}(r)$ for the cat orientation preference maps classified by individual. Each pair of lines represents the results of both hemispheres from one cat. The dashed black line at 45° shows the expected mean value when no co-circularity is present. There is statistically significant grouping between pairs at $r = 0.5, 1.75, 2.25\lambda$ ($p = 0.02, 0.01$ and 0.01 respectively).
B Maximum co-circularity strength $D_{\mathrm{best}}(r)$ after rotation by τ_{best} for all hemispheres. The line colours are labeled as in A and again, the dashed black line at 45° shows the expected mean value when no co-circularity is present. There is no significant grouping effect between pairs at any r except at $r = 0.5$ ($p = 0.03$).
C Polar histogram of τ_{best}, the angle of rotation of each orientation map required to bring it into maximal agreement with co-circularity for $r = \lambda$. Each pair of symbol type and colour represents the results of both hemispheres from one cat. Although some pairs are grouped, overall there is no statistically significant grouping ($p > 0.13$).
All statistical tests as described in Methods section.

smooth and isotropic retinotopy in the small region of cortex we are analysing (see Methods and Discussion), since retinotopy was not directly measured for the animals used in the present study, we cannot definitively eliminate the possibility that there existed a global anisotropy and/or jitter in receptive field locations. To investigate whether such uncorrected anisotropy or jitter might significantly influence our calculations of co-circularity statistics we therefore introduced them explicitly into the animal maps (see Methods). Figure 5.5 shows that in neither case were the overall form of co-circularity statistics strongly affected. This argues that, even if small uncorrected distortions in retinotopy exist in our maps, this is not the reason we did not find co-circularity with zero offset.

Secondly, we examined whether reduced symmetry in the maps arises due to other geometrical constraints on maps of this size. This was done by analysing reduced symmetry in the ECP model of map development, a simple analytical model based on map symmetries. This model is not driven by specific patterns of input activity and is not Hebbian. Nonetheless, it reproduces many qualitative and quantitative features of map layout (Wolf, 2005). The maps it generates are a useful control for testing whether the co-circularity with offset is a geometric feature of all maps of this size. Figure 5.6 shows three example ECP maps and corresponding $P(\beta, \alpha|r)$ distributions. The distributions (figure 5.6G–I) and measurements of D_{best} (not shown) show minimal co-circularity with offset for the maps, as expected. This lack of co-circularity occurs despite the maps sharing many other important properties with the animal maps. This provides evidence that geometry constraints alone cannot account for the reduced symmetry we find.

5.4   Discussion

Natural images display co-circularity, a particular form of reduced symmetry (Sigman et al., 2001; Geisler et al., 2001). Previous analysis of orientation maps from a cat and a monkey showed that reduced symmetry existed in these maps (Lee et al., 2003; Lee and Kardar, 2006) and in tree shrew (Schnabel, 2008). Lee et al. (2003) proposed that co-circularity in the visual world was the likely source, but a lack of knowledge of the absolute axes of visual space in the maps they analyzed prevented these authors confirming that the

A

B
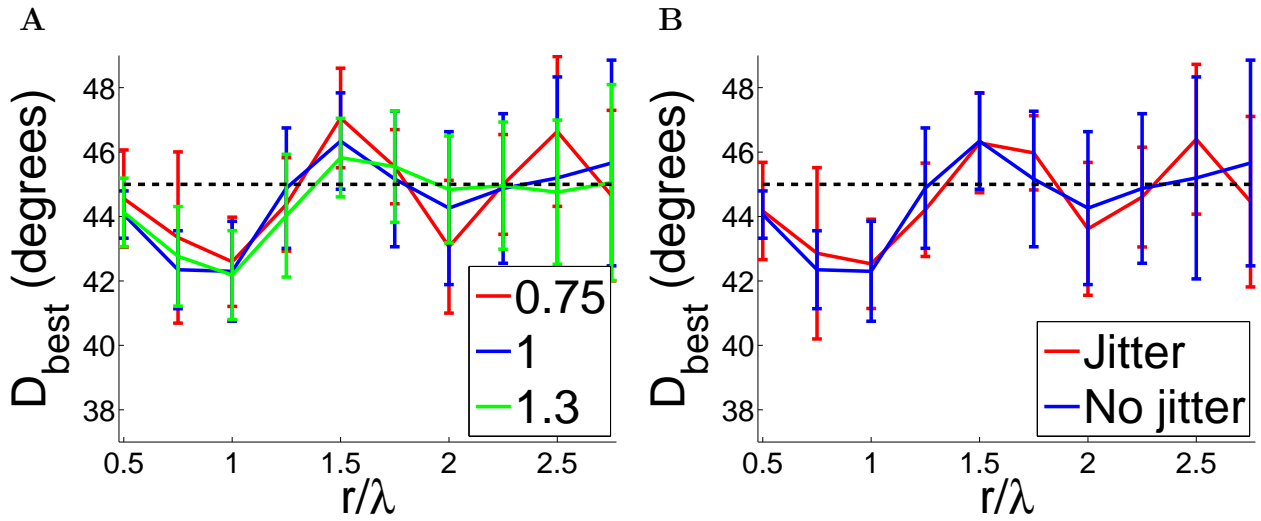
FIGURE 5.5: **The effect of anisotropy and jitter on cat map co-circularity with offset**

A We calculated co-circularity with offset D_{best} in the normal cat maps when anisotropy was added by stretching the x direction by a factor ρ. If there is existing anisotropy in the map we would expect this to correct for it. We examined $\rho = 0.75$ and $\rho = 1.3$. We found no significant changes in the strength of co-circularity with offset D_{best} or τ_{best}, indicating that our results are robust to these perturbations.
B Similarly, we added jitter consisting of uniform random noise between $[-1, 1)$ pixels, to the normal cat maps. There was no significant change to co-circularity with offset D_{best} or τ_{best}, indicating the measurements are robust to jitter.

reduced symmetry they observed was indeed co-circularity. We analysed reduced symmetry in several cat orientation maps with known retinotopic axes. We found that, while reduced symmetry of the form of co-circularity with offset ($\tau_{\text{best}} \neq 0$) was present in each case, both its strength and the angular offset required to bring it into alignment with co-circularity differed apparently randomly between individuals. We also compared the reduced symmetry in maps from normal, stripe-reared and dark-reared cats. We found no strong relationships between rearing condition and the strength or co-circularity offset τ_{best}, both of which again appeared to be randomly distributed across rearing conditions and hemispheres. These results lead us to conclude that the co-circular coupling between position and orientation in natural scenes is not the cause of coupling between these dimensions in orientation maps in cat cortex.

5.4.1 Assumptions about the retinotopic map

As in previous work (Lee et al., 2003), we have assumed that the mapping from visual space to the cortex is isotropic, smooth and continuous over the lengths we calculate correlations. We imaged only the small exposed dorsal surface of V1 approximately 5 mm $\times$ 5 mm in size, spanning both hemispheres. This area covers approximately $-5°$ to $+5°$ of elevation and from the vertical meridian to $+5°(-5°)$ of azimuth in the left(right) hemisphere (Tusa et al., 1978; Payne and Peters, 2001). However, within this area, we calculate correlations over a range of only 0 to 2.75 wavelengths, which corresponds to a cortical distance of up to 1.3 mm. We will first discuss the validity of our assumption of isotropy, followed by smoothness, over these length scales in this area of cat V1.

Evidence for isotropy in cat V1 comes from microelectrode recordings by Tusa et al. (1978) which showed that in this small central portion of the visual field, the cortical magnification factor is close to constant and isotropic. Meanwhile, evidence from primates is conflicting. Tootell et al. (1988) and Blasdel and Campbell (2001) found an anisotropic magnification factor in macaque, with the vertical magnification factor up to 1.5 times the horizontal magnification factor. However, the magnification factor in marmoset monkey is reported to be remarkably isotropic (Schiessl and McLoughlin, 2003).

Evidence for smoothness in mammalian V1 in general comes from several studies using optical imaging. In ferret, Yu et al. (2005) found a smooth map with no evidence of local distortions due to the orientation map. In tree shrew, Bosking et al. (2002) found that the spatial map was smooth down to the limit of resolution of their optical imaging technique, about 0.5° or 100 µm, and again they found no evidence for disruption of local retinotopy due to the orientation map. A similar result was found in marmoset monkeys (which have greater visual acuity) by McLoughlin et al. (2005), except that these authors reported smoothness down to a remarkable scale of 0.03°.

In contrast, Das and Gilbert (1997) suggested that cat maps are locally non-uniform and reported systematic distortions in local map retinotopy that were correlated with changes in orientation preference. However, more recent studies have found results more consistent with the data for other species described above. Using tetrode recordings, Hetherington and

Swindale (1999) showed that receptive field scatter in cats was of the order of 0.5°, and that there was no correlation between this and orientation preferences. Using both optical imaging and electrophysiological recording in cat, Buzas et al. (2003) concluded that "our findings are compatible with a locally smooth and linear representation of visual space that is not coupled to the representation of stimulus orientation." Thus, although there may be some variations between species, the weight of evidence supports our assumption of a locally smooth and continuous retinotopy in cat, at least for the small region of cortex on which our analysis is based.

However, in order to investigate the sensitivity of our results to potential uncorrected anisotropy or jitter that might be present, we added artificially generated anisotropy and jitter to the maps. We found that this did not qualitatively change the co-circularity statistics. Thus, even if anisotropy or jitter is present in the experimental maps, it is unlikely to significantly affect our results and conclusions.

Our analysis relies on an estimate of the vertical retinotopic axis, based on an assumption of the existence of a stereotypical retinotopic map in cat V1, as suggested by the literature (Payne and Peters, 2001). We found that the coupling between orientation and retinotopic maps indicated by τ_{best} was independent in different individuals. This result could only occur with co-circular coupling of the retinotopic and orientation maps if the vertical axis of the retinotopic map was randomly oriented in each hemisphere. Recordings in cat (Albus, 1975; Tusa et al., 1978; Diao et al., 1990) and other animals (Tootell et al., 1988; Schuett et al., 2002; Adams and Horton, 2003) have never found inter-individual variability to be this great. We estimate that errors of up to 10° exist in our estimations of the vertical axes, these could introduce systematic errors in each calculation of D_{best} and the consequent selection of τ_{best}. However, the distribution of τ_{best} would still be clustered around zero degrees if co-circular coupling was occurring between the retinotopic and orientation maps. This finding is robust even if there was anisotropy in the maps that was unaccounted for.

5.4.2   Genetic influences on map structure

Kaschube et al. (2002) found an activity-independent contribution to the specification of

column layout by comparing orientation maps in cats in the two hemispheres of individual brains, in littermates and in unrelated colony-mates. In particular, they found significant similarities in column size and shape in the two hemispheres of individual brains, and in littermates. However, we found that the rotation of the cortical coordinates that best aligns our experimental maps with co-circularity is not correlated between hemispheres in individuals. This suggests that, unlike low-order properties, higher-order details of column layout, such as reduced symmetry, are not influenced by genetics.

Due to both practical and ethical considerations, experiments involving rearing cats with controlled visual stimulation are generally restricted to small n-values, this limits the strength of some of our conclusions, particularly those concerning genetic influences. A related issue is the difficulty of defining wavelengths for the orientation maps in our dark-reared cats. For instance, Figure 5.3C might be interpreted as showing a subtle stretching of the length-scale of reduced symmetry in dark-reared cats compared to the normal and stripe-reared cases. However, given the ambiguity in defining an appropriate length scale over which to normalize the variation of reduced symmetry with distance for maps without a sharp peak in their Fourier spectra, in conjunction with small n values, we suggest that it would be unwise to attach strong weight to such subtle effects. While there are clearly differences between the general structure of the maps in dark-reared versus normal animals, a thorough analysis of these differences beyond their potential effects on co-circularity statistics is outside the scope of this paper.

5.4.3   Lateral connections and map stability

Our results concern only the geometric layout of orientation columns rather than patterns of lateral connections. Initial work examining how natural scene statistics influenced cortical maps was stimulated by the experimental findings in cat that, while short-range lateral connections are isotropic, long-range connections occur preferentially between co-oriented columns whose receptive fields are aligned along the axis of that orientation preference (Nelson and Frost, 1985; Bosking et al., 1997). Sigman et al. (2001) speculated that V1 may use such mechanisms to either decorrelate co-circular input, or oppositely, reinforce

responses to probable configurations. This latter postulate is consistent with a Hebb-type learning rule encoding the high probability of co-circular edge segments in natural scenes into the lateral connections of maps (Ben-Shahar and Zucker, 2004). Subsequent modelling work has included re-creating these patchy lateral connections in networks with suitable training (Choe and Miikkulainen, 2004), creating orientation maps using models that incorporate co-circular lateral connections (Lee et al., 2003; Lee and Kardar, 2006; Schnabel et al., 2007a), and analysing what the structure of visual hallucinations reveals about the symmetry of long-range intracortical connections (Bressloff et al., 2002). However, since we focused entirely on the geometric layout of columns, our results do not conflict with the evidence that lateral connections in V1 have reduced symmetry, or that their structure may be caused by the co-circularity of natural images. The next chapter is devoted to an examination of co-circularity in lateral connections.

We turned to theoretical modelling to better understand how the reduced symmetry might arise in the maps if it was not driven by visual input or genetically encoded. We used a simple analytical model, the ECP model, of map development. We found that the ECP maps did not contain significant co-circularity or co-circularity with offset.

A collaborator used a more realistic, Hebbian model of map development, the elastic net (EN) model (Durbin and Willshaw, 1987; Carreira-Perpiñán et al., 2005), to generate artificial maps trained on either co-circular or symmetric input. The details of these results are described in the paper associated with this chapter (Hunt et al., 2009). This work demonstrated that maps containing co-circularity with offset could arise even when the model was trained on symmetric input as a form of spontaneous symmetry breaking. In agreement with the animal results, such maps had a spread of τ_{best} values.

Lee et al. (2003) and Lee and Kardar (2006) (see also Schnabel et al., 2007b) also utilised theoretical modelling. They aimed to demonstrate that the coupling of orientation and position is necessary for the development of stable maps containing pinwheels. In their model, maps with isotropic lateral connections produce a stable rainbow pattern containing no pinwheels, while maps in which orientation and position are coupled through anisotropic lateral connections evolve to a stable pattern containing pinwheels. Analysis of these two types of simulated maps revealed that maps with pinwheels and anisotropic connections had

reduced symmetry, while rainbow maps with isotropic connections had full symmetry. The authors concluded that the existence of pinwheels in a stable map implies that orientation and position are coupled in the inputs. They acknowledge, though, that stable maps containing pinwheels may occur biologically without coupled inputs due to other mechanisms, such as developmental freezing. However, they did not look for reduced symmetry in a simulated map with isotropic connections that was frozen so that it still contained pinwheels.

More recently, Kaschube et al. (2010) have provided strong evidence that invariance constraints alone may explain many of the important statistics of orientation preference maps. They demonstrated that all pattern formation processes which preserve 4 postulated invariance relationships belong to the same universality set. The existence of such universality may demystify the similarities in maps across phylogenetically distant species. Kaschube et al. showed that such invariances, coupled with the assumption of inhibitory long-range interactions, results in stable solutions containing pinwheels without requiring position-orientation coupling. However, the solutions to such a universality class are ECP maps. Our results here demonstrate that, in at least one way, ECP maps differ from real maps, they lack the co-circularity with offsets we find in real maps. This may indicate that an additional invariance is needed in their model.

An issue related to map stability is the fact that we examined only maps from young cats (ages 5–6 weeks) rather than adults. It is thus conceivable that the reason we did not find co-circular structure consistent with that in natural images was simply that it had not developed yet and that adult animals would show such structure. However, imaging of orientation maps in both cat (Gödecke et al., 1997) and ferret (Chapman et al., 1996) at several developmental time-points in individual animals has shown that once the map is detectable, it is highly stable and does not change with time. Thus we think it is unlikely that co-circular structure would only emerge later in development. In addition, one of the normal reared cats was optically imaged at two time points (5 and 6 weeks) and we found no significant differences in reduced symmetry statistics between these two time points (data not shown). Furthermore, Sengpiel et al. (1999) and White et al. (2001) previously demonstrated that dark- and stripe-rearing have definite first-order effects on map structure. Therefore, it is reasonable to look for higher-order effects in these animals.

In summary, our results suggest that some statistical properties of cortical map structure are determined neither by genetics, nor by the statistics of visual input during the critical period. Rather, they may arise spontaneously, as a result of more general rules for self-organising cortical development. Consequently, these statistical properties may not have a significant functional role in adults. This is similar to the situation suggested for the cortical column (Swindale, 1990, 1998; Horton and Adams, 2005), which has much greater appeal as a candidate fundamental principle of cortical organisation. However, even without a functional role, each demonstrated statistical feature of cortical maps provides a useful constraint for developmental mechanisms.

Figure 5.6: **ECP maps**
A–C Sample ECP orientation maps.
D–F The associated $P(\beta, \alpha | r = 0)$ distributions for cortical points with small separation do not have reduced symmetry since pairs tend to be oriented similarly regardless of direction β.
G–I The $P(\beta, \alpha | r = \lambda)$ distributions for cortical points separated by one wavelength also do not show any clear co-circularity.

It is the stupidest children who are the
most childish and the stupidest
grown-ups who are the most grown-up.

CS Lewis

6

Co-circularity in lateral connections: neuronal dichotomies?

Synopsis

The statistical structure of the visual world offers many useful clues for understanding biological visual systems. One important process in visual object recognition is that of grouping together edges which belong to the same contour. As previously discussed, edges in natural scenes tend to lie on a common circle, they are "co-circular." Previous work has shown that co-circularity can predict human performance at contour grouping, something the primary visual cortex is known to be involved in. Here, we test the hypothesis that long-range excitatory lateral connections in the primary visual cortex, which are believed to be important for contour grouping, are co-circularly connected.

By analysing data from tree shrews, where information on both lateral connectivity and

the overall structure of the orientation map was available, we found a surprising diversity in the relevant statistical structure of the connections. In particular, the extent to which connections were co-circular varied significantly between injection sites.

Overall, these data suggest the intriguing possibility that V1 may contain both co-circular and anti-cocircular connections.

6.1   Introduction

Grouping edges which belong to the same object or contour is a vital part of object recognition. Biological vision systems excel at this task, yet it is still extremely challenging for artificial vision systems (Chellapilla et al., 2005). In humans, contour detection abilities develop after birth (Kovács et al., 1999), suggesting that learning from visual experience plays an important role.

As discussed in chapter 4, edge pairs in natural scenes have correlations between their relative orientations and positions. In particular, edges in natural scenes are co-circular (Sigman et al., 2001). Co-circularity is a generalisation of the well-known property of co-linearity. Co-linear edges, edges aligned along their long axis, are easier for human to group and discriminate in texture recognition (Schmidt et al., 1997; Roelfsema, 2006; Schwarzkopf and Kourtzi, 2008). Further, it has been shown that co-circularity is an important predictor of contour detection performance in human subjects (Geisler et al., 2001). The importance of co-circularity in contour grouping raises the question: what biological substrate underlies this effect?

An obvious candidate is the excitatory lateral connections in the primary visual cortex (V1). These connections are primarily found in layer 2/3 and are longer range than the inhibitory connections found predominately in other layers (Kisvárday et al., 1997). In humans, these connections develop after birth and their development coincides with improvement in contour detection (Kovács et al., 1999; Burkhalter et al., 1993; Hou et al., 2003). Further, there is compelling psychophysical evidence that they play a role in co-linear facilitation (Polat and Sagi, 1993; Cass and Spehar, 2005; Cass and Alais, 2006; Field and Hayes, 2004).

