

**Perceptual Grouping in a Self-Organizing
Map of Spiking Neurons**

Yoonsuck Choe

August 2001 AI01-292

Department of Computer Sciences
The University of Texas at Austin
Austin, TX 78712

This document is a reformatted version of the original dissertation.

Copyright

by

Yoonsuck Choe

2001

The Dissertation Committee for Yoonsuck Choe certifies that this is the approved version of the following dissertation:

**Perceptual Grouping in a Self-Organizing Map of Spiking
Neurons**

Committee:

Risto Miikkulainen, Supervisor

Benjamin J. Kuipers

Raymond J. Mooney

Wilson S. Geisler

Joydeep Ghosh

**Perceptual Grouping in a Self-Organizing Map of Spiking
Neurons**

by

Yoonsuck Choe, B.S., M.A.

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

August 2001

Dedicated to my parents, Woojin Choe and Mija Kim, and to my family.

Acknowledgments

First of all, I would like to thank my adviser Risto Miikkulainen for his continuous interest, support, and encouragement throughout this research. He taught me how to conduct high-quality research, and provided timely advice at moments when I needed it most. I would also like to thank Wilson Geisler, for the intriguing discussions through which many of the ideas presented in this thesis materialized. The committee members Benjamin Kuipers, Raymond Mooney, and Joydeep Ghosh deserve great thanks for their continuous interest and support, and for carefully reading and commenting on the manuscript. I am also grateful of DeLiang Wang, Xiuwen Liu, David Horn, and Lawrence Cormack for their interest and for the invigorating discussions. I am grateful of Jim Bednar for his piercing comments that helped make this thesis more accurate and readable, for the numerous discussions through which I learned a lot more about relevant neuroscience facts, and also for his friendship. I would like to thank Joseph Sirosh, for laying out the foundation for this thesis, and for providing me with an invaluable experience at HNC in San Diego during the Summer of 1999, for which I am also thankful toward Marc Ilgen. I am very grateful of Wonhui Cho, Un Yong Nahm, Sang Kyu Shin, Chan-Gun Lee, and other Korean friends for their friendship and support. I would also like to thank Marty and Coquis Mayberry, Tino Gomez, John Gunnels, and other friends at UTCS for their warm friendship, and for the intriguing discussions on topics ranging from Russian to pet care. I would also like to thank the fine folks at the UTCS Neural Networks lab, for helping me broaden my perspective. Chris Edmondson-Yurkanan also deserves special thanks for treating me as a colleague in the courses that I worked as a teaching assistant for her, and for her keen interest in my research. I would like to thank the developers of GNU software (<http://www.gnu.org>), the Linux operating system, and VNC (Virtual Network Computing; <http://www.uk.research.att.com/vnc>), without which this research would not have been possible. My sincerest thanks go to my parents, who have been loving, caring, patient, supportive, and understanding throughout my whole life, and to my wife Ahran for her love and patience, and to Dana (Yonsue), for her energy and courage. I would like to thank my brother and his family for their love and support, and the rest of my (extended) family, for gladly sharing their heart-felt emotions in times of happiness and sorrow. This research was supported by the National Science Foundation (grants #IIS-9811478 and #IRI-940004P), the University of Texas at Austin (High Performance Computer Time Grant, 1995), and the Schlumberger Graduate Fellowship (1999).

YOONSUCK CHOE

The University of Texas at Austin
August 2001

Perceptual Grouping in a Self-Organizing Map of Spiking Neurons

Publication No. _____

Yoonsuck Choe, Ph.D.
The University of Texas at Austin, 2001

Supervisor: Risto Miikkulainen

Perceptual grouping is the process of identifying the constituents in the visual scene that together form a coherent object. The goal of this thesis is to understand the neural mechanisms of perceptual grouping. The hypotheses are that (1) perceptual grouping is carried out through synchronized firing of neurons representing the same object, and that (2) self-organized lateral connections encoding statistical regularities of the visual environment mediate such a synchronization. A self-organizing neural network of spiking neurons was developed to test these hypotheses in the perceptual grouping task of contour integration. The network self-organized orientation maps and patchy lateral connections similar to those found in the visual cortex, and the contour integration, segmentation, and completion performance measured by the degree of synchrony in neural populations accurately predicted human performance. Such results suggest that synchronized activity can represent perceptual events, and statistical properties of the input can shape the structure of the cortex and the perceptual performance. By providing a computational framework where perceptual performance and neural structure can be compared, the model helps us understand the neural mechanisms of perceptual grouping.

