Correlations, Competition, and Optimality: Modelling the Development of Topography and Ocular Dominance

Geoffrey Goodhill

April 1992

Cognitive Science Research Paper
Serial No. CSRP 226

The University of Sussex
School of Cognitive and Computing Sciences
Falmer
Brighton BN1 9QH
GREAT BRITAIN

Abstract

There is strong biological evidence that the same mechanisms underly the formation of both topography and ocular dominance in the visual system. However, previous computational models of visual development do not satisfactorily address both of these phenomena simultaneously. In this thesis we discuss in detail several models of visual development, focussing particularly on the form of correlations within and between eyes.

Firstly, we analyse the "correlational" model for ocular dominance development recently proposed in [Miller et al 1989]. This model was originally presented for the case of identical correlations within each eye and zero correlations between the eyes. We relax these assumptions by introducing perturbative correlations within and between eyes, and show that (a) the system is unstable to non-identical perturbations in each eye, and (b) the addition of small positive correlations between the eyes, or small negative correlations within an eye, can cause binocular solutions to be favoured over monocular solutions.

Secondly, we extend the elastic net model of [Goodhill 1988, Goodhill & Willshaw 1990] for the development of topography and ocular dominance, in particular considering its behaviour in the two-dimensional case. We give both qualitative and quantitative comparisons with the performance of an algorithm based on the self-organizing feature map of Kohonen, and show that in general the elastic net performs better. In addition we show that (a) both algorithms can reproduce the effects of monocular deprivation, and (b) that a global orientation for ocular dominance stripes in the elastic net case can be produced by anisotropic boundary conditions in the cortex.

Thirdly, we introduce a new model that accounts for the development of topography and ocular dominance when distributed patterns of activity are presented simultaneously in both eyes, with significant correlations both within and between eyes. We show that stripe width in this model can be influenced by two factors: the extent of lateral interactions in the postsynaptic sheet, and the degree to which the two eyes are correlated. An important aspect of this model is the form of the normalization rule to limit synaptic strengths: we analyse this for a simple case.

The principal conclusions of this work are as follows:

- It is possible to formulate computational models that account for (a) both topography and stripe formation, and (b) ocular dominance segregation in the presence of *positive* correlations between the two eyes.
- Correlations can be used as a "currency" with which to compare locality within an eye with correspondence between eyes. This leads to the novel prediction that stripe width can be influenced by the degree of correlation between the two eyes.

Acknowledgments

I would firstly like to thank my supervisor Harry Barrow, for his help, guidance, and the large number of ideas he contributed which helped greatly to shape and direct the work described here.

I am also deeply indebted to both David Willshaw and Peter Dayan. David was a great source of inspiration and encouragement to me throughout the thesis, and I am particularly grateful to him for his hospitality during my occasional visits to Edinburgh. Peter is a very special person who combines both vast technical expertise and deep insight with great generosity and good-naturedness. His contribution to seeing the project through, especially in the closing stages, is immesurable. Chapter 4 of this thesis describes work that was carried out jointly with Peter.

I was also extremely fortunate to have occasional but important feedback on some specific aspects of the work from Geoff Hinton and Christopher Longuet-Higgins. Geoff has helped me in many ways both before and during this project, for which I am deeply grateful. I greatly appreciate both the trouble Christopher has taken to try to instil in me academic standards as rigourous as his own, and our very enjoyable music-making. In addition, I am grateful to him for pointing me in the direction of perturbation analysis in Chapter 4.

Ken Miller has provided help in many ways to this project, for which I am grateful. He willingly answered my many questions about his model, provided kind hospitality in San Francisco, and in particular emphasised to me the properties of subtractive normalization. I also much appreciated occasional conversations with Jack Cowan and Klaus Obermayer, and extremely useful feedback on the thesis from my examiners, Graeme Mitchison and David Young.

Many people at Sussex have helped me in many ways: in particular I am grateful for the useful conversations I've had with Julian Budd, Dave Cliff, Antonio Da Silva Costa, Inman Harvey, Jim Stone, and particularly Alistair Bray, whose patience in helping me with various matters went far beyond what could be reasonably expected. I also thank Richard Dallaway for much practical assistance with LaTeX. More recently in Edinburgh I have gained much from discussions of biological matters with David Price, and also Richard Ribchester.