It has been previously demonstrated that long-range lateral connections connect preferentially to patches of similar orientation and primarily along the axis of orientation in V1 (Ts'o et al., 1986; Mitchison and Crick, 1982; Malach et al., 1993; Bosking et al., 1997; Stettler et al., 2002). This is consistent with their role in co-linear facilitation. Although previous work has demonstrated that the variance in some of these data is not adequately explained by co-linearity alone (Ben-Shahar and Zucker, 2004), whether there is a more general tendency towards co-circularity in these connections has not been investigated.

Here we reanalysed previously published functional anatomical data regarding lateral connections in tree-shrew (Bosking et al., 1997), using a noise-resistant measure of co-circularity that we introduced in chapter 4. Surprisingly, we found a large variation in the statistics of lateral connections between animals. One unifying explanation for both this and previous results is that lateral connections in V1 are connected in a variety of ways, both co-circularly and anti-cocircularly. We suggest reasons why this diverse strategy may be a desirable arrangement for making sense of visual scenes.

6.2 Methods

We reanalysed data from 4 tree shrews previously published by Bosking et al. (1997). The experimental methods used are described in detail in that paper and we provide only a brief overview here. Tree shrew orientation preference maps were obtained using optical imaging. Additionally, 540nm light was used to map surface blood vessels used for alignment. Biocytin was then injected into a specific site in V1 and the animal was sacrificed 16 hours later. Slices of V1 were imaged to locate the biocytin bouton and the surface blood vessels. The blood vessel information was then used to align the orientation preference maps with the bouton images, giving overlaid information on the underlying connectivity from the injection site on the animal. The original experiment used a total of 10 cases, however, due to problems recovering the archival data, we were only able to recover the data for 4 cases.

FIGURE 6.1: **Topography correction of tree shrew orientation preference maps**
A shows a schematic of the topography of the tree shrew. The area shaded grey is a compressed topography of the ipsilateral hemisphere and was excluded from analysis due to its very different magnification factor. To facilitate analysis the maps were rotated to align the vertical axis **B** and then flipped to give a right-handed co-ordinate system **C**.

6.2.1   Topography

In order to accurately quantify the co-circularity present in these results it is important to know the underlying topography. Fortunately, the topography of the tree-shrew is consistent between individuals and well characterised (Bosking et al., 2000). To align the maps (and overlaid boutons) with visual space they were first rotated so the V1/V2 border was vertical, then flipped along the vertical axis to give a right-handed co-ordinate system (shown schematically in figure 6.1). The tree-shrew V1 has a compressed representation of the ipsilateral visual field near the V1/V2 border with a very different magnification factor from the rest of V1. Therefore, we drew a line representing the edge of the contralateral visual field (the vertical meridian) and eliminated the ipsilateral portion of the map. The map was then rotated slightly to ensure the vertical meridian was represented vertically.

This gave us an orientation preference map of V1 with an aligned injection site and bouton sites which indicate the lateral projections from the injection site. The rotating and flipping ensured that points on the map corresponded to points in visual space within a scaling factor. After truncation of the ipsilateral hemisphere, the magnification factor was approximately constant for the region of V1 imaged, and cortical distance and direction

could be used as a reliable proxy for visual field distance and direction. The aligned maps with overlaid bouton and injection sites are shown in figure 6.2.

In the original analysis (Bosking et al., 1997), electrophysiological recordings taken during the injection were used to determine orientation. Originally, the results from the electrophysiological recordings could not be recovered. Therefore, we determined the orientation of the injection site using the orientation values of the underlying pixels that were measured using optical imaging. Later, the electrophysiological recordings were located and were found to differ from the optical imaging by an average of $15°$ (range $0 - 23°$).

6.2.2   Characterising lateral connectivity

To test for the presence of co-circularity in the lateral connections we used D_{diff}, the noise-resistent measure of co-circularity introduced in chapter 4. However, now we only considered pairs of orientations that were joined by the lateral connections. Analogous to previous definitions, for each pair of injection/bouton sites on the orientation preference maps we defined the following:

θ_i	orientation preference of injection site i;
θ_b	orientation preference of bouton site b;
r	cortical distance between injection site and bouton site;
ϕ	angle of the line from injection site i to bouton site b;
$\alpha = \theta_b - \theta_i \pmod{180°}$	relative orientation preference of the bouton b relative to injection site i;
$\beta = \phi - \theta_i \pmod{180°}$	angular position of the bouton site b relative to the orientation preference of the injection site i.

The orientation preferences θ_i and θ_b were defined in real space where the stimulus is presented, relative to zero degrees along a horizontal x-axis. They were always taken on the y-positive side of the horizontal and thus range from 0 to $180°$. The relative position of the cortical points ϕ was defined relative to the horizontal axis of the retinotopic map and was calculated modulo 180. Except where noted, calculations of orientations and angular positions were done using modulo $180°$ arithmetic.

The wavelength λ of each map was calculated as the mean of the Fourier spectrum averaged over all directions as previously described (Carreira-Perpiñán and Goodhill, 2004). As there is variation in column spacing between individuals (Kaschube et al., 2002), and the imaging of each map may be at different magnifications, the use of λ as the unit of length ensured that the quantification of reduced symmetry was not affected by overall changes in map scaling. The wavelength gives a rough measure of the size of iso-orientation patches in the map.

We calculated a probability distribution for the connection statistics for each animal $p(\alpha, \beta, r)$. However, as outlined in chapter 5, if ϕ is not uniformly distributed this can introduce spurious correlations. We corrected for this by calculating a 4 dimensional probability $p(\alpha, \beta, r, \phi)$ and then marginalised along ϕ. This was done by discretising both distance r and orientations α, β and ϕ and binning. Given the full distribution it was then possible to calculate the relevant marginal distributions such as $p(\alpha|\beta, r)$ and $p(\beta|\alpha, r)$.

6.2.3   Searching for co-circularity

We calculated both co-circularity $D_{\text{diff}}(r)$ and co-circularity with offset $D_{\text{diff}}(r, \tau)$ using the same methods as described in chapter 5. However, now we only considered ordered pairs of injection sites and bouton sites. To calculate $D_{\text{diff}}(r)$ we binned r in intervals of 1.5λ. The first bin, referred to hereafter as $r = 0$, contained pairs where $0 \leq r < 0.75\lambda$, the $r = 1.5\lambda$ bin contained the pairs where $0.75\lambda \leq r < 2.25\lambda$ and so on. This bin spacing balances the requirements of having a reasonable number of pairs in each bin to reduce noise, while allowing fine enough resolution to investigate trends with r. For co-circularity with offset we used bin sizes of $10°$ for τ. As before, we also calculated $\tau_{\text{best}}(r)$ for each value of r,

In addition, the distribution of co-circularity from each injection/bouton pair can provide useful information since various distributions of d_{diff}, the co-circularity of individual bouton sites would result in the same mean $D_{\text{diff}}(r)$. Therefore, in addition to calculating the average co-circularity $D_{\text{diff}}(r)$, the values of $d_{\text{diff}} = |\theta_i + \theta_b - 2\phi|_{\text{mod } 180°}$ for each pair at distance r were binned to create a histogram $H(r, d)$. The r bins were the same as used for calculating $D_{\text{diff}}(r)$ and the d_{diff} bins were spaced $15°$ apart. This allowed a basic characterisation of the

distribution of co-circularity values.

In the original analysis of the tree shrew data it was demonstrated that lateral connections are denser along the axis of injection site orientation and connect preferentially to sites of similar orientation. This is evidence that lateral connections are likely to be co-linear. In order to test whether the co-circularity we found in the horizontal connections was simply due to this co-linearity, we recalculated $D_{\text{diff}}(r)$ excluding all points lying near the axis of the injection site. This was done by excluding all injection/bouton pairs with $\beta < 30°$ or $\beta > 150°$.

6.2.4   Controls and statistics

It was important to have a reliable control to assess the significance of any co-circularity measured in the maps. We are interested in the significance of any deviation from $45°$ that occurs in $D_{\text{diff}}(r)$. One comparison was simply examining the values of $D_{\text{diff}}(r, \tau)$ at $\tau \neq 0$ since these are expected by our hypothesis to be less co-circular then with the true origin ($\tau = 0$). This means that we would expect τ_{best} to be near 0 when connections are the most co-circular.

We also created 99 control cases ($N = 99$) for each animal by adding a large amount of Gaussian noise ($\sigma^2 = 100$ pixels ≈ 1.5mm) to the bouton positions and calculated $D_{\text{diff}}(r)_j$ for each of these controls (j denotes the index of the control). The added noise is high enough to overwhelm any co-circular preferences in the connections since it is much larger than the wavelength of the maps, which means that it is larger than the size of iso-orientation patches in the maps.

A permutation test was used to test if the value $D_{\text{diff}}(r)$ for each animal differed significantly from the control values. We tested $D_{\text{diff}}(r)$ independently at each value of r. Permutation tests were used because we did not know what distribution the control values might be from and because we only had one true value to compare with the control distribution. We calculated the difference between the mean of the control values $D_{\text{diff}}(r)_j$ and the animal case $D_{\text{diff}}(r)$.

$$\Delta Q(r) = \frac{\sum_j D_{\text{diff}}(r)_j}{N} - D_{\text{diff}}(r) \tag{6.1}$$

Our null hypothesis was that $D_{\text{diff}}(r)$ was from the same distribution as the controls $D_{\text{diff}}(r)_j$. If the null hypothesis is true then the value of $D_{\text{diff}}(r)$ can be exchanged with one of the controls $D_{\text{diff}}(r)_j$ without affecting the expected value of $\Delta Q(r)$. There are $N + 1 = 100$ possible permutations under these exchanges (since we also include the original permutation). We calculated $\Delta Q(r)_i$, the value of $\Delta Q(r)$ for each of these permutations; using i as an index over all the permutations. The significance of the true $\Delta Q(r)$ was measured by finding the likelihood of this difference arising under the null hypothesis

$$p(r) = \frac{\sum_{i=1}^{N+1} H(|\Delta Q(r)_i| - |\Delta Q(r)|)}{N + 1} \tag{6.2}$$

where $H(.)$ is the Heaviside function ($H(x) = 0$ for $x < 0$, $H(x) = 1$ for $x \geq 0$). We considered $p < 0.05$ to indicate significance.

6.2.5 Natural scenes

We used the natural scene analysis from chapter 4 to compare the statistics of the lateral connectivity with the statistics of edges in natural scenes. In particular, we compared the distribution of d_{diff} values in the lateral connectivity results with the analogous calculation in natural scenes.

6.3 Results

6.3.1 Alignment and topography

Figure 6.2 shows the aligned orientation preference maps with overlaid bouton and injection sites for all animals. In each case the vertical meridian was selected using previously published results of tree shrew topography (Bosking et al., 2000) and the map rotated to make it vertical. All pixels representing receptive field positions in the ipsilateral visual field (i.e. those to the right of the vertical meridian) were removed and these sections were not used in subsequent analysis due to the very different magnification factor in this region of the cortex. For analysis, the remaining region of V1 was treated as isotropic with a constant

Animal	No.
ts9509	5411
ts9514	4974
ts9529	10708
ts9531	4821

TABLE 6.1: **Number of included boutons for each animal**

magnification factor. We have previously demonstrated that this approximation does not significantly affect our measures of co-circularity (chapter 5, Hunt et al. (2009)).

There was a varying number of boutons sampled from each animal. Boutons which were not positioned on the orientation map or on the excluded side of the vertical meridian were ignored. The number of included boutons for each animal is listed in table 6.1. Although there was a large number of boutons for each case, each bouton cannot be considered an independent measurement as many boutons are clustered together near a single region of the map.

6.3.2 Co-circularity

We quantified the co-circularity of the lateral connections for each animal separately. Figure 6.3 shows the calculated $D_{\text{diff}}(r, \tau)$ for each animal using the connections shown in figure 6.2. Each animal was treated separately because there were substantial differences between individuals. To establish whether variations in D_{diff} were significant, each animal was compared with 100 control cases generated by adding noise to the bouton positions. The noise was much larger than the wavelength of the maps (≈ 35 pixels), ensuring that boutons were moved significantly outside their original iso-orientation patch. Comparisons were made using nonparametric statistics (see Methods).

Co-circularity with offset τ is a generalization of co-circularity. Edges are co-circular when tangent to a common circle, and are co-circular with offset τ when edges occur at an orientation τ to the tangent of a common circle (so $\tau = 0$ is plain co-circularity). Co-circularity with non-zero τ is not found in natural scenes and thus we would predict that if co-circularity in the lateral connections is significant, then it should be present more strongly than co-circularity with any non-zero offset τ. We measured co-circularity with offset at a

range of non-zero offsets to test this hypothesis (figure 6.3).

These results indicate that the degree of co-circularity varies significantly between animals. However, previous work has demonstrated that lateral connections are elongated along their visual axis, indicating likely co-linearity (Bosking et al., 1997). Since co-circularity is a generalisation of co-linearity, any co-linearity present contributes significantly to the degree of co-circularity. In order to clearly establish the presence of co-circularity rather than just that of co-linearity in the lateral connection preferences, we re-calculated our results with all co-linear connections excluded. Figure 6.4 shows the same measurements as the previous figure, but without co-linear connections.

Case ts9509 is strongly anti-cocircular at longer ranges (figure 6.3) and this tendency remains even if co-linear connections are removed (figure 6.4). Case ts9514 is strongly co-circular, although, when co-linearity is excluded, this tendency becomes less strong (and it is anti-cocircular at large r). The two other cases show inconsistent results at different wavelengths, although case ts9531 shows strong co-circularity for the longest connections.

Examining the results in the 4 cases at different values of τ confirms that there is significant variation in the statistics of wiring between the 4 animals. Case 9509 has values of $\tau_{\text{best}}(r)$ which are near $45°$, while case 9514 shows $\tau_{\text{best}}(r)$ near $0°$(the other two cases show inconclusive results). These results demonstrate that the variation between co-circularity and anti-cocircularity between the cases is significant and confirms that the wiring of these connections is not simply co-circular as our initial hypothesis predicted. As we discuss later, these findings indicate that the results cannot be explained by minor variations in topographical alignment between the animals.

In order to better understand the statistics of edge arrangements we also considered some other ways of examining edge statistics. Firstly, because multiple d_{diff} distributions would result in the same mean value, we plotted the distributions of d_{diff} and compared the distributions with those found in natural scenes (figure 6.5). In natural scenes the distribution of d_{diff} has a small upward slope towards $d_{\text{diff}} = 0$. Although the slope of this distribution is small, this represents a sample across many scenes and within a small region of a scene steeper gradients are often found. For the lateral connections, the distributions appeared to be symmetric distributions centred (by definition) on the mean, D_{diff}. This

indicates that for each individual injection site the connectivity at a given wavelength tends to be centered on a particular value of co-circularity (i.e. being either predominately co-circular or anti-cocircular or neutral).

A ts9509

B ts9514

C ts9529

D ts9531

FIGURE 6.2: **Orientation preference maps with lateral connections overlaid**
The orientation preference map for all cases with the injection site (white) and bouton traces (black) overlaid. The orientation of the white injection site markers indicates the preferred orientation at these points. The boutons denote lateral connection terminations. The orientation preference maps have been rotated, flipped and truncated (see methods) to create right-handed co-ordinate systems aligned with visual space and the overlaid boutons and injection sites have been similarly transformed after alignment. Scale bar is 1mm.

FIGURE 6.3: **Quantification of co-circularity**

D_{diff} was calculated for all cases as a function of cortical distance, indicated in both millimeters and wavelengths λ. The dotted black line marks $D_{\mathrm{diff}}(r) = 45°$, deviations below this line indicate co-circularity. Significance ($p \leq 0.05$) compared with the control cases is m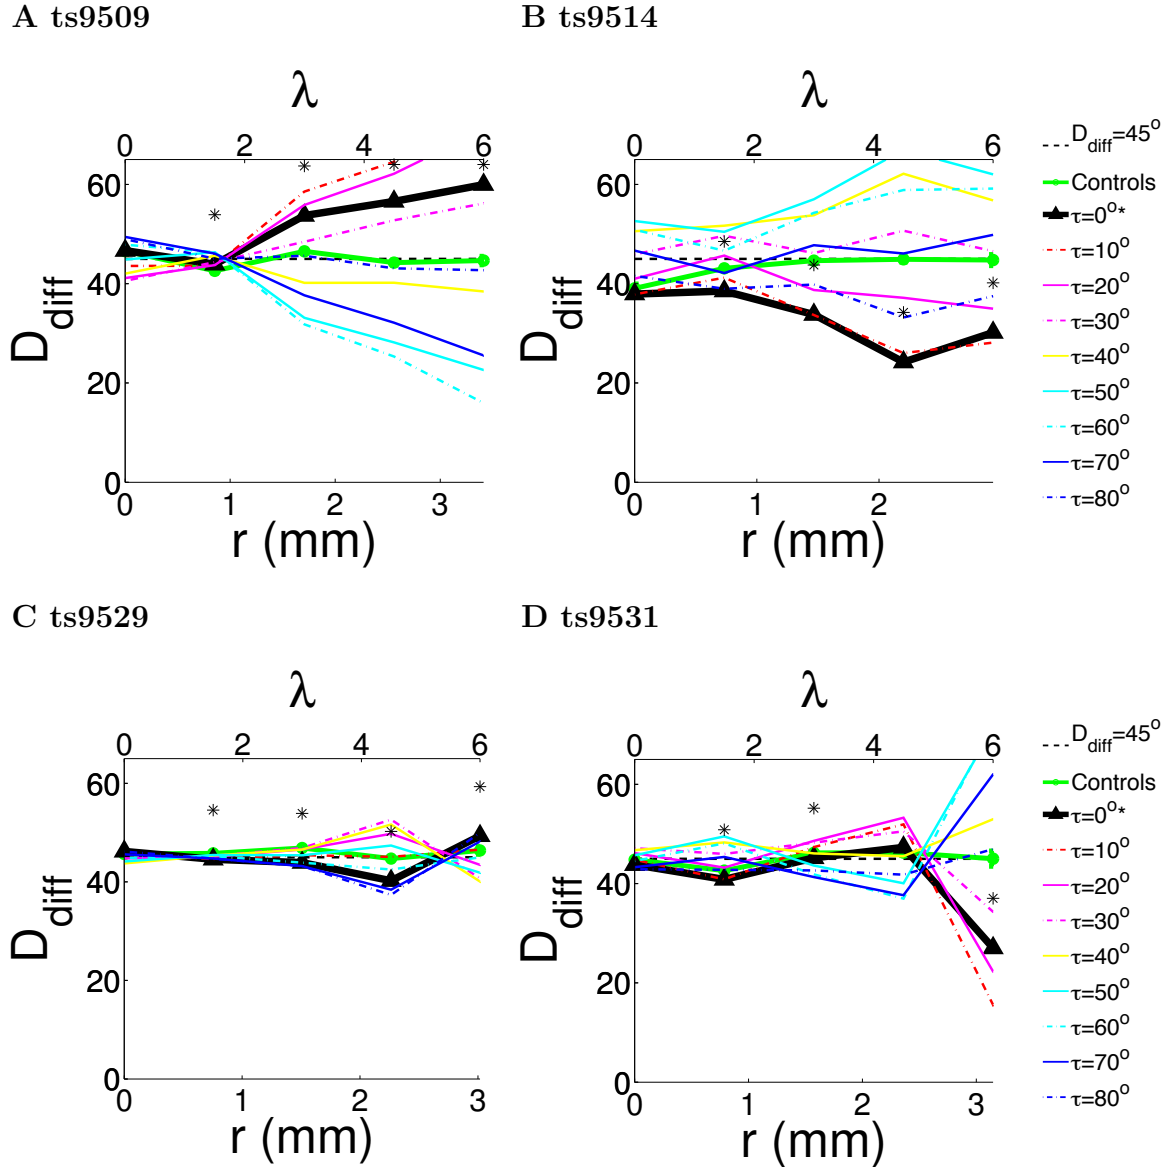arked as '*' on figure. We also report $\tau_{\mathrm{best}}(r)$ and the difference $D_{\mathrm{diff}}(r, 0) - D_{\mathrm{diff}}(r, \tau_{\mathrm{best}}(r))$.