Contents

Acknowledgments	v
Abstract	vii
List of Figures	xi
Chapter 1 Introduction	1
1.1 The Perceptual Grouping Task	1
1.2 Special Case: Contour Integration	2
1.3 Neural Mechanisms of Perceptual Grouping	5
1.4 Approach	6
1.5 Outline of the Dissertation	6
Chapter 2 Background	8
2.1 Temporal Coding	8
2.1.1 Why Temporal Coding?	8
2.1.2 Experimental Evidence for Temporal Coding	9
2.1.3 Modes of Synchronization for Temporal Coding	11
2.1.4 Computational Models of Temporal Coding	12
2.1.5 Conditions of Synchronization	14
2.1.6 Summary and Discussion	15
2.2 Self-Organization in the Visual Cortex	15
2.2.1 Structure of the Visual Cortex	15
2.2.2 Evidence of Input-Driven Self-Organization	16
2.2.3 Computational Models of Self-Organization	18
2.2.4 Summary and Discussion	20
2.3 Computational Models of Contour Integration	20
2.4 Open Issues in Contour Integration	22
2.5 Conclusion	23

Chapter 3	The PGLISSOM Model	24
3.1	Motivation	24
3.2	Network Organization	25
3.3	Neuron Model	27
3.3.1	Leaky Synapse	27
3.3.2	Calculating the Input Activation	29
3.3.3	Threshold Mechanism	30
3.4	Self-Organization	31
3.5	Summary	33
Chapter 4	Temporal Coding	34
4.1	Synchronization	34
4.1.1	Synaptic Decay Rate	34
4.1.2	Effect of the Connection Range	36
4.2	Desynchronization	38
4.2.1	Connection Types	38
4.2.2	Effect of Noise	40
4.3	Robustness of Binding and Segmentation	40
4.3.1	Effect of Relative Input Size on Synchronization	42
4.3.2	Overcoming Noise with Strong Excitation	43
4.3.3	Overcoming Noise with a Longer Refractory Period	44
4.4	Conclusion	45
Chapter 5	Self-Organization	46
5.1	Simulation Setup	46
5.2	Retinotopic Organization	48
5.3	Orientation Map	49
5.3.1	Receptive Fields, Global Order and Map Features	49
5.3.2	Orientation Preference Histogram	51
5.3.3	2-D Fourier Power Spectrum	52
5.3.4	Orientation Gradient and Orientation Selectivity	54
5.4	Lateral Connections	55
5.5	Conclusion	59
Chapter 6	Contour Integration and Segmentation	61
6.1	Simulation Setup	61
6.2	Contour Integration Performance	63
6.3	The Role of Lateral Connections	65
6.4	Contour Segmentation Performance	68
6.5	Conclusion	70

Chapter 7 Hemifield Differences in Anatomy and Performance	71
7.1 Motivation	71
7.2 Simulation Setup	72
7.3 Differences in Afferent and Lateral Connections	73
7.4 Differences in Contour Integration Performance	76
7.5 Conclusion	78
Chapter 8 Illusory Contours and Contour Completion	79
8.1 Motivation	79
8.2 Contour Completion Performance	81
8.3 Effects of Afferent and Excitatory Lateral Connections	86
8.4 Saliency of Closed vs. Open Contours	87
8.5 Conclusion	89
Chapter 9 Discussion	91
9.1 Temporal Coding in Neural Systems	91
9.1.1 Synchrony as a Perceptual Representation	91
9.1.2 Interpretation of Temporal Codes	93
9.2 Functional Divisions in the Visual Cortex	94
9.2.1 Layers of the Visual Cortex	94
9.2.2 Fovea vs. Periphery	95
9.3 High-Level Influence on Perceptual Grouping	95
9.4 Predictions of the PGLISSOM Model	96
9.5 Summary	97
Chapter 10 Future Work	98
10.1 Psychophysics	98
10.2 Neuroscience	99
10.3 Computational Sciences	100
10.4 Artificial Vision	101
10.5 Summary	102
Chapter 11 Conclusion	103
Bibliography	104
Vita	120