Deep thanks are due to all those who provided useful comments on earlier drafts of this thesis: David Willshaw, Peter Dayan and Harry Barrow, who greatly improved several versions, and also Inman Harvey and Antonio Da Silva Costa. Any errors that remain are, of course, my own.

I would also like to mention David Griffel, both for his excellent teaching and kind words while I was a mathematics undergraduate, and more recently for some insightful feedback on particular aspects of this work.

I wish to express my very great thanks to Isabelle Desjeux, both for various practical assistance, but more importantly for providing much-appreciated emotional support in some of my weakest moments.

I am grateful to the SERC and SD-Scicon for their financial support, and to the Nuffield Foundation, from whom a Science Travel Grant supported a useful period of study in Edinburgh.

Contents

Chapter 1

Introduction

1.1 Development of brain connections

Brains are mostly wire. The human brain contains about 10^{11} neurons, but roughly 1000-10,000 times as many connections between them. These connections vary in length between very short range (less than a millimetre) linking small patches of brain, to quite long range (several centimetres or more) connecting different regions of the brain, or connecting sensory structures to more central brain structures.

There is strong evidence that to a large extent these connections are "specified" during ontogenesis: fibres growing from a cell are very specific for which other cells they connect to, both in terms of general regions and particular areas within a region. It is clear that, if the connection matrix of the brain were random, the amount of information available in DNA would be grossly insufficient to specify each of these connections explicitly.

However, the overwhelming biological evidence is that connection patterns in the brain are *not* random. A number of particular types of pattern occur frequently throughout the brain, the general properties of which are remarkably similar over a wide range of animal species. Patterns, even if made up of millions of connections, can be specified using a much smaller amount of information than if these connections were random. It thus becomes feasible to imagine that connection strengths are specified in DNA, but in the form of rules that generate patterns of connections rather than an explicit connection matrix (for discussion see [Gierer 1988]). One part of understanding the development of brain connections is therefore to try to discover these rules. Note that we are not claiming that these rules necessarily exist for the purpose of generating patterns: Pattern-formation is simply one of their outcomes in the biological system.

We must also consider the following. Although in general outline connection patterns in the brains of two individuals of the same species are very similar, on a fine scale there can be large variations between them. This is true even for genetically identical individuals (see e.g. [Changeux & Danchin 1976]). Thus besides finding the rules that generate overall patterns, we also have to account for the individual variation in a way that does not imply large amounts of genetic specification. Part of this variation may be due to entirely random processes. However, the experimental data shows that the detailed patterns of connections between sensory structures and more central brain structures can be systematically influenced by the *environment* in which that particular pattern develops. Different environments, both in terms of sense data from the outside world and the biochemical environment inside the brain itself, can lead to differences in the detailed pattern of connections (see e.g. [Blakemore 1978]). It is clear that the outcome of the rules that generate patterns is bound up with the context in which those rules operate (see e.g. [Changeux & Danchin 1976]).

In this thesis we examine a number of different models for the development of connections in the visual system, and in particular study the way in which the outcome of these models is influenced by the form of their inputs.

1.1.1 The importance of development

Understanding of the development of the brain may also provide insight into how *learning* takes place. Rules that account for developmental processes may also have relevance to learning: for instance, issues of normalization may play a similar role in both (for further discussion see [Gierer 1988]). One of the advantages of studying development is that it involves substantial changes in connectivity and the response properties of individual cells that can be observed using standard neuroanatomical and neurophysiological techniques. This is in contrast to the changes that take place during learning: these can be hard to measure, especially given that the overall properties of a network of neurons can be altered significantly by the cumulative effects of just very small changes at a large number of synapses. Thus there is the possibility of more biological information being available to constrain models of development than of learning.

1.2 Topographic maps and striped projections

We now focus on one particular type of connection in the brain: that between individual structures such as lateral geniculate nucleus and V1, or V1 and V2. One of the most prevalent features of these connection patterns is that they are very often arranged in the form of a *topographic map*. We will return later to the issue of a precise mathematical definition of this term: for the present, we take it to mean that neighbouring points in the first structure are connected to neighbouring points in the second. For instance, for mappings from sensory structures this means that a small movement on the sensory surface corresponds to a small movement in the region of brain representing that sensory surface. There are a very large number of such maps in the brain:

for example over 200 in the macaque monkey, where in some cases the entire visual field is represented in an ordered manner across a surface 1mm square (M. Sereno, Personal Communication). In general we will refer to the first structure (e.g. the eye) as the input structure, or input space if we are considering more general mappings (see below), and the second structure as the output or target structure.

