

Model of all known spatial maps in primary visual cortex

Tobias Fischer



Master of Science
Artificial Intelligence
School of Informatics
University of Edinburgh
2014

Abstract

The primary visual cortex of mammals is the most extensively studied area in the visual system. The first studies discovered that there is a retinotopic mapping from the retina to the primary visual cortex. Retinotopic mapping is where neighboring neurons in the cortex respond to neighboring locations on the retina. In further research various other cortical maps such as orientation maps and color maps were discovered. All the cortical maps are overlaid onto the same set of neurons, and there is evidence that they interact with each other.

There are a variety of models aiming to replicate the properties of neurons in the primary visual cortex. The majority of these focus on a small subset of all known spatial cortical maps. For this thesis, an all maps model based on the Gain Control, Adaptation, Laterally (GCAL) model has been developed. There, it has been suggested that the underlying principles of firing rate neurons, arranged in two dimensional sheets, using Hebbian learning to adapt to either artificial input patterns or natural images, can account for a variety of different maps, as well as their combination. This required substantial work on the software package in use, Topographica, which led to a superior system to define models. It is likely this will be used by almost all users of Topographica in the future.

This thesis is a small step towards the goal of gaining an understanding of why V1 is wired as it is in mammals, and eventually how the whole visual system works. The improvements in Topographica which have been made in this project have resulted in the production of maintainable, modular models, which will hopefully lead to more exciting research in the area of computational neuroscience of vision. The model which has been built will help in gaining insights to the interplay of the various cortical maps in the primary visual cortex.

Acknowledgements

I am very grateful to my supervisor, Jim Bednar, who has always encouraged my work, and has provided valuable advice. He answered all my questions promptly, and his knowledge and ideas brought this project much further than I could have done by myself.

I also thank the other members of the Computational Systems Neuroscience Group, especially Jean-Luc Stevens and Philipp Rudiger. It was an excellent working atmosphere, and their suggestions and assistance improved the quality of my code considerably.

A special thanks goes to John K. Tsotsos, my former supervisor, for sharing his passion about visual systems; he is an inspiring example for a scientist. Our conversations and his talks have influenced my research interests substantially. Also thank you to his research associate, Eugene Simine, as he did not hesitate to allow me to remotely access one of their machines, which greatly helped in the building of resource demanding models.

I would also like to express my gratitude to Jill Dyne, my girlfriend, for her continued support. She always listened to all the issues I encountered, and showed interest in the work I have done.

I am indebted to my parents, who have been supportive all my life, always showed their love, and believed in me, and whatever I have done.

Finally, I would like to thank the German National Academic Foundation for their financial support. Without them, it would not have been possible for me to study at this wonderful university. Over the last years, they have provided me the opportunity to meet many interesting people, who have broadened my horizon.

Declaration

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

(Tobias Fischer)

Table of Contents

1	Introduction	1
2	Biological background	3
2.1	Visual system	3
2.2	Cortical maps found in animals	5
2.2.1	Measurement methods	6
2.2.2	Retinotopic mapping	7
2.2.3	Orientation map	7
2.2.4	Ocular dominance map	9
2.2.5	Disparity map	10
2.2.6	Direction map	11
2.2.7	Spatial frequency map	12
2.2.8	Hue map	13
2.3	Map interaction	13
2.4	Map development	14
3	Model background	17
3.1	Sparse coding	17
3.2	Elastic net models	19
3.3	Models based on self-organizing maps	20
3.3.1	Self-Organizing Map	20
3.3.2	Gain Control, Adaptation, Laterally connected model	22
3.3.3	Multiple maps using GCAL	24
3.4	Open questions	29
4	Architecture	31
4.1	Pattern generation	32
4.1.1	Feature coordinators	32

4.1.2	Pattern coordinator	34
4.2	Submodels	36
4.2.1	Previous implementation	36
4.2.2	Using a hierarchy of submodels	37
4.2.3	Automatically wiring up sheets	38
4.2.4	Deferred model instantiation	39
4.2.5	Summary submodels	40
4.3	Motion model	41
4.3.1	Previous proposals	41
4.3.2	New approach: Multiple projections from LGN sheets to V1	42
4.4	Discussion	43
4.4.1	Feature coordinators	43
4.4.2	Pattern coordinator	44
4.4.3	Submodels	45
4.5	Architectural summary	45
5	Methods	47
5.1	Orientation preference	47
5.2	Ocular preference	49
5.3	Disparity preference	51
5.4	Ocular and disparity preference	52
5.5	Direction preference	52
5.6	Spatial frequency preference	54
5.7	Color preference	55
5.8	Color and ocular preference	56
5.9	Combined model	58
5.10	Summary	60
6	Results	63
6.1	Orientation preference	64
6.2	Ocular preference	65
6.3	Disparity preference	67
6.4	Ocular and disparity preference	68
6.5	Direction preference	70
6.6	Color preference	73
6.7	Color and ocular preference	74

6.8	Spatial frequency preference	75
6.9	Combined model	76
7	Discussion and Future Work	79
7.1	Discussion of the results	79
7.1.1	Orientation preference	79
7.1.2	Ocular dominance preference	80
7.1.3	Disparity preference	80
7.1.4	Ocular and disparity preference	80
7.1.5	Direction preference	80
7.1.6	Color preference	81
7.1.7	Color and ocular dominance preference	81
7.1.8	Spatial frequency preference	81
7.1.9	Combined model	81
7.1.10	Summary	82
7.2	Future work scientifically	83
7.2.1	Interdependence of maps	83
7.2.2	Uniform coverage	83
7.2.3	Relative order of map development	83
7.2.4	Gaussian stimuli for color and spatial frequency	84
7.2.5	Lateral connections between spatial frequency channels	84
7.3	Future work for Topographica	84
7.3.1	Improved visualization	85
7.3.2	Stereo camera input streams	85
7.3.3	Component based submodels	85
8	Conclusions	87
	Bibliography	89

List of Figures

2.1	Mapping from the visual field to the LGN and V1	7
2.2	Orientation map in V1 of a macaque monkey	8
2.3	Orientation preference map overlaid by ocular dominance stripes . . .	9
2.4	Micro-architecture disparity map in cat visual cortex	10
2.5	Direction map in V1 of a ferret	11
2.6	Spatial frequency map in V1 of a macaque monkey	12
2.7	Hue map in area V1 of a macaque monkey	13
3.1	Comparison of a sparse variable and a Gaussian variable.	18
3.2	Basic GCAL model with resulting orientation map	24
3.3	LISSOM color model and resulting maps	25
3.4	Orientation map and spatial frequency map in a LISSOM model . . .	26
4.1	Previous proposals to model motion selectivity	42
5.1	Architecture of a GCAL model with neurons selective for orientation .	48
5.2	Oriented Gaussian input pattern	49
5.3	Architecture of a GCAL model with neurons selective for orientation and ocular dominance	50
5.4	Oriented Gaussian input patterns with brightness differences	50
5.5	Oriented Gaussian input pattern with position offset	51
5.6	Oriented Gaussian input pattern with position offset and brightness dif- ference	52
5.7	Oriented Gaussian input pattern moving over time	53
5.8	Architecture of a GCAL model with neurons selective for orientation and spatial frequency	54
5.9	Architecture of a GCAL model with neurons selective for orientation and color	56

5.10	Example natural image input pattern	57
5.11	Visualization of receptive fields in color simulations	57
5.12	Architecture of a GCAL model with neurons selective for orientation, ocular dominance and color	58
5.13	Example of natural image patterns in the case of a two-retina simulation	59
5.14	Architecture of a unified GCAL model with neurons selective for ori- entation, ocular dominance, disparity, direction, spatial frequency and color	61
5.15	Example natural image patterns for the combined model	62
6.1	Orientation map for the basic GCAL model	64
6.2	Orientation map with overlaid ocular dominance boundaries	65
6.3	Ocular dominance preference and selectivity	66
6.4	Ocular preference map of a LISSOM model for comparison	66
6.5	Orientation map with overlaid ocular dominance boundaries for a dis- parity simulation	67
6.6	Disparity preference and selectivity + ocular selectivity for a disparity simulation	68
6.7	Orientation map with overlaid ocular dominance boundaries for a bright- ness difference+disparity simulation	69
6.8	Disparity preference and selectivity + ocular selectivity for a disparity simulation	69
6.9	Orientation map for the motion simulation	70
6.10	Example connection fields of an direction selective neuron	71
6.11	Direction preference map	71
6.12	Orientation map with overlaid direction arrows	72
6.13	Orientation map for the color simulation	73
6.14	Overlaid preference+selectivity maps for orientation and color in a color simulation	74
6.15	Orientation and color maps with overlaid ocular dominance boundaries for a color+ocular dominance simulation	75
6.16	Orientation map for the spatial frequency simulation	76
6.17	Spatial frequency map	77

Chapter 1

Introduction

The human visual system works remarkably well, processing an enormous amount of complex information in real time. Visual processing areas make up approximately half of the brain, and yet we do not have a good understanding of how we carry out everyday tasks, such as face recognition, object detection and obstacle avoidance.

The primary visual cortex (also: striate cortex, V1, area 17) is the first cortical area that processes visual input. Compared to higher visual areas, it is widely studied, and many experiments have been carried out focusing on this area. There, matching qualitatively across species, cortical maps were found. For example, nearby neurons in V1 respond to nearby areas on the retina (Connolly and Van Essen, 1984), and nearby neurons in V1 also prefer similar orientations of the input stimuli (Blasdel, 1992b). These maps were found for a variety of pattern features, i.e. preferences for position, orientation, direction, spatial frequency and color, as well as the eye the pattern was presented on, and the difference of pattern locations in the two eyes. Each map has distinct features, and they are all overlaid on the same set of neurons.

Many computational models which aim to replicate these maps, in order to improve our understanding of the visual system, have been created. In this thesis, the Gain Control, Adaptation, Laterally (GCAL) model Stevens et al. (2013) is extended to cover all known spatial feature maps. The original model produces biologically realistic position preference and orientation maps, using a set of two-dimensional sheets of firing-rate neurons representing different areas of the visual pathway. By adding more feature specific sheets, and presenting more complex input patterns, the same basic principles can account for a variety of cortical maps. The initially unspecific neurons in V1 become selective for input features using normalized Hebbian learning.

In chapter 2 an introduction to the visual system is given. This is of importance

as the models discussed in later chapters are based on these findings. Furthermore, all known spatial maps in the visual cortex are presented and their properties are discussed. As one of the aims of this thesis is to build a model of all known spatial maps, it is important to be able to compare the maps resulting in the model with maps found experimentally in animals (Bednar, 2012). There is clear evidence that the cortical maps are not independent of each other, i.e. they are related to each other (Blasdel, 1992a; Müller et al., 2000; Yu et al., 2005; Landisman and Ts'o, 2002). Support for this hypothesis as well as various suggestions of how these maps emerge are also shown.

In chapter 3 various models of the primary visual cortex are introduced. They are classified into three main groups, each with distinct properties: models based on the elastic net, sparse coding models and self-organizing models. The focus is on self-organizing maps, as these are the origins for the model presented in this thesis. In addition, various papers which report multiple maps are discussed. The chapter ends with open questions which might be answered by an all maps model.

In order to build an all maps model which is maintainable and understandable, various changes to the software package in use, Topographica, are necessary. These have been implemented for this thesis and are described in detail in chapter 4. Because the underlying GCAL model is input-driven, it is essential to be able to create input patterns covering various features. For example, an orientation map only emerges if the input patterns have varying orientations. The input patterns for an all maps model have to cover a number of different features, and modeling these features was standardized and simplified. Furthermore, a superior way of defining models has been implemented. A class-based system allows the splitting up of the model into levels, and reusing those levels in multiple models. A novel way of modeling motion, based on multiple delayed projections to V1, is also shown in this chapter.

The structure of an all maps model is described in chapter 5. Starting with a model which only covers two dimensions, positional preference and orientation preference, it is shown how other dimensions can be added to this model. Then the all maps model is introduced, which is a combination of the other models.

Corresponding to the models discussed in chapter 5, the resulting maps are shown and compared to previously published maps in chapter 6. There, it emerges that all individual models produce maps which are comparable to previous publications, however the combined model does not yet produce cortical maps due to issues in determining the correct parameters in a large parameter space.

In chapter 7, a discussion of the architecture, methods and results is presented. This also includes suggestions for future research.

Chapter 2

Biological background

This chapter first provides a brief introduction to the visual system of mammals. Then, experimental findings are reviewed in regards to spatial feature maps in the visual cortex of various species. The individual maps are first investigated separately, and then their interaction is discussed. Furthermore, an overview how these maps evolve during the development of animals is given.

Studying the visual system has several advantages over other sensory systems. Firstly, the visual system covers a large area in the brain, with many neurons involved. The early stages of the visual system are well understood and outlined below. From an experimental viewpoint, it is relatively easy to conduct experiments, as the input can be controlled easily. Also, vision is important for the animal behavior in general.

2.1 Visual system

In this section, an introduction to the visual system of mammals and various reasons why the visual system is studied in this thesis are given. Also, it is important to understand the various stages of visual processing, as models of the primary visual cortex build upon these principles.

Visual input is processed in the visual system, which is part of the central nervous system. Photons emitted by light sources, and reflected by objects, hit the eyes. There, in the retina, the visible range of wavelengths is sampled by two or three different cone types. Most mammals are dichromatic, and therefore possessing only two cone types: Short and Medium-Long. Humans possess three cone types, activated by light with Short (blue), Medium (green) and Long (red) wavelength. Furthermore, there are rods which are responsible for night-vision (Bear et al., 2006).

The retinal ganglion cells (RGC) in the inner layer of the retina encode the light level at a given location. Furthermore, a pre-processing step is done: RGC only get activated when there is a difference of light exposure of their center receptive field and surround receptive field. On-center cells are excited by light in the center of their receptive field, and inhibited by light in the surround. Off-center cells show the opposite behavior (Famiglietti and Kolb, 1976). Therefore, RGC are not activated in uniform areas. Furthermore, there is no orientation preference in these cells. The sampling in animals having a fovea is not uniform, instead there is a higher sampling rate of the fovea compared to the periphery. There, the cones are both: smaller and higher in their density (Schein, 1988). There are only very few isolated rod cells found in the fovea (Wässle et al., 1995).

The activation is transmitted along the optic nerve to the optic chiasm, where the nerves partially cross. The images of the inner side of each retina cross to the other side of the brain, whereas the images of the outer side of the retinas are kept on the same side. Then, input from the right visual field goes further along the left optic tract to the left lateral geniculate nucleus (LGN) and vice versa (Crick and Asanuma, 1986). The two lateral geniculate nuclei are located in the thalamus, where also other nuclei are located, which process input from other sensory systems.

The geniculostriate pathway connects the LGN with the primary visual cortex. V1 is a slightly folded sheet of cells and is one of the most widely studied areas in the brain. It consists of six horizontal layers, with incoming connections from the LGN terminating in layer 4. Outgoing connections to higher areas depart from layers 2 and 3, and feedback connections to the LGN depart from layer 6 (Adler et al., 2011). For this thesis, the cortical maps found in V1 are of particular importance and are presented in more detail in section 2.2. For the interested reader, the paper by Crick and Asanuma (1986) is recommended to gain insights into the anatomy and physiology of the cortex.

Interestingly, cortical maps are found for most sensory systems, not just the visual system. For example, the auditory cortex and somatosensory cortex both show a topographic mapping (Udin and Fawcett, 1988). It even has been shown that the auditory cortex develops similarly to the visual cortex if the retinal projections are rewired in a way that the visual input arrives in the auditory cortex. This suggests that the underlying principles of development are shared across different brain areas, and that this development is input-driven (Sharma et al., 2000).

In most mammals, V1 connects to several other higher visual areas. In monkeys and humans, it is thought that two streams emerge: the ventral stream which projects to

the inferior temporal cortex (IT) and the dorsal stream which projects to the posterior parietal (PP) cortex (Mishkin et al., 1983). The ventral stream is thought to be related with object recognition and identification, whereas the dorsal stream is thought to be involved with the localization of these objects (Milner and Goodale, 1998). However, these areas are not nearly as well understood as V1 (Kaas, 2001). Furthermore, in some species like the least shrew, which is one of the smallest mammals, V1 directly connects to the motor system with no further visual areas in between (Catania et al., 1999).

2.2 Cortical maps found in animals

This section reviews a variety of cortical maps which are found in animals. Experiments showing which cortical maps exist in animals, which properties they have, and differences in maps between species are crucial in evaluating models of the visual system. The maps presented in this section act as a reference to the maps which emerge from the models built in chapter 5.

Neurons in V1 are organized in a retinotopic arrangement, i.e. nearby neurons in V1 respond to nearby areas of the retina (Van Essen et al., 1984). However, they are also selective for a variety of other dimensions. For example, it has been found that neurons prefer stimuli which have a certain orientation, and most neurons prefer one eye over the other (Löwel et al., 1988).

It is important to point out that a single neuron has preferences for many dimensions, i.e. a neuron has a preferred position on the retina, a preferred stimuli orientation and a preferred eye. Although neurons respond the strongest to inputs which fulfill all preferences, they usually also respond to stimuli which have similar features. In this section, the individual maps which represent preferences for a certain feature are presented, and their interaction is discussed.

2.2.1 Measurement methods

Most of the studies mentioned below use optical imaging to measure cortical maps in animals. First, the skull of the animal is removed, which exposes the surface of the visual cortex. Then, visual patterns are presented and the response of neurons is recorded. Usually sine gratings are used as patterns. Depending on the study, either the absorbed light of the cortex is measured, or chemicals are applied to the cortex which emit light when nearby neurons are active (Blasdel and Salama, 1986).

The preference of neurons is typically computed with differential imaging, i.e. two images for complementary stimuli are recorded, and the difference of these two images reveals areas which prefer one of the stimuli. For example, when measuring orientation maps, one first captures an image for a particular orientation, and then subtracts the image captured for the orthogonal orientation (Blasdel, 1992b). These steps are usually repeated and averaged. Then, the responses for different stimuli are compared, and the stimulus which is most effective to drive particular neurons emerges.

In the 1990's, an alternative method to measure maps in a finer scale emerged. Using two-photon microscopy, individual cells can be tagged with their feature preference (Swindale, 2006). The basic principle is that two photons have to hit a chemical from different directions before light is emitted. This allows a non-blurry measurement of single cells, in 2D as well as 3D (going into depth) (Denk et al., 1990). This method is much more precise than optical imaging, however the area which can be measured is smaller.

Early studies were made using microelectrodes, which were placed directly in cells (Hubel and Wiesel, 1959). Although very precise measurements were possible, the responses of only very few cells could be recorded. It did not allow the measuring of cortical maps, but only preferences of single neurons.

2.2.2 Retinotopic mapping

Retinotopy refers to the organized mapping of a receptive field position in the retina to retinotopic coordinates in a brain area, usually the LGN or an area of the visual cortex. This first has been discovered in wounded soldiers, where the area of blindness of the visual field could be predicted from the brain area which was damaged (Holmes, 1918).

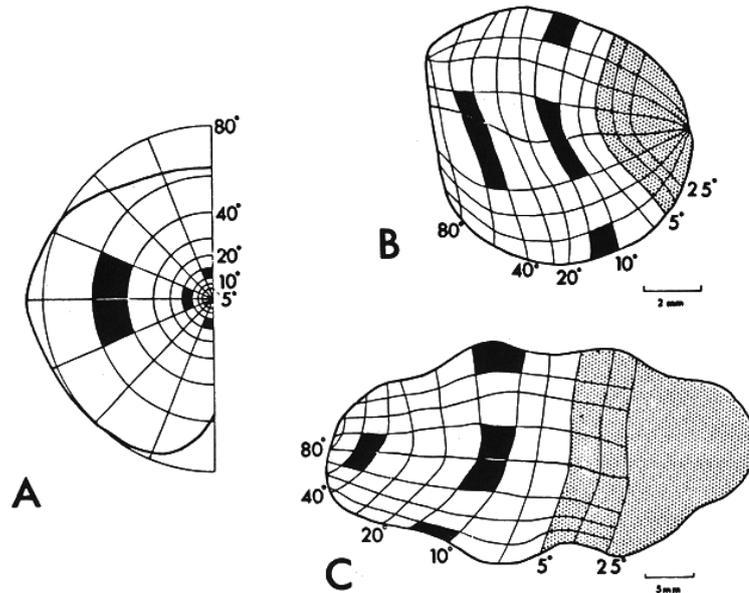


Figure 2.1: Mapping from the visual field (A) to the LGN (B) and V1 (C) in a macaque monkey. The central 5 degrees are overrepresented. Figure reprinted from Connolly and Van Essen (1984).

This was later verified using micro-electrodes. For example, retinotopic mapping was shown in cats (Tusa et al., 1978; Tusa and Palmer, 1980), rats (Espinoza and Thomas, 1983) and macaque monkeys (Van Essen et al., 1984; Connolly and Van Essen, 1984). In macaque monkeys, due to the fovea, the central 5 degrees in the visual field occupy approximately 40% of the cortex.

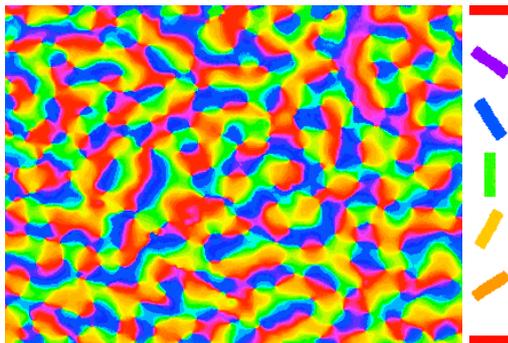
2.2.3 Orientation map

Neurons in V1 are not only arranged in a retinotopic manner, but nearby neurons are also often selective for similar orientations of stimuli. Nearby neurons with a similar orientation preference form iso-orientation patches. One can observe linear zones, where the orientation preference changes slowly and continuously. Pinwheels are neurons near areas of all possible orientation preferences. The orientation selectivity of

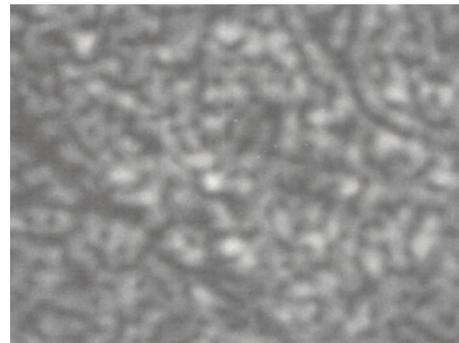
these neurons is very low. Saddle points lie between pinwheel neurons, and show local minima of orientation preference in orthogonal direction. At fracture lines, the orientation preference changes rapidly from one orientation to another (Blasdel, 1992b).

Orientation maps have been found in the primary visual cortex of animals such as macaque monkeys (Blasdel, 1992b), tree shrews (Bosking et al., 1997), ferrets (Rao et al., 1997) and cats (Löwel et al., 1988; Ohki et al., 2005). There, the general structure of orientation maps is similar, although for example there is a higher density of pinwheels in ferrets compared to cats (Müller et al., 2000). In a rat V1 it has been shown that although neurons are selective for orientation, nearby neurons do not have preferences for nearby orientations (Ohki et al., 2005).

Using two-photon microscopy, Ohki et al. (2005) have validated the presence of fractures, as well as pinwheel neurons with a low selectivity for orientation in cats.



(a) Orientation preference map found in V1 of a macaque monkey. An area of 8x6mm is shown. One can observe iso-orientation patches, which result in pinwheels where several iso-orientation patches come together. Also, fractures can be seen, which are areas where orientation preference changes rapidly from one orientation to a very different orientation.



(b) Selectivity preference map in V1 of the same macaque monkey. Darker areas correspond to less selective neurons, typically near fractures and pinwheels.

Figure 2.2: Orientation preference and selectivity maps of a macaque monkey primary visual cortex. Both figures reprinted from Blasdel (1992b).

2.2.4 Ocular dominance map

Neurons typically prefer inputs from one eye over input from the other eye. At the same time, most neurons are binocular, i.e. they respond to inputs of both eyes. The eye preference alternates in stripes (Blasdel, 1992a) or patches (Crair et al., 1997). Features which are found in orientation maps, i.e. pinwheels, fractures and saddle points, are not found in ocular dominance maps.

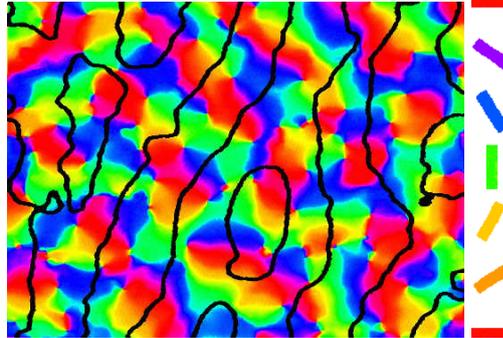


Figure 2.3: Orientation preference map found in V1 of a macaque monkey overlaid by ocular dominance stripes. An area of 4x3mm is shown. One can see that ocular dominance stripes and iso-orientation patches tend to intersect orthogonally. The same neurons have preferences for orientation as well as ocular dominance. Figure from Blasdel (1992a), reprinted modified as in Miikkulainen et al. (2005).

Ocular dominance maps are found in several species, such as macaque monkeys (Hubel et al., 1977; Blasdel, 1992a; Blasdel et al., 1995; Horton and Hocking, 1996), cats (Shatz and Stryker, 1978; Löwel et al., 1988; Crair et al., 1997; Müller et al., 2000) and ferrets (Müller et al., 2000). In cats, ocular dominance maps are more patchy compared to the stripy maps in macaque monkeys. Orientation patches and ocular dominance patches usually intersect orthogonally in macaque monkeys (Blasdel, 1992a), and are less pronounced in cats (Müller et al., 2000). Ferrets usually show a weaker relation between orientation and ocular dominance maps (Yu et al., 2005). Pinwheels usually lie in the center of ocular dominance stripes (Müller et al., 2000).