List of Figures

1.1	The Examples of Perceptual Grouping	2
1.2	Association Fields for Contour Integration	3
2.1	Superposition Catastrophe and its Solution	9
2.2	Synchronization Signaling Global Stimulus Configuration	10
2.3	Layered Structure and Long-Range Lateral Connections in the Visual Cortex	16
2.4	Orientation Map and Lateral Connections in Tree Shrews (<i>color figure</i>)	17
2.5	The RF-LISSOM Architecture	18
2.6	Orientation map and lateral connections in RF-LISSOM (<i>color figure</i>)	19
2.7	Lateral Connection Profiles	21
3.1	The Overall Organization of the PGLISSOM Network	26
3.2	The Neuron Model	28
4.1	Effect of Connection Type and Decay Rate on Synchronization	36
4.2	Effect of Connection Extent on Synchronization	37
4.3	Binding and Segmentation with Different Connection Types	39
4.4	Effect of Noise on Desynchronization	41
4.5	Effect of Relative Input Size on Synchronization	42
4.6	Overcoming Noise with Strong Excitation	43
4.7	Overcoming Noise with a Longer Refractory Period	44
5.1	Retinotopic Organization	48
5.2	Oriented Receptive Fields	50
5.3	Orientation Maps (<i>color figure</i>)	51
5.4	Characteristic Features of the Orientation Map (<i>color figure</i>)	52
5.5	Orientation Preference Histograms	53
5.6	Fourier Power Spectrum of a Cortical Orientation Map	53
5.7	2-D Fourier Power Spectra of PGLISSOM Orientation Maps	54
5.8	Relationship between Orientation Gradient and Orientation Selectivity	55
5.9	Gradient of a Cortical Orientation Map	56
5.10	Lateral Connections in Three Shrew	57

5.11	MAP2 Excitatory Lateral Connections (<i>color figure</i>)	58
5.12	Simultaneous Activation of Neurons	58
5.13	Difference in Orientation Preference vs. the Number of MAP2 Excitatory Lateral Connections	60
6.1	Inputs for the Contour Integration Experiments	62
6.2	Multi-Unit Activities in the Contour Integration Experiments	64
6.3	Contour Integration Performance	65
6.4	Spatial Relationship between Receptive Fields	66
6.5	Lateral Connection Statistics and Edge Co-Occurrence Statistics (<i>color figure</i>) . . .	67
6.6	Input for Contour Segmentation	69
6.7	Multi-Unit Activities for the Three Contours	69
6.8	Contour Segmentation Performance	70
7.1	Orientation Maps in Lower vs. Upper Half of PGLISSOM (<i>color figure</i>)	73
7.2	Orientation Selectivity in Lower vs. Upper Half of MAP1	74
7.3	Excitatory Lateral Connection Statistics in Upper vs. Lower Half of MAP2 (<i>color figure</i>)	75
7.4	Inputs in the Lower and Upper Hemifields	76
7.5	Multi-Unit Activities in the Lower vs. Upper Half of MAP2	77
7.6	Contour Integration Performance in Lower vs. Upper Hemifield	78
8.1	Two Types of Illusory Contours	80
8.2	Contour Completion across Edge-Inducers	81
8.3	Cooperative Bipole Cells	82
8.4	Inputs for the Contour Completion Experiments	82
8.5	Multi-Unit Activities for Contour Completion Experiments	83
8.6	Illusory Triangle Inputs	84
8.7	Multi-Unit Activities for the Illusory Triangle Inputs	84
8.8	Within and Across Contour Correlation	85
8.9	Afferent Contributions in MAP2	87
8.10	Multi-Unit Activities for Contour Completion with Removed Connections	88
8.11	Contour Completion Performance with Removed Connections	88
8.12	Saliency of Closed vs. Open Contours	89
8.13	Saliency of Kanizsa Triangles	90