Many regions of brain that receive such connections have a columnar structure: moving in a direction orthogonal to the surface, cells with similar response properties are encountered. These regions can almost always be regarded as two-dimensional from the point of view of understanding the abstract properties of the map. This is relatively unproblematic for "spatial" maps, where the mapping is from some sensory surface or other brain surface that is also two-dimensional. However, it can be the case that variables other than, or in addition to, position are represented on the target surface: so called "computational maps" [Knudsen et al 1987]. It is a general property of such maps in the brain that small movements in the input space again lead to small movements on the target surface, but now with some discontinuous jumps that are inevitably produced by mapping to a space of lower dimensionality. The simplest case of this is where one variable additional to position is represented, and this variable can take only two discrete values: in other words, a mapping from two and (exactly) a bit dimensions to two dimensions. In cases such as these in the brain, the map is often laid out in a series of alternating regions, each representing one of the two values of the binary variable. These maps are commonly "striped": that is, the alternating regions are a series of bands which are much longer in one direction than the other. This is well described for the case of sensory inputs from the monkey hand in [Constantine-Paton & Law 1982]:

> The bands separate sensory inputs from the hand according to the kinds of information the input carries. In some bands the neurons respond only to the onset of a touch; in the intervening bands the neurons have a more prolonged response. It is as if the map of the hand in the somatic sensory cortex of the monkey has been constructed by alternating stripes cut out of two distinct maps of the hand.

The representation of position in such maps commonly has two important properties. Firstly, the map of each value of the binary variable (i.e. the map that would be formed if all the stripes representing the other value were removed), is topographic in the above sense. Secondly, these maps are "aligned": the two values of the binary variable at each position on the input surface are mapped to nearby positions in the target structure, given the constraint that they are represented in different stripes.

1.3 Topography and ocular dominance in the visual system

In the course of vertebrate ontogenesis, the ganglion cells of the retina send axons through the eye stalk to the brain, where they project to cells of central structures. In mammals the retinae project to the lateral geniculate nucleus (LGN), which then project to primary visual cortex. In lower vertebrates such as amphibians and fish there are direct projections from retina to optic tectum. In both cases, the mapping is topographic, or as sometimes referred to in this context, "retinotopic". However, in mammals the primary visual cortex is naturally binocularly innervated, i.e. the same part of brain receives inputs from both eyes. In this case, in the early stages of development, the retinotopic projections from the two eyes to the common target structure are uniformly intermingled. As development proceeds, each part of the target gradually becomes more densely innervated by one eye and less densely innervated by the other. Eventually, a striped pattern of innervation can be observed that is reminiscent of the pattern of zebra stripes: so-called "ocular dominance stripes" (see e.g. [Hubel & Wiesel 1977]). The extent to which this happens varies between species. In animals such as cats and monkeys segregation is almost total, and the thickness of these stripes is approximately 400μm. Ocular dominance stripes are also present in humans, with width approximately 500μm to 1000μm [Horton et al 1990]. In lower vertebrates, similar striped projections do not occur naturally but can be induced experimentally (see e.g. [Constantine-Paton & Law 1978]). One very important feature of all these mappings is that they are *plastic*: their detailed form is influenced by the visual experience of the animal, and other procedures which alter the internal environment of the brain.

In one sense there is difference between the case of the striped mapping from the hand described earlier and ocular dominance stripes. In the former case, both values of the binary variable exist at the same spatial location in the input space, whereas in the latter the two values exist at different spatial locations. However, as we argue further later, the crucial point is that, even though cells at similar locations in the two eyes are spatially separated, they will be *correlated*, and it is this fact which defines the correspondance between the two eyes, rather than spatial relationships.