2.2.5 Disparity map

Several studies have shown that the response of neurons in the visual cortex depends on the disparity, i.e. the positional offset of the stimuli from one eye to the other. In both, macaque monkeys (Poggio and Fischer, 1977) and cats visual cortex (Barlow et al., 1967) the response of these neurons were reported. This is thought to be important for stereo vision, which can be used to detect obstacles or grasp objects.

Barlow et al. (1967) measured vertical as well as horizontal disparity, and found that the response of neurons is more modulated by a horizontal offset. Some neurons had a preferred disparity of up to 6.5 degrees of the visual field, whereas most neurons preferred an offset between 0 and 3.5 degrees.

Unfortunately, to the authors knowledge there has not been an optical image study measuring disparity, and therefore there is no large-scale disparity map for animals recorded as yet. A micro-architecture for disparity recorded by two-photon imaging is shown in figure 2.4.

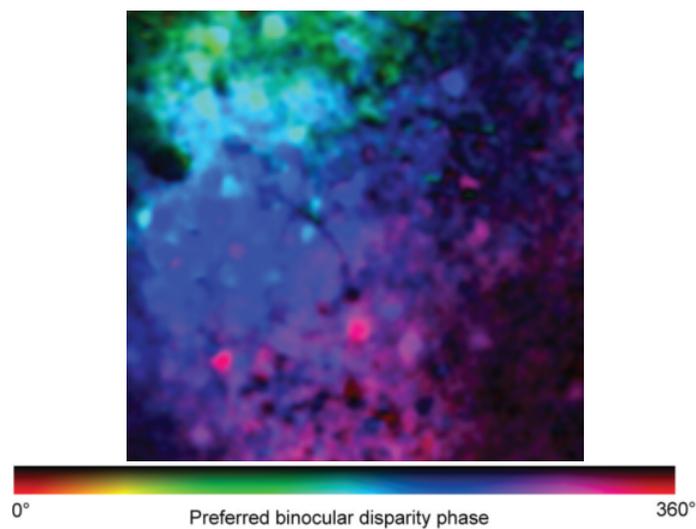


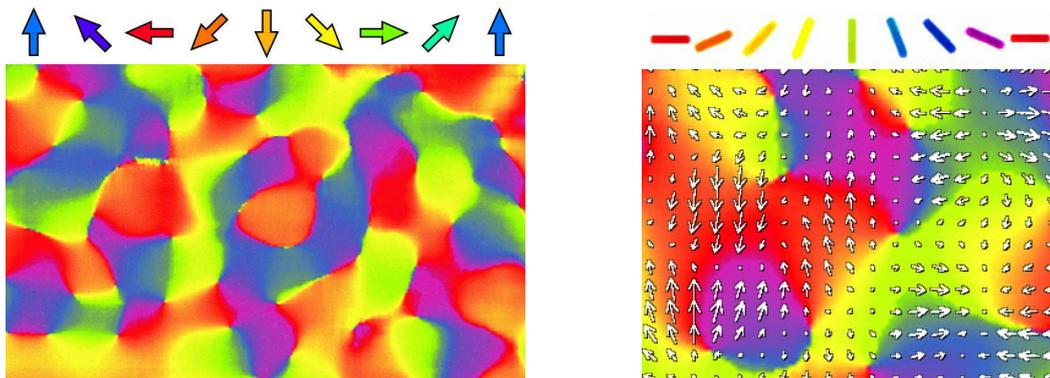
Figure 2.4: Micro-architecture disparity map in a cat visual cortex. An area of 0.3x0.3mm is shown (note that this is ≈ 100 times smaller than most other maps presented in this chapter). There, it emerges that on a small scale, disparity preference is clearly organized. Figure reprinted from Kara and Boyd (2009).

2.2.6 Direction map

Stimuli can be moved in various directions, and the preferred stimulus direction is shown in direction maps. These have been found in cats (Ohki et al., 2005) and ferrets (Weliky et al., 1996), but not in macaque monkeys (Weliky et al., 1996). Direction maps are similar to orientation maps, and show many fractures where the preferred direction changes by 180 degrees.

Weliky et al. (1996) show that orientation and direction maps are closely related. They find two opposite direction patches within a single orientation patch, whereas the preferred directions are the two directions orthogonal to the preferred orientation. They suggest that this is due to the need to represent all orientations and directions in the cortex. As with orientation maps, direction maps show high continuity. However, continuity in orientation maps seems more important than continuity in direction maps.

The direction preference may arise from lagged cells in the LGN. Some cells in the LGN respond to the onset of a stimulus only after a certain delay, which can be 100ms up to 1s. Nonlagged cells typically respond after 30-80ms (Saul and Humphrey, 1990).



(a) Direction map in V1 of a ferret. An area of 3.2x2mm is shown. One can see the abrupt changes from one direction preference to another preference, often by 180 degrees.

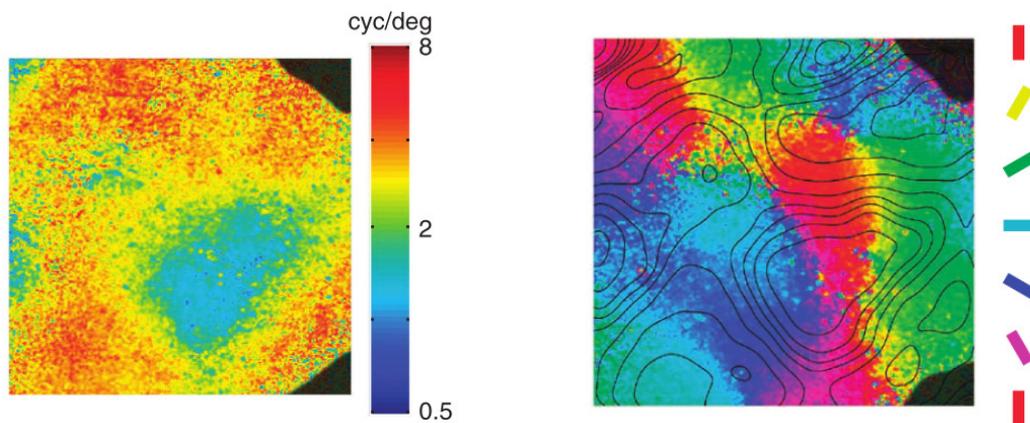
(b) Orientation map in V1 of a ferret overlaid by direction arrows. An area of 1.4mmx1mm is shown. One can see that within one orientation patch, often two direction patches with opposite direction preference emerge. Typically neurons with a low orientation selectivity, such as pinwheels, also show a low direction selectivity. The direction preference is usually orthogonal to the orientation preference.

Figure 2.5: Direction preference map and orientation preference map with overlaid direction preference arrays. Both figures reprinted from Weliky et al. (1996).

2.2.7 Spatial frequency map

Spatial frequency refers to the number of sine gratings within a given distance. There are more bars within a unit distance at a high spatial frequency compared to a lower spatial frequency. As with orientation maps and direction maps, spatial frequency maps tend to change smoothly and continuously. Most of the studies about spatial frequency have been made in cats (Issa et al., 2001; Sirovich and Uglesich, 2004; Ribot et al., 2013; Tani et al., 2012). In Nauhaus et al. (2012) spatial frequency maps of macaque monkeys are computed.

In the macaque monkey, spatial frequency maps tend to intersect orthogonally with orientation maps. Just like the relationship between orientation and ocular dominance maps in macaque, Nauhaus et al. (2012) show that the relation of orientation and spatial frequency maps is stronger in macaque compared to cats.



(a) Spatial frequency map in V1 of a macaque monkey. An area of 0.7×0.7 mm is shown. Similarly to orientation maps, the spatial frequency preference changes smoothly. However, no pinwheels or other features commonly found in orientation maps can be found in spatial frequency maps.

(b) Orientation map in V1 of a macaque monkey. The spatial frequency contours are overlaid in black. An area of 0.7×0.7 mm is shown. One can clearly see the orthogonal intersections of spatial frequency preferences and orientation preferences.

Figure 2.6: Spatial frequency map and orientation map overlaid with spatial frequency contours of a macaque monkey. Both figures reprinted from Nauhaus et al. (2012).

2.2.8 Hue map

There is evidence for color selective cells in V1, where color blobs emerge which respond best to a particular color (Landisman and Ts'o, 2002). Most of the studies related to color have been made in macaque monkey, see e.g. Landisman and Ts'o (2002); Xiao et al. (2003); Xiao (2014). Xiao (2014) found overlapping responses in the visual cortex when stimulated with different colors, whereas “the response peaks shifted systematically as a function of the stimulus color” (Xiao, 2014).

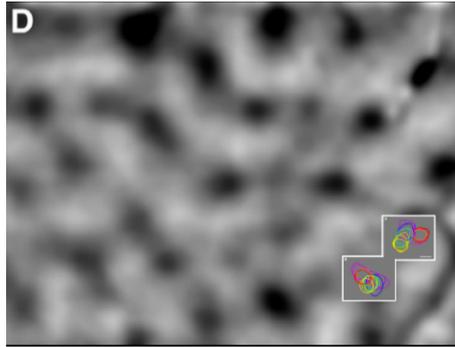


Figure 2.7: Color selective cells found in area V1 of a macaque monkey. An area of 3.75x3mm is shown. Figure reprinted from Xiao et al. (2007).

2.3 Map interaction

The last section presented all known cortical maps in animals. Here, the interaction of these maps is discussed. This is important, as the aim of this thesis is to build an all maps model of the primary visual cortex. Not only should the individual maps match those found in animals, but also the interaction between these maps.

Selectiveness for a particular feature, such as ocular dominance, is likely to affect the selectiveness for another feature, such as orientation. This section discusses some of the experimental findings which support this hypothesis. Furthermore, in regards to the thesis' aim to build an all maps model, it is important to note that the previously mentioned experiments were made in various species, and it is not known whether a single species is selective for all the spatial features, or just a subset of them.

In the late 1960's, orientation columns and ocular-dominance columns were found in macaque monkeys using electrode measurements Hubel and Wiesel (1968). Using optical imaging, it emerged that they usually intersect orthogonally in macaque monkeys Blasdel (1992a). However, in cats and ferrets this relationship is not nearly as

strong (Müller et al., 2000; Yu et al., 2005).

Furthermore, there seems to be agreement that pinwheels found in the orientation map usually lie in the center of ocular dominance stripes (Müller et al., 2000). This suggests a relationship where neurons which are highly selective for ocular dominance are less selective for orientation. Computational models might help to propose similar relationships between other spatial maps, which then could be validated in animal studies.

A clear relationship is also seen between orientation and direction maps. In ferrets, it has been shown that within one iso-orientation patch there are often two direction patches, with an opposite direction preference. Also, the direction preference is orthogonal to the orientation preference. In Miikkulainen et al. (2005), it is suggested that in theory direction patches rather than orientation patches could become the largest-scale organization, in the case animals are raised in an environment with high speeds of visual motion.

Landisman and Ts'o (2002) show that color patches tend to cross ocular dominance columns. Also, it seems as if the color patches contain binocular color-selective cells. The same experiments also suggest that color blobs exclude orientation maps from forming in the vicinity.

2.4 Map development

This section outlines the cortical map development in animals. There is clear evidence that cortical maps change considerably in the first weeks and months after birth, and that visual input partly drives these alterations. The cortical maps in the all maps model built for this thesis also gradually evolve over time, and therefore discussing which principles might be responsible for the map development is of great value.

In a comparison of normal cats and cats which lids were sutured before eye opening, it was found that orientation and ocular dominance maps in these cats developed identically within the first three weeks after eye opening. However, the maps developed differently after this time. In normal cats, the orientation selectivity remains high, whereas in binocular deprived cats the selectivity as well as responses to visual input in general decrease. That suggests that the basic structure of cortical maps exists at birth, but needs visual experience to stay intact (Crair et al., 1998). Similarly, macaque monkeys raised in darkness also form ocular dominance columns (Horton and Hocking, 1996).

Retinal waves are thought to be responsible for the basic structure of cortical maps at birth. A retinal wave is spontaneous activity on the retina, and retinal waves are drifting across the retina. By blocking spontaneous activity, the anatomical organization of ocular dominance columns is destroyed irreversibly (Huberman et al., 2006). Therefore, retinal waves might be “an evolved adaptation (...) imparting an informational robustness and redundancy guide” (Ackman and Crair, 2014).

Chapter 3

Model background

A computational model of the visual cortex is an implementation of a theory of how the visual system works. The theory is formed by the experimental findings presented in chapter 2, and these findings are used to validate the theory. It should also be possible to make predictions using these models, which then can be verified in animals. The visual system should be modeled in enough detail so it can be compared to animal studies, however it should omit everything beyond that so it is simple enough to understand it (Miikkulainen et al., 2005).

In this chapter, various models resulting in biologically realistic maps are presented. The models are categorized into three main groups. Each group is presented individually, and their similarities and differences are pointed out. The chapter also provides a brief overview of contributions that these models have made for today's understanding of the visual cortex.

3.1 Sparse coding

Sparse coding models are closely related to the area of computer vision. There, the independent component analysis, which is a special case of sparse coding, is widely used. Although there is no known model covering particularly many input dimensions simultaneously, the concept of sparse codes are also important for understanding the GCAL model, and therefore they are presented in this thesis at some length.

Sparse coding is based on the finding that only a small subset of all neurons is active at any time t , and the pattern of activation represents information in the brain (Rolls and Tovee, 1995). For each input pattern, a different subset of neurons is activated. This system can be built in a way so it is optimal for input patterns with known sta-

tistical properties, e.g. natural images (Hyvärinen and Hoyer, 2001). This is thought to have different advantages in respect to the energy consumption, storage capacity, and data representation. Furthermore, the coding is simple to decode at subsequent stages (Olshausen and Field, 2004).

Sparse codes provide a compromise between dense and local codes. Imagine a set of N binary neurons, i.e. they are either active or inactive. In local codes, each neuron is only active for exactly one input pattern. Therefore N different patterns can be represented, which is likely to be a too small number, even for billions of neurons in a brain. Dense codes are the other extreme, where on average $0.5 \times N$ neurons are active at any time. This would allow the representation of 2^N patterns, a far too large number given the number of neurons in the brain. Therefore, most capacity would be redundant, and decoding patterns would be extremely difficult as all neurons have to be taken into account. Sparse codes are in between these two extremes, using a small fraction of neurons to represent patterns (Földiák, 2002).

Mathematically, an input image $I(x,y)$ is modeled as a superposition of n basis vectors $a_i(x,y)$:

$$I(x,y) = \sum_{i=1}^n a_i(x,y)s_i \quad (3.1)$$

The s_i are coefficients which have to be calculated for each image, such that on average the activation is sparse. This is achieved by fixing the expectation $E\{s_i^2\}$ to a desired activation value, and then using a convex function G to measure the sparse-

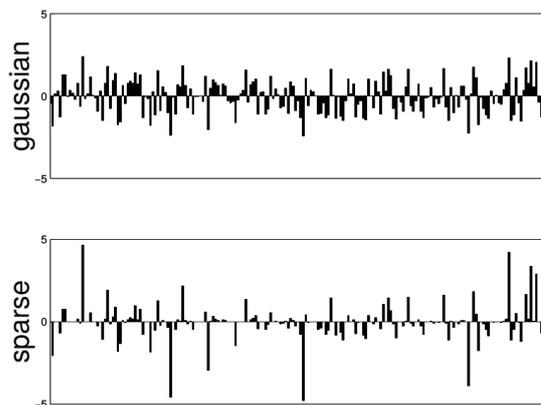


Figure 3.1: This figure compares random samples from a Gaussian variable (top) and sparse variable (bottom). One can see that the sparse variable is close to zero at most times, with some large values in between. The Gaussian variable has more non-zero values, but less extremes. The random observations were made using the same variance for both variables. Figure reprinted from Hyvärinen et al. (2009).

ness: $E\{G(s_i^2)\}$ (Hyvärinen and Hoyer, 2001).

In Hyvärinen and Hoyer (2001) it has been shown that a sparse coding algorithm can account for position, orientation and frequency preference. Using independent component analysis, it has also been shown how color and disparity preference can emerge (Hyvärinen et al., 2009). Most of the recent work focuses on application of sparse representation in more complex visual tasks, such as visual tracking (Zhang et al., 2013) and image annotation (Liu et al., 2014).

3.2 Elastic net models

The elastic net is a regularization method known from statistics and machine learning. It has also been employed to build models which result in maps similar to those in animals, where the focus is trying to determine why cortical maps in biology have the properties observed in studies, rather than building a visual system (which is the aim of GCAL).

The elastic net combines the ridge regression and LASSO regularization methods. Both methods aim to minimize the coefficients of the variables, whereas ridge regression uses the L2 norm and LASSO regularization uses the L1 norm. In practice, LASSO leads to sparse solutions, i.e. most coefficients are nearly zero, and ridge regression leads to dense solutions, i.e. most coefficients are non-zero. The elastic net algorithm also produces sparse solutions, but in comparison to LASSO highly correlated variables all have non-zero coefficients, rather than picking one of the correlated variables.

In Carreira-Perpiñán et al. (2005) a model based on the elastic net resulting in maps for positional preference, orientation, ocular dominance and spatial frequency is presented. There, the elastic net is used to optimize a trade-off of the coverage of the stimulus space, i.e. “any combination of stimuli values is represented somewhere in cortex” (Carreira-Perpiñán and Goodhill, 2002), and continuity, for example, retinotopy or iso-orientation patches. The stimuli are represented as vectors, whereas each vector entry represents one feature dimension. As in animal studies, pinwheels which lie in the center of ocular dominance stripes emerge, and maps have the tendency to intersect orthogonally.

The elastic net also has been used to investigate the differences in maps of various species. As discussed in 2.2.4, ocular dominance maps are patchier in cats compared to stripy maps in the macaque monkey. Using the elastic net, it has been suggested

that this is due to the relative order of the map development. It seems that orientation maps develop before ocular dominance maps in cats, and vice versa in the macaque monkey (Goodhill and Cimponeriu, 2000).

3.3 Models based on self-organizing maps

The GCAL models built for this thesis have their origin in the self-organizing map by Kohonen (1982). First, the original algorithm is presented. GCAL as an extension to the SOM which can account for more developmental and functional phenomena is described after, whereas the GCAL model has less artificial limitations, and a higher degree of biological realism.

3.3.1 Self-Organizing Map

A self-organizing map is a network of neurons, typically arranged in a two-dimensional grid. Each neuron has a weight vector associated, which maps the usually high-dimensional input space to the two-dimensional grid. A neighborhood function is employed so that inputs which only differ slightly are mapped to nearby areas of the network (Kohonen, 1990). In case of a cortical map, this neighborhood function can be seen as lateral interaction between neighboring neurons. These have been found experimentally in animals, see for example Gilbert et al. (1990).

During training, the network learns as follows: First, the activity of all neurons is calculated as the weighted sum of the input vector $V = \{v_1, v_2, \dots, v_n\}$, $v \geq 0$ and weight vector $W_{ij} = \{w_{1,ij}, w_{2,ij}, \dots, w_{n,ij}\}$:

$$\eta_{ij} = \|\vec{V} \cdot \vec{W}_{ij}\| \quad (3.2)$$

Then, the neuron (r, s) with the highest activation is determined. This neuron implicitly gets the maximal activity η_{max} assigned in equation 3.3. This step is important so that the whole input space is covered, even for an initially bad match. As the neighborhood function, a Gaussian can be used and the activity for all other neurons is calculated as follows, whereas (i, j) is the grid position of the neuron:

$$h_{rs,ij} = \eta_{max} \exp\left(-\frac{(r-i)^2 + (s-j)^2}{\sigma_h^2}\right) \quad (3.3)$$

The parameter σ_h controls the width of the Gaussian function, which is gradually decreased over time. It has to start out large as the initial weights are random, and

therefore the activity patterns are random as well. Using the learning function in equation 3.4, the activity pattern is concentrated around the winning neuron after a certain training period, which then needs a Gaussian which is narrower (Kohonen, 2001).

As learning rule, the mechanism proposed by Hebb is employed, which states that a connection is strengthened when two neurons fire at the same time (Hebb, 1949):

$$w'_{k,ij} = w_{k,ij} + \alpha v_k h_{rs,ij} \quad (3.4)$$

Where k is the k -th component of the weight/input vector and α is the learning rate. Often the normalized version of the Hebbian learning rule is used, where the sum of the weights is constant. This is necessary so that the weights do not become instable due to ever-increasing values (Rochester et al., 1956).

Using this version of the self-organized map, a retinotopic mapping emerges when feeding the network with Gaussian input patterns located at a random position. As in animals (see section 2.2.2), neurons in the center of the network respond to centrally located input patterns, neurons at the edge of the network respond to inputs at the edge of the receptor; and nearby neurons respond to input patterns with nearby location. If there is a higher probability for certain locations in the input space, there will be more neurons selective for these locations compared to under-represented locations (Ritter, 1991). To rephrase this: the resources (neurons) are allocated according to the input distribution, and therefore the statistics of the environment are encoded.

This version of a self-organizing map has certain drawbacks. Firstly, picking a winner is biologically implausible. There is no known mechanism which could supervise this behavior. Secondly, a full connectivity between the neurons is required, which is again not biologically realistic. Furthermore, the lateral connections are isotropic, which contradicts known animal studies, where lateral connections are mainly existent between neurons preferring the same orientation. It is also unclear how the radius σ_h could shrink over time in an animal (Miikkulainen et al., 2005). In addition, the initial connections on the visual pathway are roughly retinotopic, as axons follow signaling gradients (Tessier-Lavigne and Goodman, 1996).

The self-organizing map implicitly makes use of two concepts which allow mapping of a higher dimensional space to a low dimensional space. The first concept is the principal surface, which takes advantage of the fact that input dimensions do not vary independently. For example, the direction of a pattern must be roughly orthogonal to its orientation so a movement can be seen. As this is often not enough to map the input into two dimensions, folding is another technique that can be used. There, a curve first

stretches along the dimensions with the highest variance, and then makes tight turns across the width of the area to cover the remaining dimensions (Miikkulainen et al., 2005).

3.3.2 Gain Control, Adaptation, Laterally connected model

The standard GCAL model consists of four two-dimensional sheets:

1. The photoreceptor sheet, which can be fed with any two-dimensional input pattern, e.g. Gaussian patterns or natural images. The activation $\Psi_i(t)$ of each unit i at time t typically ranges between 0.0 and 1.0.
2. Two sheets representing the preprocessing done in the retinal ganglion cells (and transmitted to the lateral ganglion nucleus): One On and one Off sheet, modeling the difference of light exposure between the center and the surround of retinal ganglion cells, as described in section 2.1. This is abstracted by Difference-of-Gaussian afferent connections from the photoreceptor sheet. Each unit in the On/Off sheets connects to multiple units in the photoreceptor sheet, forming a receptive field. Neighboring units have different, but overlapping receptive fields. Units in the On sheet respond to bright areas surrounded by dark areas, and vice versa. Furthermore, gain control is employed to allow a wide range of input contrasts.
3. The V1 sheet, which is connected to the On/Off sheets. In comparison to the self-organizing model, each unit only connects to a small area of the On/Off sheets. Units also connect laterally to surrounding units in V1, with short-range excitatory as well as long-range inhibitory connections. These lateral connections sharpen the response. The inhibitory connections are non-isotropic, allowing representation of long-term correlations of the input patterns. Furthermore, the encoding is more efficient using lateral connections, as redundancy in the input is suppressed, which then allows detecting changes in the input more efficiently (Sirosh et al., 1996).

For the LGN sheets, gain control is employed, which ensures that the model is robust for varying input contrast. The activation $\eta_{j,O}(t + \delta t)$ of a LGN unit at position j on sheet O for time $t + \delta t$ is calculated as follows:

$$\eta_{j,O}(t + \delta t) = f \left(\frac{\gamma_O \sum_{i \in F_{j,P}} \Psi_i(t) \omega_{ij}}{k + \gamma_S \sum_{i \in F_{j,S}} \eta_{i,O}(t) \omega_{ij,S}} \right) \quad (3.5)$$

Where f is a half rectangle, ensuring that the response is always positive, γ_O is a boost factor chosen so that the unit activation is normally between 0.0 and 1.0. F denotes the projection fields, F_P for the afferent projection and F_S for the lateral connections, which suppress highly active units. k ensures that the denominator is always positive, and γ_S is a rescaling factor. The weights ω_{ij} from unit i in the photoreceptor sheet to unit j in the On/Off sheets is given by a Difference of Gaussians, with a smaller center and wider surround. For the On sheet, the surround is subtracted from the center and vice versa. For the exact values of the constants, see Stevens et al. (2013).