Chapter 1

Introduction

The goal of this thesis is to understand the neural mechanisms of perceptual grouping through computational modeling of the visual cortex. Perceptual grouping can happen at many levels, and to establish a solid link between perceptual events and neural mechanisms, it is necessary to define the model at an appropriate level. Low-level phenomena such as contour integration are well-suited for this purpose because the necessary anatomical and neurophysiological data is well established. In this chapter, the problem of perceptual grouping in general and contour integration in particular are introduced. Psychophysical results are reviewed, and possible neural mechanisms will be outlined. Based on these facts, I will motivate the approach and conclude with an overview of the the dissertation.

1.1 The Perceptual Grouping Task

Perceptual grouping is the process of identifying the constituents in the visual scene that together form a coherent object. Examples of perceptual grouping are shown in figure 1.1. The complexity of such tasks varies widely, and they can take place at various levels of the visual processing hierarchy. The human visual system employs a large number of grouping rules at multiple levels, based on spatial, temporal, and chromatic relationships (Geisler and Super 2000). Understanding the neural mechanisms of perceptual grouping of such a broad range is a daunting task, but by focusing our attention on the low-level phenomena first, it may be possible to make progress. At that level, we have a large body of rich neurophysiological data that correspond to equally rich psychophysical data that can be used to constrain, validate, and test the models.

For this reason, in this thesis, I will focus on the low-level perceptual grouping task of contour integration. The next section defines the task, explains why it is important and appropriate for the current study, and reviews the known psychophysical results.

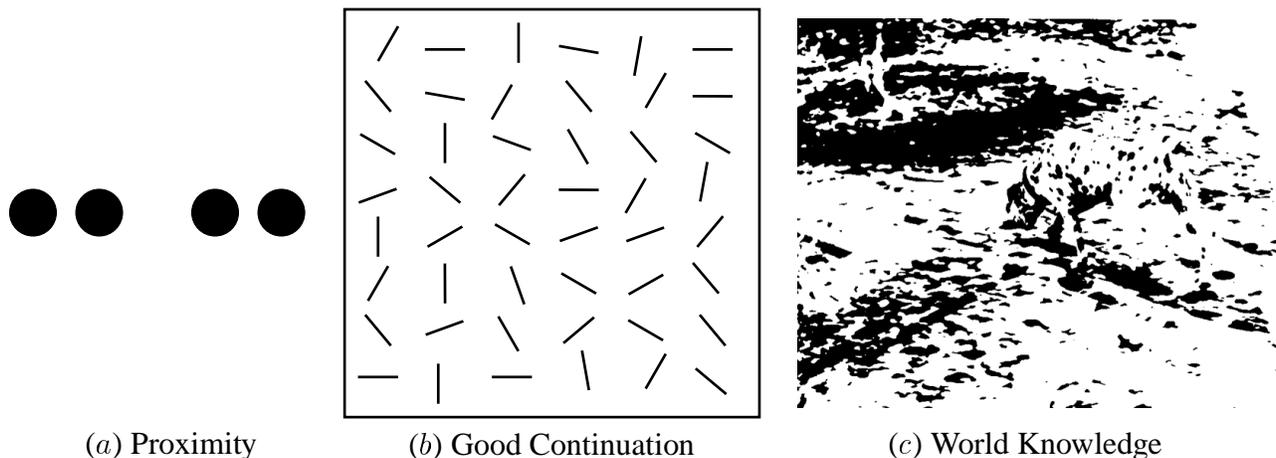


Figure 1.1: The Examples of Perceptual Grouping. Perceptual grouping is the process of identifying constituents in the visual scene that together form a coherent object. Perceptual grouping can take place at many different levels, from the very low level (a), to the very high level (c). (a) Grouping by proximity. We tend to group the two black disks that are close to each other to form a unit. Thus we perceive two groups, one on the left and one on the right. (b) Grouping by good continuation. In the random background of oriented edges (or contour elements), we can easily notice the long, straight series of contour elements that runs from the mid-left of the frame to the bottom right corner. The task is known as *contour integration*. (c) Grouping requiring world knowledge. In this seemingly unintelligible image lurks a Dalmatian dog sniffing on the pavement (a photo by R. C. James; the dog is in the top right the image, facing left). Without knowledge of the world, especially that of the dog, it would be impossible to group together the dots that form the Dalmatian.