A Case ts9509 had D_{diff} significantly different from the control case at $r = 1.5, 3, 4.5, 6\lambda(p < 0.01)$. The values of τ which minimized the D_{diff} at $r = 1.5, 3, 4.5, 6\lambda$ were $10°, 60°, 60°, 60°$ respectively, with differences $D_{\mathrm{diff}}(r, 0) - D_{\mathrm{diff}}(r, \tau_{\mathrm{best}}(r))$ of $0.2°, 21.9°, 31.2°, 44.0°$ respectively.

B Case ts9514 had D_{diff} significantly different from the control case at $r = 1.5, 3, 4.5, 6\lambda(p < 0.01)$. The values of τ which minimized the D_{diff} at $r = 1.5, 3, 4.5, 6\lambda$ were $0°, 10°, 0°, 10°$ respectively, with differences $D_{\mathrm{diff}}(r, 0) - D_{\mathrm{diff}}(r, \tau_{\mathrm{best}}(r))$ of $0°, 0.08°, 0°, 2°$ respectively.

C Case ts9529 had D_{diff} significantly different from the control case at $r = 1.5, 3, 4.5, 6\lambda(p < 0.01)$. The values of τ which minimized the D_{diff} at $r = 1.5, 3, 4.5, 6\lambda$ were $0°, 80°, 80°, 40°$ respectively, with differences $D_{\mathrm{diff}}(r, 0) - D_{\mathrm{diff}}(r, \tau_{\mathrm{best}}(r))$ of $0°, 0.7°, 2.9°, 9.3°$ respectively.

D Case ts9531 had D_{diff} significantly different from the control case at $r = 1.5, 3, 6\lambda(p < 0.04)$. The values of τ which minimized the D_{diff} at $r = 1.5, 3, 4.5, 6\lambda$ were $0°, 70°, 60°, 10°$ respectively, with differences $D_{\mathrm{diff}}(r, 0) - D_{\mathrm{diff}}(r, \tau_{\mathrm{best}}(r))$ of $0°, 3.9°, 10.3°, 11.5°$ respectively.
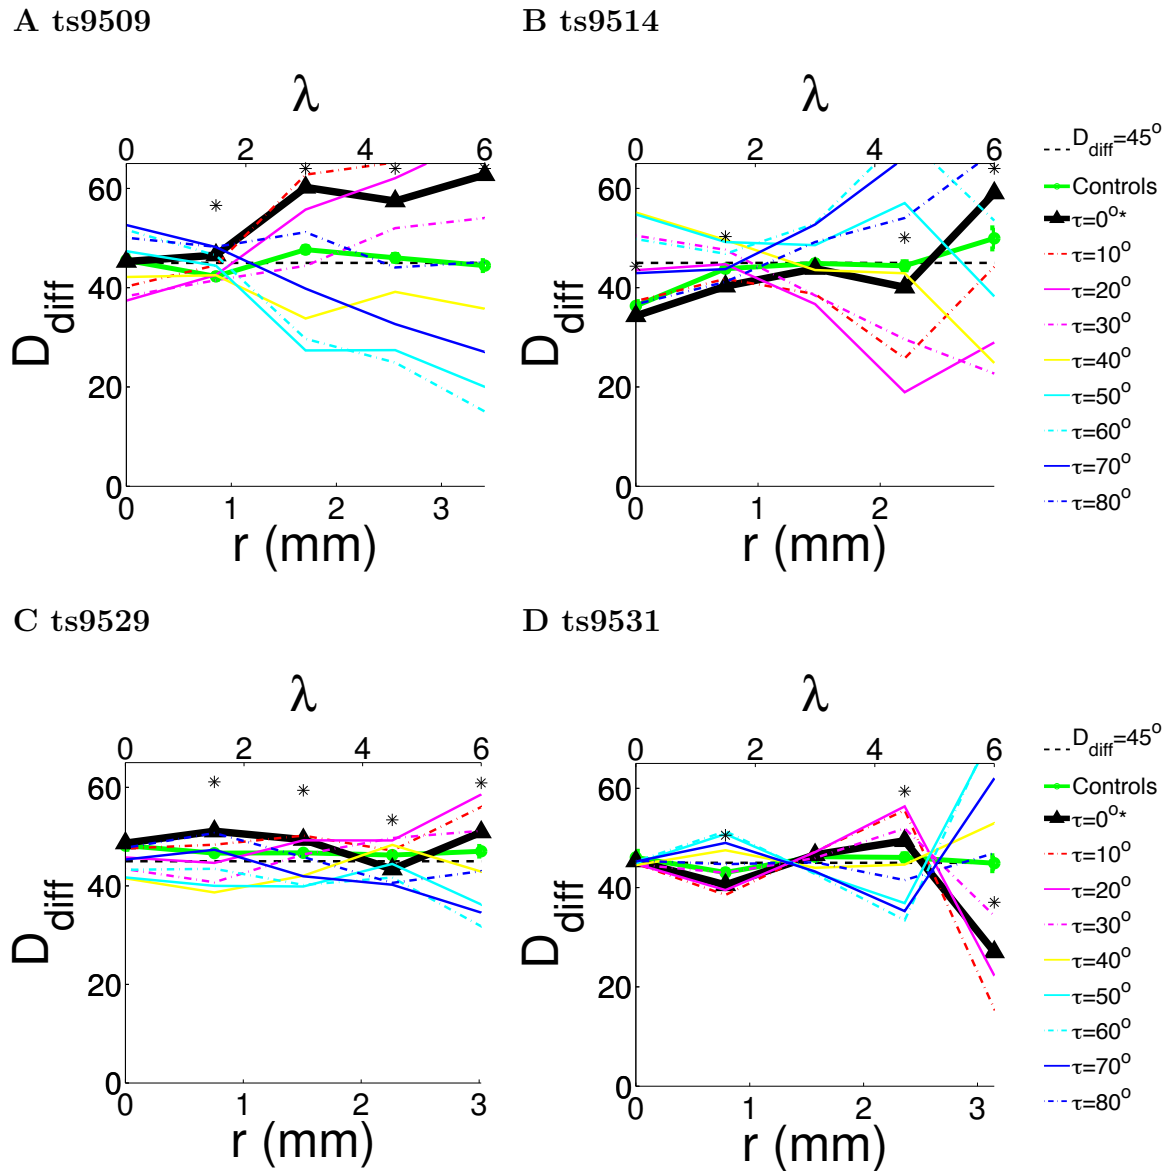
FIGURE 6.4: **Quantification of co-circularity excluding co-linearity**

$D_{\mathrm{diff}}(r)$ was calculated with all bouton sites which are co-linear with the injection site excluded $(30° \geq \beta \leq 150°)$. All cases where the true value was significantly different from controls were marked with '*'.

A Case ts9509 had D_{diff} significantly different from the control case at $r = 1.5, 3, 4.5, 6\lambda (p < 0.01)$. The values of τ which minimised the D_{diff} at $r = 1.5, 3, 4.5, 6\lambda$ were $30°, 50°, 60°, 60°$ respectively, with differences $D_{\mathrm{diff}}(r, 0) - D_{\mathrm{diff}}(r, \tau_{\mathrm{best}}(r))$ of $5°, 32.8°, 32.5°, 47.7°$ respectively.

B Case ts9514 had D_{diff} significantly different from the control case at $r = 0, 1.5, 4.5, 6\lambda (p < 0.02)$. The values of τ which minimised the D_{diff} at $r = 1.5, 3, 4.5, 6\lambda$ were $0°, 20°, 20°, 30°$ respectively, with differences $D_{\mathrm{diff}}(r, 0) - D_{\mathrm{diff}}(r, \tau_{\mathrm{best}}(r))$ of $0°, 7.1°, 21.1°, 36.4°$ respectively.

C Case ts9529 had D_{diff} significantly different from the control case at $r = 1.5, 3, 4.5\lambda (p < 0.01)$. The values of τ which minimised the D_{diff} at $r = 1.5, 3, 4.5, 6\lambda$ were $40°, 50°, 70°, 60°$ respectively, with differences $D_{\mathrm{diff}}(r, 0) - D_{\mathrm{diff}}(r, \tau_{\mathrm{best}}(r))$ of $12.4°, 9.5°, 3.1°, 19.1°$ respectively.

D Case ts9531 had D_{diff} significantly different from the control case at $r = 1.5, 4.5, 6\lambda (p < 0.02)$. The values of τ which minimised the D_{diff} at $r = 1.5, 3, 4.5, 6\lambda$ were $10°, 60°, 60°, 10°$ respectively, with differences $D_{\mathrm{diff}}(r, 0) - D_{\mathrm{diff}}(r, \tau_{\mathrm{best}}(r))$ of $2.0°, 3.9°, 16.0°, 11.5°$ respectively.
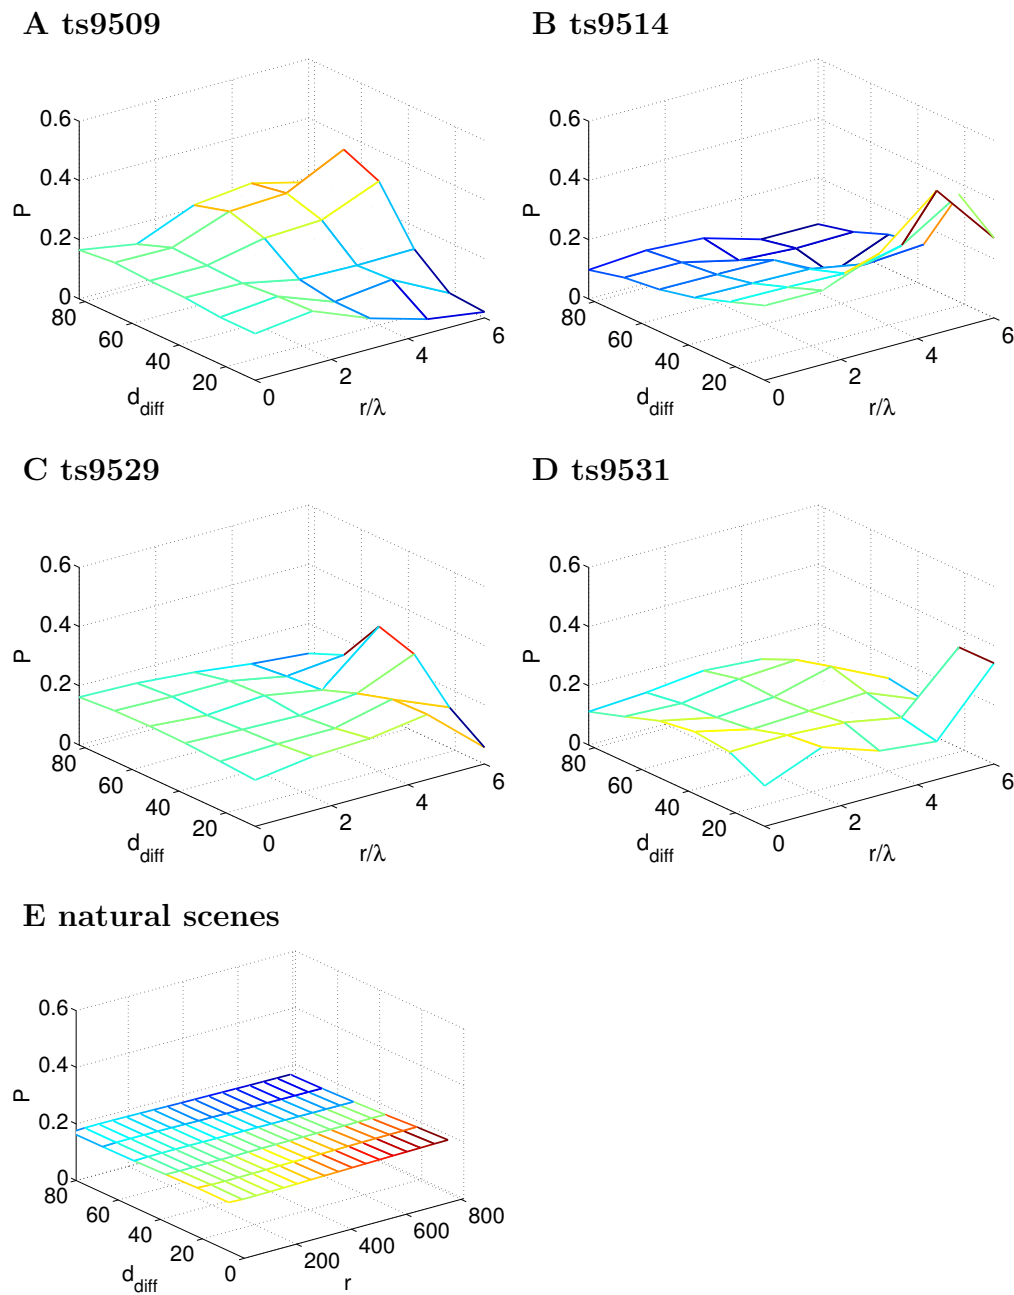
FIGURE 6.5: **Distribution of D_{diff} values for the 4 cases compared with natural scene statistics**

A-D For each animal the probability distribution $P(d_{\mathrm{diff}}|r)$ was calculated by binning into bins of $20°$ in d_{diff} and 1.5λ in r. Because of the low number of connections the results have a significant amount of noise.

E An analogous statistic was calculated for natural scenes (replication of data shown in chapter 4). Although difficult to see at this scale (which matches that in panels A-D), there is a small decline in probability with increasing d_{diff} so the mean value $D_{\mathrm{diff}}(r) < 45°$. Because a large number of natural scenes could be counted and averaged, there was very little noise in these results.

6.4   Discussion

We examined lateral connections in V1 and compared the statistics of their connections with the edge arrangement of natural scenes. These connections develop after eye-opening (Burkhalter et al., 1993; Kovács et al., 1999) and therefore may be driven by natural scene input. As found previously, at short distances the lateral connectivity is isotropic (and probably mostly inhibitory (Kisvárday et al., 1997)). However, long-range lateral connections are clearly anisotropic. We hypothesised that these connections might display similar connectivity statistics to edges in natural scenes. However, we found that the connectivity statistics showed a range of behaviour across the individual animals and both co-circularity and anti-cocircularity were present in the wiring. This study is, as far as we know, the first to directly examine lateral connectivity in V1 for evidence of co-circular connections.

Some limitations of our results are as follows. Firstly, we only had access to connectivity data from 4 animals, which makes drawing conclusions challenging. Secondly, the experimental data has a significant amount of noise due to error in the alignment of optical imaging results with the histology. Because co-circularity is a 2nd order property these errors are compounded when calculating co-circularity statistics. An additional possible source of noise is based on more theoretical considerations which are inherent to any experimental study of these connections. Suppose the brain has some desired function for specifying lateral connectivity (for instance some spread of co-circularity values), it is unlikely this function is realised without a significant level of noise in the connectivity. Because of this, even without any experimental uncertainty it may be difficult to recover a good approximation to the original generating function. Assuming the noise in the connectivity is independent between connections, this issue could be overcome by increasing the number of injection sites. Another complication is that the restricted captive environment in which the tree shrews were raised was probably not as rich as a truly natural environment, however, at the level of statistical characterisation considered here, these differences are unlikely to be consequential. Because of these limitations, it is difficult to make definitive conclusions from our results. However, this data is still able to provide some important insights into functional connectivity in V1.

The variations between the 4 animals are not explainable by minor variations in topography. We examined the value of co-circularity for each of the animals for various values of τ, which is the value of co-circularity that would be present if the assigned topography was rotated by τ. In order for, for instance, case ts9509 to be in agreement with case ts9514, they would have to contain a combined error of $45°$ in the assigned orientation of the topography. Other work has demonstrated that the topography of tree shrews has little inter-individual variability and consistent local structure Bosking et al. (2002). Additionally, our previous work has shown that the co-circularity statistics we used are robust to other perturbations in topography such as anisotropy.

Previous work (Ben-Shahar and Zucker, 2004), has shown that co-linearity alone does not fully explain the anisotropy in the original paper (Bosking et al., 1997) and some broader form of connectivity is required. Here, we postulated that co-circularity similar to that found in natural scenes would account for this. However, we found that there is a large variation in co-circularity statistics between different injections sites (and animals). In particular, case ts9509 shows strong anti-cocircularity, while case ts9514 shows strong co-circularity.

Several other studies have examined horizontal connections in V1 in a variety of animals (Ts'o et al., 1986; Stettler et al., 2002). Early work in tree shrew (Rockland and Lund, 1982) found long-range (around 2.5mm) lateral connectivity and some indications of orientation selectivity in the connections, as did similar work in cats (Gilbert and Wiesel, 1989). There is conflicting information on whether connections in V1 are strongly iso-linear (Bosking et al., 1997) (tree shrew) or mostly isotropic (Kisvárday et al., 1997) (cat). However, there is significant evidence from psychophysics (Polat and Sagi, 1993; Cass and Spehar, 2005) that lateral connections play a role in some forms of co-linear facilitation. It has also been demonstrated that connections are refined during development (Schmidt et al., 1997) and that there are orientation specific interactions in neuron spiking (Ts'o et al., 1986; Das and Gilbert, 1999; Kinoshita et al., 2009). In summary, there is strong evidence for co-linear lateral connections, but also indications that other types of connections are present.

Previous work has shown that an important functional role of lateral connections may be inhibitory, such as iso-orientation suppression Li (2002); Kinoshita et al. (2009). However, the majority of connections we have considered here, particularly at longer distances

($r > 500\mu$m), are believed to be excitatory (Kisvárday et al., 1997; Bosking et al., 1997). Excitatory connections are known to be connected co-linearly and to be important for co-linear facilitation (Polat and Sagi, 1993; Cass and Spehar, 2005; Cass and Alais, 2006). There is some evidence that V1 cells in some species may also be involved in border-ownership computation (Tamura et al., 1996), however, since we are not aware of any characterisation of the type of connectivity necessary for these computations, we did not test the proposition here.

6.5   Conclusions

We suggest that there may be a role for both anti-cocircular and co-circular lateral connections in visual processing. Edges in natural scenes are arranged co-circularly (with a small but consistent bias). For tasks such as co-linear facilitation (Cass and Spehar, 2005), excitatory co-linear lateral connections are necessary. This facilitation is thought to be involved in contour integration Geisler et al. (2001). However, beyond facilitation, it is precisely when input conflicts with what is expected, in this case anti-cocircular edges, that it is providing novel and important information. Indeed, it has been shown that high entropy areas of visual scenes are usually at the center of gaze (Reinagel and Zador, 1999). From this point of view, edges that are anti-cocircular are the most salient, and excitatory connections to facilitate their detection are a reasonable postulate. It is possible that the brain uses a combination of these strategies to both detect the unexpected (anti-cocircular edges) while facilitating detection of the expected (anti-cocircular). Our findings are consistent with such a bifurcated connection strategy.

An interesting experimental followup to this work would be to examine the lateral connectivity of animals reared under unusual rearing conditions (as done for orientation maps in chapter 5). Since these connections develop after eye opening it is likely that these rearing conditions would have an effect on lateral connectivity (indeed previous work has indicated that strabismus during development has an effect on connectivity (Schmidt et al., 1997)). In chapter 5, such results allowed us to make strong conclusions about the source of variation in the orientation preference map statistics. However, examining lateral connectivity

in this way would require significant experimental investment, particularly if, as we found here, the lateral connection properties vary dramatically between individuals. If available, such results would provide further insight into the developmental mechanisms and statistics underlying lateral connections and the degree to which these connections are influenced by natural visual scenes during development.

I like nonsense, it wakes up the brain
cells.

Dr Seuss

7

On the saliency of naturalistic and entropic
edge arrangements

Synopsis

Previous work (Baker and Graf, 2009) showed that visual noise with a naturalistic spatial power spectrum dominated over noise with different power spectra in binocular rivalry experiments.