Units in V1 receive afferent connections A from the On/Off sheets, lateral excitatory connections E from surrounding sheets, and longer range lateral inhibitory connections I . The initial activation is calculated purely from the afferent connections, which have an initially Gaussian shape. Then, there are several settling steps which purely depend on the activation of the surrounding sheets. The activation $\eta_{j,V}(t)$ of V1 unit j at time t is the half-rectified sum of the activity of the projections $p = (A, E, I)$:

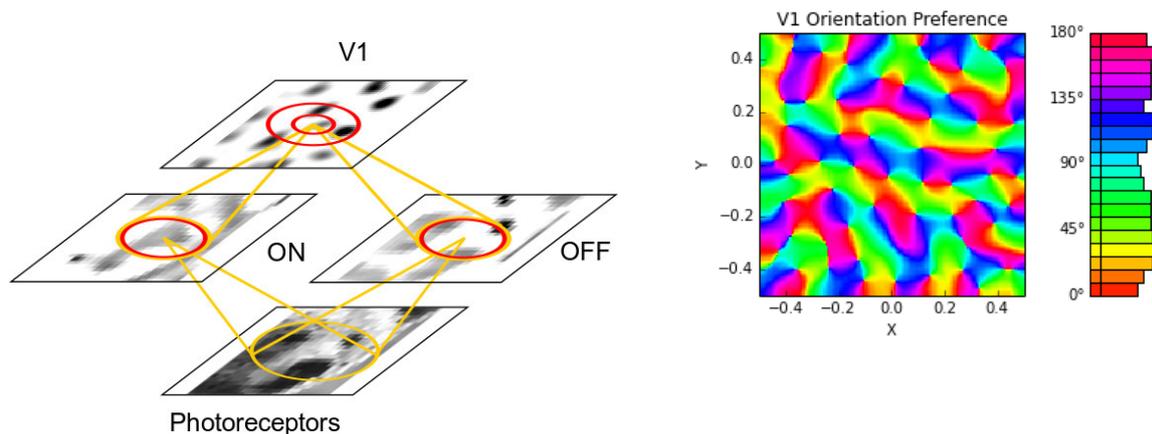
$$\eta_{j,V}(t) = f \left(\sum_p \gamma_p \sum_{i \in F_{j,p}} \eta_{i,p}(t) \omega_{ij,p} \right) \quad (3.6)$$

With $\eta_{i,p}(t)$ as the activation of the afferent/laterally connected units, $\omega_{ij,p}$ as the connection strength from unit i to unit j for projection p and γ_p to weight the different connection types. Here, f has a variable threshold aiming to bring the average of the activity values of a unit to a desired value. For the detailed adaptation process, see Stevens et al. (2013).

Then, a normalized Hebbian learning rule with learning rate α is employed to update the weight $\omega_{ij,p}$ from unit i to unit j of projection p :

$$\omega_{ij,p}(t) = \frac{\omega_{ij,p}(t-1) + \alpha \eta_j \eta_i}{\sum_k (\omega_{kj,p}(t-1) + \alpha \eta_j \eta_k)} \quad (3.7)$$

Using this basic model, biologically realistic orientation maps result for a wide range of input contrasts. The map development is stable, i.e. only after a relatively short training, the maps have similar properties to those found in animals. Furthermore, statistics in the input patterns are reflected on the resulting maps. This is demonstrated by simulating a goggle-reared environment, where one orientation is dominating the visual scene (Stevens et al., 2013).



(a) Network structure of the basic GCAL model, with one photoreceptor sheet which projects into one On sheet and one Off sheet. The extent of the connection fields are shown with orange cones. Red circles represent lateral connections. For the On/Off sheets, these are as wide as the connection field projection down from a single V1 unit. For the V1 sheet, the lateral excitatory connections are shown with the smaller circle, and the inhibitory connections are shown with a wider circle. Figure reprinted from Stevens et al. (2013).

(b) Resulting orientation map in GCAL. The histogram shows that all orientations are equally represented.

Figure 3.2: Structure of the basic GCAL model, and resulting orientation map when Gaussians are used as input patterns.

3.3.3 Multiple maps using GCAL

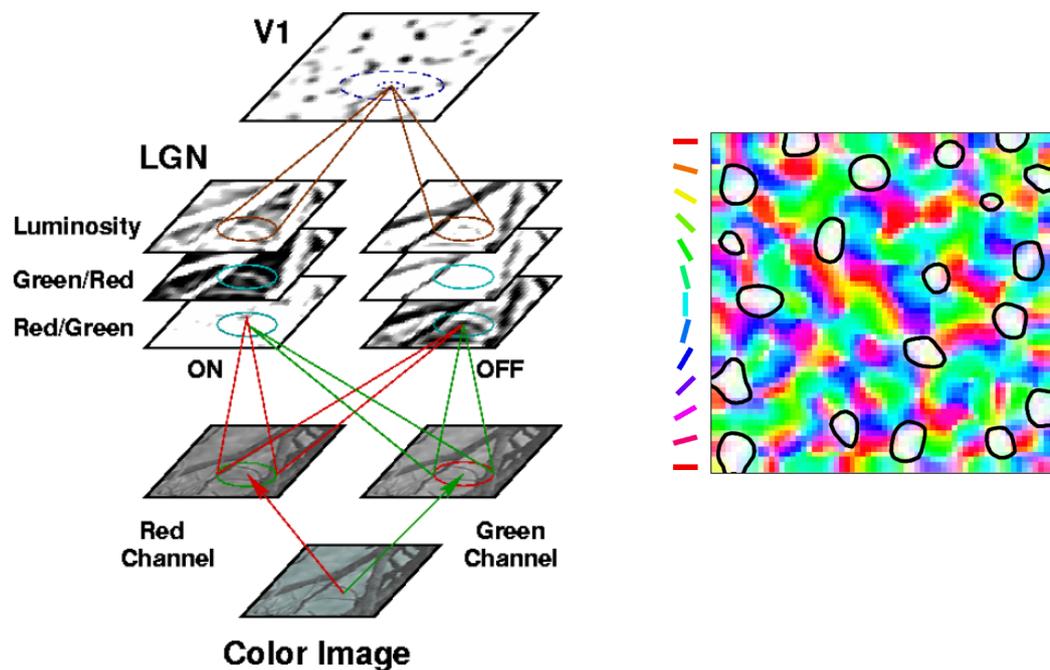
Earlier studies using the predecessor of GCAL, the Laterally Interconnected Synergetically Self-Organizing Map model (LISSOM), it has been shown that using the same principles as introduced above, it is possible to explain a wide variety of cortical maps. Here, the methods which were used to extend the basic LISSOM model to model multiple maps are presented together with the corresponding papers.

Orientation + Color

For color simulations, multiple retina sheets are used, each of them representing a particular cone type. For a dichromatic simulation, one sheet for medium wavelength (green) and another sheet for long (red) wavelength are used. Rather than just one pair

of On/Off sheets, three pairs of these sheets are required. One pair with units which get their center afferent input from the green photoreceptor sheet, and their surround from the red photoreceptor, and another pair where the center and surround is swapped, as well as a pair of luminosity sheets which is connected to both photoreceptor sheets. The units in V1 connect to all sheets (Bednar et al., 2005).

Units in V1 typically become selective for either orientation, or one of the color channels. That is, if a neuron is highly selective for orientation, it has nearly no preference for color. Similarly, if a neuron is selective for a specific color, it is neither selective for other colors nor for orientation.



(a) Required sheets and projections for a color simulation. In comparison to the basic LISSOM/GCAL model, there are several photoreceptor sheets representing the various cone types found in animals. They project into two pairs of LGN sheets which are color-selective and one pair of luminosity sheets which equally receives input from all photoreceptor sheets.

(b) Resulting orientation map, whereas color blobs are surrounded by black lines. Color blobs have a low orientation selectivity, which is denoted by brighter colors.

Figure 3.3: LISSOM model for a dichromatic color simulation and the resulting maps. Both figures reprinted from Bednar et al. (2005).

Orientation + Spatial frequency

To model spatial frequency, the visual pathway again has to be modified. Similarly to the color model, further LGN sheets have to be introduced. For spatial frequency, each spatial frequency channel (consisting of an On/Off sheet pair) has a different size of receptive fields in the photoreceptor sheet. Therefore, the channels have varying selectivity for different frequencies of the input pattern. This also implies that the sizes of the sheets differ between the channels because of the required edge buffering.

It has been found that the spatial frequency range widens if a unit in V1 only connects to one of the spatial frequency channels rather than all of them (which is achieved by pruning the weak connections after an initial training), however this comes with the cost of distorted orientation maps (Palmer, 2009).

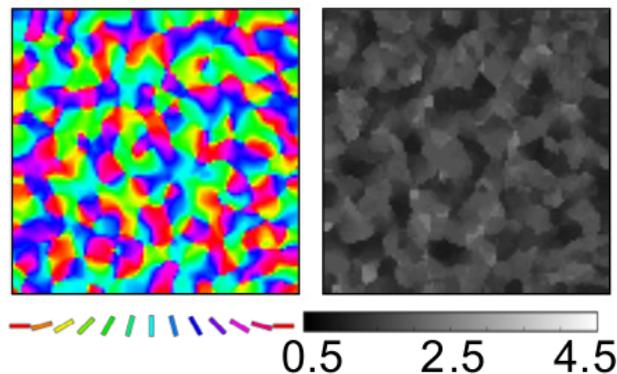


Figure 3.4: Resulting spatial frequency map and orientation map of a LISSOM simulation with multiple LGN sheets selective for different spatial frequencies of the photoreceptor sheet. Figure reprinted from Palmer (2009).

Orientation + Ocular dominance + Direction

The model presented in Bednar and Miikkulainen (2006) extends the basic LISSOM model by two dimensions, ocular dominance and direction. To model ocular dominance, two retinas are used, corresponding to the two eyes of mammals. Ocular dominance maps result from brightness differences, i.e. the input pattern is brighter in one of the two eyes. The brightness difference is randomly distributed, in a way that the activities sum up to 2.0.

Direction preference is achieved with multiple lagged LGN sheets. The input pattern first arrives at lag $N = 3$ at time t , and is then moved slightly before lag $N - 1$ is activated at time $t + 1$ and so forth. These differences in time result in spatiotemporal

fields. This is described in more detail in section 4.3, where also other ways to model motion are presented. Also, in figure 4.1a the model structure is shown.

In Miikkulainen et al. (2005), it has been shown that differences in the pattern location between the eyes, also known as disparity, does not result in biologically realistic ocular dominance maps. Disparity would be a more intuitive way to model ocular dominance, as animals encounter this every day due to the displacement of the eyes, as opposed to non-intuitive dimming differences.

In Burger and Lang (1999) and Burger and Lang (2001) a model which has a very similar structure to this model is presented, however does not implement direction maps. Using random stimulus patterns (in contrast to Gaussian input patterns / natural images used in LISSOM/GCAL), orientation and ocular dominance maps emerged.

All maps model

In a previous project, LISSOM was used to build an all maps model. This model shows that in principle it seems possible to combine all cortical maps found in animals in a single simulation. Basically, the presented models of section 3.3.3 are all combined into one model. One exception is a different motion model which was used in this simulation, which uses several lagged retina sheets rather than lagged LGN sheets (see 4.3). This is biologically not plausible and computationally expensive, as per lagged retina the corresponding lagged LGN sheets have to be added to the simulation as well (Gerasymova, 2009).

When simulating all dimensions, the model complexity is tremendous. There are 24 retina sheets, resulting of 2 eyes * 4 time lags * 3 cones. Per time lag and eye, there are eight LGN sheets for color processing and two LGN sheets modeling spatial frequency, so in total $2 * 4 * (8 + 2) = 80$ LGN sheets. The model structure is not shown in a figure due to the high complexity, which can be reduced significantly, as shown in section 5.9.

Besides the high complexity and the biologically unrealistic lagged retina sheets, the model has certain other drawbacks. The combination of the ocular dominance and color dimensions is not realistic, as the brightness difference to model ocular dominance is applied per color channel in the RGB color space. This means that if the activation of the red channel in the left eye is R_{Left} , then the activation in the right eye

is calculated as $R_{Right} = 2.0 - R_{Left}$. As this is done for each channel, the hue¹ for a specific unit differs between the left and right eye. This leads to a strong change in the color map, where the preferred hue is shown.

From a software point of view, the model is highly inflexible and difficult to maintain. The various components cannot be easily replaced, for example, it is not possible to use GCAL instead of LISSOM.

A unified model which is much more comprehensive and also models of all known cortical maps has been suggested in Bednar (2012), however this paper reports results from different models covering subsets of all dimensions. An implementation is described as “very much a work in progress” (Bednar, 2012) and is aim of this thesis. The proposed model contains only three sheets per eye, one per cone type, compared to the 12 retina sheets in the work from Gerasymova (2009). Because motion is modeled with delayed projections from the LGN to V1, only 20 LGN sheets are required. The model structure is not shown in a figure as it is very similar to the combined model presented in section 5.9.

In this thesis, the focus is on replicating maps which are similar to those found in animals. The paper by Bednar (2012) shows how a unified model could also explain a variety of other findings, such as contrast-invariant tuning, the properties of simple and complex cells in the cortex and aftereffects. However, these explanations are beyond the scope of this thesis.

Input patterns for multiple map simulations

In the previous section, it has been shown how the basic GCAL model can be extended to cover various dimensions. This usually requires the extension of the model with retina and/or LGN sheets. However, besides adjusting the model structure, the input patterns have to be adjusted as well. For example, if the input patterns of the left and right eye are exactly the same, no ocular dominance map will develop. Even in the basic GCAL model an orientation map only emerges because the orientation of the Gaussian input patterns is random (Miikkulainen et al., 2005).

For ocular dominance simulations, the brightness of the input patterns has to vary. A related but different input dimension is known as disparity. A disparity map emerges only in the case the pattern locations in the left and right eye differ. For the motion

¹The hue value of the color is a measure where the color is found in the color spectrum, independently of its brightness which is encoded in the value or lightness and the pureness which is described by the saturation in the HSL/HSV color spaces (Joblove and Greenberg, 1978).

model, the input pattern has to be moved over time. Similarly, the spatial frequency of input patterns has to vary so that a spatial frequency map emerges. As for color, these dimensions are inherent to natural images.

Here, the fact that the input statistics are reflected on the resulting maps has to be stressed. For example, if cardinal orientations are overrepresented in the input, as they usually are in natural images, then more neurons become selective for these orientations, compared to underrepresented orientations (Stevens et al., 2013).

For the GCAL model, any artificial two-dimensional input patterns such as two-dimensional Gaussians or natural images can be used. Using Gaussian patterns has the advantage that the parameters can be controlled more easily.

3.4 Open questions

Although computational modeling has successfully been used to provide various hypotheses of how the visual system in general, and cortical maps in particular, might evolve and what their function might be (see Goodhill (2007); White and Fitzpatrick (2007); van Ooyen (2011) for reviews), there are still many open questions. Here, two of them have been chosen which are particularly interesting because it may be possible to provide an answer using an all maps model.

Using vector-like inputs similar to the ones which are used in the elastic net, the self-organizing map was applied without modifications to model multiple maps. In the paper by Swindale (2000), the question is raised how many maps can be represented in the two-dimensional visual cortex. This is important, as there might be undiscovered cortical maps. Furthermore, this would allow statements about the interaction of maps. With theoretical assumptions beyond the scope of this thesis, the author suggests that there might be an upper limit of nine or ten maps. The interested reader might also want to read a follow-up study in Swindale (2004).

The same author proposes that maps might be optimized for uniform coverage, i.e. whatever combination of input features one presents to the cortex, the total neural response is independent of this combination (Swindale et al., 2000, 2002). However, another study doubts that the way coverage was measured was correct (Carreira-Perpiñán and Goodhill, 2002). Whether the assumption of uniform coverage is correct or not remains an open question, and a model of all known spatial maps could be used to develop this further.

Chapter 4

Architecture

This chapter introduces three main concepts necessary to build an all maps model in Topographica, a software for simulating cortical maps (Bednar, 2009).

It is crucial to generate patterns which cover all dimensions one wants to model, due to the input-driven self-organization of the model. With the former implementation it was difficult to generate complex patterns, and hard to extend more dimensions. Here, the concept of feature coordinators is introduced, whereby each coordinator adds one more dimension. The pattern coordinator acts as a convenient wrapper to use the feature coordinators.

The submodels machinery is one of the main contributions of this thesis. Automatic wiring of sheets, easier setting of parameters and the ability to modify and inspect the model structure before instantiation are a few of the advantages gained by using this new concept. The superiority of this implementation over the previous is especially obvious for models with many sheets and projections between them. The old results can be replicated if needed, but it becomes much more clear how the model structure emerges in the new system.

The last concept discussed is a new motion model. Direction maps can emerge in many types of temporal delays within the model. Previously, this has been done with a set of lagged retina sheets or a set of lagged LGN sheets. Here, following the suggestion in Bednar (2012), multiple projections with increasing delays from LGN sheets to V1 are used to model motion.

4.1 Pattern generation

As described in section 3.3.3, it is crucial that input patterns cover the feature dimensions one wants to model. In the case of Gaussian input patterns, for example, the model of V1 will not develop orientation-specific neurons if all Gaussian patterns are oriented the same way. On the other hand, if one only wants to model position preference (without orientation preference), it is important that all Gaussians have the same orientation. Using oriented Gaussian patterns does, however, alter the position preference map. In order to build a model that covers various feature dimensions, it is therefore crucial to have a mechanism that allows the ability to specify the needed feature dimensions of the input patterns, and the synthesis of the feature dimensions (if needed). For this thesis, two concepts called feature coordinators and pattern coordinator have been implemented, and they are described in detail in this section.

4.1.1 Feature coordinators

A feature coordinator modifies an input pattern in some way, and returns the modified input pattern. A feature is something coordinated between input patterns. This could be either a parameter of an input pattern (such as the size), or a variable from which values for an existing parameter can be calculated (such as a position offset). Each input to the model can consist of one or more overlaid input patterns.

Each feature coordinator synthesizes exactly one feature dimension. This allows any arbitrary combination of feature dimensions by applying the desired feature coordinators in a serial manner (see section 4.1.2). By supplying the name of the pattern generator to the feature coordinators, it is possible to coordinate a set of patterns across multiple retinas.

Position feature coordinators

The position feature coordinators are necessary in order to develop a position preference map. There are two position feature coordinators, one for modifying the x coordinate of a supplied pattern and the other for modifying the y coordinate. By default, the position is altered around the point of origin by a uniform random number (within set bounds) provided as a parameter.

Orientation feature coordinator

The orientation feature coordinator modifies the orientation of a supplied pattern, by default to a uniform random number between $-\pi$ and $+\pi$. This allows development of an orientation map.

Spatial frequency coordinator

The spatial frequency coordinator varies the sizes of the input patterns in order to develop a spatial frequency map. Two parameters are supplied to this coordinator: the spatial frequency spacing Γ and the highest spatial frequency channel ν_{max} . The spatial frequency spacing corresponds to the factor by which successive spatial frequency channels increase in size. The size of the supplied input pattern is multiplied by a uniform random number between one and $\Gamma^{\nu_{max}}$. This corresponds to sizes which suit the receptive fields of the spatial frequency channel with the smallest receptive fields up to sizes which suit the channel with the biggest receptive fields.

Motion feature coordinator

The motion coordinator sweeps the supplied pattern orthogonal to its orientation. At each time step, the pattern is moved further. The speed of the sweeping can be adjusted by a parameter, as can the amount of time required before showing a new pattern to the model. Further details are explained in section 4.3.

Dimming feature coordinator

In order to get realistic ocular dominance maps in GCAL based models, it is necessary to introduce dimming differences of patterns in the left and right retina. This will be discussed in more detail in section 6.3. Dimming differences can be achieved with the dimming feature coordinator. The coordinator can only be used with two retinas, and is implemented such that the brightnesses of the two input patterns add up to 2, i.e. if the brightness of the left input pattern is 0.7, the brightness of the right input pattern must be 1.3. The fraction by which the pattern brightness varies between the two eyes can be supplied as parameter, and is again created within a uniform random number generator.

Disparity feature coordinator

The disparity (difference of pattern locations) of two input patterns is coordinated by the disparity feature coordinator. Currently, this only works for horizontal offsets, but could be easily extended to allow vertical offsets as well. In practice, the x -coordinate of the input pattern is shifted to the right in one retina, and by the same amount to the left in the other. The amount of shift is controlled by a uniform number generated between zero and a parameter ξ_{max} .

Summary feature coordinators

This section introduced the concept of feature coordinators and described each feature coordinator in detail. The coordinators can be applied in any arbitrary order because the parameters of the supplied patterns are not overridden, but modified. The only exception is the motion coordinator, which has to be applied after the orientation coordinator, because the motion is calculated orthogonal to the orientation of the pattern.

The implemented coordinators can synthesize most input dimensions, except for color. However, adding more feature coordinators is straightforward due to the common interface and atomic operation of each individual feature coordinator. A mathematically more precise definition of the feature coordinators can be found in the corresponding sections of chapter 5.

4.1.2 Pattern coordinator

The pattern coordinator is a newly introduced concept which allows using the feature coordinators presented earlier in a user-friendly way. The pattern coordinator is a class which can be instantiated, and the only required parameter on instantiation is a list of the desired dimensions to be contained in the returned input patterns. Internally, there is a mapping to a particular feature coordinator for each dimension. Furthermore, one can specify the names of the returned input patterns, which implicitly allows setting of the number of input patterns.

Initially, each input pattern is created with default parameters; and all parameters set to their default values. E.g. the x , y and orientation parameters are all set to 0. The appropriate feature coordinators are then applied to these input patterns, which gradually synthesizes more and more dimensions of the patterns. Finally, the individual patterns are passed to a composite generator, which returns the input pattern. It is worth mentioning that this concept is general with respect to the number of input

patterns returned, i.e. more than two input patterns can be created, thus allowing the modeling of animals with more than two eyes, e.g. three-eyed *Rana pipiens* (Reh and Constantine-Paton, 1985).

Pattern coordinator for images

In order to use images as inputs, a special class has been defined. This class allows the reading of all images from a specified folder, and the building of input patterns from them. Because some of the dimensions are inherent to images, the corresponding feature coordinators are skipped - namely, the orientation coordinator, color coordinator and spatial frequency coordinator are skipped by default.

The pattern coordinator for images also allows the filename of the images to be specified with templates. This is useful in the case that even more dimensions are inherent to the image dataset. For example, the scenario in which one folder is not enough to ensure the relationship of the corresponding left and right images in a stereo image dataset is maintained. This can be circumvented by using filename templates with patterns that are replaceable depending on whether the pattern is used as left or right input pattern. Because the concept is very general, this not only applies to stereo image datasets, but also to videos supplied to a model as successive images.

Summary pattern coordinator

The pattern coordinator provides default implementations with a standard set of parameters for common manipulations of input patterns. This has the advantage that the modeler does not need to specify them in detail. However, if needed, the modeler can override the behavior of certain feature coordinators by setting their parameters appropriately, or even decide to implement a new feature coordinator for a certain feature. In general, this is not needed, and therefore a high-level interface to input patterns was created.

This also reduces the duplication of code considerably where, previously, input patterns were created within the model files. When creating a new model, the input pattern generation was usually copied from another file for convenience and then altered to ensure the desired dimensions were covered. Therefore, the input pattern generation was duplicated across models, but contained many small differences depending on the used image dataset, the number of eyes, and whether, for example, motion was modeled. This approach was very fault prone.

Using the pattern coordinator class, a new instance of the class can be created, and the desired dimension for modeling is the only required parameter. Switching to an image dataset from an artificial stimuli such as Gaussians is also straightforward, as the only change necessary being the instantiation of an image pattern coordinator class rather than a normal pattern coordinator class. The parameters do not need to be changed. This also allows the comparison of the resulting maps of different models, as one can enforce the use of exact same input patterns across models.

4.2 Submodels

In section 4.1, a method for unifying the input patterns across model files was presented. Here, a mechanism to reuse so-called submodels across various models is described. This reduces the code duplication to a minimum while, at the same time, making differences across models easily visible. Furthermore, it allows for the parameters of the model to be changed before instantiation, and the structure of the model can be explored before model instantiation. This is important, as instantiating a model is expensive in terms of both time and memory consumption.

4.2.1 Previous implementation

So far, a model was created in a single file typically consisting of the following sections:

1. Import declarations
2. Input pattern generation
3. Sheet instantiation
4. Connecting sheets
5. Setup of analysis methods

The first section was typically shared across all model files, and imported various classes from *Topographica*.

The second section was similar across the model files, but with slight differences depending on the dimension one wants to model, as described in 4.1.2.

In the third section, the modeler had to create sheet instances. This involved determining the number of sheets, the particular sheet types to use, and passing the desired parameters to the sheets. This is a relatively simple procedure when the model consists of just a few sheets, but becomes very complex for a model covering many input

dimensions, as the number of sheets increases quickly. Typically this has been done with nested loops - once dimensions could be enabled/disabled with parameters, the increased complexity necessitated the addition of many conditional statements.

The created sheets were connected in the fourth section. Again, this required a complex set of operations containing various nested loops and conditional statements in more complex models. The parameters for the connection had to be passed to each connection statement, which resulted in duplicated code even within a model file. For example, a spatial frequency sheet was connected to the V1 sheet with exactly the same parameters as any other LGN sheet, the only difference being the size.

The fifth and final section involved setting up the measurement process, and was usually rather short. This, however, differed between the model files, as models covering more dimensions also have a more complex measurement setup.

4.2.2 Using a hierarchy of submodels

To avoid code duplication across models as described previously, the concept of submodels is introduced. The most general submodel, **Model**, provides a structure which has to be fulfilled by all subclasses, as well as a set of helper functions. These loosely correspond to the sections of the old model files as described in 4.2.1. Subclasses of **Model** then gradually become more specific through the implementation of more and more methods and the addition of more parameters. For example, in **VisualInput-Model**, parameters shared by models that receive visual input are added. This includes the input contrast and the area modeled by the cortex, amongst many others. Also, this class calls the pattern coordinator for input pattern generation.

Starting with an early vision model, the concept of levels is used. A level is a set of sheets in which the sheets within a level share most parameters and match conditions which are described later. In practice, the early vision model consists of two levels: **Retina** consisting of one or two retina sheets, and **LGN** containing at least two sheets, LGN On and LGN Off, but likely containing more sheets in complex models.

Rather than using nested loops, a product is used to specify the number of sheets. This might look as follows:

```
'Retina':args['eyes'],
'LGN': args['eyes'] * args['polarities'] * args['SFs']
```

The first line creates a level “Retina”, containing as many sheets as the list “args[‘eyes’]” contains entries, and labeled with the entries of this list. The second line creates the

“LGN” level. In the simplest case, where “args[’eyes’]” and “args[’SFs’]” are lists with just one entry and “args[’polarities’]” is a list with the entries “On” and “Off”, this statement creates only two sheets. However, if “args[’eyes’]” contains the entries “Left” and “Right”, “args[’polarities’]” is a list as before, and “args[’SFs’]” contains the entries “1” and “2”, eight sheets are created.