1.3.1 Why study this phenomenon?

There are a number of reasons why topography and ocular dominance are useful developmental phenomena to study for both neuroscientists and theoreticians. Firstly, the visual system is perhaps the most extensively researched and well-understood part of the brain. The form of the map from the eye (via the LGN) to the brain is a macroscopic property which develops as a result of microscopic changes in synaptic strengths. It therefore provides a way of magnifying alterations to various parameters of the system so that their effects can be clearly seen. In addition, ocular dominance is a relatively straightforward property of neurons in the visual system to measure: it is comparatively easy to determine by anatomical and physiological techniques to what extent cells have become monocular. Alterations in the parameters can be seen to have a clear effect on the monocularity of cells and on the global properties of stripe morphology: for instance, stripe width. For these reasons (among others - see chapter 2) there exists a substantial amount of experimental data.

The motivation for studying both phenomena together comes largely from experiments on frogs and goldfish performed by Constantine-Paton and others. Here, stripes were produced in normally monocular tectum by procedures such as diversion of both optic nerves into the same tectum, and implant of a third eye (reviewed in chapter 2). The importance of these experiments is summed up in [Constantine-Paton & Law 1982]:

> Questions about the functional or evolutionary significance of stripes in the brain of three-eyed frogs are clearly irrelevant. The third eye is abnormal, and in the absence of substantial input from both of the eyes to a single optic tectum the normal frog would get no benefit from a mechanism that evolved specifically to segregate tectal inputs into stripes. On the other hand, the survival of free-living frogs, and in particular their ability to catch the insects on which they feed, depends critically on a robust mechanism to ensure that a precise map of the contralateral retinal surface develops in each tectum.[1] We began, therefore, to consider the possibility that stripes might arise from the same developmental mechanism that generates maps.

The hypothesis that the same developmental mechanisms underly both topography and ocular dominance stripes is the main motivation that drives this thesis. As we will see, there are very many models that address either the formation of topographic maps *or* of ocular dominance stripes, but very few that satisfactorily address both. Considering both simultaneously allows a wider range of biological knowledge to be applied to constrain models.

1.4 Levels, models and mathematics

1.4.1 Levels

[Marr 1982] identified three levels of analysis for understanding "....any machine carrying out an information-processing task": in our case, the brain. The highest of these he called the computational level, which is the abstract definition of the problem: the goal of the computation, and why it is appropriate. The second is the algorithmic level, at which an abstract procedure is defined that can solve the problem as formulated at the computational level. Lastly, there is the implementational level, where the algorithm is actually carried out in some piece of hardware: at this level we might have for instance a computer program, or a demonstration of how the algorithm could be realised in a certain piece of brain tissue. Working down through these levels has sometimes been taken as a methodology for research: *first* formulate the problem, *then* find an algorithm, *then* show how it can be implemented.

This thesis does not strictly follow this "top-down" approach. We present algorithms without necessarily defining a computational problem, and use certain implementation issues to constrain theorizing at the algorithmic level. Our reasons for this are as follows [Churchland & Sejnowski 1988, Sejnowski et al 1988, Sejnowski & Tesauro 1989]:

[1] The contralateral eye is on the opposite side of the brain from the tectum under consideration: the ipsilateral eye is on the same side.

- It is often difficult, and sometimes misleading, to attempt to define "the computational problem". Brains evolved rather than being designed by an engineer. It may be the case that a certain mechanism in the brain solves a number of computational problems, or that the phenomenon under investigation is a byproduct of mechanisms for solving a quite different computational problem from that which our intuitions might first suggest.

- Even if we can accurately define a particular computational problem, to paraphrase [Churchland & Sejnowski 1988], "the space of algorithmic possibilities is consumately vast": there may be many possible algorithms that solve the problem. However, by considering the implementational level in conjunction with the algorithmic level, that is restricting attention to those algorithms which satisfy certain constraints arising from consideration of the hardware in which they are implemented (i.e. nervous tissue), one can proceed much more quickly to discovering the algorithm the brain actually employs. Furthermore, the implementational level can often provide inspiration to the algorithmic level.

- There may not be just one computational level and one algorithmic level. The brain contains many levels of structure, and "....it is likely that there is a corresponding multiplicity of algorithmic and computational levels as well" [Sejnowski & Tesauro 1989].

We now discuss the *types* of models presented in this thesis.