We hypothesised that, similarly, edges with naturalistic arrangements might dominate in rivalry with random edge arrangements. We constructed stimuli with different types of edge statistics: co-linear, co-circular, "naturalistic" and random. We used binocular rivalry to measure the salience of edge arrangements with statistical structure versus randomly arranged edges.

In contradiction of our initial hypothesis, we found that randomly oriented edges dominated

over all other types of edge arrangements. We argue that these results, in conjunction with previous work considering gaze direction and saccades, indicate that entropy reduction may be a key consideration in early visual processing. We show that this is consistent with the evidence in chapter 6 that V1 may have multiple wiring strategies.

7.1   Introduction

In previous chapters, we have examined the role of natural scene statistics in the development and wiring of mammalian early visual processing. In this chapter, we examine the role such statistics may play in the subjective experience. This psychophysical approach is complementary to the earlier examination of biological substrates. It allows us to measure the impact the natural scene statistics have on the function of visual processing in a fully-developed, complete brain. As earlier in this thesis, here we focus on the statistics of edge arrangements.

As discussed previously, natural scenes contain many statistical regularities. Both intuitively, and from the principled statistical approach to receptive field formation discussed in chapter 3, we know that edges are the basic components of natural scenes. For instance edge arrangements define the contours at the boundaries of objects (Geisler, 2008; Ledgeway et al., 2005; Geisler et al., 2001) and are utilised in texture recognition (Grossberg and Pessoa, 1998; Wolfson et al., 1995; Ben-Shahar, 2006). We have already extensively considered the prevalence of co-circularity in natural scene edge arrangements. The previous chapter showed that there appears to be significant variation in the co-circularity of the lateral connectivity between individual V1 neurons. Because of this, we wanted to examine how edge statistics might modify the salience of a stimulus.

Binocular rivalry provided a psychophysics paradigm which allowed us to measure low-level salience while minimising subjective bias or the need for complex and potentially confusing feedback from the subjects. Binocular rivarly occurs when two markedly different stimuli are presented, one in each eye. The subjective perception is usually that only one stimulus is perceived at a time, and there is stochastic switching between the two stimuli (Walker, 1975; Blake et al., 1990; Lehky, 1995; Brascamp et al., 2006; Tong et al., 2006; Kim

et al., 2006; Krips and Furst, 2009; Clifford, 2009). Binocular rivalry has been systematically studied for nearly two centuries. Sometimes one stimulus is subjectively perceived more of the time than the other, in this case we say that stimulus dominates the rivalry. Others have demonstrated that more salient (Paffen et al., 2006; Paffen and Verstraten, 2008; Yu and Blake, 1992; Koch, 2004), contrasting or luminescent stimuli are more dominant (Blake, 2001). In this chapter, we used an approach similar to Baker and Graf (2009). Baker and Graf used binocular rivalry to examine the salience of noise with differing power spectrums. Since their work provides a model of both the experiment and interpretation, we start with a brief summary of the most relevant part of their work.

Natural scenes have a characteristic power spectrum (distribution of spatial frequency amplitudes). Spatial frequencies in natural scenes have an amplitude of $A(f) \propto 1/f$ (Field, 1987; Dong and Atick, 1995), at least when measured on large enough scales (Langer, 2000). Baker and Graf examined whether subjects found visual noise more salient if it was generated with a naturalistic power spectrum. To test this, Baker and Graf used a binocular rivalry setup; subjects were presented with two different noise stimuli. The stimuli were tinted two different colours and subjects were instructed to indicate through a button press which colour they were experiencing at any given time. Baker and Graf used this feedback to measure the dominance of different types of noise stimuli. By the use of counterbalanced presentations they could eliminate any effect due to colour or eye preference. Figure 7.1 shows some example stimuli and a summary of their results, where they showed that noise with a naturalistic power spectrum dominated over noise with different power spectra whether $\alpha > 1$ or $\alpha < 1$.

The attraction of using binocular dominance to meaure salience of different noise stimuli is that subjects need only report a very basic observation (which colour is dominating their perception) regardless of the complexity of the underlying statistics of interest. The simplicity of this task ensures that individual differences in task strategies or preferences do not have a chance to intrude. Additionally, this simplicity allows changes in salience to be measured accurately and in an unbiased way. We used binocular rivalry to measure the change in salience of edges which contain differing types of edge statistics.

There are many possible ways of imposing statistical structure on edge arrangements, here
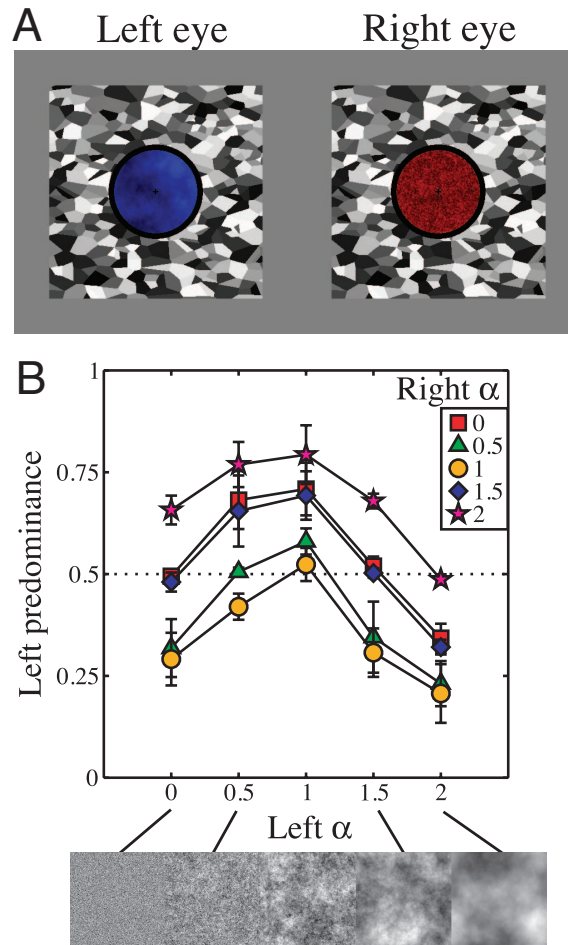
FIGURE 7.1: **Summary of Baker and Graf's experiment**
A Example of the stimuli Baker and Graf used. Each eye was presented static noise with different amplitude spectrums, all of the form $A(f) \propto 1/f^{-\alpha}$. The stimuli were tinted and observers were asked to indicate which tint they were perceiving over time. **B** Baker and Graf found that noise stimuli with $\alpha = 1$ dominated in the rivalries over noise generated with both larger and smaller values of α. Natural scenes have previously been shown to have an amplitude spectrum of this form.
Figure from Baker and Graf (2009).

we restrict ourselves to considering a small number of arrangements known to be present in natural scenes and relevant to grouping edges. We created four types of edge stimuli: co-linear, co-circular, "natural" and random. The details of how each type was created is in the methods. However, we will highlight that the "natural" stimuli were created by directly extracting the arrangement of edges in natural scenes, as opposed to introducing the desired structure artificially. This was used to ensure that the results we found were not due to the

artificial nature of our stimuli.

We used binocular rivalry to find which types of edge arrangements were the most salient. Our initial hypothesis was that, much like in Baker and Graf (2009), edge arrangements similar to natural scene arrangements would dominate. Surprisingly, we found that randomly oriented edges dominated over all other types of stimuli. We discuss later why this provides support for entropy reduction as an important component of visual salience.

7.2   Methods

We used a standard approach to testing the binocular rivalry dominance of the different stimuli similar to that used in Baker and Graf (2009). The primary challenge was to generate well-defined stimuli with edge arrangements containing the appropriate statistics.

7.2.1   Gabor fields

Many important results in visual psychology have been found using simple stimuli. These stimuli have the advantage of being well-characterized and easy to create, but have the disadvantage of being very different from the natural world and therefore possibly not indicative of normal visual processing. At the other extreme are natural scenes, which contain, by definition, natural statistics but it can be difficult to characterise or to modify their statistics in well-defined ways.

We created "Gabor fields" as a useful intermediary between simple stimuli and the full complexity of natural scenes. Each field consists of 300 Gabor patches, placed randomly with the field region. The orientation of the Gabor's could be modified to introduce well-defined statistical structure. Figure 7.4 shows some examples of Gabor fields.

Gabors were used as the basic building block because they are both a good description of V1 simple cell receptive fields (Jones and Palmer, 1987) and, not coincidently, the sparse casual components of natural images are Gabor-like (see chapter 3). A symmetric envelope

was used for the Gabors so that each Gabor was parametrized as:

$$g_{x_0, y_0, \theta, \sigma}(x, y) = \exp\left[\frac{-(x - x_0)^2 - (y - y_0)^2}{2\sigma^2}\right] \cos\left[2\pi f\left(x \cos\theta + y \sin\theta\right)\right] \qquad (7.1)$$

described by a spatial position x_0, y_0, orientation θ and with a Gaussian envelope with variance σ^2. The spatial frequency of the Gabors were chosen to be near the expected peak spatial frequency of the majority of V1 simple cells (Valois et al., 1982): $f = 4\,\text{cycles}/^\circ$. The Gaussian envelope was fixed at $\sigma = 1/f$. The total size of the Gabor field was a circle of diameter $10\,^\circ$, this was chosen to accomodate all the Gabor patches without overlaps. The Gabor positions were chosen stochastically with constraints to avoid Gabors overlapping. This was achieved by placing the first Gabor anywhere in the field with uniform probability. Each subsequent Gabor position was chosen by stochastically selecting a new position with a weighting inversely proportional to the distance of the nearest existing Gabor. This ensured that few Gabors overlapped, while avoiding introducing regularity into the arrangement of the Gabors.

We then created sets of Gabor fields with varying types of statistics necessary for our rivalry displays. The random stimuli set was the simplest to generate. After the Gabors were positioned as described above, each Gabor was assigned an orientation θ chosen uniformly between $[0, \pi)$. Co-linear stimuli were generated in a similar fashion, except each Gabor in a field was assigned the same orientation. The creation of co-circular and natural stimuli was more involved as we explain next.

Co-circular Gabor fields were created using simulated annealing (figure 7.2). In order to use simulated annealing, a quantification of co-circularity was needed. We used the previously introduced measure of co-circularity, $D_{\text{diff}}(r)$ which measures the mean deviation from completely co-circular of each pair of edges at distance r. This means that a set of randomly oriented edges should have D_{diff} near 45°, with lower values indicating increasing co-circularity. In order to remove the dependence on r and get a scalar quantity for annealing, each pair of Gabors were weighted using a Gaussian weighting with a mean of $5\,^\circ$ and a standard deviation of $1\,^\circ$, to calculate a scalar summation of D_{diff}. This weighting introduces little co-circularity at short separations in order to avoid creating nearby edges with a large

FIGURE 7.2: **Generating co-circular edges using simulated annealing**
Co-circular edge arrangements were created using simulated annealing. **A** We started with a set
of randomly oriented edges. **B** Each step of the annealing process, one of the edges was chosen at
random and its orientation was rotated slightly. The direction of the rotation was also random.
The co-circularity between all the edges was calculated both before and after the edge was rotated.
If the proposed random change increased the co-circularity of the edges it was accepted. The
changed edge set was then used as the starting point for the next step of the annealing. If the
change decreased the co-circularity it was rejected probabilistically. The reason for this is that
occasionally accepting "backward" steps helps avoid getting stuck in poor local minima.

degree of co-linearity, while maximising the amount of co-circularity introduced into the
arrangement. Simulated annealing was used to anneal randomly oriented fields with starting
$D_{\text{diff}} \approx 45°$ until $D_{\text{diff}} \leq 30$ was achieved. New states were generated by adjusting a randomly
chosen Gabor patch's orientation by $\pm 1°$. State changes during annealing were accepted with
a probability:

$$p = \min\left(\exp((D_{\text{old}} - D_{\text{new}})/k_b), 1\right) \tag{7.2}$$

where D_x denotes the co-circularity of the old and new states. The annealing temperature
was adjusted after each round of annealing to be:

$$k_b(n) = 3\exp\left[-n/3000\right] \tag{7.3}$$

where n is the number of annealing steps taken. This simulated annealing reliably generated
scenes with D_{diff} values of 30° from randomly oriented scenes.

In order to ensure that the results we found were not due to some statistical artifact
from the artificial generation of co-circular scenes, we also created Gabor fields using edge

arrangements extracted directly from natural images. This had the disadvantage of introducing wide variation in the statistics of the edge arrangements, but ensured that the statistics were naturalistic. The reason natural scenes could not be used directly is that we needed to separate the statistics of the edge arrangements from other confounding factors, such as changes in power spectrum or high-level salience. To achieve this, we picked a random scene fragment from the van Hateren image dataset (van Hateren and van der Schaaf, 1998) for each stimuli. Figure 7.3A shows an example. Edge positions were chosen stochastically as described earlier. The scene fragments were filtered, using 3-level steerable pyramids (Simoncelli and Freeman, 1995) to find the power in each orientation at each pixel. The use of steerable filters allowed a multiscale filter to be calculated efficiently so that orientation was considered at multiple length scales in the original scene. Each pixel was then assigned the orientation of the filter with the strongest power at that position. Figure 7.3B shows an example orientation image. Finally, each edge was assigned the orientation of the underlying pixel. Figure 7.3C shows the Gabor field overlaid on its source image. This method allowed us to source the underlying image edge orientation statistics while maintaing stochastic edge arrangments and fixed spatial frequencies.

All Gabor fields contained the same number of Gabors, with their positions chosen stochastically. The only variation between each class of stimuli was in the statistics of the edge orientations. Figure 7.4 shows an example of each type of stimuli (including tinting, which was independently varied as described below). The creation of these synthetic stimuli ensured that only the variable of interest was modified when comparing different stimuli types; there were no correlations with spatial frequency, contrast or contours which could confound our results.

7.2.2   Binocular rivalry displays

Once the stimulus sets were created a standard binocular rivalry setup was used to measure dominance. Stimuli were displayed on an Dell 2407 WFP LCD monitor (1920x1200 pixels, 32-bit colour) running at 60 Hz driven by a ATI Radeon 2400 Pro graphics card[1] controlled

---

[1]Advanced Micro Devices, Inc., Sunnyvale, CA. USA

Figure 7.3: **Creating natural edge fields**
Natural scene edge arrangements were created by randomly selecting a region of a picture from the van Hateren dataset **A** The orientation of each pixel was then determined using a multiscale filter (described in the text). Note: This image has been lightened compared with the original image to improve the clarity in the manuscript. The images used in the analysis were not modified. **B** Gabor patches were randomly overlayed and the orientation of the underlying pixel was assigned to the Gabor **C**. The Gabor patches (without the background picture) were used as the natural edge arrangement stimuli for the rivalry experiments. Note: In order to facilitate comprehension, the image and overlayed Gabor's have been contrast-enhanced. The low contrast version was used in the experiment.

by a Dell PC running Windows XP[2]. The MATLAB Psychophysics toolkit (Brainard, 1997) under MATLAB[3] was used to display the stimuli. Subjects viewed the display through a standard mirror stereoscope 600 mm from the display monitor. Subjects performed the experiment in a dark room. All subjects had normal, or corrected to normal vision.

Each experiment consisted of 20 trials of a non-random stimulus type rivaled with a random stimulus. One stimulus was tinted red and the other blue for each trial (tinting was performed by using only the corresponding pixel colour of the RGB display). Each stimulus was used only once per subject. Subjects reported which colour they were experiencing by holding down one of two mouse buttons. They were instructed to click neither button if they were unsure which colour was dominant. The trial was counterbalanced across eye and tint condition. Each trial had consisted of a 40 second rivalry followed by a 20 second blank screen and then the next trial. The 40 second trial length was chosen because it was long enough to allow a number of rivalry switches but minimized the amount of "bleaching"

---

[2]Microsoft, Redmond, WA, USA
[3]version 2009b, The Mathworks, Natick MA, USA
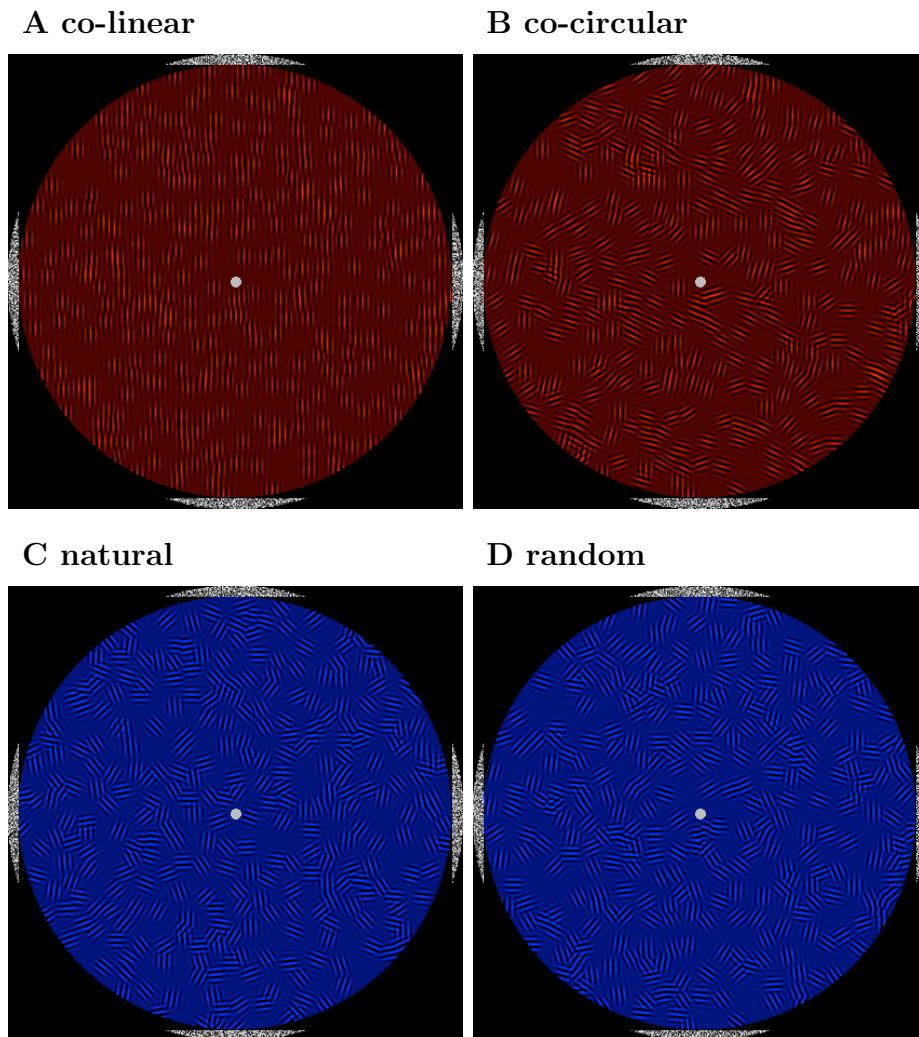
FIGURE 7.4: **Example stimuli**
Four stimuli types were generated for the rivalry experiment. **A** co-linear stimuli containing edges of a single orientation, **B** co-circular stimuli with strong co-circularity ($D_{\mathrm{diff}} = 30°$), **C** natural edge arrangements (described earlier) **D** and randomly oriented stimuli.
Note that each non-random stimuli was only ever paired during rivalry with the random stimuli. The arrangement of the stimuli here does not indicate pairing. Examples of the two tint colours are also shown here for the stimuli. Both eye and tint colour were interleaved for each stimulus type so the tint used for a particular stimulus here has no significance.

that occurred due to visual adaptation. Subjects performed blocks of 5 trials and then the experiment paused until subjects were ready to continue. This allowed subjects to stretch and minimized fatigue and eye strain.