This is a huge improvement over using nested loops, which could sometimes reach up to 5 levels deep. For each of the levels, one method specifying the parameters of the level has to be implemented. Internally, this method can slightly adapt the parameters depending on the properties of the actual sheet. This avoids duplicating code. Here, the most general version of the **Retina** and **LGN** levels has been implemented. By setting the dimensions parameter of the early vision model accordingly, a very general version with over 20 sheets in total (see section 6.9), or a subset of this model with the basic GCAL model as the simplest model (see section 3.3.2) with just one sheet on the **Retina** level and two sheets on the **LGN** level, can be created using the same class.

The early vision model can be used by any model of the cortex. In this thesis, GCAL has been set up on top of the early vision model. However, the concept is general, and any other visual cortex model could also use the **EarlyVisionModel**. For GCAL, another level called **V1**, which consists only of one sheet with the same name, is introduced. Because the early vision model is used, code duplication across model files is avoided. Import declarations, training patterns and sheet instantiation are now taking place within one specific class shared across all models rather than being performed individually per model file. The following section shows how creating connections between sheets can be considerably simplified.

4.2.3 Automatically wiring up sheets

Traditionally, connections between sheets were explicitly stated using a **connect** command with the source sheet, the destination sheet, and the projection parameters as arguments. This again was usually done with nested loops. Using this machinery, it was hard to understand which sheets got connected, and extending this procedure to more sheets was fault prone. Here, match conditions set per level are proposed to specify the connections.

Before describing the match conditions, another concept of sheet properties needs to be introduced. This is a generalized version of a sheet name. A typical sheet name is **LeftLGNOn**, whereas **Left** refers to an eye, **LGN** refers to the sheet level and **On**

refers to the polarity. A sheet property allows each of those parts to be set and accessed individually.

Match conditions are defined per level. Each match condition consists of a set of criteria which must be fulfilled by an incoming sheet. As a simple example, V1 connects to all LGN sheets. Therefore, the only criteria is:

```
{'level':'LGN'}
```

To specify the match conditions, the properties of the sheet for which the match condition is specified can be used, e.g. connecting LGN sheets with a property **eye** to retina sheets of the same eye is implemented as follows:

```
{'level': 'Retina', 'eye': properties['eye']}
```

There, `properties['eye']` refers to the eye of the LGN sheet. This can easily be extended to also work in the case that only one eye is used, and therefore `properties['eye']` does not exist. Furthermore, it is possible to specify more than one match condition. This allows very specific connections, as well as connecting to different levels, e.g. there is usually one match condition for afferent connections and another match condition for lateral connections.

Within the **Model** class, a method that automatically checks whether a match condition holds and connects sheets has been implemented. For each pair of sheets, there is a test whether the source sheet fulfills all criteria in one of the match conditions of the destination sheet. If so, a connection between the two sheets is created. If none, or just a subset of the criteria are fulfilled, no connection is established. Because this method is implemented in a general manner, rather than explicitly connect sheets, all a modeler need do is create match conditions.

There is one method per match condition to specify the parameters of the projection in the case of a match. To adapt the parameters to the specific connection, the properties of the source as well as the destination sheet are passed to these methods.

4.2.4 Deferred model instantiation

By introducing specification classes for sheets and projections, which act as templates for actual sheet and projection instances, it is possible to investigate the network structure as well as the parameters of the sheets and projections, before instantiation. This is desirable when e.g. running many simulations in a batch mode with varying network structure. This is useful in order to be able to record the responses of all sheets, as the

names of the sheets have to be known before they get instantiated. This only becomes feasible using specification classes.

Even if the network structure across various simulations stays constant, one might want to change a specific parameter of a projection. Again, previously this was not possible, and only became feasible by introducing the specification classes. A resolve method allows accessing the actual sheet/projection after the specification object has been instantiated.

Models are now classes which can be instantiated. To instantiate a model, the model object is called. This internally calls all the specification classes which automatically instantiate actual sheets and projections between them.

4.2.5 Summary submodels

The introduction of submodels has greatly improved the work-flow of a modeler using Topographica. Modelers now have access to an early vision model and can build their model of the primary vision model on top of it. Alternatively, the GCAL model has been implemented and can be modified. The code duplication will be cut down to a minimum once all models have been changed to use the new system. Comparing models is now much easier than before, as the model files no longer need to be compared as a whole. Using the class-based approach, instead, a subclass should explicitly list all the modifications to the base class.

The machinery is set up in such a way that various cortex models can be used on top of the same retina and LGN sheets. Also, one can imagine that the set of LGN sheets is replaced by a single sheet with neurons which are randomly generated by choosing a certain number of units in the center, along with a non-specific surround (see Kneisel (2013)). Such a model could build upon the visual input model, replace the early vision model, and **ModelGCAL** could inherit from the new model, but would not need to be changed in other ways.

The implementation of submodels is currently based on inheritance. That is, a cortex model is inherited from the early vision model, and extends its functionality by adding a cortex level. In theory, this also allows the addition of higher visual areas on top of V1, by inheriting from the submodel that includes V1.

A limitation of the inheritance-based approach is that multiple models, all inheriting from the same submodel, cannot currently be combined. Imagine a model which adds V2 on top of V1, and another model which adds the visual area MT on top of

V1. These models are currently incompatible, and therefore cannot be combined. This could be circumvented by connecting submodels not via inheritance but, rather, in the actual model file as components. This would require substantial changes, as attributes would not be shared between the classes anymore. Instead, they would need to be communicated across the components.

The match conditions are of great importance for models covering many input dimensions, and hence a large number of sheets with complex wiring. These allow the properties of sheets to be specified such that each sheet allows incoming connections from other sheets. This not only simplifies the wiring up of the sheets, but also allows a single set of parameters associated with a match condition, rather than having the parameters spread over the model file.

Because of the helper functions introduced in the **Model** class, the modeler's main focus is shifted to deciding which sheets the model should include, which parameters these sheets should have, and which parameters the projection connecting the sheets should have. Previously, all those decisions still had to be made, but they were mixed with code. The submodels machinery allows defining parameters more explicitly, as parameters are now class attributes.

4.3 Motion model

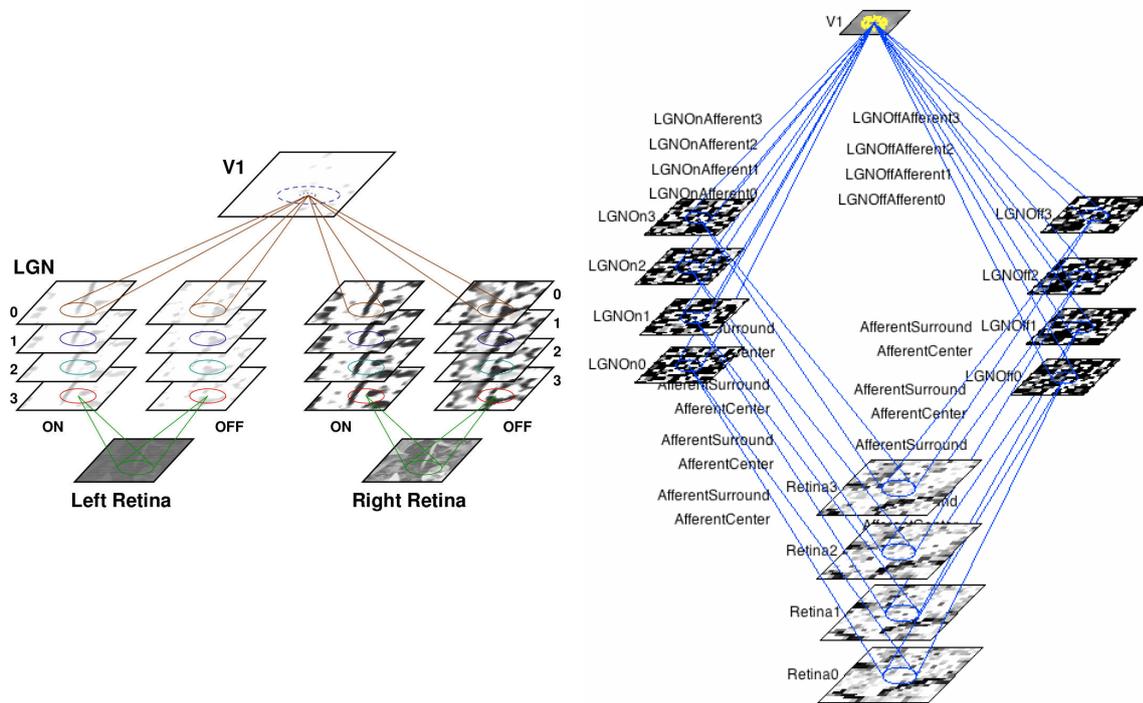
Neurons only become selective for motion if temporal delays are introduced to the model. These delays correspond to the lagged cells found in cat LGN (see section 2.2.6). There are various ways to implement this in a model.

4.3.1 Previous proposals

In Topographica, a motion concept based on lagged retina sheets has been implemented (Gerasymova, 2009). For each retina sheet, the input is swept further. For each retina, there is one LGN On sheet and one LGN Off sheet. For example, if one wants to model four motion lags, four retina sheets and eight LGN sheets are needed.

In the paper from Bednar and Miikkulainen (2006), it is shown how lagged retina sheets can be avoided. There are only two retina sheets (modeling the left and right eye), however they change their pattern over time. At time zero, the first LGN sheet with the highest lag computes its activation. At each further time step, the input pattern on the retina sheets is swept in a certain direction, and the activation of subsequent

LGN sheets with lower delay is computed. Once all LGN sheets have been activated, the responses are propagated to the V1 sheet. This is much more realistic than having multiple retina sheets representing motion.



(a) Model with two retina sheets, one for the left eye and one for the right eye. Temporal delays are added at the LGN level. Reprinted from Bednar and Miikkulainen (2006).

(b) Model with several lagged retina sheets, which introduce temporal delays that are propagated to corresponding LGN On and Off sheets. Reprinted from Gerasymova (2009).

Figure 4.1: Previous proposals to model motion selectivity.

4.3.2 New approach: Multiple projections from LGN sheets to V1

The idea from Bednar and Miikkulainen (2006) has been taken further for this thesis. Rather than modeling multiple LGN sheets representing lagged cells, just one LGN On and one LGN Off sheet is contained in the model. Each sheet has multiple projections to V1, and each projection has a different delay. That means that the activation in V1 is gradually built up over time. This concept had first been proposed in Bednar (2012), but was not yet implemented.

Although a prototype for this new representation of temporal delays existed previously, it had severe issues. As described in sections 4.1.1 and 4.1.2, there are several

advantages to sequentially applying feature coordinators in order to generate input patterns. This was not possible with the prototype, as all parameters needed to be set on the class handling motion. It was not possible to pass an input pattern to this class, and get a moved pattern returned.

In this thesis, the prototype has been improved to work properly with any kind of input patterns. Usually, the pattern is swept orthogonally to its orientation, but in theory any motion direction can be applied on arbitrary input patterns. This is done by generating a new input pattern on the supplied pattern generator only every n -th time step, rather than for each time step. At the first time step, the input pattern is returned as is. At the second time step, it is swept aside in the desired direction, and so on.

This implementation has the advantage that the model structure does not need to be changed in order to represent temporal delays. It looks exactly like the basic GCAL model described in 3.3.2; however, it has been extended with multiple projections from LGN sheets to V1, which allow modeling motion. In section 6.5, it is shown that this representation of temporal delays results in qualitatively similar results compared to the previous proposals.

4.4 Discussion

This section critically analyzes the feature coordinators, the pattern coordinator, and the class-based submodels system introduced in this chapter. Strengths and limitations are highlighted, and suggestions for improvements and extensions are given.

4.4.1 Feature coordinators

Feature coordinators are a very general concept. A feature coordinator is supplied an input pattern and expected to modify this pattern in a certain way, returning a new pattern based on what was supplied. This concept is very powerful, and has proven to successfully synthesize the following dimensions: position, orientation, direction, spatial frequency, dimming and disparity. A challenge for the future will be synthesizing color for mathematically defined inputs such as Gaussians - this is challenging due to the complex nature of the task, not because a feature coordinator is too general or too limited in its capabilities (see section 7.2.4).

From an implementation point of view, it is challenging to transfer this concept to measure maps. This is the inverse problem to the generation of training input patterns.

As part of Topographica, the feature mapper concept exists; its job is to coordinate the presentation of patterns and measure of the responses. Within feature mapper, a similar hierarchy of methods exists which creates patterns that vary in a certain way. It is desirable that feature coordinators are used within feature mapper in the future, although this would require substantial work. Rather than randomly varying feature dimensions for training inputs, the dimensions need to be varied in a deterministic way to measure responses.

In summary, feature coordinators have improved the input pattern generation considerably. A clear, extensible hierarchy of feature coordinators evolved from an ad-hoc way of defining input patterns with a lot of random number generators bundled together. Generalizing this concept even further to also measure maps is future work.

4.4.2 Pattern coordinator

The pattern coordinator class provides a convenient interface to the feature coordinators described above. Furthermore, it is designed to work with artificial stimuli such as Gaussians as well as images to be read from a folder. This allows modelers to use predefined input patterns rather than having to define patterns themselves. They can choose from a wide range of dimensions that the pattern should have. Furthermore, changing from artificial stimuli to natural image inputs is now straightforward, as all that needs to be changed is using the subclass **PatternCoordinatorImages** rather than the superclass **PatternCoordinator** directly. The parameters can be kept the same, except that a folder for the image dataset needs to be supplied.

This is superior to the previous definition of input patterns, where specifying the dimensions to be covered had to be done with conditional statements. Also, the input patterns often varied between models even if the same dimensions were covered, because the seeds for the random number generators were set differently. The pattern coordinator provides a standardized interface, and random streams are guaranteed to be the same even if used in different models.

One drawback is that only the individual patterns, but not the composite input pattern generated of the individual patterns, can be modified, e.g. it is not possible to sweep the whole pattern in a certain direction, but only the individual patterns. This could be addressed in future extensions.

4.4.3 Submodels

The class-based modeling approach using submodels is one of the main contributions of this thesis. Here, the strengths and limitations of this new machinery are discussed.

One of the main advantages is that model files are now just instantiating a class rather than defining the whole model. There is a hierarchy of classes, and a modeler can choose to use only the very basic **Model** class, for example in the case a model of the auditory cortex is to be created. At the other extreme, a modeler can use **ModelGCAL** to model the primary visual cortex, and only modify a certain method in a subclass. Alternatively, a modeler can choose to build up on any level of the hierarchy in between these extremes. Whatever decision is made, the submodels can be easily reused.

The **Model** class provides a set of helper functions which simplify the workflow considerably. It is much easier, for example, to define how many sheets and projections a model should have. The projections between sheets can be easily set up using simple rules called match conditions, which avoids using nested loops. Also, previously a subset of dimensions within a model was made possible with conditional statements; this is now mostly needless, as the helper functions skip parameters which are not set because the dimension they belong to is not modeled.

The biggest limitation of this approach is that it is mostly restricted to feed-forward models. Each level in the hierarchy of classes adds more sheets which are to be connected to sheets of lower levels. Therefore, the higher levels in the hierarchy have to know which sheets in the lower levels exist, and how to connect to them. It would be desirable that, rather than having a hierarchy of classes, one would have components which are interchangeable.

This component-based approach would be a major extension to the current implementation. One of the biggest issues is that attributes are currently shared across the different levels, and would need to be communicated appropriately across the levels instead. In addition, it is not clear how feedback connections could be specified.

4.5 Architectural summary

This chapter introduced a set of tools which are necessary in order to build an all maps model. The input pattern generation was modularized, and the addition of dimensions was simplified. Each feature coordinator synthesizes exactly one dimension, and applying them sequentially on a template pattern results in input patterns that cover all

dimensions modeled in this thesis.

The class-based submodels system is superior to the previous way of defining models. For a modeler, it is more obvious how sheets and projections between sheets are created. The code duplication within one file and across files was cut to a minimum. Submodels can be reused and/or modified depending on the requirement. This makes differences between related models much more visible than before. It is expected that new models in Topographica will consistently use this new system.

Multiple projections with increasing delays from LGN sheets to V1 is a novel way of modeling motion. To the author's knowledge, it is the first time that this has been addressed in a publication, although originally proposed in Bednar (2012). The prototype has been extended considerably, and biologically realistic direction maps are resulting, as shown in 6.5.

In summary, a set of tools necessary to model an all maps model of the primary visual cortex has been developed. Furthermore, a new way of modeling motion has been presented. It is expected that the pattern coordinator using the feature coordinators and the submodels system will be used by the majority of Topographica users in the future.

Chapter 5

Methods

Starting with the basic GCAL model, which was shown to result in biologically realistic orientation maps in section 3.3.2, this chapter demonstrates how this model can be extended to model all other known cortical maps. First, it is shown which extensions are needed to model the other cortical maps individually. This includes the changes in the model structure, as well as the modifications of the input patterns using feature coordinators. For each model, a graphical representation of the model structure as well as sample input patterns are printed. Note that all input patterns are created using the pattern coordinator, which itself makes use of the various feature coordinators (see section 4.1).

As a sample of how individual models can be combined, a model for color preference and ocular dominance is presented. This example has been chosen because this combination of dimensions was known not to work in Gerasymova (2009). Therefore, in the corresponding section, the required changes are highlighted.

A combination of the various individual models, resulting in a combined model, is described at the end of this chapter. This model covers the orientation, ocular dominance, disparity, direction, spatial frequency and color dimensions simultaneously. Although the complexity of this model is still very high, the model is much simpler compared to the work in Gerasymova (2009).

5.1 Orientation preference

To model orientation preference, the basic GCAL model as described in section 3.3.2 is used as is. Figure 5.2 shows the Gaussian input patterns which are used. The orientation of the Gaussians are distributed uniformly random between 0 and π . Alternatively,

natural images can be used (not shown here, see sections 5.6 and 5.7 for reference).

The structure of this model is very simple. The retina sheet outputs its activity to the LGN sheets, which perform a Difference of Gaussian operation as well as gain-control, which enables the model to work for a wide range of input contrasts. The V1 sheet is connected to both LGN sheets. Furthermore, short-range lateral excitatory and long-range lateral inhibitory connections between neurons in V1 exist. The afferent connections from the LGN sheets, as well as the inhibitory connections, are adapting using the Hebbian learning rule.

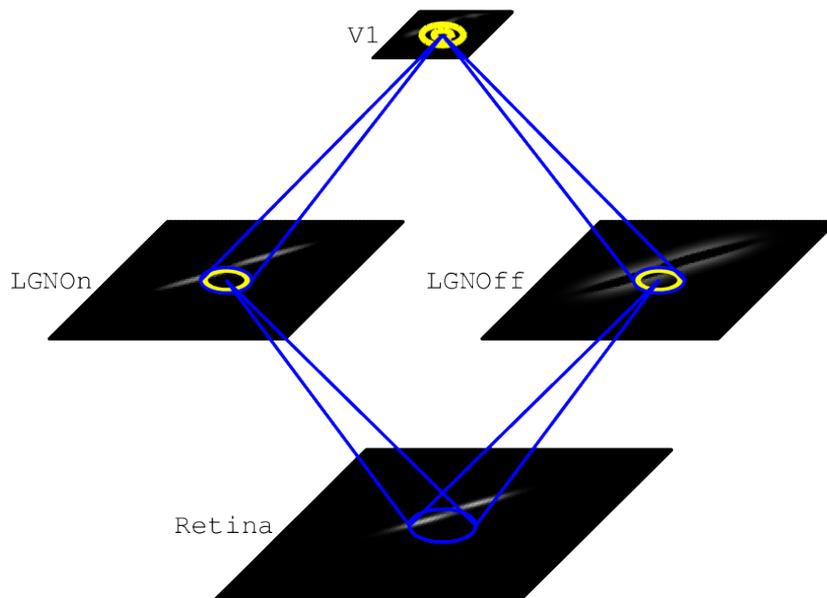


Figure 5.1: Architecture of a GCAL model with neurons selective for orientation. A Gaussian input pattern is shown on the retina sheet. Blue cones from the LGN sheets to the retina sheet visualize the receptive field of a single LGN unit. Yellow circles at the LGN level represent the lateral connections within a LGN sheet due to gain control. The activity at the LGN level is shown as well. There, one can see that units on the LGN On sheet are active in regions where the center of the receptive field is brighter than the surround of the receptive field. For the LGN Off sheet, units are active in regions where the surround of the receptive field is brighter than the center. Again, blue circles visualize the receptive fields from V1 to the LGN sheets. Two yellow circles on the V1 sheet represent the excitatory lateral connection (smaller circle) and the inhibitory lateral connection (bigger circle). Whiteish areas on the V1 sheet indicate a slight activation of V1.

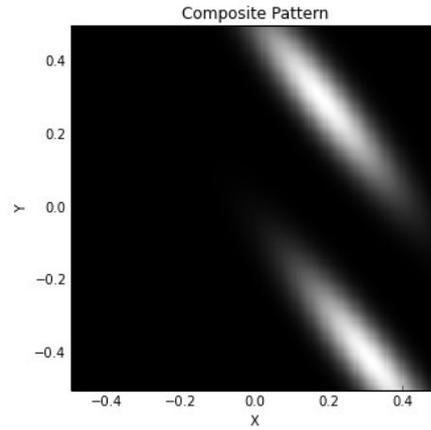


Figure 5.2: Oriented Gaussian input pattern for the basic GCAL model.

5.2 Ocular preference

As ocular dominance expresses the preference for one eye over another eye, a second retina sheet representing a second eye needs to be added for ocular preference simulations. Therefore, the simulation contains two retina sheets representing the left and right eye. For each retina sheet, there is one LGN On and one LGN Off sheet. As before, all LGN sheets connect to V1.

In GCAL (and its predecessor LISSOM), biologically realistic ocular dominance maps only emerge due to brightness differences in the input. The fraction by which the pattern brightness varies between the two eyes is set to $\varepsilon = 0.7$. The brightnesses β_{left} and β_{right} of the left and right eyes are then calculated as follows, where ξ is a random number between 0 and 2 (see a sample pair of input patterns in figure 5.4):

$$\beta_{left} = (1 - \varepsilon) + \varepsilon \cdot (2.0 - \xi)$$

$$\beta_{right} = (1 - \varepsilon) + \varepsilon \cdot \xi$$

In comparison to previous simulations using LISSOM, additional connections between the LGN sheets of the two eyes are required. This is due to the gain control at the LGN level. To recall, gain control allows using a wide variety of input contrasts, by dividing each units activity by the sum of the activity of surrounding units. Therefore, a brightness difference is eliminated by gain control, as the ratio of numerator and denominator in equation 3.5 is approximately constant for different brightnesses (as long as the sum of activations is sufficiently bigger than k).

To circumvent this behavior, lateral connections between the two LGN On sheets as well as between the two LGN Off sheets are introduced. Then, the denominator in

equation 3.5 is the same for the units at position j of both sheets (in contrast to adding up the activations of just one sheet, which results in different denominators each sheet), but the activation $\Psi_i(t)$ of the retina sheet is higher for either the left or the right eye, and therefore the activation $\eta_{j,O}$ differs between the LGN sheets.

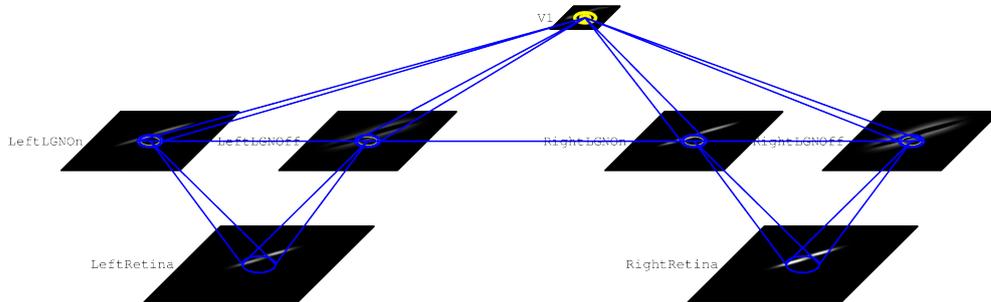


Figure 5.3: Architecture of a GCAL model with neurons selective for orientation and ocular dominance. Compared to the orientation-only model, a second retina sheet with corresponding LGN sheets is introduced. The blue lines across the LGN sheets indicate lateral connections due to gain control of sheets with the same polarity. Please accept apologies for the small font in this and some of the following figures showing model architectures.

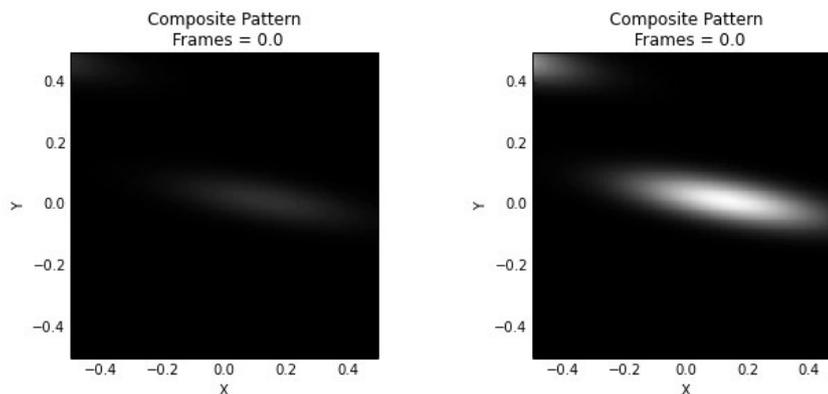


Figure 5.4: Gaussian input for a GCAL model with neurons selective for orientation and ocular dominance. One can clearly see the brightness differences of the input patterns. In this example, the right pattern is much brighter than the left pattern.