1.4.2 Models

[Sejnowski et al 1988] identify two types of brain models which they refer to as "realistic" and "simplifying": these can be seen as extremes on a continuum. A realistic model of a particular brain system attempts to take into account all that is known about the system, down to the level of the biophysics of individual cells. A simplifying model on the other hand tries to extract and work with only those properties of the system that are relevant to the phenomenon under investigation. [Sejnowski et al 1988] draw attention to a number of general problems with investigating realistic models, which also apply directly to the topography and ocular dominance problem. Firstly there is the issue that to formulate a realistic model with confidence requires the support of a huge amount of experimental data, which may not be available in sufficient detail to adequately constrain the parameters of the model. Secondly, the danger arises with realistic models that they may contain so many parameters and low-level details that very little higher-level understanding of the system can be obtained from the model. Thirdly, there is the practical problem that the computational requirements for simulating the vast detail of a realistic model may restrict consideration to an uninterestingly small fragment of the system.

The models discussed in this thesis tend, to differing degrees, towards the simplifying end of the spectrum. That is, they are formulated in terms of just a few key parameters that are assumed to be important in map formation. For instance, we study the effects of correlations in activity between cells in the retina or LGN, and the lateral connectivity in the cortex/tectum. These are simply not known at a detailed enough level in the natural system to be investigated in the context of a realistic model of map formation. We prefer therefore to examine the influence of general correlation and interaction parameters, abstracted away from cellular details. Our motivation is the hypothesis that these are the key parameters that determine the properties of the map, in particular the stripe width. Lastly, regarding the issue of computational resources, a general problem with studying map formation is that it is a property of many rather than just a few cells. Therefore it is not practically possible to devote large computational resources to simulating each individual cell.

1.4.3 Mathematics

What mathematical statements can be made about the topography and ocular dominance problem, and to what extent are these useful?

We note firstly that the term "topographic" does not have a well-defined mathematical meaning. It is used in the biological and theoretical literature pertaining to this problem to mean loosely that, in general, either points close together in the target structure are mapped to points close together in the input structure, or that points close together in the input structure map to points that are close together in the target structure (we return to issues of the difference between these two statements in chapter 5). Here "close together" is usually a purely qualitative measure, perhaps based on the visual appearance of the map as realised in some biological experiment or as a picture on a computer screen.

Although topography is not well-defined, unfortunately some authors appear to use this word interchangeably with "topology", which does have a precise meaning. A topological map, or homeomorphism, is a map that is one-to-one and bicontinuous (i.e. continuous in both directions) [Alexandroff 1961]. We argue that it may not be useful to apply such a definition to brain maps. Firstly, biological maps are virtually never one-to-one: each fibre from an input cell may contact thousands of cells in the target structure. Secondly, given a mapping between many thousands or even millions of cells in the brain, it would be most unlikely that there was not even one deviation from perfect ordering, and one deviation is all that is required for a map to fail to be topological. Another (ab)use of terminology is to drop the word "topographic" altogether, and simply use the word "map" to mean topographic map (as in the quote from [Constantine-Paton & Law 1982] above), Here "map" is referring to a much more restricted class of entities than is defined by the mathematical meaning of the word.

It seems clear that what is needed is a measure of the *degree* of topography of a particular map, that is, the degree to which it approaches some ideal of perfect continuity (or bicontinuity). Indeed, this is the way in which the word is often used informally in the literature. However, these distinctions are again qualitative. A number of quantitative measures of topography have been proposed (see e.g. [Rankin & Cook 1986, Roe et al, 1990, Cottrell & Fort 1986, Érdi & Barna, 1984, Kohonen 1988]). Unfortunately each makes particular assumptions with respect to an individual system, and cannot be applied in all situations.

Another important reason for not using the mathematical concept of a topological map here is that the ocular dominance case involves mapping between spaces of different dimensionalities, and such maps (if they are assumed 1-1) cannot preserve all neighbourhoods of the higher dimensional space in the lower dimensional space. However, the topography and ocular dominance map still has some property in common with a "perfectly" topographic map as could exist between spaces of the same dimensionality, and we will tend, in common with most others, to refer to this case as being topographic with occasional discontinuities.