It was important to ensure the binocular rivalry was occurring for each subject and that confounding factors were minimized. Fusion was aided by surrounding the stimuli in each with a small ring of white noise, identical in both eyes. Subjects placed their heads on a chin-rest during the experiment and before each experiment subjects used a mouse to select the display position which maximised fusion. Additionally, since rivalry dominance is affected by perceived contrast, each subject did calibration trials where two random stimuli were rivalled and the contrast of the red stimuli was adjusted until the stimuli were reported as dominating approximately equally.

7.2.3   Statistical analysis

There is a large amount of noise in any binocular rivalry experiment due to the inherent stochasticity of rivalry switching and accumulated errors in subject's reporting. Additionally, we a priori expected that changes in edge statistics would cause only small perturbations in the dominance of our stimuli.

For each rivalry experiment we used a two-sided, one-sample t-test to test if the fraction of time random stimuli dominated the rivalry was significantly different from 0.5 (the amount expected if the edge arrangements did not affect rivalry). This test was performed using the dominance of all subjects pooled together for each stimuli type and also for each subject individually.

We also wanted to determine whether the fraction of dominance of random stimuli was significantly different between the different experiments. One-way ANOVA tests were used to find if there was any difference in the fraction of random dominance between the 3 experiments. Since the ANOVA indicated significance, we used Tukey's honest significance test (Kramer, 1956; Braun, 1994; Sonnemann, 2008) to perform a 3-way comparison to find which experiments had a significantly different fraction of random dominance. Both tests were done using all the subjects results pooled together.

Stimuli set	n naive	n total subjects
Co-linear	2	4
Co-circular	2	4
Natural	2	5

TABLE 7.1: **Number of subjects per experiment**
Each subject performed 20 binocular rivalry trials. All results were used in the analysis. Most subjects participated in all 3 experiments.

Additionally, we fitted a 2-way ANOVA to test whether eye and tint colour had any affect on our results. This was to measure the noise that these confounding factors contributed to the results. It is important to note that even if these factors are significant, they cannot contribute to a non-null result due to the counterbalancing in our experimental design. We fit the ANOVA across all subject's results together. We report both the p-value and the effect size η^2 (the effect size measures the amount of variance in the results explained by a variable).

For all tests $p < 0.05$ was considered to indicate significance. All statistics were calculated using MATLAB built-in functions.

7.3   Results

Contrasting with our original hypothesis, we found randomly arranged edges consistently dominated rivalries. This dominance was statistically significant when results were pooled across subjects for all pairings we tested. Figure 7.5 summarises these results. We also analysed subjects individual results separately, in every experiment in which the subject had a statistically significant deviation from 50% randomly oriented edges dominated (some subjects did not have a statistically significant deviation from 50% dominance, this is not surprising as each subject only performed 20 trials of each stimulus type). Table 7.1 shows the number of subjects used for each stimulus type. Every trial had at least 2 naive subjects. The strength of the dominance was lower for the natural stimuli compared with co-linear and co-circular stimuli, but still statistically significant. These results demonstrate that randomly oriented edges dominate in rivalry.
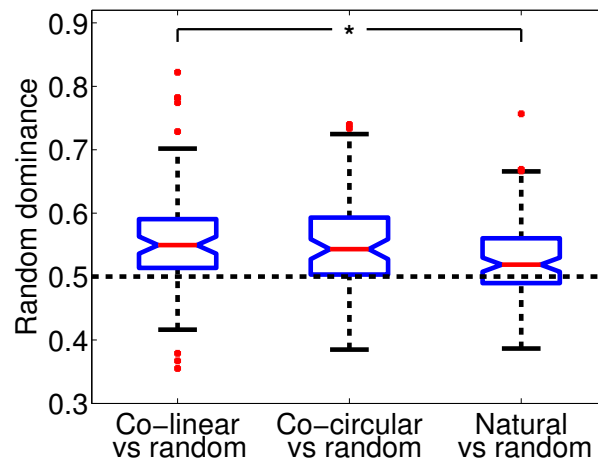
FIGURE 7.5: **Binocular rivalry of Gabor fields**
Box and whisker plot showing the fraction of a trial that the random stimuli were dominant, as reported by all subject's feedback. Notches show 95% confidence intervals. The red dots indicate outliers which are defined as points $> q_3 + 1.5(q_3 - q_1)$, where q_1 and q_3 represent the lower and upper quartiles respectively, mutatis mutandis for the low outliers. In all stimuli conditions a 2-sided, 1-sample Student's t-test found that the random stimuli dominated ($p > 0.0005$). The mean dominance of random stimuli was 0.56, 0.55, 0.53 for co-linear, co-circular and natural stimuli respectively. The ANOVA showed a significant difference between the 3 experimental conditions $p = 0.007$. A Tukey honest-difference tests showed the only significant difference occurred in the dominance of natural scenes compared with co-linear (marked with the '*' on the graph).

Although random edges consistently dominated the rivalries, they did not do so by a large amount. Given the relatively small number of trials and the many sources of noise, this small change in rivalry dominance is important and robust. As we will demonstrate shortly, the dominance of the random stimuli is comparable to the dominance of naturalistic spatial noise we found we when replicated Baker and Graf's results. We expand on this issue in the Discussion.

We used ANOVA to measure the contribution of noise from confounding factors such as eye and tint colour. Our counterbalanced design ensures that these factors cannot contribute to a non-null result but they may add noise to the results. A 2-way ANOVA was performed, with the fraction of time subjects indicated the random stimuli was dominant as the result variable and stimuli position and tint as treatment conditions. In the co-linear experiment the ANOVA showed small dependence on tint $p = 0.008, \eta^2 = 0.09$, and in the natural stimuli experiment a small dependence on position $p = 0.01$, $\eta^2 = 0.07$. No other trials had
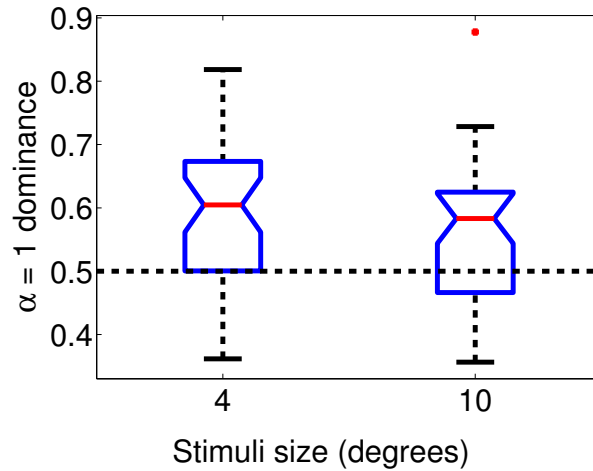
Figure 7.6: **Stimulus size does not significantly affect binocular rivalry**
To test if the size of the stimulus was likely have played a role in our findings we repeated the Baker and Graf (2009) experiment on spatial noise but using two different sizes of stimuli. The box and whisker plot shows the dominance of naturalistic spatial noise ($\alpha = 1$) over white noise ($\alpha = 0$). For both the small (4°) and large stimuli (10°) naturalistic noise dominated the rivalry $p < 0.006$ with a dominance of 0.59 and 0.57 respectively. There was no significant difference $p = 0.29$ between the two stimulus sizes.

a statistically significant deviation on either tint or stimulus position. This indicates that the majority of the variance in the trials is not due to eye position or tint colour, but in two experiments these factors contributed a small amount of noise. These factors are not an exhaustive list of noise sources.

One issue with our stimuli is that, in order to contain a large number of edges, they were significantly larger than the stimuli used in Baker and Graf (2009). This means that, as expected (Kang, 2009), rivalry switching was sometimes incomplete. Subjects sometimes perceived patches with both colours simultaneously. We instructed subjects to only press a button when one colour dominated clearly so these periods of partial switching should not confuse the results. Additionally, any such confusion should not contribute to a non-null result as it would be distributed equally between stimulus types. To check that stimulus size was not playing a role in our results we repeated the Baker and Graf experiment demonstrating the dominance of naturalistic spatial noise using two different stimulus sizes. Using both small (4°) stimuli similar to their original experiment and larger stimuli (10°) we found naturalistic spatial noise dominated (figure 7.6). There was some reduction in dominance

of the larger stimuli but it was not statistically significant. This indicates that, although larger stimuli may have longer switching transitions, this was not an important factor in our results. Additionally, it demonstrates the mean dominance of the random Gabor fields are comparable to the dominance of naturalistic spatial noise.

7.4 Discussion

The results discussed here demonstrate that edge arrangments with high entropy in their arrangement are consistently considered more salient than edges containing statistical structure commonly encountered in natural scenes. This reinforces our finding throughout this thesis that early visual processing is tuned to utilise to the statistical regularities found in natural scenes, although in this case in a surprising way. It provides evidence that early visual processing may use higher-level statistical properties to process visual input efficiently.

Although our rivalry tests showed only small changes in binocular dominance based on edge statistics, these small changes may indicate important differences in early visual processing of the edges. The small increase in the dominance of random edges over naturalistic edge arrangments does not necessarily mean the effect of edge arrangment on visual processing is small. The fraction of time random edges dominated was comparable to the dominance of naturalistic spatial frequencies. Binocular rivalry experiments are inherently noisy due the stochastic nature of binocular switching (Walker, 1975; Blake et al., 1990; Lehky, 1995; Brascamp et al., 2006; Tong et al., 2006; Kim et al., 2006; Krips and Furst, 2009) in addition to noise in the subject's reporting. This is exacerbated because the Gabor fields are necessarily large (to incorporate a large number of edges), which means that binocular switching is not instantaneous. Another consideration is that the changes in edge statistics used in our stimuli are relatively subtle and, in the case of the co-circular stimuli, only present over large distances. Therefore, despite the small change in magnitude, the dominance of random edges may indicate importance differences in low-level salience.

It is intriguing that, opposite to our initial hypothesis, randomly oriented edges dominated over more ordered stimuli. This indicates that visual processing may discount easily predictable input and instead focus attention on the unexpected. There is support for this

idea in several recent models of predictive coding (Schultz and Dickinson, 2000; Hosoya et al., 2005; Zhaoping, 2006; Schwartz et al., 2007; Lesica et al., 2007). We showed in chapter 6 evidence for both anti and co-circular lateral connections in V1. These seemingly opposite wiring strategies may be unified by understanding them to be performing two distinct tasks: matching predicted statistical structure efficiently and increasing the salience of unexpected edge positions.

In contradistinction to our finding that randomly oriented edges dominated binocular rivalry, Baker and Graf (2009) found that stimuli with naturalistic power spectrums dominated over other stimuli, including the most entropic choice: white noise. However, it is not surprising that naturalistic power spectrums may be treated differently from naturalistic edge arrangements during rivalry. Edges arise from appropriate combinations of differing spatial frequencies, in this sense, edge arrangements are higher-level statistics than the power spectrum. The dominance of stimuli with naturalistic power spectrums could be understood as a very narrow prior on the input power spectrum, which causes input with non-naturalistic spatial frequencies to be discounted. The different treatment of naturalistic edge statistics may reflect a higher-variance prior belief about edge statistics so that non-naturalistic edge arrangements attract attention rather than becoming discounted as nonsense. This difference in processing is reasonable since a close match to the $1/f$ power spectrums is found universally in natural scenes (Field, 1987; Dong and Atick, 1995), while edge arrangements in natural scenes are more variable (see chapter 4 and Sigman et al. (2001)). Regarding possible mechanisms, if the differences in binocular dominance are arising in V1, both lateral connections and predictive feedback from other visual areas provide scope to explain the dominance of stimuli with naturalistic power spectrum and entropic edge arrangements.

Other groups have used eye-tracking as a method for deducing low-level salience. This approach has the advantage of corresponding more closely to natural viewing conditions (whereas binocular rivalry with no correlation between eyes is unusual) but the disadvantage of providing more scope for individual idiosyncrasies (Leonards and Scott-Samuel, 2005; Lorigo et al., 2006), habitual search patterns and task specific behaviours (Parkhurst et al., 2002; Najemnik and Geisler, 2005; Williams et al., 1997). Additionally, high-level visual features may play a role (Cerf et al., 2008; Johnson et al., 1991). Reinagel and Zador (1999);

Renninger et al. (2007); Itti and Baldi (2009) found that observers tend to cover regions of natural scenes with less spatial correlation (more entropic) regions of natural scenes. Bruce and Tsotsos (2006); Renninger et al. (2007) found that eye-movements were preferentially in directions which maximised the reduction in entropy. Although these groups considered local spatial correlations rather than edge arrangements, both of these findings are consistent with the notion that entropy is a key issue in assigning input salience.

Motoyoshi and Kingdom (2010) recently examined the role of co-circularity in texture discrimination. They created textures containing pairs of edges, visually similar to our Gabor fields. Subjects were asked to distinguish between textures. They found that subjects could distinguish textures containing co-circularity from random textures relatively easily compared to other textures. This result is complementary to our finding that randomly arranged edges dominate during rivalry. Motoyoshi and Kingdom (2010) measured discriminatory ability to show that co-circularity is a key statistic in texture discrimination. However, our results show that the salience of edge arrangements is increased by removing statistical regularities, thus increasing entropy. Both these finding are compatible with the view that the brain is tuned to deal quickly with co-circular structure, while devoting attention to less-predictable structure.

Binocular rivalry provides a straightforward method for estimating low-level visual salience accurately. Here we have used the dominance of edge arrangments to indicate their salience. An avenue for future work that may elucidate the role lateral connections play in edge salience would be to examine asymmetries in the spread of rivalry switching, as has been done for other types of stimuli (Knapen et al., 2007; Arnold et al., 2009). Such asymmetries may be indicative of underlying asymmetries in the neural circuitry. There is also wide scope for further work examining visual processing of other types of edge statistics using binocular rivalry. We have shown here that entropic edge arrangements dominate binocular rivalry. This is in agreement with previous work showing that entropy reduction is a key consideration in early visual processing.

Then my daughter was about seven
years old, she asked me one day what I
did at work. I told her I worked at the
college - that my job was to teach
people how to draw. She stared at me,
incredulous, and said, "You mean they
forget?"

Howard Ikemoto

8

Conclusions

The aim of the thesis was to improve our understanding of the function and development and the mammalian primary visual cortex. There is already a huge body of work on this topic. Our approach has been to use the statistics of natural scenes as a guide to the function of the primary visual cortex. As we discussed earlier in the thesis, the hypothesis that early visual cortex is optimally encoding its input has a long history, dating back to at least Barlow (1961), with significant advances in the 1990s (Olshausen and Field, 1996; van Hateren and van der Schaaf, 1998; van Hateren and Ruderman, 1998; Bell and Sejnowski, 1997). This earlier work demonstrated that elucidating primary visual function by considering the statistics of its input and relevant computational functions was a promising avenue. Here, we provide a summary of the findings in this thesis and some discussion of future directions. We restrict our discussion here to less detailed summaries of the findings in the thesis, individual chapters contain more detailed discussion of the specific results.

8.1 Summary of thesis

In chapter 3 we expanded on previous coding models of receptive field development to consider the development of binocular receptive fields. Since all visual mammals have binocular visual systems (Pettigrew, 1986), this is an important extension which allows for better comparison with animal experience. We found that, rather than being a trivial extension to previous findings, binocular sparse coding introduces some new questions about coding in the primary visual cortex. In particular, we found that the asymmetry due to the horizontal separation of animal eyes induces stronger inter-ocular correlations between horizontal edges than vertical ones. Our sparse coding model predicts that this asymmetry should lead to a significant over-representation of vertically-oriented neurons. This over-representation does not appear to exist in real animals which may indicate other constraints at work in receptive field development. The extension to binocular field development also allowed us to demonstrate that the changes in receptive field development observed in animals with modified visual input are qualitatively similar to those predicted by our model. This indicates the receptive fields are, at least partially, learned from visual input during development. It also demonstrates that sparse coding models, despite the mismatch with vertical over-representation, capture many important aspects of biological receptive field development. In agreement with previous work, we found that the sparse features of natural scenes are mostly edges. Because edges are such important features, the remainder of the thesis focussed on the statistics of edge arrangements in the visual world and the cortex.

8.1.1 A world of edges

The first step to understanding the influence naturalistic edge arrangements have on the development and wiring of V1 was to create a good model of natural edge statistics. In chapter 4, we extended a previous analysis of edge statistics in natural scenes (Sigman et al., 2001). Sigman et al. generalised the observation that co-linear edges are common in images by showing that pairs of edges in natural scenes tend to be tangent to a common circle, termed co-circular. We showed that this co-circularity in edge arrangements does not significantly diminish in natural scenes over substantial length scales ($> 14°$), much greater

distances than pixel intensities remain correlated. We also showed some preliminary work on grouping different scene types on the basis of their co-circularity statistics and found that the majority of natural scenes have a similar, strong degree of co-circularity. Most of the variance is due to a small fraction of outlier scenes. The ubiquity of long-range orientation correlations in natural scenes potentially make them an important statistic for understanding early visual processing. The remainder of the thesis considered the influence of naturalistic edge statistics on the structure and function of the primary visual cortex.

In chapter 5 we considered the statistics of V1 orientation preference maps and found, surprisingly, that while orientation preference maps contain statistical structure, it is not the same as that found in natural scenes and is not affected by modified visual input during development. This is in agreement with recent evidence that orientation map structure may be driven more by intrinsic pattern formation processes rather than visual input (Kaschube et al., 2010).

However, in chapter 6, we found that lateral connections between V1 columns may reflect a dualistic strategy to utilising edge correlations. We presented evidence that lateral connectivity in V1 may have both anti-cocircular and co-circular components. This dualistic approach may facilitate the encoding of expected edge arrangements, while raising the salience of unexpected scene properties.

This dual strategy hypothesis was reinforced by the functional results we found in chapter 7. We showed that randomly-oriented edge arrangements were more salient to human observers than naturalistic edge arrangements, such as co-linear or co-circular edges. Others have shown that co-linear edges are easier to detect (Polat and Sagi, 1993), but this result indicates that entropic arrangements may be more eye-catching. Additionally, this result provides further evidence that points to V1 lateral connectivity as an important brain substrate for contour detection.

Overall, these findings suggest that there are two underlying challenges encountered in early visual processing. Visual input containing common statistical structure needs to be processed efficiently – this is a restatement of optimal coding principles. However, visual input containing unexpected (less probable) arrangements cannot be discarded, rather it needs

to attract attention[1]. These, almost opposite, principles both postulate a strong connection between natural scene statistics and early sensory processing. We have demonstrated that both principles may be important at many different levels, from dealing with spatial correlations (receptive field development) to V1 lateral connectivity and the salience of edge arrangements measured in adult humans. These ideas bear similarity to predictive coding which proposes that the brain is constantly comparing a model of the world with observation and that errors in prediction are the most salient features of an observation (Schultz et al., 1997; Schultz and Dickinson, 2000; O'Doherty et al., 2003; Doherty et al., 2004).

8.2   Future directions

The central nervous system is among the most complex networks we have ever tried to unravel (Koch and Laurent, 1999; Strogatz, 2001; Bullmore and Sporns, 2009). Hidden in its non-linear complexities are clues to our past, our future, magnificent machines and most of all, us (Koch, 2004). There is now cause for cautious optimism that progress is being made in our understanding. Modern electronics are now capable of creating machines with comparable numbers of components to small regions of the cortex (Markram, 2006). Such advances in technology seem bound to unravel mysteries. Of course, this optimism must be tempered by history. Previous generations have fallen sway to the belief that the latest technology, be that vacuum diodes or centrifugal steam engines would resolve the mysteries of the brain (Brooks, 2001). Even so, there is a common theme that links these purported neural analogies: the brain as a network.