5.3 Disparity preference

For modeling disparity preference, the same model as for the ocular preference simulation is used, but the input differs. Rather than using brightness differences between the two retina sheets, in the disparity simulation the pattern location is slightly offset between the eyes, as shown in figure 5.5. The position of the patterns is calculated as follows, whereas ξ is a uniformly distributed number between 0 and $\xi_{max} = 0.08333$, and x is also uniformly distributed between 0 and $0.75 - \xi_{max}$:

$$x_{left} = x + \xi$$

$$x_{right} = x - \xi$$

Here, it has to be stressed that ocular preference and disparity are two separate dimensions, although modeled in the same way. A neuron in V1 can prefer one eye over the other, while keeping the activity level constant regardless of the input pattern disparity. On the other hand, another neuron might be highly selective for disparity, only reacting to input patterns with a specific offset, while being selective to input from both eyes.

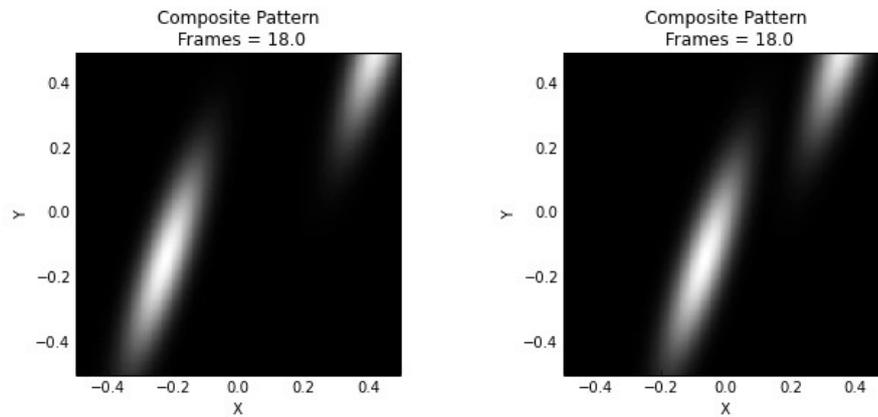


Figure 5.5: Oriented Gaussian input pattern with position offset. In this example, the bottom pattern is clearly offset to the right in the right retina (hitting the x-axis at ≈ -0.35 in the left retina, and at ≈ -0.2 in the right retina). The pattern at the top is slightly offset to the left in the right retina.

5.4 Ocular and disparity preference

When modeling ocular dominance and disparity simultaneously, the input patterns of the two eyes have to vary in brightness as well as their position offset. This undertaking is straightforward using the pattern coordinator. The model structure itself is not altered, compared to the individual ocular dominance/disparity simulations.

This simulation is important, as earlier studies have shown that disparity by itself cannot explain biologically realistic ocular dominance maps (Miikkulainen et al., 2005), and it is clear that a disparity map can only emerge if there is an offset between the pattern locations of the two eyes. Therefore, only if the input pattern show both, brightness differences and position offsets, ocular dominance maps and disparity maps become apparent. Figure 5.6 shows a pair of sample inputs where both features are visible.

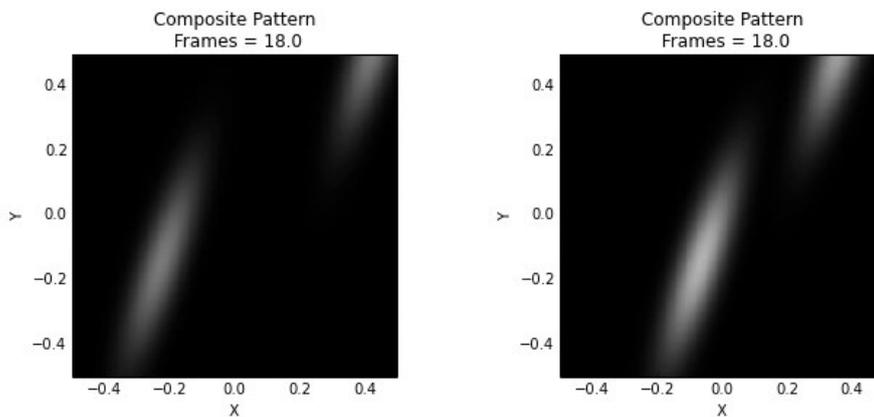


Figure 5.6: Oriented Gaussian input pattern with position offset and brightness offset. The same inputs as in figure 5.5 are shown, but now the pattern on the right retina is slightly brighter than the pattern on the left retina.

5.5 Direction preference

A new motion model was introduced in section 4.3.2. The model structure compared to the basic GCAL model consists of the same sheets (a figure of the model structure is therefore omitted), however multiple delayed projections (here: $n = 4$) from the LGN sheets to V1 are introduced. To account for these additional projections, the learning

rate from neurons in V1 is divided by the number of projections. Furthermore, the strength for these projections is increased by a factor of 1.5 to push the activity level of neurons in V1 closer to their desired average activity.

The main contribution in this thesis is the introduction of a feature coordinator for motion, which moves the input pattern orthogonally to its orientation in the case of Gaussian patterns, where the orientation ranges between 0 and 2π (in contrast to 0 and π which would only allow upwards/leftwards motion rather than motion in all directions). For natural images, the direction is chosen randomly between 0 and 2π as it is hard to determine the orientation of a natural image pattern. At each time step, the pattern is swept by 3.0/24.0 units.

After n time steps, a new pattern is propagated from the retina sheets to V1. A sample series of Gaussian patterns for time $t = 0$ until time $t = n$ (with sweeping of one pattern from time $t = 0$ until time $t = n - 1$ and a subsequent onset of a new pattern at time $t = n$) with $n = 4$ is shown in figure 5.7.

Visual inspection has raised the issue that the connection fields of neurons show unusual behavior at the edges, if run with standard parameters. That is, especially for the projections which are activated at $t = 0$ and $t = n - 1$, neurons at the LGN layer which are at the very edge of the connection field of neurons in V1 show an unusually high weight. To circumvent this behavior, the position boundaries of the Gaussian input patterns are increased. Therefore, the patterns can be “tracked” over time (as in, they are seen by neurons in V1 from time $t = 0$ until time $t = n - 1$). As a rule of thumb, the boundaries are increased by $n - 1$ times the amount of the pattern is swept per time step.

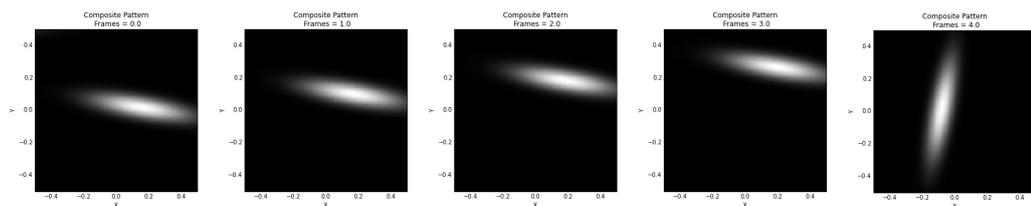


Figure 5.7: Oriented Gaussian input pattern moving over time. The leftmost pattern shows the input at $t = 0$. Then, at $t = 1$, the pattern is moved orthogonally to its orientation slightly further to the top-right. This sweeping continues until $t = n - 1$. At $t = n = 4$, a new pattern with changed orientation is onset. This pattern again is swept for n time steps, before another pattern serves as an input model input (not shown here).

5.6 Spatial frequency preference

Following the proposal in Palmer (2009), multiple LGN sheets selective for different spatial frequencies in the input are extended to model spatial frequency. Here, there are $\nu = 3$ different spatial frequency channels with different sizes of receptive fields in the photoreceptor sheet. The receptive fields for successive spatial frequency channels increase by a factor of $\Gamma = 2.5$. Both, ν and Γ are represented as variables in the source code and can easily be changed. However, more than $\nu = 4$ spatial frequency channels are computationally not feasible.

Similarly to the ocularity simulation, additional lateral connections between the spatial frequencies are added. This allows the different spatial frequency channels to “compete” with each other. A more detailed investigation of this principle would be desirable, but is left for future work.

Figure 5.8 shows the extended model, which also displays the different activity values of the spatial frequency channels depending on the present frequencies in the

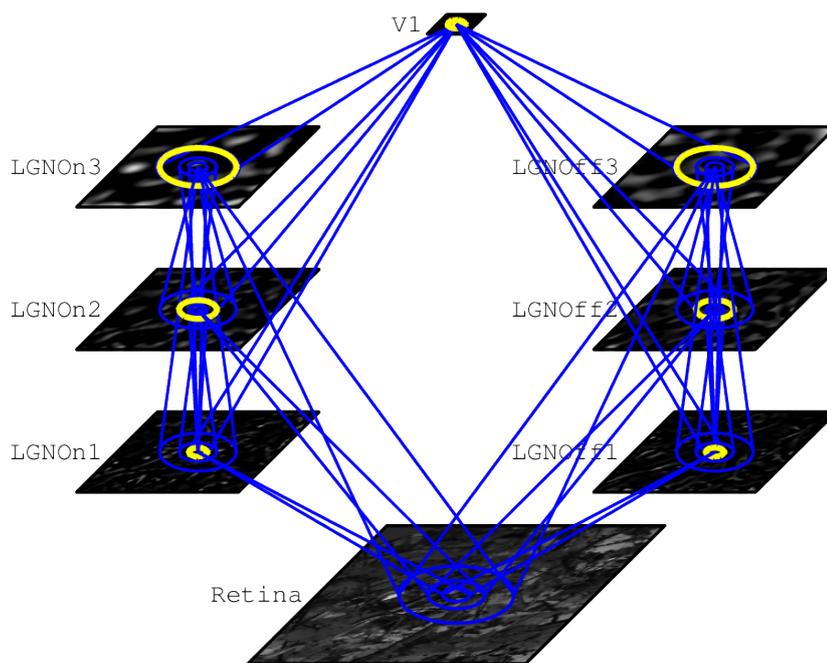


Figure 5.8: Architecture of a GCAL model with neurons selective for orientation and spatial frequency. Compared to the orientation-only model, a second and third spatial frequency channel with each one LGN On and one LGN Off sheet is introduced. Blue cones between the spatial frequency channels show connections due to gain control. The size of these connections has to be investigated further, as discussed in section 7.2.

input image.

As input patterns, (gray-scale) natural image patches from the McGill dataset (Olmos and Kingdom, 2004) shown in figure 5.10 are used. This has the advantage over Gaussian input patterns in that different spatial frequencies are inherent to natural images, whereas they would need to be synthesized for Gaussian patterns. This is still work in progress, see section 7.2.4 for further discussion.

5.7 Color preference

In comparison to Bednar et al. (2005), a new color model by Ball (2014) (improved and updated to work with the most recent version of Topographica by Spigler (2014)) is used. Separate retina sheets for each color are replaced by a single retina sheet, with the possibility to access different color channels. This further reduces the number of retina sheets, as now just like in biology just one sheet per eye is needed. Each color channel within the retina can be seen as one type of cone cells.

Following the opponent-process theory (Hurvich and Jameson, 1957), four pairs of sheets at the LGN level modeling the following opponent channels are present: red-green, green-red, blue-yellow and black-white (also known as luminance). For the red-green pair, the center of the receptive field is connected to the red channel and the surround is connected to the green channel. For the green-red pair, the center and surround are inter-converted. The blue-yellow pair is a special case, as it is modeled spatially co-extensive as proposed in Dacey (2000). Therefore, the receptive fields in the center from the blue channel and the surround which is connected to the red and green channels are equally wide. The luminosity pair receives input from all color channels, in the center as well as the surround. The LGN On sheets receive excitation in the center and inhibition in the surround and vice versa for the LGN Off sheets. A visualization can be found in figure 5.11.

The model was built in a way that also a dichromatic simulation is feasible, however this is not shown here. In contrast to the ocular dominance and spatial frequency simulations, gain control is not used for color selective channels (red-green, green-red and blue-yellow) due to the lack of gain control in the color sensitive parvo cells (Solomon and Lennie, 2005; Pitkow and Meister, 2012). Therefore, gain control is only used for the luminosity sheets.

As for the spatial frequency simulation, natural images from the McGill dataset are used as input patterns (see figure 5.10). The ratio between strengths of LGN color

versus luminosity channel connections to V1 has been chosen as 0.1, i.e. the strength of the connections from luminosity sheets to V1 is much higher compared to the strength of color connections.

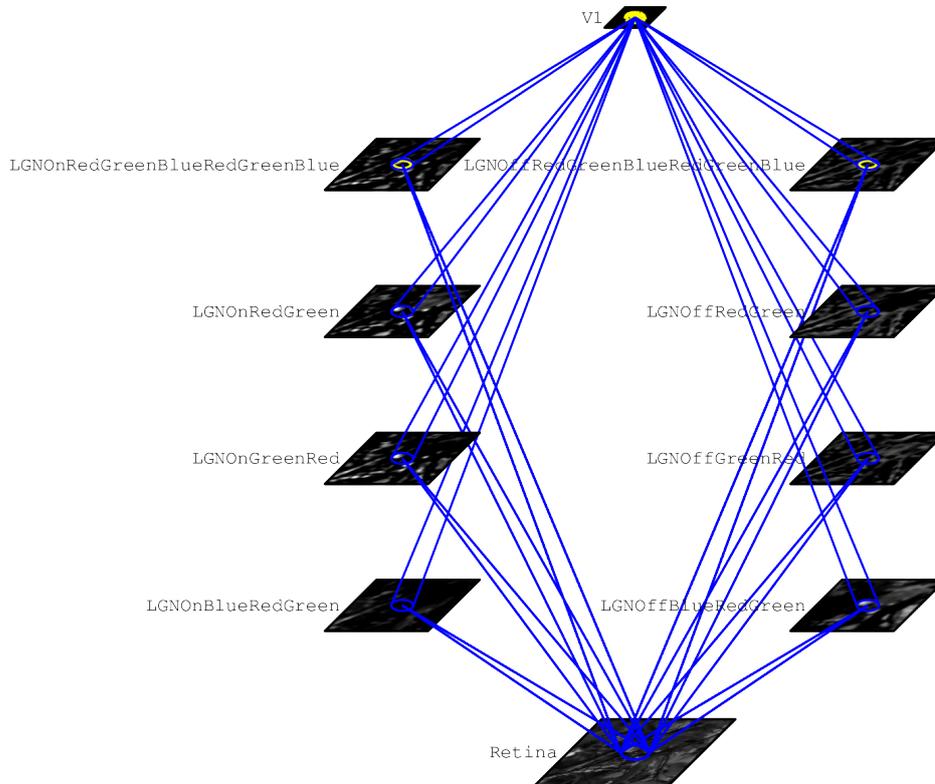


Figure 5.9: Architecture of a GCAL model with neurons selective for orientation and color. Compared to the orientation-only model, color-selective LGN sheets are introduced. Here, “RedGreenBlue” stands for luminosity, as the luminosity channel gets input from all three cone types, here modeled within a single retina. “RedGreen” sheets are wired internally to get their center input from the red cones, and surround input from the green cones. For the “GreenRed” sheets, this is interchanged. The “BlueRed-Green” sheets get their center input from the blue cones, and their surround input from both, red and green cones. These sheets are modeled spatially co-extensive, i.e. the receptive fields of center and surround are of equal size. Please accept apologies that the pattern on the retina sheet is drawn in gray-scale rather than color.

5.8 Color and ocular preference

In comparison to the color preference simulation above, a second eye with exactly the same setup as the first eye is introduced. This then also requires four more pairs of

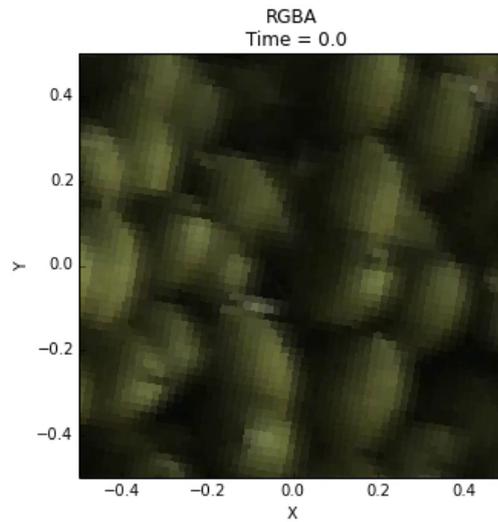


Figure 5.10: Example natural image input pattern

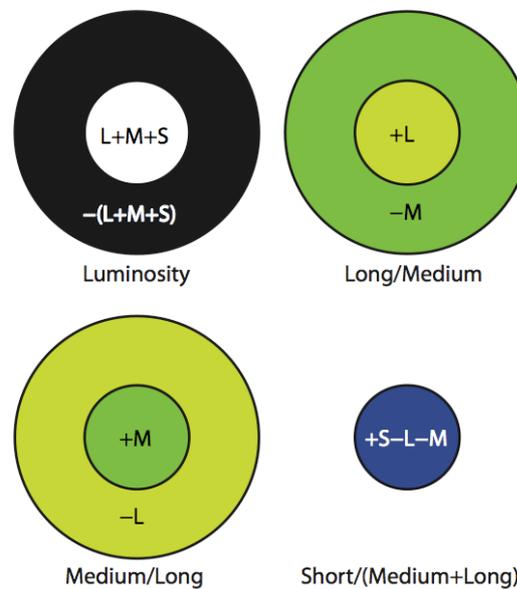


Figure 5.11: Visualization of receptive fields in color simulations. Here, Long (L) stands for the red cone type, Medium (M) for the green cone type and Short (S) for the blue cone type. Long/Medium/Short are labels according to the wavelengths the cones are responsive to. Figure reprinted from Paula (2007).

opponent channels at the LGN level for the second eye. Otherwise, the model structure is the same, as can be seen by comparing figures 5.9 and 5.12.

In comparison to the previous all maps model (see 3.3.3), the input patterns between the eyes actually differ in their brightness/value in the HSV color space rather than subtracting the same constant from the R, G, B values which results in a changed

hue value.

The implementation is as follows: each pixel i is read from the image file, resulting in the RGB values r_i, g_i, b_i . Then, these values are converted into the HSV color space, resulting in values h_i, s_i and v_i . As for the pure ocular dominance simulation, a variable $\varepsilon = 0.7$ is introduced describing the brightness difference between the eyes. Here, the random variable ξ takes values between $(-1 + \varepsilon)/2.0$ and $(1 - \varepsilon)/2.0$. Then, the HSV values for the left and right eye are calculated as follows:

$$h_{i,left} = h_{i,right} = h_i$$

$$s_{i,left} = s_{i,right} = s_i$$

$$v_{i,left} = v_i + \xi$$

$$v_{i,right} = v_i - \xi$$

Converting these HSV values back to RGB leads to image pairs where one of the patterns has a higher brightness than the other, while still having the same hue for each pixel. Note that in some cases a cropping is necessary, where the values $v_{i,left}$ and $v_{i,right}$ above 1.0 are cropped to 1.0. A sample image pair is shown in figure 5.13.

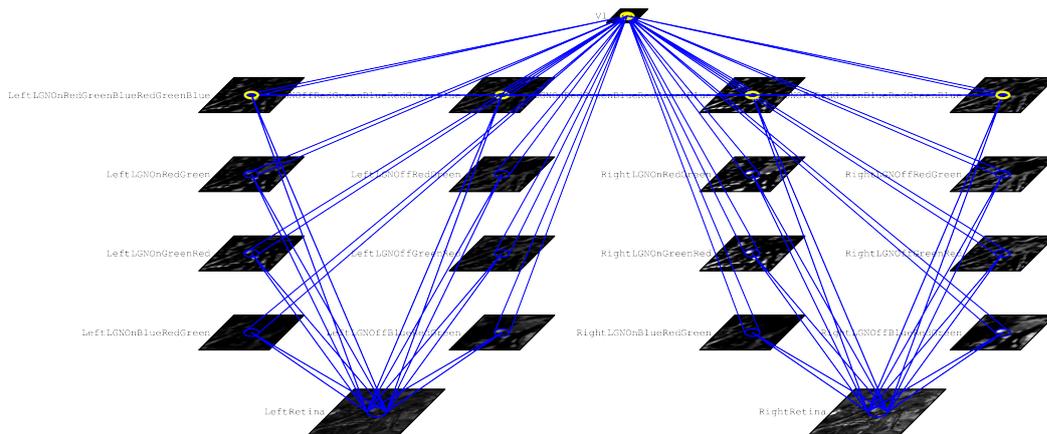


Figure 5.12: Architecture of a GCAL model with neurons selective for orientation, ocular dominance and color. Compared to the color model, a second retina sheet and corresponding LGN sheets are introduced. Furthermore, lateral connections between the luminosity sheets are added.

5.9 Combined model

Here, it is shown how the individual models presented above are combined to build a unified model with neurons selective for the following feature dimensions: positional

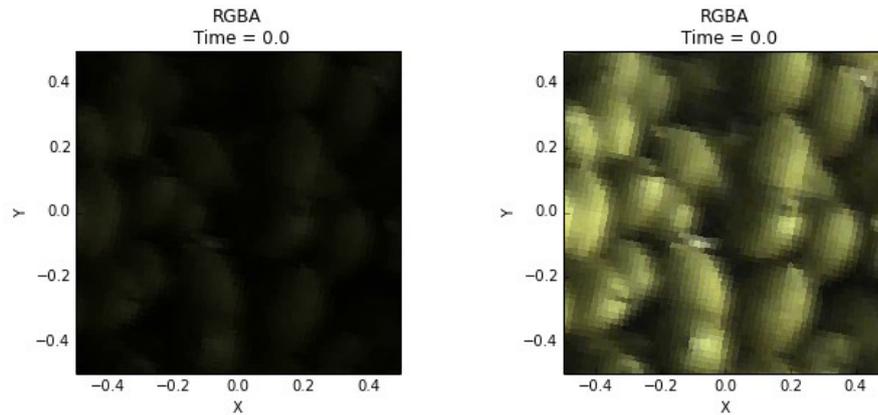


Figure 5.13: Example of natural image patterns in the case of a two-retina simulation. There, one can see clear brightness differences between the left and right input images. However, the corresponding hue values in the two images are the same.

preference, orientation, eye of origin, phase disparity, direction, color and spatial frequency. In order to build a unified model, the concepts of all individual models have to be combined. Furthermore, the input patterns which are presented to the model have to cover all input dimensions so that feature preferences of neurons in V1 emerge.

The basic GCAL model with one retina sheet, one pair of sheets at the LGN level and one sheet representing V1 is extended as follows: To model color, the one pair of LGN On/Off sheets is replaced by four opponent channels as described in section 5.7. Extending this model to also include spatial frequency is possible by introducing $\nu - 1 = 2$ spatial frequency sheets, as the first spatial frequency channel is implied by the luminosity sheets. As the input, colored natural images are used, which do not need any adjustments at this stage as color and spatial frequency are both inherent.

By mirroring this pathway, i.e. introducing a second retina sheet and building the same structure of LGN sheets for this sheet, the ocular dominance and disparity dimensions are covered. Unfortunately, there is no color calibrated stereo dataset, and therefore disparity has to be synthesized as described in section 5.3. Also, brightness differences are introduced (see section 5.2).

Introducing temporal delays from every single LGN sheet to V1 results in a preference for motion. As in section 5.5, $n = 4$ projections per sheet with increasing delays are used. The motion of the image patches is synthesized.

Figure 5.14 is visualizing the architecture of the all maps model. In total, there are

27 sheets, two of them on the retina level, 24 LGN sheets and the V1 sheet. These are connected by in total 158 projections, thereof 60 from the retina sheets to the LGN sheets (including one projections per cone per color selective sheet), 72 gain control connections at the LGN level, 26 connections from the LGN sheets to V1, as well as the two lateral projections within V1. The input for the combined model for both retina sheets at different times is shown in figure 5.15.

5.10 Summary

Starting with the basic GCAL model, this chapter first introduced the required changes to this model to cover the other feature dimensions individually. Changes made are the extension of the model by more sheets, the extension by more projections, as well as the alteration of the input patterns to cover the required features, and in most of the cases, a combination of these.

Then, the ocular dominance and color model have been merged to show the combination of two previously separate models. In this particular case, the four opponent channels were duplicated, whereas the duplication is wired to the second retina sheet introduced by the ocularity model. One could also think of it the other way around. Starting with the ocularity model, the previously two pairs of LGN sheets were replaced by four opponent channels. This clarifies that there is no “order” of dimensions, and any subset of dimensions can be modeled.

At the end of the chapter, a unified model resulting in selectivity for all known spatial dimensions in the primary visual cortex was presented. This is still a very complex model, however following the proposals in Bednar (2012) is much more comprehensible than the previous attempt to build an all maps model in Gerasymova (2009).

In the next chapter, the resulting maps of these models are shown and discussed.