1.5 Mechanisms of map formation

1.5.1 Cooperation and competition

Two important concepts for understanding pattern formation in physical (and more particularly biological) systems are cooperation and competition. To explain these we will refer to an abstract system, initially disordered, which organizes itself due to the growth of various "modes" (patterns) of the system. *Cooperation* is the process whereby different modes can enhance each other's strength (ability to dominate in the final system). Balanced against this are often processes of *competition*. This has been defined by [Keddy 1989] as

> the negative effects which one mode[2] has upon another by consuming, or controlling access to, a resource that is limited in availability.

Here a resource is something required for the continued growth or maintenance of modes. Thus cooperation encourages mutual support between modes, while competition selects between modes. Both can be important for the development of interesting structure. Without cooperation, modes other than the very strongest have no influence on development, while without competition all modes can grow unrestrained, possibly obscuring the interesting structure contained within each.

Classic examples of physical systems which can be described by these concepts are the Rayleigh-Bénard instability, and the organization of crystal lattices and magnetic domains (for further discussion see [von der Malsburg & Singer 1988]). One of the first applications to biological systems was made by [Turing 1952], who formulated mathematically mechanisms by which an array of cells can become chemically differentiated from an almost uniform initial state. For other examples of biological interest see [Meinhardt 1982]. We now focus on the relevance of the concepts of cooperation and competition to the formation of brain maps, within the context of an extremely influential theory of the mechanism underlying changes of synaptic strength in the nervous system.

[2] The word appearing here was originally "organism", since [Keddy 1989] is primarily concerned with Population Biology.

1.5.2 Hebbian rules

[Hebb 1949] made the following hypothesis regarding the circumstances under which the connection between two cells in the brain is increased in strength:

> *When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.*

This has become known as the Hebb rule, and the term "Hebbian" is generally applied to any postulated synaptic learning rule that implements some version of this idea (see [Stent 1973] for discussion of a possible physiological implementation). Direct evidence for the existence of such a rule comes from, among others, various experiments in the visual system (reviewed in [Rauschecker 1991]).

1.5.3 Some consequences of Hebbian rules

We now relate this to the concepts of cooperation and competition described above. There are two ways in which the Hebb rule can be seen as implementing a form of cooperation. Firstly, the postsynaptic cell cooperates with the presynaptic cells with which it simultaneously fires: these connection strengths are increased, making it yet more likely that the postsynaptic cell will fire when these presynaptic cells fire, thus further increasing the connection strength, and so on. Secondly, presynaptic cells whose activities are *correlated* will tend to cooperate in forming strong connections with the same postsynaptic cell. Their joint firing makes it more likely that the postsynaptic cell will fire, thus increasing all their connection strengths, and so on. This form of cooperation between correlated presynaptic cells is crucially important in the models of map formation discussed in this thesis.

The Hebb rule as expressed above prescribes only increases in synaptic strengths (which we will also refer to as "weights"). Applied just in this form, all weights would grow without bounds, however small the correlation between the postsynaptic cell and each presynaptic cell. One way of limiting synaptic strengths is to specify that weights are reduced under certain pairings of pre- and postsynaptic activity. Recent (though controversial) evidence that such a process may be operating in the brain comes from the discovery of long term depression in the hippocampus [Stanton & Sejnowski 1989]. Here, the pairing of presynaptic activity with postsynaptic inactivity produces a decrease in synaptic strength (see [Willshaw & Dayan 1990] for a discussion of some of the computational issues this raises). An alternative way of preventing unbounded growth is to introduce some form of competition, where the resource is some substance required for the maintenance of synaptic strengths. This is often implemented in theoretical models by means of weight normalization: each pre- or postsynaptic cell (or sometimes both) is assumed to have a fixed amount of "weight substance" that it can share out between the connections to or from that cell. This normalization can take various forms, and the form can greatly influence the final pattern of weights, as we will see later.

1.6 Introduction to the thesis

In this thesis we investigate three different models of map formation in the visual system, and compare them along a number of dimensions that speak to their adequacy as models of the biological system. The most general considerations are as follows:

- Their ability to form maps and stripes from the same mechanism.
- The degree to which their mechanisms and parameters can be interpreted biologically.
- Their ability to form appropriate mappings in the presence of positive correlations between the two eyes (the significance of this will become apparent in chapters 2 and 3).
- The degree to which the morphology of stripes produced by the model matches that of natural stripes, and the effect on model results of monocular deprivation.
- The factors that determine the width of the stripes.

We now briefly summarize the content and main results of the thesis.