New techniques such as connectomics (Yap et al., 2010) which allow large connectivity diagrams to be elucidated, automatic behaviour screening (Bejsovec et al., 2004; Bang et al., 2002; Hicks et al., 2006), and of course, the family of -omics (proteomics, genomics) are already in various stages of revolutionising neurobiology (Collins et al., 2003; Joyce and Palsson, 2006; Geschwind and Konopka, 2009). These changes in the volume of data gathering are not just providing new evidence, but also changing the way we do science. Such advances

---

[1]It seems our brains may have predated Mr Cheney in realising that "unknown unknowns" are the items of input that are necessarily the most salient.
`http://www.defense.gov/transcripts/transcript.aspx?transcriptid=2636`

in the speed of data acquisition are desperately needed. The human central nervous system is composed of $\sim 10^9$ neurons with complex connectivity. Each neuron is itself a complex network with $\sim 10^5$ interacting genes. With these numbers, even pairwise correlations are exceedingly difficult to elucidate. Such high-throughput work may make more principled consideration of some of these network dynamics tractable. Ethical concerns will limit some of these high-throughput techniques to simpler species such as *C. elegans*. Functional learning models may be critical for interpreting the connectivity and developmental expressions in the cortex. There is a long way to go but we now have some right to look at the complex networks we are studying with some degree of optimism.

Since others have already considered the future of neuroscience more comprehensively (Adams, 1998; Bell, 1999; Akil et al., 2011), here we will focus on the more manageable topic of the future of naturalistic statistics in understanding brain development, particularly the primary visual cortex.

The explosion in computational capabilities will, of course, have an impact on the study of natural scene statistics. Artificial sensors have already equalled biological vision systems at many aspects of light detection (Vollmer, 2011). However, artificial, active vision systems with wide fields of view remain computationally challenging. The feasibility of such systems, particularly for technological applications with strict size and power constraints, will require improvements in artificial attention (Ishii et al., 2009). This allows devices, as in the brain, to discard the irrelevant quickly with little computational cost[2] (Aziz and Mertsching, 2008).

Another area where the primary visual cortex may act as a inspirational guide is in dealing with noise. Advances in technology relegated noise to the sensing periphery, most modern digital computers attempt to provide guaranteed accuracy, masking the underlying noisiness of their substrate (Shepard and Narayanan, 1996). Brains, on the other hand, are noisy throughout (Shadlen and Newsome, 1994; London et al., 2010; Averbeck et al., 2006). Despite this, the brain is capable of efficient, complex, accurate computation (Ma et al., 2006). Attempts to fabricate ever smaller, faster computers are increasingly finding it difficult to continue to hide the inherent noisiness of the underlying substrate (Sadek et al., 2004). Efficient use of such noisy computers will require a revision of the "Turing model" deterministic

---

[2]Readers are encouraged to deny that such ability was required for processing this thesis.

approach to computation. Mammalian vision with its noisy computation (Zhaoping, 2006) may be one source of inspiration for such approaches.

Increasing computational power will make more complex models of natural scene statistics feasible (Felsen and Dan, 2005). These improved models will lead to both improved image compression (Simoncelli, 1998) and artificial vision algorithms (Jurie and Triggs, 2005), and a better understanding of biological visual systems (Simoncelli and Olshausen, 2001; Geisler and Diehl, 2002; Srivastava et al., 2003; Geisler, 2008). Nonetheless, the complexity of these statistical models scale horrifically with dimension (Simoncelli, 2005). This means that clever approaches which capture the complexity of natural scenes while remaining tractable will continue to be crucial (Perlovsky, 1998). Biological visual systems may aid us in learning how to construct such concise models of the world.

Much like onions, biological visions systems have many layers. In this thesis, we have concentrated on the encoding of static images by V1 simple cells. This constraint has allowed us to enjoy the advantage of well-defined input, single-level learning models along with a well-characterised biological substrate. However, multi-layer learning models which incorporate temporal input and feedback from higher levels are necessary to describe the function of the full mammalian visual system (Olshausen, 2003). Hyvärinen and Hoyer (2001) have already shown that V2 response can be modelled by multi-layer unsupervised learning models. Extending such work to incorporate higher visual areas and richer dimensions will require new advances in our learning models, but are vital for improving our understanding of the visual system as a whole. Such multi-layer models are becoming increasing computationally tractable, but they may require accurate reproductions of natural scene input to match biological systems. For this reason, robotic acquisition systems which have similar properties and movements to animals may be needed (Schneider et al., 2005). As we have shown in chapter 4, using a more complete model of visual input can have important effects that may not be obvious in simplified models. In chapter 7, we used binocular rivalry to demonstrate functional differences in processing of different edge arrangements. Visual psychophysics may play an important role in testing functional predictions about visual processing made by multi-layer visual models.

As introduced above, mammalian visual systems are active visual systems. The eyes are

not static receptors but constantly on the move, driven by feedback and attentional demands. Understanding this active vision where feedback can be used to gather new information in a task-dependent manner is another area where optimal encoding models can make a significant contribution (Geisler and Cormack, 2011). Several groups have demonstrated that saccades are dependent on both task and spatial frequencies (Najemnik and Geisler, 2005, 2008) and may be utilised in a task-dependent optimal manner (Rucci, 2008; Rucci et al., 2007; Ko et al., 2010; Geisler et al., 2009). Better understanding of active vision will require models incorporating feedback and detailed charaterization of the statistics of visual input that are relevant to particular tasks (Schneider et al., 2005; Rucci et al., 2007; Rucci, 2008; Ko et al., 2010).

As the sophistication of theoretical models increase it will be vital to continue to compare them with experiment. One important source of comparison for cortical development is with animals which are reared with deprived or modified visual input. In chapter 5, this type of experiment played a key role in eliminating visual input as a source of reduced symmetry in orientation preference maps. Unfortunately, there have been few experiments which have examined functional connectivity in such modified reared animals. Such results could help refine theories of cortical development. For both practical and ethical reasons, the number of animals in these modified rearing experiments is not likely to ever be large (Festing and Wilkinson, 2007). This means that theoretical models are necessary to make the best use of such valuable results. However, since the amount of data needed to detect higher-order statistical changes grows rapidly with increasing order, some tests will likely continue to remain unfeasible. Data sharing, particularly in standardised formats (Liu and Ascoli, 2007; Gardner et al., 2003), will help ensure the maximum public benefit is derived from these experiments (Russell and Burch, 1959; Gluck, 2003) and encourage detailed comparisons between species and rearing conditions.

Even with improved data sharing, the comparison between models and experimental results are not straightforward. Detailed quantitative comparison, particularly of higher-order statistics, requires an appreciation for the minutiae of the experimental approach. However, such comparisons are vital. This will require people from different expertise domains to collaborate closely; not only should experiments guide theoretical models, good models ought

to lead to fruitful experiments. In this thesis, we had to good fortune to collaborate with two groups with unique, relevant experimental data. A range of experimental insights, from histology in visual cortex to binocular rivalry in human subjects, were used. By definition, comprehensive models of cortical function and development will need to incorporate a wide array of experimental insight.

Vision has provided humans with many of our most important observations about the world, from the microscopic to the telescopic. Improved artificial visual systems are going to be essential for designing machines which can sense and interact with our world and us. We have shown here that better models of visual statistics can improve our understanding of biological vision. In return, better biological understanding has led to improved models of visual processing and encoding. There are significant challenges remaining in our understanding of both artificial and biological visual systems, however, it is reasonable to be optimistic that we may be gathering an almost adequate arsenal to combat our ignorance. Optimal encoding models have proved useful in making sense of early visual encoding and their hierarchical versions may play a key role in elucidating further levels of biological visual processing.

The more that you read,

The more things you will know.

The more that you learn,

The more places you'll go.

Dr. Seuss

References

Adams DL, Horton JC (2003) A precise retinotopic map of primate striate cortex generated from the representation of angioscotomas. *J Neurosci* 23:3771–3789.

Adams S (1998) *The Dilbert Future: Thriving on Business Stupidity in the 21st Century* Harper Paperbacks.

Akil H, Martone ME, Van Essen DC (2011) Challenges and opportunities in mining neuroscience data. *Science* 331:708–12.

Akmajian A, Demers RA, Farmer AK, Harnish RM (2010) *Linguistics: An Introduction to Language and Communication* The MIT Press.

Albert MV, Schnabel A, Field DJ (2008) Innate Visual Learning through Spontaneous Activity Patterns. *PLoS Computational Biology* 4:e1000137.

Albus K (1975) A quantitative study of the projection area of the central and the paracentral visual field in area 17 of the cat. I. The precision of the topography. *Exp Brain Res* 24:159–179.

Andrieu C, De Freitas N, Doucet A, Jordan M (2003) An introduction to MCMC for machine learning. *Machine learning* 50:5–43.

Arnold D, James B, Roseboom W (2009) Binocular rivalry: Spreading dominance through complex images. *Journal of Vision* 9:1–9.

Atick J, Redlich A (1992) What does the retina know about natural scenes? *Neural Computation* 4:196–210.

Attias H, Schreiner C (1997) Temporal low-order statistics of natural sounds. *Advances in neural information processing systems* 9:27–33.

Attneave F (1954) Some informational aspects of visual perception. *Psychological Review* .

Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nature reviews. Neuroscience* 7:358–66.

Aziz MZ, Mertsching B (2008) Fast and robust generation of feature maps for region-based visual attention. *IEEE transactions on image processing* 17:633–44.

Baker DH, Graf EW (2009) Natural images dominate in binocular rivalry. *Proc Natl Acad Sci USA* 106:5436–5441.

Bang PI, Yelick PC, Malicki JJ, Sewell WF (2002) High-throughput behavioral screening method for detecting auditory response defects in zebrafish. *Journal of neuroscience methods* 118:177–87.

Barlow HB (1961) Possible principles underlying the transformations of sensory messages. *Sensory Communication* pp. 217–234.

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

Beck J, Rosenfeld A, Ivry R (1989) Line segregation. *Spatial Vision* 4:75–101.

Bejsovec A, Lecuit T, Modolell J (2004) The fly Olympics: faster, higher and stronger answers to developmental questions. *EMBO reports* 5:1037–40.

Bell AJ, Sejnowski TJ (1997) The "independent components" of natural scenes are edge filters. *Vision Res* 37:3327–3338.

Bell A (1999) Levels and loops : the future of artificial intelligence and neuroscience. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* 354:2013.

Ben-Shahar O, Zucker S (2004) Geometrical Computations Explain Projection Patterns of Long-Range Horizontal Connections in Visual Cortex. *Neural Networks* 16:445–476.

Ben-Shahar O (2006) Visual saliency and texture segregation without feature gradient. *Proc Natl Acad Sci USA* 103:15704–15709.

Betsch BY, Einhäuser W, Körding K, König P (2004) The world from a cat's perspective - statistics of natural videos. *Biological Cybernetics* 90:41–50.

Blake R, Westendorf D, Fox R (1990) Temporal perturbations of binocular rivalry. *Perception and psychophysics* 48:593–602.

Blake R (2001) A primer on binocular rivalry, including current controversies. *Brain and Mind* 2:5–38.

Blakemore C, Cooper GF (1970) Development of the brain depends on the visual environment. *Nature* 228:477–478.

Blasdel G, Campbell D (2001) Functional retinotopy of monkey visual cortex. *J Neurosci* 21:8286–8301.

Bonhoeffer T, Grinvald A (1991) Iso-orientation domains in cat visual cortex are arranged in pinwheel-like patterns. *Nature* 353:429–431.

Bonhoeffer T, Grinvald A (1993) The layout of iso-orientation domains in area 18 of cat visual cortex: optical imaging reveals a pinwheel-like organization. *J Neurosci* 13:4157–4180.

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.

Bosking WH, Kretz R, Pucak ML, Fitzpatrick D (2000) Functional specificity of callosal connections in tree shrew striate cortex. *J Neurosci* 20:2346–2359.

Bosking WH, Crowley JC, Fitzpatrick D (2002) Spatial coding of position and orientation in primary visual cortex. *Nat Neurosci* 5:874–882.

Bosking WH, Zhang Y, Schofield B, Fitzpatrick D (1997) Orientation Selectivity and the Arrangement of Horizontal Connections in Tree Shrew Striate Cortex. *Journal of Neuroscience* 17:2112–2127.

Brainard DH (1997) The Psychophysics Toolbox. *Spatial Vision* 10:433–436.

Brascamp JW, van Ee R, Noest AJ, Jacobs RHaH, van den Berg AV (2006) The time course of binocular rivalry reveals a fundamental role of noise. *Journal of vision* 6:1244–56.

Braun H (1994) *The Collected Works of John W. Tukey: Multiple Comparions, Volume VIII* Chapman and Hall/CRC.

Brent RP (2002) *Algorithms for Minimization Without Derivatives* Dover Publications.

Bressloff PC, Cowan JD, Golubitsky M, Thomas PJ, Wiener MC (2002) What geometric visual hallucinations tell us about the visual cortex. *Neural Computation* 14:473–491.

Brooks R (2001) Steps towards living machines In Gomi T, editor, *Evolutionary Robotics. From Intelligent Robotics to Artificial Life*, pp. 72–93. Springer.

Bruce N, Tsotsos J (2006) Saliency based on information maximization. *Advances in neural information processing systems* 18:155.

Bullmore E, Sporns O (2009) Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews. Neuroscience* 10:186–98.

Burkhalter A, Bernardo KL, Charles V (1993) Development of local circuits in human visual cortex. *Journal of Neuroscience* 13:1916.

Burton GJ, Moorhead IR (1987) Color and spatial structure in natural scenes. *Applied optics* 26:157–70.

Butts DA, Rokhsar DS (2001) The information content of spontaneous retinal waves. *J Neurosci* 21:961–973.

Buzas P, Volgushev M, Eysel UT, Kisvarday ZF, Buzás P, Kisvárday ZF (2003) Independence of visuotopic representation and orientation map in the visual cortex of the cat. *European Journal of Neuroscience* 18:957–968.

Cang J, Niell C, Liu X, Pfeiffenberger C, Feldheim D, Stryker M (2008) Selective Disruption of One Cartesian Axis of Cortical Maps and Receptive Fields by Deficiency in Ephrin-As and Structured Activity. *Neuron* 57:511–523.

Cang J, Kaneko M, Yamada J, Woods G, Stryker MP, Feldheim DA (2005a) Ephrin-as guide the formation of functional maps in the visual cortex. *Neuron* 48:577–589.

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

Carreira-Perpiñán MA, Goodhill GJ (2004) Influence of lateral connections on the structure of cortical maps. *J Neurophysiol* 92:2947–2959.

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.

Cass J, Alais D (2006) The mechanisms of collinear integration. *Journal of Vision* 6:915–922.

Cass JR, Spehar B (2005) Dynamics of collinear contrast facilitation are consistent with long-range horizontal striate transmission. *Vision Res* 45:2728–2739.

Caywood MS, Willmore B, Tolhurst DJ (2004) Independent components of color natural scenes resemble V1 neurons in their spatial and color tuning. *J Neurophysiol* 91:2859–2873.

Cerf M, Harel J, Einhäuser W, Koch C (2008) Predicting human gaze using low-level saliency combined with face detection. *Advances in neural information processing systems* 20:241–248.

Chalupa LM (2007) A reassessment of the role of activity in the formation of eye-specific retinogeniculate projections. *Brain research reviews* 55:228–236.

Chapman B, Stryker MP, Bonhoeffer T (1996) Development of orientation preference maps in ferret primary visual cortex. *J Neurosci* 16:6443–6453.

Chellapilla K, Larson K, Simard PY, Czerwinski M (2005) Building segmentation based human-friendly human interactive proofs (HIPs). *Proceedings of the Second International Workshop on Human Interactive Proofs* 2:1–26.

Chen LM, Friedman RM, Roe AW (2005) Optical imaging of SI topography in anesthetized and awake squirrel monkeys. *J Neurosci* 25:7648–59.

Chen LM, Heider B, Williams GV, Healy FL, Ramsden BM, Roe AW (2002) A chamber and artificial dura method for long-term optical imaging in the monkey. *Journal of neuroscience methods* 113:41–9.

Chen Y, Martinez-Conde S, Macknik SL, Bereshpolova Y, Swadlow HA, Alonso JM (2008) Task difficulty modulates the activity of specific neuronal populations in primary visual cortex. *Nat Neurosci* 11:974–982.

Choe Y, Miikkulainen R (2004) Contour integration and segmentation in a self-organizing map of spiking neurons. *Biol Cybern* 90:75–88.

Chow CC, Jin DZ, Treves A (2002) Is the world full of circles? *Journal of Vision* 2:571–576.

Christian B (2011) *The Most Human Human: What Talking with Computers Teaches Us About What It Means to Be Alive* Doubleday.

Clifford CWG (2009) Quick guide: Binocular rivalry. *Current Biology* 19:1022–1023.

Collins FS, Green ED, Guttmacher AE, Guyer MS (2003) A vision for the future of genomics research. *Nature* 422:835–47.

Coppola DM, Purves HR, McCoy AN, Purves D (1998a) The distribution of oriented contours in the real world. *Proc Natl Acad Sci USA* 95:4002.

Coppola DM, White LE, Fitzpatrick D, Purves D (1998b) Unequal representation of cardinal and oblique contours in ferret visual cortex. *Proc Natl Acad Sci USA* 95:2621–2623.

Crist R, Li W, Gilbert C (2001) Learning to see: experience and attention in primary visual cortex. *Nature Neuroscience* 4:519–525.

Crook JM, Kisvárday ZF, Eysel UT (1998) Evidence for a contribution of lateral inhibition to orientation tuning and direction selectivity in cat visual cortex: reversible inactivation of functionally characterized sites combined with neuroanatomical tracing techniques. *The European journal of neuroscience* 10:2056–75.

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

Das A, Gilbert CD (1999) Topography of contextual modulations mediated by short-range interactions in primary visual cortex. *Nature* 399:655–61.

Diamond ME, Petersen RS, Harris JA (1999) Learning through maps: Functional significance of topographic organization in primary sensory cortex. *J Neurobiol* 41:64–68.

Diao YC, Jia WG, Swindale NV, Cynader MS (1990) Functional organization of the cortical 17/18 border region in the cat. *Exp Brain Res* 79:271–282.

Doherty JO, Dayan P, Schultz J, Deichmann R, Friston K, Dolan RJ (2004) Dissociable Roles of Ventral and Dorsal Striatum in Instrumental Conditioning. *Science* 304:452.

Doi E, Inui T, Lee TW, Wachtler T, Sejnowski TJ (2003) Spatiochromatic receptive field properties derived from information-theoretic analyses of cone mosaic responses to natural scenes. *Neural Computation* 15:397–417.

Dong DW, Atick JJ (1995) Statistics of natural time-varying images. *Network: Computation in Neural Systems* 6:345–358.

Dragoi V, Turcu CM, Sur M (2001) Stability of cortical responses and the statistics of natural scenes. *Neuron* 32:1181–1192.

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

Elder JH, Goldberg RM (2002) Ecological statistics of Gestalt laws for the perceptual organization of contours. *Journal of Vision* 2:324–353.

Elstrott J, Anishchenko A, Greschner M, Sher A, Litke A, Chichilnisky E, Feller M (2008) Direction Selectivity in the Retina Is Established Independent of Visual Experience and Cholinergic Retinal Waves. *Neuron* 58:499–506.

Fagiolini M, Pizzorusso T, Berardi N, Domenici L, Maffei L (1994) Functional postnatal development of the rat primary visual cortex and the role of visual experience: dark rearing and monocular deprivation. *Vision research* 34:709–20.

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.

Feller MB (2009) Retinal waves are likely to instruct the formation of eye-specific retino-geniculate projections. *Neural Development* 4:24.

Felsen G, Dan Y (2005) A natural approach to studying vision. *Nat Neurosci* 8:1643–1646.

Felsen G, Touryan J, Han F, Dan Y (2005) Cortical sensitivity to visual features in natural scenes. *PLoS Biol* 3:e342.

Festing S, Wilkinson R (2007) The ethics of animal research. *EMBO Rep* 8:526–530.

Field DJ (1987) Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A* 4:2379–2394.