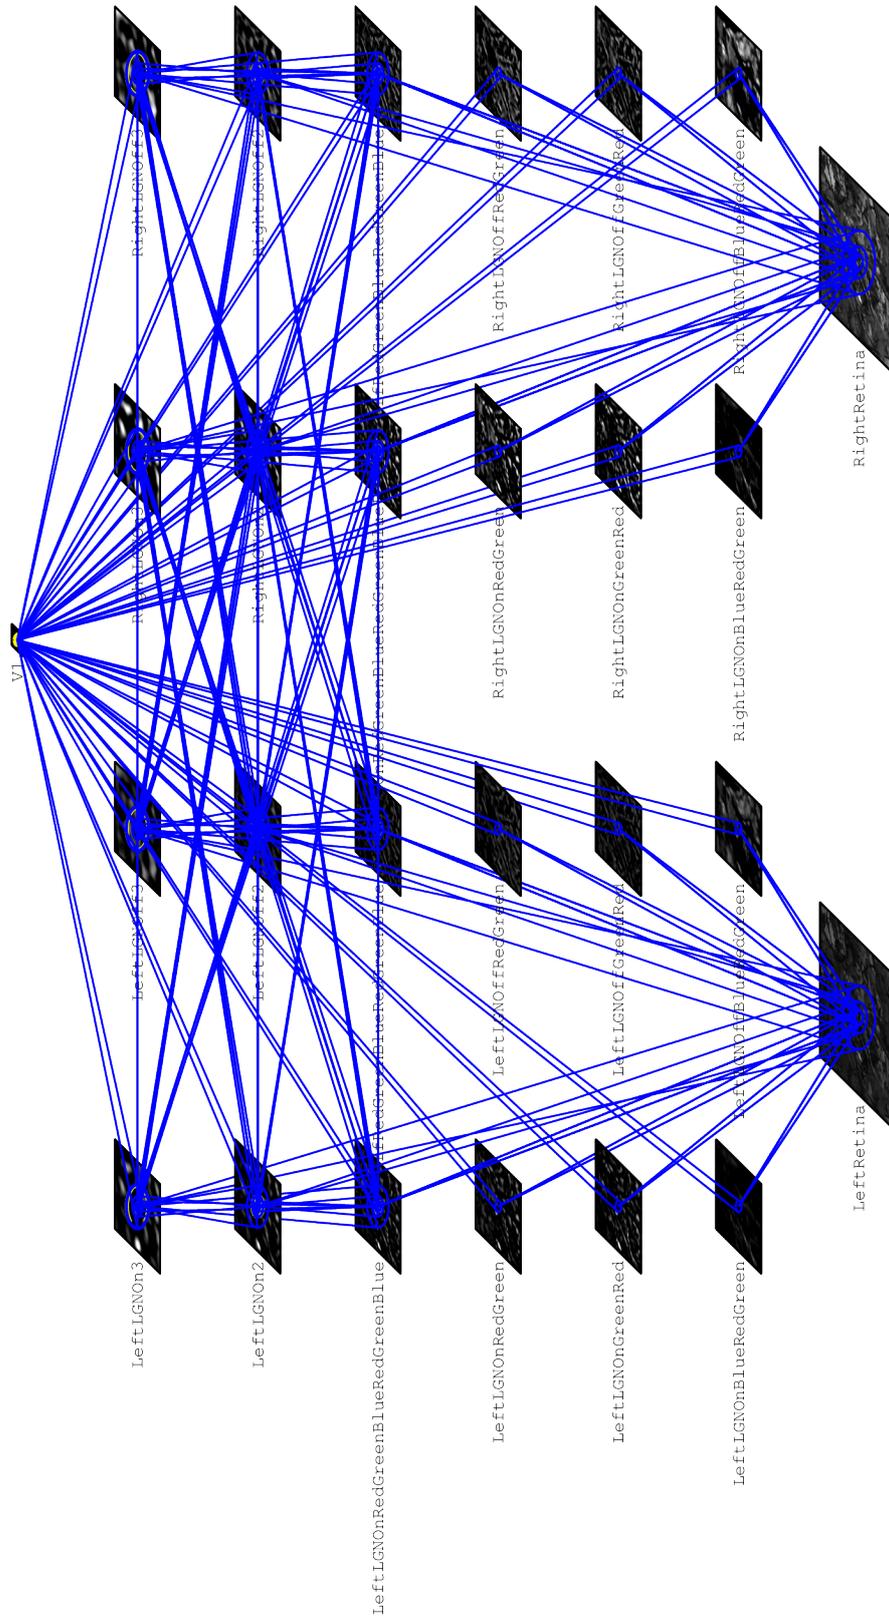


Figure 5.14: Architecture of a unified GCAL model with neurons selective for orientation, ocular dominance, disparity, direction, spatial frequency and color. Separate pathways for the left and right eye allow neurons to become selective for ocularity as well as disparity. Color sensitive LGN sheets allow the development of a hue map. Furthermore, LGN channels which are responsive to different frequencies in the input let neurons in V1 become selective for a range of spatial frequencies. Temporal delays introduced by multiple projections from every single LGN sheet to V1 model selectiveness for motion.

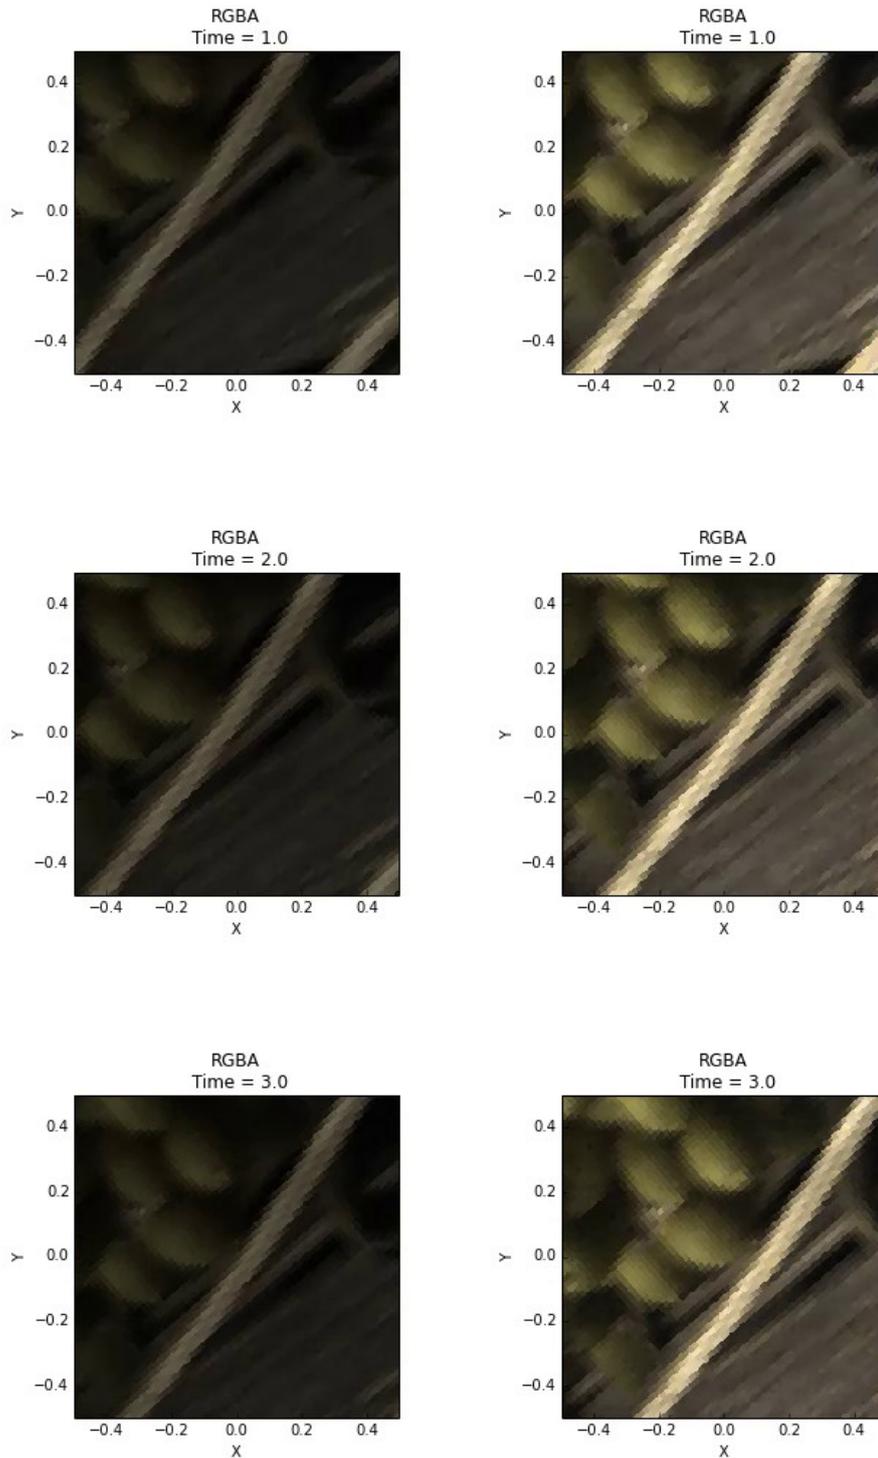


Figure 5.15: Example natural image patterns for the combined model. The brightness of the pattern varies between the eyes. Furthermore, in this example the pattern is slightly offset to the right in the right retina. Also, the pattern is swept over time, here in south-east direction.

Chapter 6

Results

In the previous chapter, a variety of models accounting for different cortical maps were introduced. Here, the resulting maps are investigated, and for the models combining feature preferences the interplay between the maps is discussed. They are also compared to the maps found in animals, as reviewed in chapter 2.

For each model, the orientation map after 20000 training iterations is shown, including an indication of pinwheels which were found. Furthermore, the Fourier power spectrum is shown, which visualizes how frequently orientation patches repeat. Biological maps usually have ring-shaped Fourier spectra (Blasdel, 1992a,b). The integral of the power spectrum over the radius is shown in the bottom-right corner of these plots. A red line indicates the estimated radius of the isotropic ring (Stevens et al., 2013).

The maps are measured by presenting sine gratings with varying features to the retina. For example, to measure orientation, sine gratings with varying orientation, phase and position are used. The phase and position parameters have to be varied so that at least one appropriate pattern has been presented for each neuron. The orientation preference is then calculated by forming a vector for each measured orientation, whereas the vector angle corresponds to the measured orientation and the vector length corresponds to the highest activity of this neuron for this orientation. Then, the sum of these vectors is calculated, which again results in a vector. The angle of this vector is the resulting orientation preference.

Furthermore, the length of the resulting vector corresponds to the selectivity of neurons, which reveals the ratio of the responses to the preferred vs non-preferred orientation. A highly selective neuron responds much stronger to the preferred orientation than to others. For weakly selective neurons, the difference of the activity values for

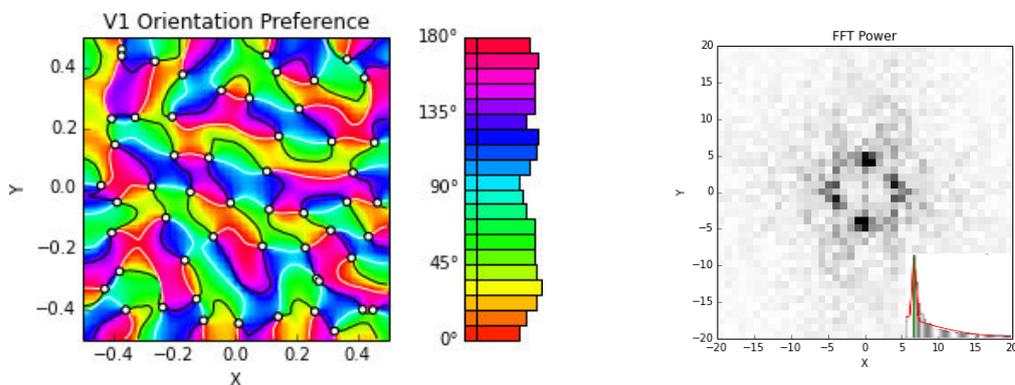
preferred vs non-preferred orientations is smaller.

Here, a series of separate maps is measured. For example, the direction preference is measured independently of the color preference, and so forth. In case of all but the phase disparity map, the corresponding feature as well as the orientation of the sine grating is varied. For the phase disparity map, the orientation of the sine grating is vertical to measure horizontal disparity. It would be desirable to measure maps as a cross product of all features, however this would be very expensive in terms of the time needed.

6.1 Orientation preference

In this section, the resulting orientation map of the basic GCAL model is briefly discussed. For the interested reader, an extensive evaluation of this model can be found in Stevens et al. (2013).

The orientation preference of the orientation map in figure 6.1 varies smoothly across the cortex. As in animals, typical characteristics such as pinwheels, saddles and fractures can be observed. The results are qualitatively the same as in the published GCAL model (Stevens et al., 2013), and an exact replicate can be created if desired (not shown here).



(a) Orientation preference map with overlaid pinwheels. The histogram to the right shows how the orientation preference is distributed, here all orientations are represented nearly uniformly.

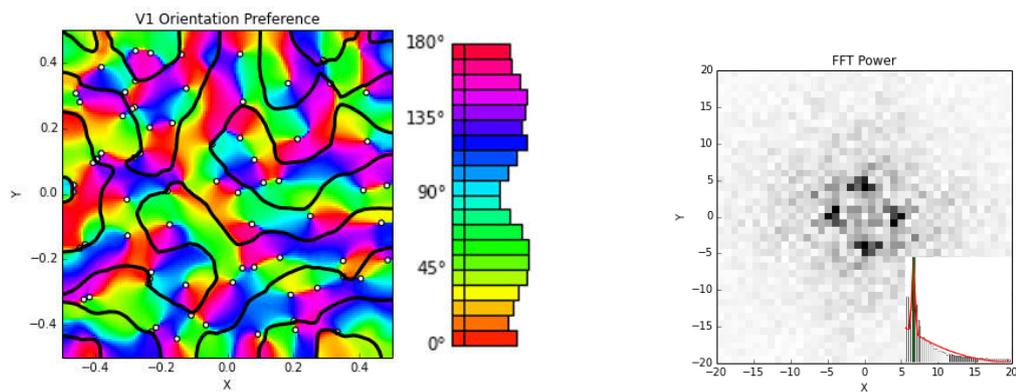
(b) Fourier power spectrum of the orientation map, which is ring-shaped and therefore indicating a high quality map.

Figure 6.1: Orientation map for the basic GCAL model

6.2 Ocular preference

Introducing a second retina and presenting patterns which vary in brightness between the retina results in ocularity maps as shown in figures 6.2 and 6.3. The orientation map is slightly altered, but still of a high quality comparable to those of animals.

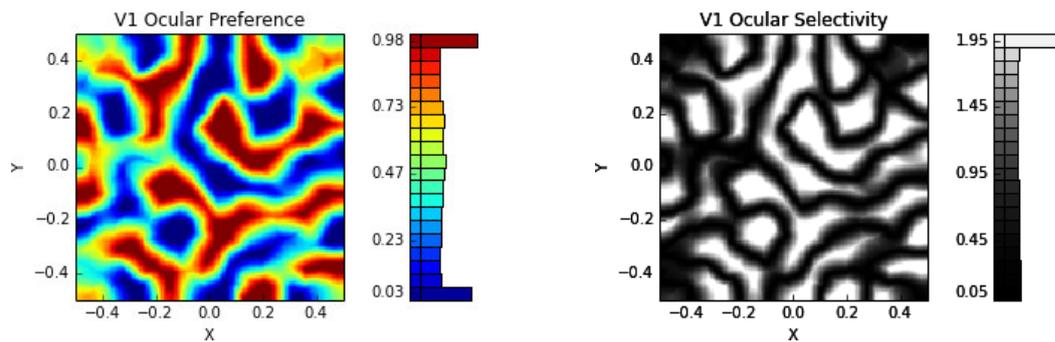
The majority of neurons are selective for one eye or the other, and there are only relatively few binocular neurons. As in animals, most neurons lie in the center of ocular dominance stripes. Some pinwheel neurons are duplicated, i.e. there are several neurons with low orientation selectivity, surrounded by other neurons of all different orientation preferences, very close to each other. These duplications occur almost always near ocular dominance boundaries, which are featured by a low selectivity for orientation. Therefore, these neurons have a low selectivity for both: orientation and ocular dominance. It would be interesting to investigate the properties of these neurons further.



(a) Orientation preference map with overlaid pinwheels and contours visualizing the ocular dominance boundaries. The ocular dominance boundaries tend to cross iso-orientation patches orthogonally, as found in animals. However, some other boundaries seem to follow fractures. Some pinwheels are duplicated, and these interestingly lie at ocular dominance boundaries. Therefore, these pinwheel neurons not only have a weak selectivity for orientation, but also a weak selectivity for ocular dominance.

(b) Fourier power spectrum of the orientation map. The spectrum is still ring-shaped, however when inspecting the histogram one can see that there is a second peak at the left, indicating that there are some pinwheels which are very close together. Therefore, the quality of the orientation map slightly decreased compared to the orientation-only simulation.

Figure 6.2: Orientation map with overlaid ocular dominance boundaries



(a) Ocular preference map. One can see two peaks, one at very low values and another one at very high values. This indicates that most neurons have a preference for one eye or the other.

(b) Ocular selectivity map. When comparing with the ocular preference map, it emerges that neurons which are weakly selective (dark areas) do not have a preference for input from a certain eye.

Figure 6.3: Ocular dominance preference and selectivity

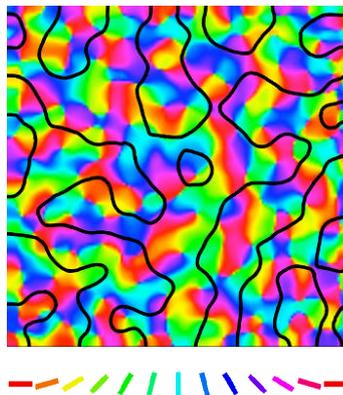
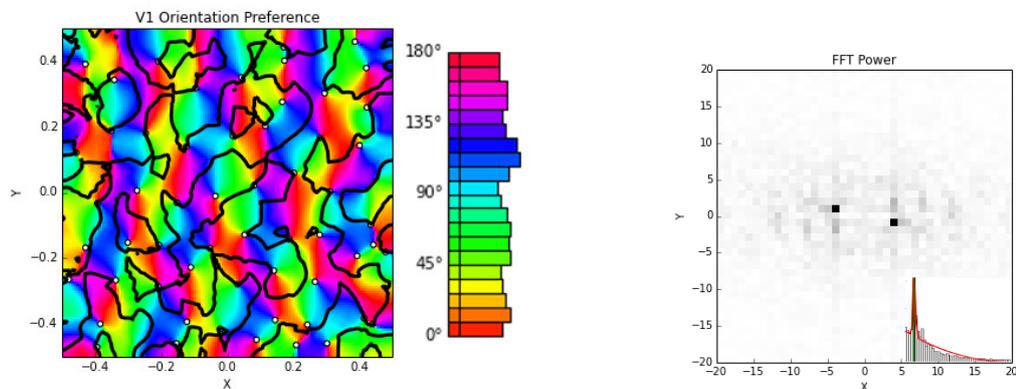


Figure 6.4: Ocular preference map of a LISSOM model for comparison. See also figure 2.3 for an ocular preference of a macaque monkey. Figure reprinted from Miikkulainen et al. (2005)

6.3 Disparity preference

A model driven by input patterns which have a position offset between the eyes does not lead to realistic ocular preference maps, as shown in figure 6.5. The ocular dominance boundaries are very irregular. The same figure shows that the orientation map is altered considerably. Iso-orientation areas are stripy rather than patchy, which results in a non ring-shaped power spectrum.

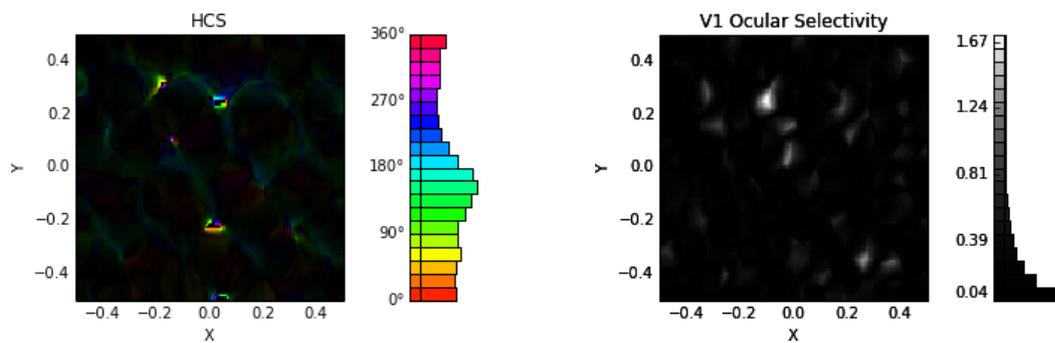
The disparity map itself indicates that most neurons are not selective for disparity. There are small areas of highly selective neurons, and it seems that there is a clear organization of disparity preference in these areas, just like in animals. That most neurons are not selective for disparity could explain the relatively few experimental findings about disparity maps in animals.



(a) Orientation preference map with overlaid pinwheels and contours visualizing the ocular dominance boundaries. The ocular dominance boundaries are very irregular, originating from the low ocular selectivity as printed in figure 6.6. Iso-orientation patches tend to run vertically.

(b) Fourier power spectrum of the orientation map. The spectrum indicates that iso-orientation patches run vertically. Usually, a ring-shaped spectrum elucidates that from any one point, the same orientation preference on average will occur again at a distance regardless of the direction.

Figure 6.5: Orientation map with overlaid ocular dominance boundaries for a disparity simulation



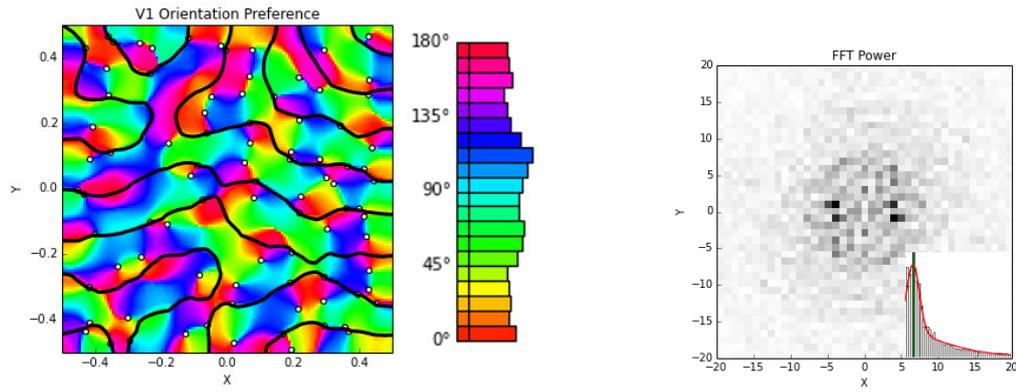
(a) Disparity preference and selectivity map. Most neurons are not selective for disparity, and it seems there is a clear organization of neurons within areas which are selective for disparity.

(b) Ocular selectivity map. Compared to the simulation with brightness differences, it emerges that most neurons are not selective for ocularity. This leads to irregular ocular dominance boundaries, as these are originating mostly from noise rather than actual ocular preferences.

Figure 6.6: Disparity preference and selectivity + ocular selectivity for a disparity simulation

6.4 Ocular and disparity preference

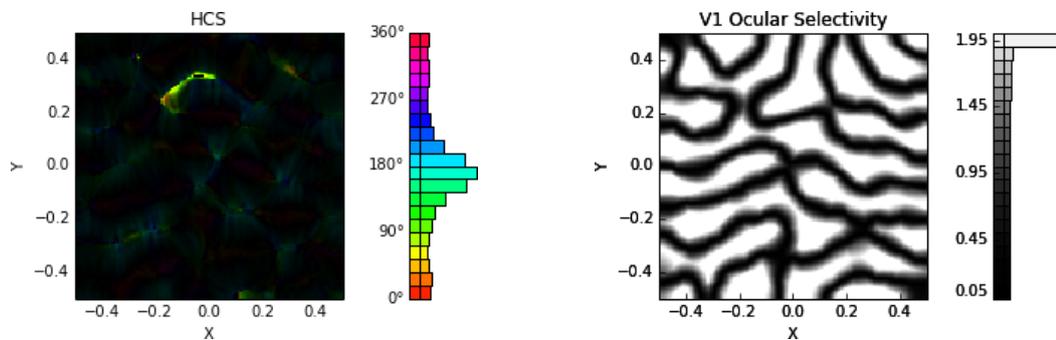
For this simulation, Gaussian patterns with brightness differences as well as position offsets between the eyes were used as training inputs. As for the ocular-only simulation, biologically realistic orientation maps and ocular dominance maps emerge. The ocular dominance boundaries tend to cross orientation patches at a right angle. Again, some duplicate pinwheels near ocular dominance boundaries are present. The disparity map looks qualitatively very similar to the one in the disparity-only simulation. The histogram records a clear preference for disparities around 180 degrees, which is reasonable as the position offset is always on the horizontal axis.



(a) Orientation preference map with overlaid pinwheels and contours visualizing the ocular dominance boundaries. Qualitatively, this map looks very similar to the ocularity-only simulation (see figure 6.2), except that there are now more ocular dominance boundaries that follow orientation borders. Please also see this figure for discussion. Compared to figure 6.5, the ocular dominance boundaries are of regular shape.

(b) Fourier power spectrum of the orientation map. Compared to the spectrum of the disparity-only simulation, the spectrum is clearly ring-shaped, indicating a higher quality orientation map, similar to those found in animals.

Figure 6.7: Orientation map with overlaid ocular dominance boundaries for a brightness difference+disparity simulation



(a) Disparity preference and selectivity map. As in the disparity-only simulation, most neurons are not selective for disparity. Here, one area with especially high disparity selectivity is obvious.

(b) Ocular selectivity map. Most neurons are highly selective for one eye. This again illustrates that brightness differences are driving the development of ocularity maps, rather than disparity in the input.