1.6.1 Chapter 2: Review of biological data

In this chapter we outline what is known experimentally about the development of topography and ocular dominance, and draw attention to the results most relevant to the thesis. We review the literature regarding both amphibians and fish (in particular frogs and goldfish), and higher mammals (in particular cats and monkeys), and argue that the development of topography and ocular dominance is similar enough in both systems for both to be modelled by the same abstract mechanisms.

1.6.2 Chapter 3: Previous theoretical work

Here we review previous theoretical work on the formation of maps and stripes in the visual system. We consider two levels: firstly, qualitative hypotheses as have commonly been used to explain various aspects of map formation, and secondly specific computational theories that have been implemented as computer programs. We conclude that most models of topography do not naturally extend to also account for ocular dominance segregation.

1.6.3 Chapter 4: Correlational models

In this chapter we present theoretical analysis of a recent and influential model of stripe formation [Miller et al 1989], related to the model of visual receptive field development of [Linsker 1986]. Both assume a pre-existing topography. Cortical unit receptive field development is driven by the correlations within and between two populations of presynaptic (LGN) units, and stripe width is determined mainly by the extent of lateral connections in the cortex. Due to the simple form of the learning rule, the behaviour of the system can be largely predicted from the eigendecomposition of the matrix of presynaptic correlations. In Miller's original analysis (see e.g. [Miller 1990(a)]) he assumed that the correlational structure of activity in both eyes is exactly the same, and that there are no correlations between the eyes. We present analysis of the effect of relaxing these two assumptions by introducing perturbating correlations, which reveals that in the first case there is an instability in the model, and in both cases monocularity can begin to break down. Generally speaking, in the perturbed case eigenvalues of binocular solutions are more likely to be larger than eigenvalues of monocular solutions, and thus there is a greater tendency for cortical units to evolve to a binocular, as opposed to monocular, final state. Some examples of numerically calculated eigenvectors and eigenvalues are presented to support these conclusions.

1.6.4 Chapter 5: Elastic and Kohonen models

This chapter investigates the properties of the elastic net model of topography and ocular dominance, and also those of a related algorithm based on Kohonen's self-organizing feature map algorithm. The elastic net is a developmental algorithm which can be interpreted as performing an optimization. It attempts to form a map, from an input space of points to a target sheet, that minimizes the separation between the representations in the target sheet of points that are neighbouring in the input space. It was first applied to the topography and ocular dominance problem in [Goodhill 1988, Goodhill & Willshaw 1990]: here we significantly extend this work by (among other things) an investigation of the effect of certain parameters on the pattern of stripes in the two dimensional case. In addition, we also investigate a version of Kohonen's algorithm for the same problem, and compare its performance with that of the elastic net. In both these models stripe width is determined by parameters which we interpret in terms of the ratio of correlations within and between eyes. As a final point we show that a global orientation for ocular dominance stripes in the elastic net case can be produced by anisotropic boundary conditions in the cortex.

1.6.5 Chapter 6: A competitive model

The final model presented considers a sheet of interacting cortical cells that compete for the right to respond to input patterns in the two eyes (similar to [Barrow 1987]). Unlike many earlier models, the mechanism we propose forms appropriate maps in the presence of distributed, simultaneous activity in the two eyes. An important difference with earlier competitive models is the form of the normalization rule used to limit synaptic strengths: we analyse this in a simple case. We show that stripe width in this model is influenced by two factors: the extent of lateral interactions in the cortex (as in Miller's model), and the degree of correlation between the two eyes (as in the elastic and Kohonen models).

1.6.6 Chapter 7: Conclusions

Here we sum up the main conclusions and contributions of the thesis. These are as follows:

- Addition of perturbing correlations in Miller's model of ocular dominance leads to an increased tendency to produce binocular rather than monocular units.
- Elastic and Kohonen models naturally produce topography and ocular dominance from the same developmental mechanisms, and can account for the effects of monocular deprivation.
- A new model is presented that for the first time accounts for the development of topography and ocular dominance in the presence of distributed activity simultaneously in the two eyes. Stripe width in this model is related to both the degree of correlation between the two retinae and the width of interactions in the cortex.

Finally we suggest possible directions for future work, and ways in which the predictions regarding the effect of input correlations on stripe width could be tested in the natural system.