Field DJ (1999) Wavelets, vision and the statistics of natural scenes. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 357:2527–2542.

Field DJ, Hayes A (2004) Contour integration and the lateral connections of V1 neurons. *The Visual Neurosciences* pp. 1069–1079.

Fischler MA, Firschein O (1987) *Readings in computer vision: issues, problems, principles, and paradigms* Morgan Kaufmann.

Fleet D, Wagner H (1996) Neural encoding of binocular disparity: energy models, position shifts and phase shifts. *Vision research* 36:1839–1857.

Florence SL, Kaas JH (1992) Ocular dominance columns in area 17 of Old World macaque and talapoin monkeys: complete reconstructions and quantitative analyses. *Visual neuroscience* 8:449–62.

Foldiak P, Endres D (2008) Sparse coding. *Scholarpedia* 3:2984.

Foldiak P, Young M (1995) Sparse coding in the primate cortex In *The handbook of brain theory and neural networks*, pp. 895–898. Citeseer.

Frostig RD, Lieke EE, Ts'o DY, Grinvald A (1990) Cortical functional architecture and local coupling between neuronal activity and the microcirculation revealed by in vivo high-resolution optical imaging of intrinsic signals. *Proc Natl Acad Sci USA* 87:6082–6086.

Fulvio J, Singh M (2009) An experimental criterion for consistency in interpolation of partly occluded contours. *Journal of Vision* 9:1–19.

Furmanski CS, Engel Sa (2000) An oblique effect in human primary visual cortex. *Nature neuroscience* 3:535–6.

Gandhi SP, Heeger DJ, Boynton GM (1999) Spatial attention affects brain activity in human primary visual cortex. *Proc Natl Acad Sci USA* 96:3314–9.

Gardner D, Toga AW, Ascoli Ga, Beatty JT, Brinkley JF, Dale AM, Fox PT, Gardner EP, George JS, Goddard N, Harris KM, Herskovits EH, Hines ML, Jacobs Ga, Jacobs RE, Jones EG, Kennedy DN, Kimberg DY, Mazziotta JC, Miller PL, Mori S, Mountain DC, Reiss AL, Rosen GD, Rottenberg Da, Shepherd GM, Smalheiser NR, Smith KP, Strachan T, Van Essen DC, Williams RW, Wong STC (2003) Towards effective and rewarding data sharing. *Neuroinformatics* 1:289–95.

Geisler WS, Perry JS, Super BJ, Gallogly DP (2001) Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res* 41:711–724.

Geisler WS (2008) Visual perception and the statistical properties of natural scenes. *Annual Review of Psychology* 59:167–192.

Geisler WS, Diehl RL (2002) Bayesian natural selection and the evolution of perceptual systems. *Philos Trans R Soc Lond, B, Biol Sci* 357:419–448.

Geisler W, Cormack L (2011) Models of overt attention In *Oxford Handbook of Eye Movements*. Oxford University Press.

Geisler W, Najemnik J, Almon D (2009) Optimal stimulus encoders for natural tasks. *Journal of vision* 9:1–16.

Geisler W, Perry J (2009) Contour statistics in natural images: Grouping across occlusions. *Visual neuroscience* 26:109–121.

Geschwind DH, Konopka G (2009) Neuroscience in the era of functional genomics and systems biology. *Nature* 461:908–15.

Ghanbari M (2003) *Standard codecs: image compression to advanced video coding* Institute of Electrical Engineers, London.

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

Gilbert C, Wiesel T (1979) Morphology and intracortical projections of functionally characterised neurones in the cat visual cortex. *Nature* 280:120–125.

Girman S, Sauvé Y, Lund R (1999) Receptive Field Properties of Single Neurons in Rat Primary Visual Cortex. *Journal of neurophysiology* 82:301.

Gleick J (2011) *The Information: A History, a Theory, a Flood* Fourth Estate.

Gluck J (2003) Ethical issues in the use of animals in biomedical and psychopharmocological research. *Psychopharmacology* 171:6–12.

Gödecke I, Kim DS, Bonhoeffer T, Singer W (1997) Development of orientation preference maps in area 18 of kitten visual cortex. *Eur J Neurosci* 9:1754–1762.

Gordon JA, Stryker MP (1996) Experience-Dependent Plasticity of Binocular responses in the primary visual cortex of the mouse. *The Journal of Neuroscience* 16:3274–3286.

Greene G, Barrett DGT, Sen K, Houghton C (2009) Sparse coding of birdsong and receptive field structure in songbirds. *Network: Computation in Neural Systems* 20:162–77.

Grinvald A, Frostig RD, Siegel RM, Bartfeld E (1991) High-resolution optical imaging of functional brain architecture in the awake monkey. *Proc Natl Acad Sci USA* 88:11559–63.

Grinvald A, Shoham D, Shmuel A, Glaser D, Vanzetta I, Shtoyerman E, Slovin H, Wijnbergen C, Hildesheim R, Sterkin A, Others (2001) In-vivo optical imaging of cortical architecture and dynamics In Windhorst U, Johansson H, editors, *Modern techniques in neuroscience research*, pp. 893–961. Springer-Verlag, Berlin Heidelberg New York.

Grinvald A, Hildesheim R (2004) VSDI: a new era in functional imaging of cortical dynamics. *Nat Rev Neurosci* 5:874–885.

Grinvald A, Slovin H, Vanzetta I (2000) Non-invasive visualization of cortical columns by fMRI. *Nat Neurosci* 3:105–107.

Grossberg S, Pessoa L (1998) Texture segregation, surface representation and figure-ground separation. *Vision Research* 38:2657–2684.

Guizzo E (2003) The essential message: Claude Shannon and the making of information theory Technical report, MIT.

Han S, Humphreys GW, Chen L (1999) Uniform connectedness and classical Gestalt principles of perceptual grouping. *Perception and Psychophysics* 61:661–74.

Hansen B, Essock E (2004) A horizontal bias in human visual processing of orientation and its correspondence to the structural components of natural scenes. *Journal of Vision* 4:1044–1060.

Hansen T, Gegenfurtner KR (2009) Independence of color and luminance edges in natural scenes. *Visual Neuroscience* 26:35–49.

Hetherington PA, Swindale NV (1999) Receptive field and orientation scatter studied by tetrode recordings in cat area 17. *Vis Neurosci* 16:637–652.

Hicks C, Sorocco D, Levin M (2006) Automated analysis of behavior: A computer-controlled system for drug screening and the investigation of learning. *Journal of neurobiology* 66:977–990.

Hinton G (2002) Training products of experts by minimizing contrastive divergence. *Neural Computation* 14:1771–1800.

Hinton GE, Salakhutdinov RR (2006) Reducing the dimensionality of data with neural networks. *Science* 313:504–7.

Hinton G (2007a) To recognize shapes, first learn to generate images. *Progress in brain research* 165:535–547.

Hinton GE (2007b) Learning multiple layers of representation. *Trends in cognitive sciences* 11:428–34.

Hirsch HVB, Spinelli DN (1970) Visual Experience Modifies Distribution of Horizontally and Vertically Oriented Receptive Fields in Cats. *Science* 168:869.

Hollander M, Wolfe DA (1999) *Nonparametric Statistical Methods, 2nd Edition* Wiley-Interscience.

Horton J, Adams D (2005) The cortical column: a structure without a function. *Phil. Trans. R. Soc. B* 360:837–862.

Horton JC, Hocking DR (1996) An adult-like pattern of ocular dominance columns in striate cortex of newborn monkeys prior to visual experience. *The Journal of Neuroscience* 16:1791–807.

Hosoya T, Baccus S, Meister M (2005) Dynamic predictive coding by the retina. *Nature* 436:71–77.

Hou C, Pettet MW, Sampath V, Candy TR, Norcia AM (2003) Development of the spatial organization and dynamics of lateral interactions in the human visual system. *J Neurosci* 23:8630–8640.

Hoyer P, Hyvärinen A (2000) Independent component analysis applied to feature extraction from colour and stereo images. *Network: Computation in Neural Systems* 11:191–210.

Hoyer PO, Hyvärinen A (2002) A multi-layer sparse coding network learns contour coding from natural images. *Vision Res* 42:1593–1605.

Hsu AS, Dayan P (2007) An unsupervised learning model of neural plasticity: Orientation selectivity in goggle-reared kittens. *Vision research* 47:2868–77.

Hubel DH (1981) Evolution of ideas on the primary visual cortex, 1955–1978: A biased historical account. *Nobel lecture* .

Hubel DH, Wiesel TN (1965) Binocular interaction in striate cortex of kittens reared with artificial squint. *Journal of Neurophysiology* 28:1041.

Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. *Journal of physiology* 195:215–43.

Hubel DH, Wiesel TN (1972) Laminar and columnar distribution of geniculo-cortical fibers in the macaque monkey. *The Journal of comparative neurology* 146:421–50.

Hubel DH, Wiesel TN (1977) Ferrier lecture. Functional architecture of macaque monkey visual cortex. *Proc R Soc Lond, B, Biol Sci* 198:1–59.

Hubel DH, Wiesel TN, LeVay S (1977) Plasticity of ocular dominance columns in monkey striate cortex. *Phil Trans R Soc Lond B* 278:377–409.

Hubel D, Wiesel T (1959) Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology* 148:574.

Hubel D, Wiesel T (1970) The period of susceptibility to the physiological effects of unilateral eye closure in kittens. *The Journal of physiology* 206:419–436.

Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology* 160:215.

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

Hunt JJ, Giacomantonio CE, Tang H, Mortimer D, Jaffer S, Vorobyov V, Ericksson G, Sengpiel F, Goodhill GJ (2009) Natural scene statistics and the structure of orientation maps in the visual cortex. *Neuroimage* 47:157–172.

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

Hyvärinen A, Hurri J, Hoyer PO (2009) *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision* Springer.

Hyvärinen A, Karhunen J, Oja E (2001) *Independent Component Analysis* Wiley-Interscience.

Ishii I, Taniguchi T, Sukenobe R, Yamamoto K (2009) Development of high-speed and real-time vision platform, H3 vision. *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems* pp. 3671–3678.

Ito M, Gilbert C (1999) Attention modulates contextual influences in the primary visual cortex of alert monkeys. *Neuron* 22:593–604.

Itti L, Baldi P (2009) Bayesian surprise attracts human attention. *Vision Res* 49:1295–306.

Johnson MH, Dziurawiec S, Ellis H, Morton J (1991) Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition* 40:1–19.

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

Joyce AR, Palsson BO (2006) The model organism as a system: integrating 'omics' data sets. *Nature reviews. Molecular cell biology* 7:198–210.

Jurie F, Triggs B (2005) Creating efficient codebooks for visual recognition In *Tenth IEEE International Conference on Computer Vision*, pp. 604–610. Ieee.

Kandel ER, Schwartz J, Jessell TM (2000) *Principles of Neural Science* McGraw-Hill Medical.

Kang Ms (2009) Size matters : A study of binocular rivalry dynamics. *Vision Research* 9:1–11.

Karklin Y, Lewicki M (2008) Emergence of complex cell properties by learning to generalize in natural scenes. *Nature* 457:83–86.

Karmarkar U, Dan Y (2006) Experience-Dependent Plasticity in Adult Visual Cortex. *Neuron* 52:577–585.

Kaschube M, Schnabel M, Löwel S, Coppola DM, White LE, Wolf F (2010) Universality in the Evolution of Orientation Columns in the Visual Cortex. *Science* 330:1113–1116.

Kaschube M, Wolf F, Geisel T, Lowel S (2002) Genetic influence of Quantitative Features of Neocortical Architecture. *Journal of Neuroscience* 22:7206–7217.

Katz LC, Shatz CJ (1996) Synaptic activity and the construction of cortical circuits. *Science* 274:1133–1138.

Kim DS, Duong TQ, Kim SG (2000) High-resolution mapping of iso-orientation columns by fMRI. *Nature neuroscience* 3:164–9.

Kim YJ, Grabowecky M, Suzuki S (2006) Stochastic resonance in binocular rivalry. *Vision research* 46:392–406.

Kind PC, Mitchell DE, Ahmed B, Blakemore C, Bonhoeffer T, Sengpiel F (2002) Correlated binocular activity guides recovery from monocular deprivation. *Nature* 416:430–433.

Kinoshita M, Gilbert C, Das A (2009) Optical Imaging of Contextual Interactions in V1 of the Behaving Monkey. *Journal of neurophysiology* 102:1930.

Kisvárday ZF, Tóth E, Rausch M, Eysel UT (1997) Orientation-specific relationship between populations of excitatory and inhibitory lateral connections in the visual cortex of the cat. *Cereb Cortex* 7:605–618.

Knapen T, van Ee R, Blake R (2007) Stimulus motion propels traveling waves in binocular rivalry. *PloS one* 2:e739.

Ko Hk, Poletti M, Rucci M (2010) Microsaccades precisely relocate gaze in a high visual acuity task. *Nature neuroscience* 13:1549–1553.

Koch C, Laurent G (1999) Complexity and the Nervous System. *Science* 284:96–98.

Koch C (2004) *The quest for consciousness: a neurobiological approach* Roberts and Company Publishers.

Köster U, Hyvärinen A (2010) A two-layer model of natural stimuli estimated with score matching. *Neural computation* 22:2308–33.

Kovács I (2000) Human development of perceptual organization. *Vision research* 40:1301–10.

Kovács I, Kozma P, Fehér A, Benedek G (1999) Late maturation of visual spatial integration in humans. *Proc Natl Acad Sci USA* 96:12204–9.

Kramer CY (1956) Extension of Multiple Range Tests to Group Means with Unequal Numbers of Replications. *International Biometric Society* 12:307–310.

Kretzmer E (1952) Statistics of television signals. *Bell Syst. Tech. J* 31:751–763.

Krips R, Furst M (2009) Stochastic properties of coincidence-detector neural cells. *Neural computation* 21:2524–53.

Langer MS (2000) Large-scale failures of $f^{-\alpha}$ scaling in natural image spectra. *Journal of the Optical Society of America A* 17:28–33.

Ledgeway T, Hess RF, Geisler WS (2005) Grouping local orientation and direction signals to extract spatial contours: Empirical tests of "association field" models of contour integration. *Vision Research* 45:2511–2522.

Lee H, Ekanadham C, Ng AY (2007) Sparse deep belief net model for visual area V2. *NIPS* p. 8.

Lee HY, Kardar M (2006) Patterns and Symmetries in the Visual Cortex and in Natural Images. *Journal of Statistical Physics* 125:1243–1266.

Lee HY, Yahyanejad M, Kardar M (2003) Symmetry considerations and development of pinwheels in visual maps. *Proc Natl Acad Sci USA* 100:16036–16040.

Lehky SR (1995) Binocular rivalry is not chaotic. *Proc. R. Soc. Lond. B* 259:71–6.

Leonards U, Scott-Samuel NE (2005) Idiosyncratic initiation of saccadic face exploration in humans. *Vision research* 45:2677–84.

Lesica NA, Jin J, Weng C, Yeh CI, Butts DA, Stanley GB, Alonso JM (2007) Adaptation to stimulus contrast and correlations during natural visual stimulation. *Neuron* 55:479–491.

Lewicki MS (2002) Efficient coding of natural sounds. *Nat Neurosci* 5:356–363.

Li B, Peterson MR, Freeman RD (2003) Oblique effect: a neural basis in the visual cortex. *Journal of neurophysiology* 90:204–17.

Li W, Piëch V, Gilbert C (2004) Perceptual learning and top-down influences in primary visual cortex. *Nature Neuroscience* 7:651.

Li W, Piëch V, Gilbert CD (2008) Learning to link visual contours. *Neuron* 57:442–451.

Li W, Piëch V, Gilbert CD, Pich V (2006) Contour saliency in primary visual cortex. *Neuron* 50:951–962.

Li Y, Hooser SV, Mazurek M, White L, Fitzpatrick D (2008) Experience with moving visual stimuli drives the early development of cortical direction selectivity. *Nature* .

Li Y, Fitzpatrick D, White LE (2006) The development of direction selectivity in ferret visual cortex requires early visual experience. *Nat Neurosci* 9:676–681.

Li Z (1998) A neural model of contour integration in the primary visual cortex. *Neural computation* 10:903–40.

Li Z (2002) A saliency map in primary visual cortex. *Trends in Cognitive Sciences* 6:9–16.

Liu Y, Bovik AC, Cormack LK (2008) Disparity statistics in natural scenes. *Journal of Vision* 8:14–19.

Liu Y, Ascoli GA (2007) *Value added by data sharing: long-term potentiation of neuroscience research* Springer.

Livingstone M, Hubel D (1984) Specificity of intrinsic connections in primate primary visual cortex. *Journal of Neuroscience* 4:2830.

Lloyd S (1982) Least squares quantization in PCM. *IEEE Transactions on Information Theory* 28:129–137.

Logothetis NK (2008) What we can do and what we cannot do with fMRI. *Nature* 453:869–878.

London M, Roth A, Beeren L, Häusser M, Latham P (2010) Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. *Nature* 466:123–127.

Long F, Yang Z, Purves D (2006) Spectral statistics in natural scenes predict hue, saturation, and brightness. *Proc Natl Acad Sci USA* 103:6013–6018.

Lorigo L, Pan B, Hembrooke H, Joachims T, Granka L, Gay G (2006) The influence of task and gender on search and evaluation behavior using Google. *Information Processing and Management* 42:1123–1131.

Lund JS, Lund RD, Hendrickson AE, Bunt AH, Fuchs AF (1975) The origin of efferent pathways from the primary visual cortex, area 17, of the macaque monkey as shown

by retrograde transport of horseradish peroxidase. *The Journal of comparative neurology* 164:287–303.

Ma W, Beck J, Latham P, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nature neuroscience* 9:1432–1438.

MacKay DJC (2003) *Information Theory, Inference and Learning Algorithms* Cambridge University Press.

MacQueen J (1967) Some methods for classification and analysis of multivariate observations. *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability* 1:14.

Mahesri A, Vardhan V (2005) Power consumption breakdown on a modern laptop. *Power-Aware Computer Systems* pp. 165–180.

Malach R, Amir Y, Harel M, Grinvald A (1993) Relationship Between Intrinsic Connections and Functional Architecture Revealed by Optical Imaging and in vivo Targeted Biocytin Injections in Primate Striate Cortex. *Proc Natl Acad Sci USA* 90:10469–10473.

Malik J (2001) Visual grouping and object recognition. *Proceedings 11th International Conference on Image Analysis and Processing* pp. 612–621.

Malonek D, Grinvald A (1996) Interactions Between Electrical Activity and Cortical Microcirculation Revealed by Imaging Spectroscopy: Implications for Functional Brain Mapping. *Science* 272:551–554.

Mannion DJ, Mcdonald JS, Clifford CWG (2010a) Orientation Anisotropies in Human Visual Cortex. *Journal of Neurophysiology* 103:3465–3471.

Mannion DJ, McDonald JS, Clifford CWG (2010b) The influence of global form on local orientation anisotropies in human visual cortex. *NeuroImage* 52:600–5.

Markram H (2006) The blue brain project. *Nature reviews. Neuroscience* 7:153–60.

Marr D, Hildreth E (1980) Theory of edge detection. *Proceedings of the Royal Society of London. Series B.* 207:187.

Marr D, Poggio T (1976) From understanding computation to understanding neural circuitry. *AI Memo* 357:1.

Marsaglia G, Tsang W, Wang J (2003) Evaluating Kolmogorov's distribution. *Journal of Statistical Software* 8:1–4.

Martin C, Berwick J, Johnston D, Zheng Y, Martindale J, Port M, Redgrave P, Mayhew J (2002) Optical imaging spectroscopy in the unanaesthetised rat. *Journal of neuroscience methods* 120:25–34.

Maurer D, Lewis TL, Brent HP, Levin AV (2011) Rapid Improvement in the Acuity of Infants After Visual Input. *Science* 286:108–110.