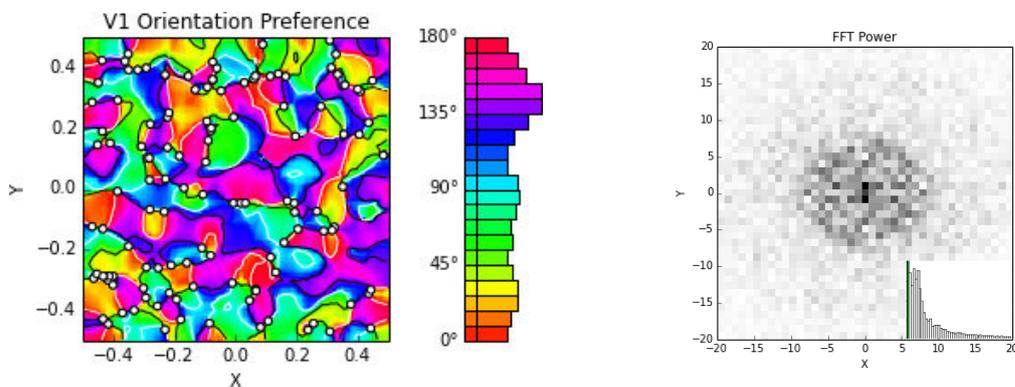
Figure 6.8: Disparity preference and selectivity + ocular selectivity for a disparity simulation

6.5 Direction preference

This section presents the resulting maps of the model in which temporal delays from the LGN sheets to V1 are introduced. As discussed in 5.5, these are especially of interest as this is a novel way of modeling motion selectiveness.

Figure 6.9 shows the resulting orientation map. Compared to the basic GCAL model, the map quality decreased, as there are several areas with a cumulation of nearby pinwheel neurons. The two dark points in the center of the Fourier plot make also reveal the same fact, and there is no clear peak in the Fourier histogram.

The produced direction map is shown in figure 6.12. There, one can see that most neurons are selective for the direction orthogonal to their orientation preference. Furthermore, there are areas where two direction patches are included in one orientation patch. The connection fields of a highly selective neuron are shown in figure 6.10.



(a) Orientation preference map with overlaid pinwheels. One can observe areas with a cumulation of nearby pinwheel neurons, which are a sign of lower map quality.

(b) Fourier power spectrum of the orientation map. The two central dark points again indicate an accumulation of nearby pinwheels. However, there is still an isotropic ring visible.

Figure 6.9: Orientation map for the motion simulation

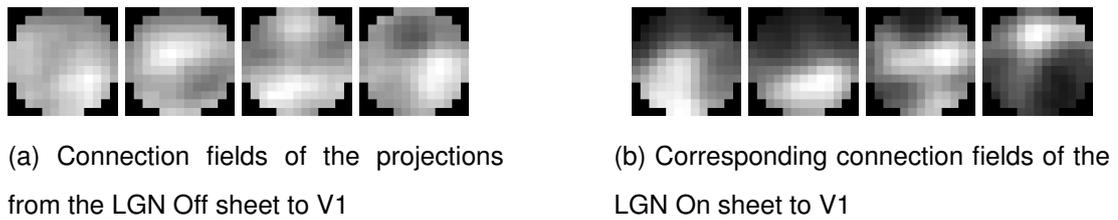


Figure 6.10: Example connection fields of a direction selective neuron. The leftmost projections are for time lag 0, and the time lag is increasing to the right. Here, the neuron is selective for a pattern which is in the bottom left of the connection field at time $t = 0$, and then moving upwards and slightly to the right. Please mind that the majority of neurons does not have such a clear indication of direction preference.

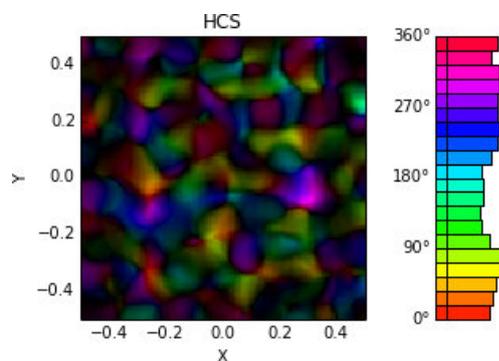


Figure 6.11: Direction preference map. Brighter areas show higher selectiveness for direction. The direction map looks similar to an orientation map, however with smaller patches. This is because orientation is the larger scale organization in this simulation. By increasing the speed of the input patterns, it is theoretically possible that direction becomes the large scale organization (Miikkulainen et al., 2005), however this is contradicting the experimental findings and not shown here.

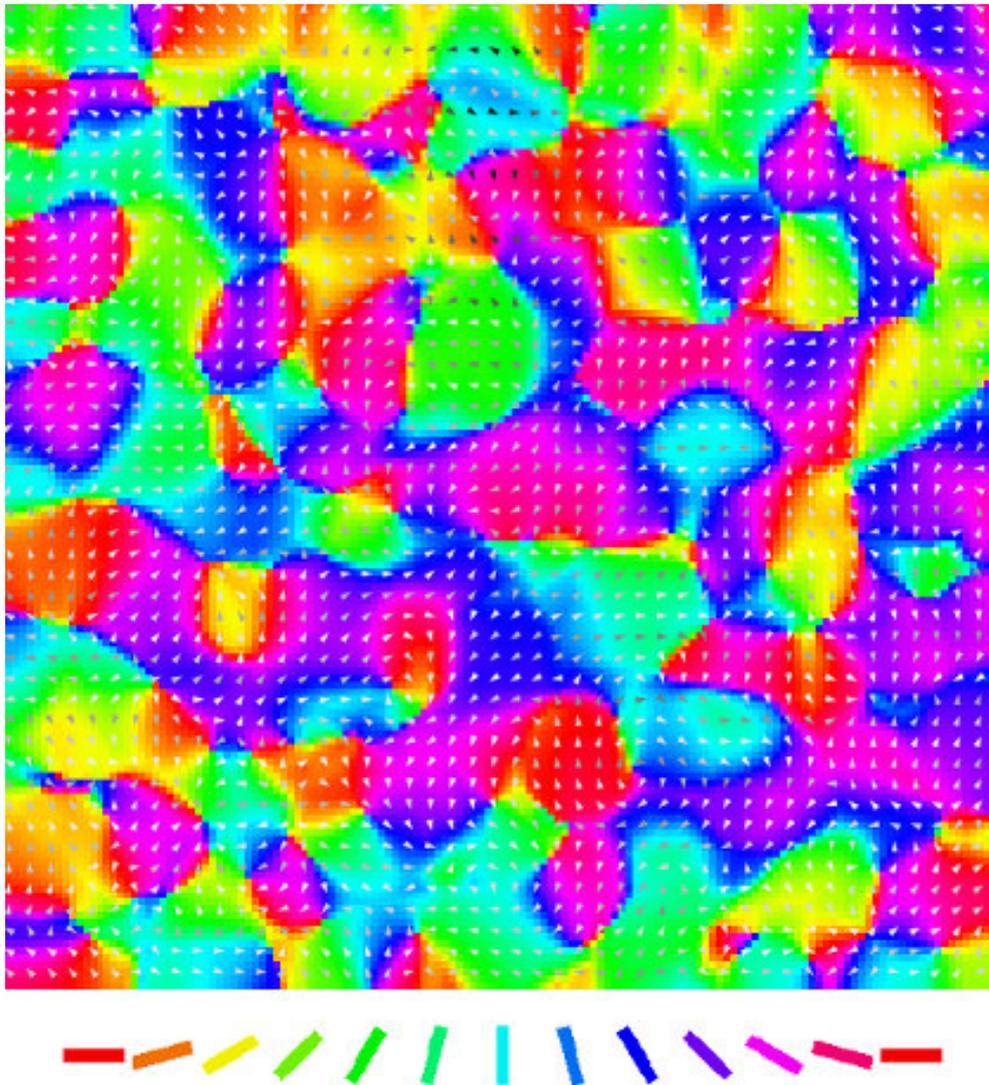
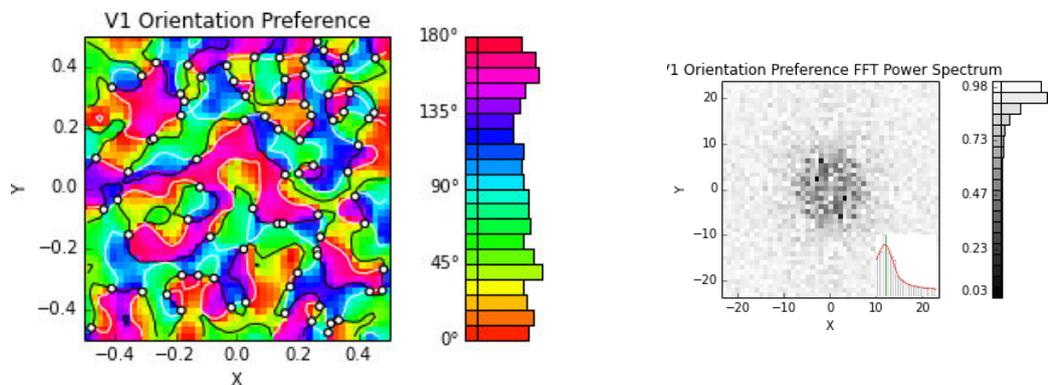


Figure 6.12: Orientation map overlaid with arrows showing the preferred direction. Darker arrows correspond to higher selectivity and vice versa. The color key for orientation preference at the bottom of the figure allows better investigation of the orthogonality between orientation and direction maps. As in animals, some orientation patches contain two direction patches of opposite direction preference. This overlaid map looks similar to an overlaid map in a LISSOM simulation with multiple lagged LGN sheets (figure 5.23 in Miikkulainen et al. (2005)).

6.6 Color preference

The resulting maps of a model using the new color model with a single retina sheet are presented here. Compared to the color map in Bednar et al. (2005), color is the largest scale organization, resulting in preferences of a certain color for nearly all neurons rather than clearly separate color blobs. In the previous model, color selective neurons had only a weak orientation selectivity. Here, all neurons are approximately equally orientation selective (see figure 6.14).

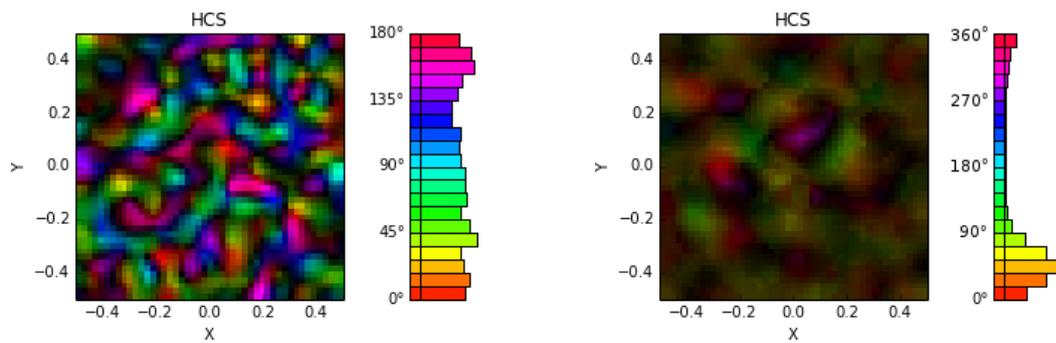
Please accept apologies that in this simulation there is a clear peak for reddish/yellowish colors, this is due to the weighting of the cone types. The color+ocular dominance preference shows a simulation with a better weighting, resulting in a wider range of color preferences.



(a) Orientation preference map with overlaid pinwheels. Compared to the basic GCAL model, the orientation map is slightly distorted, but still of a high quality.

(b) Fourier power spectrum of the orientation map, for comparison with the basic GCAL model.

Figure 6.13: Orientation map for the color simulation



(a) Overlaid preference+selectivity map for orientation. Brighter areas correspond to higher selectiveness. In comparison to previous models, all neurons are approximately equally selective for orientation, i.e. there are no weakly selective areas for orientation due to highly color selective blobs.

(b) Overlaid preference+selectivity map for color. It emerges that all neurons are selective for color, and it seems that color is the largest-scale map organization. This is in contrast to earlier studies, where only relatively few neurons were color selective organized in color blobs.

Figure 6.14: Overlaid preference+selectivity maps for orientation and color in a color simulation

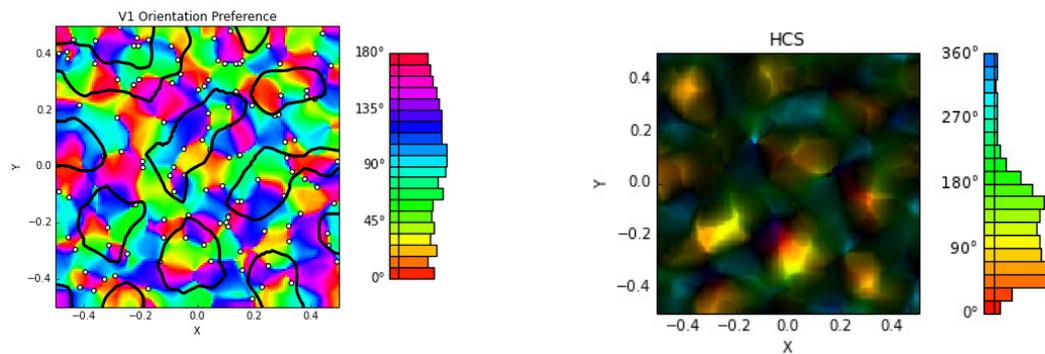
6.7 Color and ocular preference

Here, the resulting maps of a combined model for color and ocular preference are presented. These are an important milestone, as the previous implementation in Topographica based on LISSOM did not produce realistic maps for this feature combination, as discussed in section 5.8.

The orientation map is not altered considerably compared to the individual models. The ring shaped Fourier spectrum indicates a high quality orientation map (not shown here).

As for the color only model, the color map is the largest scale organization. Here, there is a wider range of color preference, due to a different weighting of the cone types. Also, the color map in figure 6.15 looks “sharper” which is due to a higher density of neurons compared to the color-only simulation. Again, please accept apologies for this.

The ocular dominance boundaries look more patchy compared to the stripy boundaries in the individual ocular dominance model. This might be due to the interplay of the ocular dominance and color map. As in animals, most pinwheels lie in the center of ocular dominance patches.



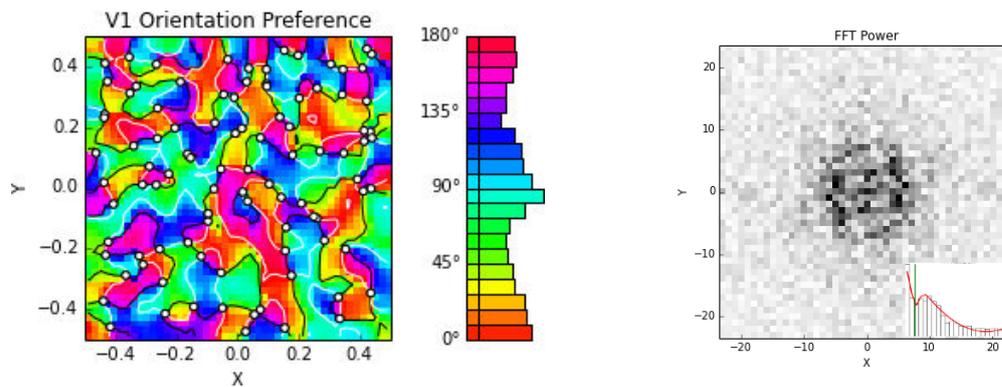
(a) Orientation preference map with overlaid pinwheels and contours visualizing the ocular dominance boundaries. As in the individual ocular dominance simulation, ocular dominance boundaries tend to cross iso-orientation patches orthogonally, but not all of them. Again, some pinwheels are duplicated, however in this simulation not at the ocular dominance boundaries. The ocular dominance boundaries are more patchy compared to the stripy boundaries in the individual simulation.

(b) Overlaid color preference+selectivity map. Compared to the color-only simulation, a wider range of colors is covered (see histograms). Also, the map looks “sharper” which is due to a higher cortical density in this simulation.

Figure 6.15: Orientation and color maps with overlaid ocular dominance boundaries for a color+ocular dominance simulation

6.8 Spatial frequency preference

Extending the basic GCAL model with spatial frequency channels allows neurons in V1 to become selective for various spatial frequencies. As with most of the other presented features, the orientation map in figure 6.16 is slightly distorted. The spatial frequency map in figure 6.17 reveals that neurons become selective for a range of different spatial frequencies. In this simulation, the range is quite small, however if a fourth spatial frequency sheet is added, the range for which neurons become selective gets wider. The connection fields of neurons in V1 show that almost all neurons have higher weights for the spatial frequency channel selective for high frequencies in the input compared to the other spatial frequency channels selective for lower frequencies in the input. This is contrary to the previous implementation in LISSOM (Palmer, 2009), where some neurons became selective for the one channel while other neurons



(a) Orientation preference map with overlaid pinwheels. The orientation map is slightly distorted, with some duplicated pinwheels. As with the other cortical maps where this occurs, this might be due to the interplay of the maps.

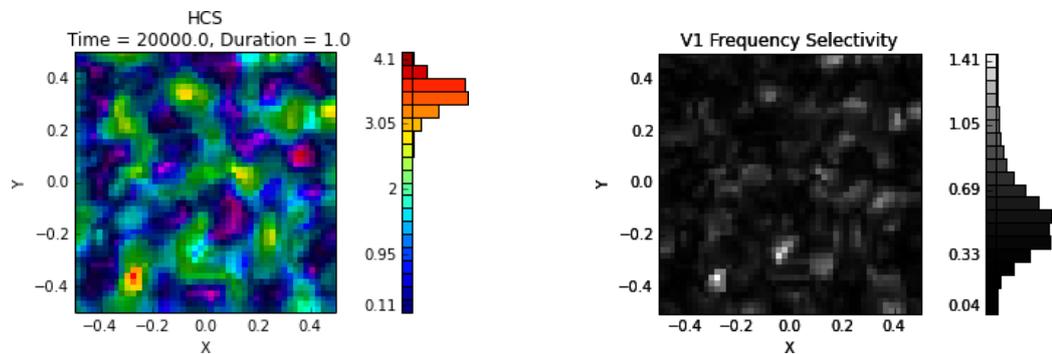
(b) Fourier power spectrum of the orientation map. The isotropic ring is clearly visible, however some pixels within the ring indicate distortions.

Figure 6.16: Orientation map for the spatial frequency simulation

became selective for others.

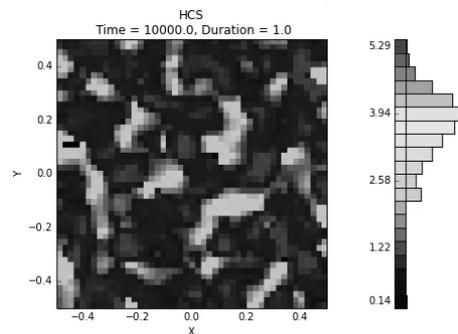
6.9 Combined model

Running the combined model is a difficult undertaking, as there are dozens of parameters which can be varied, and the computational resources required are enormous. On a 24 core Intel Xeon X5650 with 24 gigabyte RAM, simulating the combined model for 15000 training iterations and subsequent map measurement takes approximately 20 hours. At the time of writing, no parameter set was found which allowed the measurement of maps. This is likely due to activation values which are on average either too high or too low. Then, the neurons in V1 do not organize into maps, and the measurement of map fails. Possible ways to resolve this issue are discussed in the next chapter.



(a) Overlaid preference+selectivity map for spatial frequency. One can see that the preference in spatial frequency varies, however the histogram shows that the range is quite narrow.

(b) Spatial frequency selectivity map. Most neurons are somewhat selective for spatial frequency. Some smaller areas appear with a high spatial frequency selectiveness, whereas these prefer very high or very low spatial frequencies (see overlaid map on the left).



(c) Overlaid preference+selectivity map for spatial frequency for a simulation with four spatial frequency channels. In comparison to the simulation with three channels (see figure above), a wider range of spatial frequencies is covered. Accept apologies that the colors do not match between the two figures, instead please compare the histograms so that the differences become visible.

Figure 6.17: Spatial frequency map

Chapter 7

Discussion and Future Work

Building an all maps model of the primary visual cortex is a difficult undertaking. There are many scientific decisions which need to be made, for example, which sheets to connect laterally for gain control. Furthermore, there are many parameters which influence the resulting maps. Here, a discussion of the results of the previous chapter is presented. Then, future research ideas are given. The chapter ends with suggestions for improvements of Topographica, which would allow further investigation of the complex interplay of several maps.

7.1 Discussion of the results

In this section, the results of the individual simulations are discussed further. Starting with the basic GCAL model which can account for orientation preference, the other cortical maps which emerge in individual models are investigated. Then, suggestions to resolve the issue of not being able to measure maps in the all maps model are presented.

7.1.1 Orientation preference

In the simulation with the basic GCAL model, the high quality orientation maps could be reproduced qualitatively. In fact, an exact replicate of the published GCAL model (Stevens et al., 2013) can be created if desired. This is important seeing that the new class-based way of defining the model is used. Although considerably simplifying the model definition and exploration, the instantiated models still have the same properties as with the traditional way of defining models.

7.1.2 Ocular dominance preference

When adding a second retina sheet to the simulation, and changing the input patterns so they vary in their brightness between the retinas, biologically realistic ocular dominance maps emerge. Pinwheels are usually found in the center of ocular dominance stripes. Some duplicated pinwheels, which indicate small areas of low orientation selectiveness, occur near ocular dominance boundaries. The resulting maps are comparable to those in Miikkulainen et al. (2005).

7.1.3 Disparity preference

As found in previous studies, position offsets cannot account for ocular dominance stripes in simulations using LISSOM/GCAL (Miikkulainen et al., 2005). All neurons are binocular, i.e. not preferring input from one eye over the other. The disparity map shows small areas of high disparity selectiveness, which could explain why so far no large-scale disparity map was measured in animals (Kara and Boyd, 2009).

7.1.4 Ocular and disparity preference

Using inputs which have varying brightness and a position offset between eyes, biologically realistic ocular dominance and disparity maps emerge. The ocular dominance stripes tend to cross iso-orientation patches orthogonally. In this simulation, there is a clear peak for disparities around 180 degrees, which reflects the input statistics because the position offset of the patterns is in x direction.

7.1.5 Direction preference

The direction map emerging of the new motion model with multiple time delayed projections from the LGN sheets to V1 shows biologically realistic features. As in animals, most neurons are selective for the direction orthogonal to their orientation preference. Also, one iso-orientation patch often includes two smaller iso-direction patches with opposite preference for direction. This demonstrates that multiple lagged projections from the LGN level to V1 can account for biologically realistic direction maps. Future work could investigate which parameters could improve the orientation map quality. Here, the learning rate of the afferent connections to neurons in V1 was decreased, as there are more incoming connections. One could also argue that the

strength of these afferent connection should be decreased because of the same argument.

7.1.6 Color preference

The parameters of the color model in use need further investigation. Previous models have shown a clear separation of color blobs and orientation selective regions. With the new color model, there are no separate areas anymore. Even if the ratio of the input of color channels and luminosity channels to V1 is decreased, color is still the largest scale organization. Increasing the strength of the luminosity channel leads to less color selective neurons at an equally large scale. This might be due to the homeostasis process in V1 of GCAL (Ball, 2014).

7.1.7 Color and ocular dominance preference

The combination of the ocular dominance model with the color model is one of the many combinations one can model. It was shown that brightness differences introduced to natural images by converting them into HSV space and changing the value of pixels result in similar results to the individual models. Furthermore, a better choice of the cone scaling led to a hue map which shows a wider range of color preferences. This points out how small changes in the parameter space can have a quite considerable effect on the resulting maps.

7.1.8 Spatial frequency preference

The spatial frequency simulation has shown that neurons become selective for a range of spatial frequencies. However, compared to the previous work in Palmer (2009), the range is much smaller. This is due to the selectiveness of all neurons for one out of the three spatial frequency channels. The existing variety in frequency preferences emerge only because of a slightly higher weight to the other two spatial frequency channels of some neurons.

7.1.9 Combined model

The combined model covering all known spatial maps has dozens of parameters to control. Some of them are related to the input pattern generation, such as the speed

for the sweeping of patterns or the weight of the individual color channels. Other parameters can control the architecture of the model, for example the amount of spatial frequency channels or how many projections with temporal delays should be present for modeling motion. Yet others are controlling the strengths and weights for the projections between the sheets. It is expected that tweaking these parameters is necessary so that the all maps model results in useful maps.

In the trials to run the all maps model, it was not possible to generate any maps at all. This is likely due to an on average too high or too low activation of neurons in V1, which then in turn do not respond appropriately to the measurement patterns. Therefore, a parameter search over different strengths from the retina sheets to the LGN sheets, as well as for the strengths from the LGN sheets to V1, is recommended. The two parameters should be changed independently of each other.

It is expected that once a suitable parameter set is found, the all maps model will result in maps for all feature dimensions. However, in this project a parameter search was not possible due to the high computational resources which are needed; one simulation with 15000 training iterations takes about 20 hours on a fairly fast machine. It appears that using a cluster is unavoidable to run multiple simulations.

The resulting maps are likely to be distorted compared to the individual models, as a high degree of interdependence is expected. However, such maps will reveal open issues with the GCAL model, and in turn can provide hints where improvements might be necessary.

7.1.10 Summary

The resulting maps of the individual models have shown for the first time that GCAL can account for all known cortical maps in separate simulations. Previously, this was shown for LISSOM, the predecessor of GCAL. Building an all maps model is still an open issue, which is complex because of the interplay of all the different features and their corresponding model properties as well as the vast number of parameters. Compared to the previous LISSOM all maps model, it was shown that the combination of color maps and ocular dominance maps can be successfully established.

7.2 Future work scientifically

In this section, a variety of possible future research topics is presented. These are only a small subset of the many topics that can still be investigated using computational models of the visual cortex, with a focus on topics which require nearly all/all maps to be simulated simultaneously.

7.2.1 Interdependence of maps

A combined model of all known cortical maps allows the examination of various previous proposals. The investigation of the interdependence of various feature maps could be driven further by creating more models with various feature combinations. The ocular dominance+color model is a step in the right direction, but the full product of all feature combinations should be simulated, to see in detail which dimension affects which other dimension. The proposal of a very clear and significant interdependence of maps in Yu et al. (2005) could be validated using this technique.

7.2.2 Uniform coverage

One research goal is to understand the functionality of cortical maps. It has been suggested that cortical maps might be optimized for completeness and continuity, which are competing principles (see section 3.4). An investigation of completeness and continuity of an all maps model is desirable, and could give an answer to the controversial question whether cortical maps are optimized for uniform coverage.

7.2.3 Relative order of map development

Also, by varying inputs within a single simulation, one could test whether the relative order of map development corresponds to different map properties, as suggested in Goodhill and Cimponeriu (2000). As a starting point, two models of the same structure as for an ocular dominance simulation could be created. Then, in one of the models the first 10000 simulations are run with patterns of equal brightness but varying orientation. After 10000 iterations, brightness differences are supplemented. In this model, the orientation map develops before the ocular dominance map. For the second model, this procedure could be reverted, by showing Gaussians of only one orientation but varying brightness in the first 10000 iterations.

7.2.4 Gaussian stimuli for color and spatial frequency

For the color model, it would be desirable to be able to use colorized Gaussian inputs. These could be controlled in detail by the modeler, which is not the case at all for natural images. So far, there has been no way found to imitate the complex relationship between foreground and background found in natural images. In natural images, hue usually changes slowly, and often brightness changes without changes in the hue value. A starting point could be a Gaussian pattern of a random hue, set on a background of the same hue, but different brightness. However, a tentative draft of such a pattern did not lead to a meaningful color map so far.

Similarly, randomly varying the size of Gaussian patterns to simulate varying spatial frequency (as described in section 4.1) did not lead to a meaningful spatial frequency map. Compared to colorized Gaussian inputs, resolving this issue should not be as hard.

7.2.5 Lateral connections between spatial frequency channels

Also related to models including spatial frequency are the extents of the lateral connection between spatial frequency channels. Usually, the extent of the lateral connection is (approximately) as large as the receptive field of a neuron in the corresponding LGN sheet. However, the sizes of the receptive fields depends on the spatial frequency channel. Therefore, when connecting sheets across spatial frequency channels, it is not clear whether the size of the lateral connection should be as large as the receptive field of the incoming sheet, or the outgoing sheet. One could also imagine some ratio of incoming to outgoing channel. Each extend is expected to lead to different characteristics, which need to be discovered in future research.

7.3 Future work for Topographica

Topographica is a powerful simulator for computational modeling neural maps, and includes many helpful and advanced functions. However, there are still some open issues, especially when multiple maps are investigated. Here, some suggestions for future functions which would improve Topographica are presented.

7.3.1 Improved visualization

From a software point of view, the extension of the library “dataviews” which is used by Topographica to be able to overlay different cortical maps is desirable for better visualization of the interplay of maps. Thus far, it is only possible to overlay boundaries and points on a map, as done when overlaying the ocular dominance boundaries on the orientation map, and additionally inserting points indicating pinwheel neurons. One could imagine arrows as in the orientation map overlaid by direction arrows, which was done manually for this thesis. The length could not only have varying brightness to visualize selectiveness, but also varying length. This would allow more dimensions to be visualized within a single graphic.

The same library could also be extended in a way that the angles of crossings between the maps could be easily calculated. This would allow comparison to other models and experiments with animals where this property is often stated. This is more reliable than a visual inspection, and one could use them as statistical measurement to investigate the interdependence of maps.

7.3.2 Stereo camera input streams

The investigation of models which are created using a stereo camera as the input source is another interesting research idea. A stereo camera could provide input similar to those of natural images, but successive images could potentially account for motion. The input could be used directly, but possibly one wishes to introduce brightness differences so a model with realistic ocular dominance maps could emerge.

7.3.3 Component based submodels

Another important improvement would be changing the class-based submodels system to be component based rather than inheritance based, as discussed in section 4.4. This will be necessary to build models which include higher visual areas and feedback connections between the cortical areas.

Chapter 8

Conclusions

Neurons in the primary visual cortex of mammals show preferences for a variety of features, such as orientation and ocular dominance. These preferences can be visualized in cortical maps, where a regular organization of the neurons becomes visible. For example, nearby neurons in V1 have preferences for similar orientations. There is evidence that the map structure depends on the input characteristics, and therefore each feature results in a distinct map.

It is thought that cortical maps emerge due to the fact that a high dimensional input space is mapped on a two dimensional set of neurons. Then, an interdependence of all feature preferences seems unavoidable. One highly controversial suggestion is that cortical maps are organized for uniform coverage, i.e. there is a trade-off between continuity and completeness (Swindale et al., 2000).

In general, the function of cortical maps is still an open subject. It was found that there is no orientation map in rats, although neurons were found to be highly selective for orientation (Ohki et al., 2005). Also, it is unclear whether there are cortical maps in higher visual areas, where the presentation of more complex visual inputs is needed in order to get reliable responses of neurons.

In this thesis, three families of models resulting in biologically realistic cortical maps have been reviewed. Nearly all of them model only a limited, small subset of features, with the exception of Gerasymova (2009), where the predecessor of GCAL was used to implement an all maps model. However, the combination of color and ocular dominance was known not to work, and the resulting spatial frequency map was not of a high quality either. More importantly, the model by Gerasymova (2009) is highly inflexible and difficult to use.

In this thesis, a superior way of defining models was presented, where it is more

obvious how large amounts of sheets and projections can be defined, especially for complex models such as the combined all maps model. Rather than manually connecting sheets with deeply nested loops, match conditions were introduced allowing the fully automatic wiring up of sheets. Furthermore, the inspection of model properties was improved significantly, and is now possible before the (time and memory consuming) model instantiation. At the time of writing, several other users had already migrated to this new system, illustrating a considerable improvement in the software design.

The basic GCAL model with neurons selective for orientation was rebuilt using the novel way of defining models. Then, a series of models with feature preferences for orientation and one of the other input dimensions was presented. There, it emerged that the principles of GCAL can account for all other feature maps. A novel way of modeling motion using temporal delays from the LGN sheets to V1 was used, and it was shown that this results in direction maps similar to previous proposals and animal studies.

Following this, a model which accounts for orientation, color and ocular dominance preference was described, showing how multiple models can be joined to build a more complex model. This feature combination was chosen because it was known not to work in the previous attempt of building an all maps model. The combined model resulted in a set of cortical maps comparable to those of animals.

The idea of combining models was taken further and a model which in principle can account for all cortical maps simultaneously was drafted. However, the complexity of this model and the corresponding parameters which have to be carefully tuned did not allow the running of this model in a way that it results in cortical maps at this point of time.

In summary, this thesis resulted in major improvements of the modeling software *Topographica*, allowing modelers to focus on scientific questions rather than implementing code. The pattern coordinator allows the use of predefined input patterns, and the class-based submodels system improves the work-flow in various ways. A novel way of modeling motion with temporal delays from LGN sheets to V1 was defined, which is biologically plausible, in contrast to the previous implementation in *Topographica* with multiple lagged retina sheets. An all maps model was proposed, but the time-consuming investigation of the complex interplay of parameters prohibited using this model to produce cortical maps. This is not thought to be a failure of the underlying principles of GCAL, but an issue in finding a parameter set which drives the activation of neurons in V1 in a suitable way. Investigating the influence of parameters on this complex model is expected to become a research project on its own.

Bibliography

- Ackman, J. B. and Crair, M. C. (2014). Role of emergent neural activity in visual map development. *Current Opinion in Neurobiology*, 24(1):166–75.
- Adler, F. H., Kaufman, P. L., Levin, L. A., and Alm, A. (2011). *Adler's Physiology of the Eye*. Elsevier Health Sciences.
- Ball, C. E. (2014). *Modeling the development of primate color vision*. Manuscript in preparation, The University of Edinburgh.
- Barlow, H. B., Blakemore, C., and Pettigrew, J. D. (1967). The neural mechanism of binocular depth discrimination. *The Journal of Physiology*, 193(2):327–42.
- Bear, M. F., Connors, B., and Paradiso, M. (2006). *Neuroscience: Exploring the Brain*. Lippincott Williams & Wilkins, 3rd edition.
- Bednar, J. A. (2009). Topographica: building and analyzing map-level simulations from Python, C/C++, MATLAB, NEST, or NEURON components. *Frontiers in Neuroinformatics*, 3:8.
- Bednar, J. A. (2012). Building a Mechanistic Model of the Development and Function of the Primary Visual Cortex. *Journal of Physiology*, 106:194–211.
- Bednar, J. A. and Miikkulainen, R. (2006). Joint Maps for Orientation, Eye, and Direction Preference in a Self-Organizing Model of V1. *Neurocomputing*, 69:1272–1276.
- Bednar, J. A., Paula, J. B. D., and Miikkulainen, R. (2005). Self-organization of color opponent receptive fields and laterally connected orientation maps. *Neurocomputing*, 65–66:69–76.
- Blasdel, G. G. (1992a). Differential Imaging of Ocular Dominance and Orientation Selectivity in Monkey Striate Cortex. *The Journal of Neuroscience*, 12(8):3115–38.

- Blasdel, G. G. (1992b). Orientation Selectivity, Preference, and Continuity in Monkey Striate Cortex. *The Journal of Neuroscience*, 12(8):3139–3161.
- Blasdel, G. G., Obermayer, K., and Kiorpes, L. (1995). Organization of Ocular Dominance and Orientation Columns in the Striate Cortex of Neonatal Macaque Monkeys. *Visual Neuroscience*, 12(3):589–603.
- Blasdel, G. G. and Salama, G. (1986). Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321(6070):579–85.
- Bosking, W. H., Zhang, Y., Schofield, B. R., and Fitzpatrick, D. (1997). Orientation Selectivity and the Arrangement of Horizontal Connections in Tree Shrew Striate Cortex. *The Journal of Neuroscience*, 17(6):2112–2127.
- Burger, T. and Lang, E. W. (1999). An Incremental Hebbian Learning Model of the Primary Visual Cortex with Lateral Plasticity and Real Input Patterns. *Zeitschrift für Naturforschung C — A Journal of Biosciences*, 54:128–140.
- Burger, T. and Lang, E. W. (2001). Self-Organization of Local Cortical Circuits and Cortical Orientation Maps: A Nonlinear Hebbian Model of the Visual Cortex with Adaptive Lateral Couplings. *Zeitschrift für Naturforschung C — A Journal of Biosciences*, 56:464–478.
- Carreira-Perpiñán, M. a. and Goodhill, G. J. (2002). Are Visual Cortex Maps Optimized for Coverage? *Neural Computation*, 14(7):1545–60.
- Carreira-Perpiñán, M. A., Lister, R. J., and Goodhill, G. J. (2005). A Computational Model for the Development of Multiple Maps in Primary Visual Cortex. *Cerebral Cortex*, 15(8):1222–1233.
- Catania, K. C., Lyon, D. C., Mock, O. B., and Kaas, J. H. (1999). Cortical organization in shrews: evidence from five species. *The Journal of Comparative Neurology*, 410(1):55–72.
- Connolly, M. and Van Essen, D. (1984). The Representation of the Visual Field in Parvicellular and Magnocellular Layers of the Lateral Geniculate Nucleus in the Macaque Monkey. *The Journal of Comparative Neurology*, 226(4):544–64.
- Crair, M. C., Gillespie, D. C., and Stryker, M. P. (1998). The Role of Visual Experience in the Development of Columns in Cat Visual Cortex. *Science*, 279:566–570.

- Crair, M. C., Ruthazer, E. S., Gillespie, D. C., and Stryker, M. P. (1997). Relationship between the Ocular Dominance and Orientation Maps in Visual Cortex of Monocularly Deprived Cats. *Neuron*, 19(2):307–18.
- Crick, F. H. C. and Asanuma, C. (1986). Certain Aspects of the Anatomy and Physiology of the Cerebral Cortex. In McClelland, J. L., Rumelhart, D. E., and PDP Research Group, C., editors, *Parallel Distributed Processing*, pages 333–371. MIT Press, Cambridge, MA, USA.
- Dacey, D. M. (2000). Parallel Pathways for Spectral Coding in Primate Retina. *Annual Review of Neuroscience*, 23:743–775.
- Denk, W., Strickler, J. H., and Webb, W. W. (1990). Two-Photon Laser Scanning Fluorescence Microscopy. *Science*, 248(4951):73–76.
- Espinoza, S. G. and Thomas, H. C. (1983). Retinotopic organization of striate and extrastriate visual cortex in the hooded rat. *Brain Research*, 272(1):137–44.
- Famiglietti, E. and Kolb, H. (1976). Structural basis for ON-and OFF-center responses in retinal ganglion cells. *Science*, 194(4261):193–195.
- Földiák, P. (2002). Sparse coding in the primate cortex. In Arbib, M. A., editor, *The Handbook of Brain Theory and Neural Networks*, pages 1064–1068. MIT Press, 2nd edition.
- Gerasymova, K. (2009). *Computational Model of All Known Feature Maps in the Primary Visual Cortex*. Student research project, Humboldt-Universität zu Berlin, Germany.
- Gilbert, C. D., Hirsch, J. A., and Wiesel, T. N. (1990). Lateral Interactions in Visual Cortex. In *The Brain*, volume LV of *Cold Spring Harbor Symposia on Quantitative Biology*, pages 663–677, Cold Spring Harbor, NY. Cold Spring Harbor Laboratory Press.
- Goodhill, G. J. (2007). Contributions of Theoretical Modeling to the Understanding of Neural Map Development. *Neuron*, 56(2):301–311.
- Goodhill, G. J. and Cimponeriu, A. (2000). Analysis of the Elastic Net Model Applied to the Formation of Ocular Dominance and Orientation Columns. *Network: Computation in Neural Systems*, 11:153–168.

- Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory*. Wiley, Hoboken, NJ.
- Holmes, G. (1918). Disturbances of Visual Orientation. *British Journal of Ophthalmology*, 2(9):449–468.
- Horton, J. C. and Hocking, D. R. (1996). An Adult-like Pattern of Ocular Dominance Columns in Striate Cortex of Newborn Monkeys Prior to Visual Experience. *The Journal of Neuroscience*, 16(5):1791–1807.
- Hubel, D. H. and Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology*, 148:574–591.
- Hubel, D. H. and Wiesel, T. N. (1968). Receptive Fields and Functional Architecture of Monkey Striate Cortex. *The Journal of Physiology*, 195:215–243.
- Hubel, D. H., Wiesel, T. N., and LeVay, S. (1977). Plasticity of Ocular Dominance Columns in Monkey Striate Cortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 278:377–409.
- Huberman, A. D., Speer, C. M., and Chapman, B. (2006). Spontaneous Retinal Activity Mediates Development of Ocular Dominance Columns and Binocular Receptive Fields in V1. *Neuron*, 52(2):247–254.
- Hurvich, L. M. and Jameson, D. (1957). An opponent-process theory of color vision. *Psychological Review*, 64(6, Pt.1):384–404.
- Hyvärinen, A. and Hoyer, P. O. (2001). A Two-Layer Sparse Coding Model Learns Simple and Complex Cell Receptive Fields and Topography From Natural Images. *Vision Research*, 41(18):2413–2423.
- Hyvärinen, A., Hurri, J., and Hoyer, P. O. (2009). *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision*. Springer Publishing Company, Incorporated, 1st edition.
- Issa, N. P., Trepel, C., and Stryker, M. P. (2001). Spatial Frequency Maps in Cat Visual Cortex. *The Journal of Neuroscience*, 20:8504–8514.
- Joblove, G. H. and Greenberg, D. (1978). Color spaces for computer graphics. *ACM SIGGRAPH Computer Graphics*, 12(3):20–25.

- Kaas, J. (2001). The organization of sensory cortex. *Current Opinion in Neurobiology*, 11(4):498–504.
- Kara, P. and Boyd, J. D. (2009). A micro-architecture for binocular disparity and ocular dominance in visual cortex. *Nature*, 458(7238):627–31.
- Kneisel, Y. (2013). *Feasibility of Random Wiring in Primate Colour Vision*. Masters thesis, The University of Edinburgh, UK.
- Kohonen, T. (1982). Self-Organized Formation of Topologically Correct Feature Maps. *Biological Cybernetics*, 43:59–69.
- Kohonen, T. (1990). The Self-Organizing Map. *Proceedings of the IEEE*, 78(9):1464–1480.
- Kohonen, T. (2001). *Self-Organizing Maps*. Springer, Berlin, 3rd edition.
- Landisman, C. E. and Ts'o, D. Y. (2002). Color Processing in Macaque Striate Cortex: Relationships to Ocular Dominance, Cytochrome Oxidase, and Orientation. *Journal of Neurophysiology*, 87(6):3126–3137.
- Liu, W., Tao, D., Cheng, J., and Tang, Y. (2014). Multiview Hessian discriminative sparse coding for image annotation. *Computer Vision and Image Understanding*, 118:50–60.
- Löwel, S., Bischof, H.-J. J., Leutenecker, B., and Singer, W. (1988). Topographic relations between ocular dominance and orientation columns in the cat striate cortex. *Experimental Brain Research*, 71(1):33–46.
- Miikkulainen, R., Bednar, J. A., Choe, Y., and Sirosh, J. (2005). *Computational Maps in the Visual Cortex*. Springer-Verlag.
- Milner, A. D. and Goodale, M. A. (1998). *The Visual Brain in Action*. Number 27. Oxford University Press, USA.
- Mishkin, M., Ungerleider, L. G., and Macko, K. A. (1983). Object vision and spatial vision: two cortical pathways. *Trends in Neurosciences*, 6:414–417.
- Müller, T., Stetter, M., Hubener, M., Sengpiel, F., Bonhoeffer, T., Gödecke, I., Chapman, B., Löwel, S., and Obermayer, K. (2000). An Analysis of Orientation and Ocular Dominance Patterns in the Visual Cortex of Cats and Ferrets. *Neural Computation*, 12(11):2573–2595.

- Nauhaus, I., Nielsen, K. J., Disney, A. a., and Callaway, E. M. (2012). Orthogonal micro-organization of orientation and spatial frequency in primate primary visual cortex. *Nature neuroscience*, 15(12):1683–90.
- Ohki, K., Chung, S., Ch'ng, Y. H., Kara, P., and Reid, R. C. (2005). Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature*, 433(7026):597–603.
- Olmos, A. and Kingdom, F. A. A. (2004). A biologically inspired algorithm for the recovery of shading and reflectance images. *Perception*, 33(12):1463–73.
- Olshausen, B. a. and Field, D. J. (2004). Sparse coding of sensory inputs. *Current Opinion in Neurobiology*, 14(4):481–7.
- Palmer, C. M. (2009). *Topographic and Laminar Models for the Development and Organisation of Spatial Frequency and Orientation in V1*. PhD thesis, The University of Edinburgh.
- Paula, J. B. D. (2007). *Modeling the self-organization of color selectivity in the visual cortex*. PhD thesis, The University of Texas at Austin.
- Pitkow, X. and Meister, M. (2012). Decorrelation and efficient coding by retinal ganglion cells. *Nature neuroscience*, 15(4):628–35.
- Poggio, G. F. and Fischer, B. (1977). Binocular Interaction and Depth Sensitivity in Striate and Prestriate Cortex of Behaving Rhesus Monkey. *Journal of Neurophysiology*, 40(6):1392–405.
- Rao, S. C., Toth, L. J., and Sur, M. (1997). Optically Imaged Maps of Orientation Preference in Primary Visual Cortex of Cats and Ferrets. *The Journal of Comparative Neurology*, 387(3):358–370.
- Reh, T. A. and Constantine-Paton, M. (1985). Eye-specific Segregation Requires Neural Activity in Three-Eyed *Rana pipiens*. *The Journal of Neuroscience*, 5(5):1132–43.
- Ribot, J., Aushana, Y., Bui-Quoc, E., and Milleret, C. (2013). Organization and Origin of Spatial Frequency Maps in Cat Visual Cortex. *The Journal of Neuroscience*, 33(33):13326–43.

- Ritter, H. (1991). Asymptotic Level Density for a Class of Vector Quantization Processes. *IEEE Transactions on Neural Networks*, 2:173–175.
- Rochester, N., Holland, J. H., Haibt, L. H., and Duda, W. L. (1956). Tests on a Cell Assembly Theory of the Action of the Brain, Using a Large Digital Computer. *IRE Transactions on Information Theory*, 2:80–93.
- Rolls, E. T. and Tovee, M. J. (1995). Sparseness of the Neuronal Representation of Stimuli in the Primate Temporal Visual Cortex. *Journal of Neurophysiology*, 73(2):713–26.
- Saul, A. B. and Humphrey, A. L. (1990). Spatial and Temporal Response Properties of Lagged and Nonlagged Cells in Cat Lateral Geniculate Nucleus. *Journal of Neurophysiology*, 64(1):206–24.
- Schein, S. J. (1988). Anatomy of macaque fovea and spatial densities of neurons in foveal representation. *The Journal of Comparative Neurology*, 269(4):479–505.
- Sharma, J., Angelucci, A., and Sur, M. (2000). Induction of Visual Orientation Modules in Auditory Cortex. *Nature*, 404:841–847.
- Shatz, C. J. and Stryker, M. P. (1978). Ocular Dominance in Layer IV of the Cat's Visual Cortex and the Effects of Monocular Deprivation. *The Journal of Physiology*, 281:267–283.
- Sirosh, J., Miikkulainen, R., and Bednar, J. A. (1996). Self-Organization of Orientation Maps, Lateral Connections, and Dynamic Receptive Fields in the Primary Visual Cortex. In Sirosh, J., Miikkulainen, R., and Choe, Y., editors, *Lateral Interactions in the Cortex: Structure and Function*. The UTCS Neural Networks Research Group, Austin, TX.
- Sirovich, L. and Uglesich, R. (2004). The organization of orientation and spatial frequency in primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 101(48):16941–6.
- Solomon, S. G. and Lennie, P. (2005). Chromatic Gain Controls in Visual Cortical Neurons. *The Journal of Neuroscience*, 25(19):4779–92.
- Spigler, G. (2014). *Neural modelling of the McCollough Effect in color vision*. Masters thesis (in print), The University of Edinburgh.

- Stevens, J.-L. R., Law, J. S., Antolik, J., and Bednar, J. A. (2013). Mechanisms for Stable, Robust, and Adaptive Development of Orientation Maps in the Primary Visual Cortex. *The Journal of Neuroscience*, 33:15747–15766.
- Swindale, N. V. (2000). How Many Maps are there in Visual Cortex? *Cerebral Cortex*, 10(7):633–43.
- Swindale, N. V. (2004). How different feature spaces may be represented in cortical maps. *Network: Computation in Neural Systems*, 15(4):217–242.
- Swindale, N. V. (2006). Cerebral Cortex: The Singular Precision of Visual Cortex Maps. *Current Biology*, 16(23):R991–R994.
- Swindale, N. V., D, S., A, G., T, B., and M, H. (2002). Reply to Carreira-Perpiñán and Goodhill. *Neural Computation*, 14(9):2053–2056.
- Swindale, N. V., Shoham, D., Grinvald, A., Bonhoeffer, T., and Hubener, M. (2000). Visual Cortex Maps Are Optimized for Uniform Coverage. *Nature neuroscience*, 3(8):822–826.
- Tani, T., Ribot, J., O’Hashi, K., and Tanaka, S. (2012). Parallel development of orientation maps and spatial frequency selectivity in cat visual cortex. *The European Journal of Neuroscience*, 35(1):44–55.
- Tessier-Lavigne, M. and Goodman, C. S. (1996). The Molecular Biology of Axon Guidance. *Science*, 274(5290):1123–1133.
- Tusa, R. J. and Palmer, L. a. (1980). Retinotopic Organization of Areas 20 and 21 in the Cat. *The Journal of Comparative Neurology*, 193(1):147–64.
- Tusa, R. J., Palmer, L. a., and Rosenquist, a. C. (1978). The Retinotopic Organization of area 17 (Striate Cortex) in the Cat. *The Journal of Comparative Neurology*, 177(2):213–35.
- Udin, S. B. and Fawcett, J. W. (1988). Formation of topographic maps. *Annual Review of Neuroscience*, 11:289–327.
- Van Essen, D. C., Newsome, W. T., and Maunsell, J. H. (1984). The visual field representation in striate cortex of the macaque monkey: Asymmetries, anisotropies, and individual variability. *Vision Research*, 24(5):429–448.

- van Ooyen, A. (2011). Using theoretical models to analyse neural development. *Nature reviews. Neuroscience*, 12(6):311–26.
- Wässle, H., Grünert, U., Chun, M. H., and Boycott, B. B. (1995). The Rod Pathway of the Macaque Monkey Retina: Identification of AII-Amacrine Cells with Antibodies Against Calretinin. *The Journal of Comparative Neurology*, 361(3):537–51.
- Weliky, M., Bosking, W. H., and Fitzpatrick, D. (1996). A Systematic Map of Direction Preference in Primary Visual Cortex. *Nature*, 379:725–728.
- White, L. E. and Fitzpatrick, D. (2007). Vision and Cortical Map Development. *Neuron*, 56(2):327–38.
- Xiao, Y. (2014). Hierarchy of Hue Maps in the Primate Visual Cortex. *Journal of Ophthalmic & Vision Research*, 9(1):144–7.
- Xiao, Y., Casti, A., Xiao, J., and Kaplan, E. (2007). Hue Maps in Primate Striate Cortex. *Neuroimage*, 35(2):771–786.
- Xiao, Y., Wang, Y., and Felleman, D. J. (2003). A spatially organized representation of color in macaque cortical area V2. *Nature*, 421:535–539.
- Yu, H., Farley, B. J., Jin, D. Z., and Sur, M. (2005). The Coordinated Mapping of Visual Space and Response Features in Visual Cortex. *Neuron*, 47(2):267–280.
- Zhang, S., Yao, H., Sun, X., and Lu, X. (2013). Sparse coding based visual tracking: Review and experimental comparison. *Pattern Recognition*, 46(7):1772–1788.