McGuire BA, Gilbert CD, Rivlin PK, Wiesel TN (1991) Targets of horizontal connections in macaque primary visual cortex. *J Comp Neurol* 305:370–392.

McLoughlin N, Cotton P, Schiessl I (2005) A continuous smooth map of space in the primary visual cortex of the common marmoset. *Perception* 34:967–974.

Miesenböck G (2009) The optogenetic catechism. *Science* 326:395–9.

Mitchell DE, Kennie J, Schwarzkopf DS, Sengpiel F (2009) Daily mixed visual experience that prevents amblyopia in cats does not always allow the development of good binocular depth perception. *Journal of Vision* 9:7–22.

Mitchison G, Crick F (1982) Long axons within the striate cortex: their distribution, orientation, and patterns of connection. *Proc Natl Acad Sci USA* 79:3661–3665.

Motoyoshi I, Kingdom FAA (2010) The role of co-circularity of local elements in texture perception Isamu Motoyoshi. *Journal of Vision* 10:1–8.

Movshon J, Thompson I, Tolhurst D (1978) Receptive field organization of complex cells in the cat's striate cortex. *The Journal of physiology* 283:79.

Najemnik J, Geisler W (2008) Eye movement statistics in humans are consistent with an optimal search strategy. *Journal of vision* 8:1–14.

Najemnik J, Geisler WS (2005) Optimal eye movement strategies in visual search. *Nature* 434:387–91.

Nauhaus I, Benucci A, Carandini M, Ringach DL (2008) Neuronal selectivity and local map structure in visual cortex. *Neuron* 57:673–679.

Nelson JI, Frost BJ (1985) Intracortical facilitation among co-oriented, co-axially aligned simple cells in cat striate cortex. *Exp Brain Res* 61:54–61.

Niebur E, Worgotter F (1994) Design Principles of Columnar Organization in Visual Cortex. *Neural Computation* 6:602–614.

O'Doherty JP, Dayan P, Friston K, Critchley H, Dolan RJ (2003) Temporal difference models and reward-related learning in the human brain. *Neuron* 38:329–37.

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.

Ohki K, Chung S, Kara P, Hübener M, Bonhoeffer T, Reid RC, Hubener M, al. E (2006) Highly ordered arrangement of single neurons in orientation pinwheels. *Nature* 442:925–928.

Okajima K (2004) Binocular disparity encoding cells generated through an Infomax based learning algorithm. *Neural networks* 17:953–62.

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

Olshausen BA, Field DJ (1997) Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vision Res* 37:3311–3325.

Olshausen BA, Field DJ (2004) What is the other 85% of V1 doing. *Problems in Systems Neuroscience* pp. 1–29.

Olshausen B (2003) Principles of image representation in visual cortex. *The visual neurosciences* pp. 1603–1615.

Olshausen BA, Field DJ (2004) Sparse coding of sensory inputs. *Current Opinion in Neurobiology* 14:481–487.

Osindero S, Welling M, Hinton GE (2006) Topographic product models applied to natural scene statistics. *Neural Computation* 18:381–414.

Paffen CLE, Alais D, Verstraten FAJ (2006) Attention Speeds Binocular Rivalry. *Psychological Science* 17:752.

Paffen CLE, Verstraten FAJ (2008) Attention-based perceptual learning increases binocular rivalry suppression of irrelevant visual features. *Journal of Vision* 8:1–11.

Parkhurst D, Law K, Niebur E (2002) Modeling the role of salience in the allocation of overt visual attention. *Vision research* 42:107–23.

Parkinson R (1999) *Cracking Codes: The Rosetta Stone and Decipherment* University of California Press.

Párraga CA, Brelstaff G, Troscianko T, Moorehead IR (1998) Color and luminance information in natural scenes. *Journal of the Optical Society of America A* 15:563–569.

Payne BR, Peters A (2001) The concept of cat primary visual cortex. In Payne BR, Peters A, editors, *The cat primary visual cortex*, pp. 1–108. San Diego: Academic Press.

Perlovsky L (1998) Conundrum of combinatorial complexity. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20:666–670.

Perrinet L, Samuelides M, Thorpe S (2004) Sparse spike coding in an asynchronous feedforward multi-layer neural network using matching pursuit. *Neurocomputing* 57:125–134.

Pettigrew JD (1986) The evolution of binocular vision In *Visual Neuroscience*, pp. 208–22. CRC Press.

Poggio G, Fischer B (1977) Binocular interaction and depth sensitivity in striate and prestriate cortex of behaving rhesus monkey. *J Neurophysiol* 40:1392–1405.

Poghosyan V, Ioannides AA (2008) Attention modulates earliest responses in the primary auditory and visual cortices. *Neuron* 58:802–813.

Polat U, Sagi D (1993) Lateral interactions between spatial channels: suppression and facilitation revealed by lateral masking experiments. *Vision Res* 33:993–999.

Polimeni JR, Granquist-Fraser D, Wood RJ, Schwartz EL (2005) Physical limits to spatial resolution of optical recording: clarifying the spatial structure of cortical hypercolumns. *Proc Natl Acad Sci USA* 102:4158–63.

Prince S, Cumming B, Parker A (2002) Range and Mechanism of Encoding of Horizontal Disparity in Macaque V1. *Journal of Neurophysiology* 87:209.

Purves D, Augustine GJ, Fitzpatrick D, Hall WC, LaMantia AS, McNamara JO, White LE (2007) *Neuroscience* Sinauer Associates, Inc., 4 edition.

Ratzlaff EH, Grinvald a (1991) A tandem-lens epifluorescence macroscope: hundred-fold brightness advantage for wide-field imaging. *Journal of neuroscience methods* 36:127–37.

Reinagel P, Zador AM (1999) Natural scene statistics at the centre of gaze. *Network* 10:341–350.

Renninger LW, Verghese P, Coughlan J (2007) Where to look next? Eye movements reduce local uncertainty. *Journal of Vision* 7:1–17.

Ringach D, Shapley R, Hawken M (2002) Orientation Selectivity in Macaque V1: Diversity and Laminar Dependence. *Journal of Neuroscience* 22:5639.

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

Ringach D, Bredfeldt C, Shapley R, Hawken M (2002) Suppression of Neural Responses to Nonoptimal Stimuli Correlates With Tuning Selectivity in Macaque. *Journal of Neurophysiology* 87:1018.

Ringach D, Hawken M, Shapley R (2003) Dynamics of orientation tuning in macaque V1: the role of global and tuned suppression. *Journal of neurophysiology* 90:342.

Rockland KS, Lund JS (1982) Widespread periodic intrinsic connections in the tree shrew visual cortex. *Science* 215:1532–1534.

Rockland KS, Lund JS, Humphrey AL (1982) Anatomical binding of intrinsic connections in striate cortex of tree shrews (Tupaia glis). *The Journal of comparative neurology* 209:41–58.

Roelfsema PR (2006) Cortical Algorithms for Perceptual Grouping. *Annu. Rev. Neurosci.* 29:203–227.

Roelfsema P, Lamme V, Spekreijse H (1998) Object-based attention in the primary visual cortex of the macaque monkey. *Nature* 395:376–381.

Rucci M (2008) Fixational eye movements, natural image statistics, and fine spatial vision. *Network: Computation in Neural Systems* 19:253–285.

Rucci M, Iovin R, Poletti M, Santini F (2007) Miniature eye movements enhance fine spatial detail. *Nature* 447:851–854.

Ruderman D, Bialek W (1994) Statistics of natural images: scaling in the woods. *Physical Review Letters* 73:814–817.

Russell WM, Burch R (1959) *The principles of humane experimental technique* Methuen.

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.

Sadek AS, Nikoli K, Forshaw M (2004) Parallel information and computation with restitution for noise-tolerant nanoscale logic networks. *Nanotechnology* 15:192–210.

Sato M, Stryker MP (2008) Distinctive features of adult ocular dominance plasticity. *J Neurosci* 28:10278–10286.

Schiessl I, McLoughlin N (2003) Optical imaging of the retinotopic organization of V1 in the common marmoset. *Neuroimage* 20:1857–1864.

Schmidt KE, Goebel R, Löwel S, Singer W (1997) The perceptual grouping criterion of colinearity is reflected by anisotropies of connections in the primary visual cortex. *Eur J Neurosci* 9:1083–1089.

Schnabel M, Kaschube M, Löwel S, Wolf F (2007a) Random waves in the brain: Symmetries and defect generation in the visual cortex. *The European Physical Journal-Special Topics* .

Schnabel MR, Kaschube M, White LE, Coppola DM, Wolf F (2007b) Pinwheel stability, pattern selection and the geometry of visual space. In *Program No. 346.7, Abstract Viewer/Itinerary Planner*. Washington DC: Society for Neuroscience.

Schnabel M (2008) A Symmetry of the Visual World in the Architecture of the Visual Cortex Ph.D. diss., University of Göttingen, Germany.

Schneider E, Dera T, Bartl K, Boening G, Bardins S, Brandt T (2005) Eye Movement Driven Head-Mounted Camera: It Looks Where the Eyes Look In *IEEE International Conference on Systems, Man and Cybernetics*, pp. 2437–2442. Ieee.

Schuett S, Bonhoeffer T, Hübener M (2002) Mapping retinotopic structure in mouse visual cortex with optical imaging. *J Neurosci* 22:6549–6559.

Schultz W, Dickinson a (2000) Neuronal coding of prediction errors. *Annual review of neuroscience* 23:473–500.

Schultz W, Dayan P, Montague P (1997) A neural substrate of prediction and reward. *Science* 275:1593.

Schwartz O, Hsu A, Dayan P (2007) Space and time in visual context. *Nature reviews. Neuroscience* 8:522–35.

Schwarzkopf D, Kourtzi Z (2008) Experience Shapes the Utility of Natural Statistics for Perceptual Contour Integration. *Current Biology* 18:1162–1167.

Schwarzkopf DS, Vorobyov V, Mitchell DE, Sengpiel F (2007) Brief daily binocular vision prevents monocular deprivation effects in visual cortex. *Eur J Neurosci* 25:270–280.

Sengpiel F, Gödecke I, Stawinski P, Hübener M, Lowel S, Bonhoeffer T (1998) Intrinsic and environmental factors in the development of functional maps in cat visual cortex. *Neuropharmacology* 37:607–621.

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

Sengpiel F, Kind PC (2002) The role of activity in development of the visual system. *Curr Biol* 12:818–26.

Shadlen MN, Newsome WT (1994) Noise, neural codes and cortical organization. *Current opinion in neurobiology* 4:569–79.

Shalev GY, Paradiso MA (2005) The effects of natural scenes and saccades on V1 orientation selectivity. *J. Vis.* 5:596.

Shannon CE (1948) The mathematical theory of communication. *Bell Syst. Tech. J* 27:379–423.

Shatz C, Lindstrom S, Wiesel T (1977) The distribution of afferents representing the right and left eyes in the cat's visual cortex. *Brain Res* 131:103–116.

Shatz CJ, Stryker MP (1978) Ocular dominance in layer IV of the cat's visual cortex and the effects of monocular deprivation. *J Physiol* 281:267–283.

Shepard K, Narayanan V (1996) Noise in deep submicron digital design In *Proceedings of International Conference on Computer Aided Design*, pp. 524–531. IEEE Comput. Soc. Press.

Shouval H, Intrator N, Law CC, Cooper LN (1996) Effect of binocular cortical misalignment on ocular dominance and orientation selectivity. *Neural Computation* 8:1021–1040.

Sigman M, Cecchi GA, Gilbert CD, Magnasco MO (2001) On a common circle: natural scenes and Gestalt rules. *Proc Natl Acad Sci USA* 98:1935–1940.

Simoncelli EP (2005) Statistical modeling of photographic images. *Handbook of Video and Image Processing* p. 13.

Simoncelli EP, Freeman WT (1995) The steerable pyramid: A flexible architecture for multi-scale derivative computation. *2nd Annual Intl. Conf. on Image Processing* .

Simoncelli EP, Olshausen BA (2001) Natural image statistics and neural representation. *Annu Rev Neurosci* 24:1193–1216.

Simoncelli E (1998) Statistical models for images: compression, restoration and synthesis In *Asilomar Conference on Signals, Systems and Computers*, pp. 673–678. IEEE Comput. Soc.

Sirotin YB, Das A (2010) Zooming in on mouse vision. *Nature neuroscience* 13:1045–6.

Slovin H, Arieli A, Hildesheim R, Grinvald A (2002) Long-term voltage-sensitive dye imaging reveals cortical dynamics in behaving monkeys. *J Neurophysiol* 88:3421–3438.

Smith EC, Lewicki MS (2006) Efficient auditory coding. *Nature* 439:978–982.

Sonnemann E (2008) General solutions to multiple testing problems. *Biometrical Journal* 50:641–656.

Srivastava A, Lee A, Simoncelli E, Zhu S (2003) On advances in statistical modeling of natural images. *Journal of mathematical imaging and vision* 18:17–33.

Stettler DD, Das A, Bennett J, Gilbert CD (2002) Lateral connectivity and contextual interactions in macaque primary visual cortex. *Neuron* 36:739–750.

Stosiek C, Garaschuk O, Holthoff K, Konnerth A (2003) In vivo two-photon calcium imaging of neuronal networks. *Proc Natl Acad Sci USA* 100:7319–7324.

Strogatz SH (2001) Exploring complex networks. *Nature* 410:268–76.

Stryker MP (1978) Postnatal development of ocular dominance columns in layer IV of the cat's visual cortex and the effects of monocular deprivation. *Archives italiennes de biologie* 116:420–426.

Stryker MP, Sherk H, Leventhal AG, Hirsch HV (1978) Physiological consequences for the cat's visual cortex of effectively restricting early visual experience with oriented contours. *Journal of Neurophysiology* 41:896–909.

Sur M, Leamey CA (2001) Development and plasticity of cortical areas and networks. *Nat. Rev. Neurosci* 2:251–262.

Swindale N (1990) Is the cerebral cortex modular? *Trends in Neurosciences* 13:487–492.

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

Swindale NV (1998) Cortical organization: modules, polymaps and mosaics. *Curr Biol* 8:270–273.

Swindale N (2008) Visual map. *Scholarpedia* 3:4607.

Swindale NV (2007) Visual cortex: more wiggle room for the brain. *Curr Biol* 17:1055–7.

Tamura H, Sato H, Katsuyama N, Hata Y, Tsumoto T (1996) Less segregated processing of visual information in V2 than in V1 of the monkey visual cortex. *Eur J Neurosci* 8:300–309.

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.

Tanaka S, Tani T, Ribot J, O'Hashi K, Imamura K (2009) A postnatal critical period for orientation plasticity in the cat visual cortex. *PloS One* 4:e5380.

Tang H, Li H, Yi Z (2010) Online learning and stimulus-driven responses of neurons in visual cortex. *Cognitive Neurodynamics* 5:77–85.

Tani T, Tanaka S (2008) Interocular imbalance in visual experience affects orientation map formation in cat visual cortex, program no. 769.8/kk14 In *Society for Neuroscience Annual Meeting, Washington, DC, abstract only.*

Teh Y, Welling M, Osindero S, Hinton G (2003) Energy-based models for sparse overcomplete representations. *The Journal of Machine Learning Research* 4:1235–1260.

Tong F, Meng M, Blake R (2006) Neural bases of binocular rivalry. *Trends in Cognitive Sciences* 10:1364–6613.

Tootell RB, Switkes E, Silverman MS, Hamilton SL (1988) Functional anatomy of macaque striate cortex. II. Retinotopic organization. *J Neurosci* 8:1531–1568.

Ts'o DY, Gilbert CD, Wiesel TN (1986) Relationships between horizontal interactions and functional architecture in cat striate cortex as revealed by cross-correlation analysis. *J Neurosci* 6:1160–1170.

Tsodyks M, Kenet T, Grinvald A, Arieli A (1999) Linking Spontaneous Activity of Single Cortical Neurons and the Underlying Functional Architecture. *Science* 286:1943–1946.

Tusa RJ, Palmer LA, Rosenquist AC (1978) The retinotopic organization of area 17 (striate cortex) in the cat. *J Comp Neurol* 177:213–235.

Tzou KH (1992) High-order entropy coding for images. *IEEE Transactions on Circuits and Systems for Video Technology* 2:87–89.

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

van der Maaten L, Hinton G (2008) Visualizing data using t-SNE. *Journal of Machine Learning Research* 9:2579–2605.

van Hateren JH, Ruderman DL (1998) Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex. *Proc Biol Sci* 265:2315–2320.

van Hateren JH, van der Schaaf A (1998) Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc Biol Sci* 265:359–366.

Vanzetta I, Grinvald A (1999) Increased cortical oxidative metabolism due to sensory stimulation: implications for functional brain imaging. *Science* 286:1555–1558.

Vnek N, Ramsden BM, Hung CP, Goldman-Rakic PS, Roe AW (1999) Optical imaging of functional domains in the cortex of the awake and behaving monkey. *Proc Natl Acad Sci USA* 96:4057–60.

Vollmer M (2011) High speed and slow motion: the technology of modern high speed cameras. *Physics Education* 46:191.

Vorobyov V, Schwarzkopf DS, Mitchell DE, Sengpiel F (2007) Monocular deprivation reduces reliability of visual cortical responses to binocular disparity stimuli. *Eur J Neurosci* 26:3553–3563.

Walker P (1975) Stochastic properties of binocular rivalry alternations. *Perception* 18:467–473.

Webster MA (2007) Seasonal variations in the color statistics of natural images. *Network: Computation in Neural Systems* 18:213–233.

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.

Wiesel TN, Hubel DH (1965a) Comparison of the effects of unilateral and bilateral eye closure on cortical unit responses in kittens. *Journal of neurophysiology* 28:1029–40.

Wiesel T, Hubel D (1965b) Extent of recovery from the effects of visual deprivation in kittens. *J Neurophysiol* 28:1060–1072.

Williams D, Reingold E, Moscovitch M, Behrmann M (1997) Patterns of Eye Movements During Parallel and Serial Visual Search Tasks. *Canadian Journal of Experimental Psychology* 51:151–164.

Willmore B, Tolhurst DJ (2001) Characterizing the sparseness of neural codes. *Network: Computation in Neural Systems* 12:255–70.

Willmore B, Mazer J, Gallant J (2011) Sparse coding in striate and extrastriate visual cortex. *Journal of Neurophysiology* 44:in press.

Wilson RA, Keil FC (2001) *The MIT encyclopedia of the cognitive sciences* MIT Press.

Wolf F (2005) Symmetry, multistability and long-range interactions in brain development. *Physical Review Letters* 95:208701.

Wolfson SS, Orientation T, Landy MS (1995) Discrimination of orientation-defined texture edges. *Vision research* 35:2863–77.

Xiao Y, Casti A, Xiao J, Kaplan E (2007) Hue maps in primate striate cortex. *Neuroimage* 35:771–786.

Yacoub E, Harel N, Ugurbil K (2008) High-field fMRI unveils orientation columns in humans. *Proc Natl Acad Sci USA* 105:10607–10612.

Yang Z, Heeger DJ, Seidemann E (2007) Rapid and precise retinotopic mapping of the visual cortex obtained by voltage-sensitive dye imaging in the behaving monkey. *J Neurophysiol* 98:1002–1014.

Yap PT, Wu G, Shen D (2010) Human Brain Connectomics: Networks, Techniques, and Applications. *IEEE Signal Processing Magazine* 27:131–134.

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.

Yu K, Blake R (1992) Do recognizable figures enjoy an advantage in binocular rivalry? *Journal of experimental psychology* 18:1158–73.

Zepeda A, Vaca L, Arias C, Sengpiel F (2003) Reorganization of visual cortical maps after focal ischemic lesions. *J Cereb Blood Flow Metab* 23:811–820.

Zhaoping L (2006) Theoretical understanding of the early visual processes by data compression and data selection. *Network: Computation in Neural Systems* 17:301–334.