1.2 Special Case: Contour Integration

A typical visual input for the contour integration task is shown in figure 1.1b. The input consists of a series of short oriented edge segments (or *contour elements*) aligned along a continuous path, embedded in a background of randomly oriented contour elements. The task is to identify the longest continuous contour in this scene. Contour integration is an interesting problem because the relationships between constituents of the image are neither too simple (as in figure 1.1a where the distance between the centers of the disks is the only grouping criteria), nor too complex (as in figure 1.1c where complex world knowledge is required). Most importantly, contour integration is believed to occur relatively early in the visual system. The response properties and connection patterns found in the primary visual cortex seem to have exactly the right properties for explaining contour integration performance in terms of neural mechanisms.

Psychophysical experiments (Field et al. 1993a; Pettet et al. 1998; Geisler et al. 1999, 2001; McIlhagga and Mullen 1996) and computational theories and models (Geisler et al. 1999, 2001; Li 1998; Yen and Finkel 1997, 1998) suggest that there exists a highly specific pattern of interactions among the contour elements. Such interactions allow contour elements in certain orientation and position configurations to be more visible than in others.

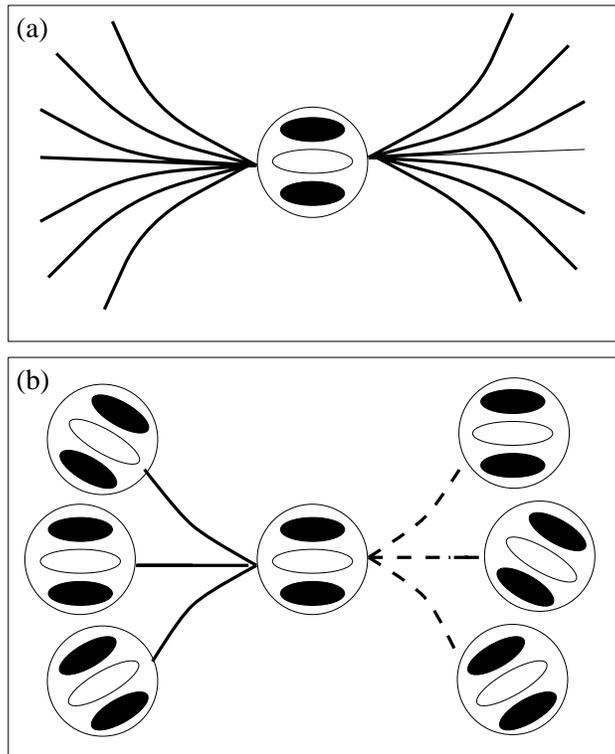


Figure 1.2: **Association Fields for Contour Integration.** The interaction patterns postulated by Field et al. (1993a) are shown. The circular disks with black and white oriented bars represent Gabor wavelets, which is the typical contour element used in the study of Field et al. (a) The contour element in the center interacts with the elements on smooth curves radiating out from the center. This plot shows the *spatial positions* where another contour element can appear relative to the orientation of the reference contour element in the center. (b) Specific rules of alignment. The *orientations* of contour elements are as important as their positions in determining if any two contour elements should enhance each other. The orientation has to be parallel (or collinear) to the smooth contour. The solid lines are where integration occurs, and the dashed lines are where it does not. Adapted from Field et al. (1993a).

Field et al. (1993a) conducted a series of experiments where each subject was told to find a contour of similarly-oriented Gabor patterns embedded among randomly-oriented Gabor patterns. Several factors affected the performance of the subjects. One important factor was the relative orientation of successive contour elements (or *orientation jitter*) along the longest contour (the path). When the orientation of successive contour elements differed more, the performance degraded. There were other factors such as inter-element distance and difference in phase of the successive Gabor patterns, but the most important factor was relative orientation along the path.

Based upon these results, Field et al. suggested that local interactions between contour elements follow specific rules and form the basis for contour integration in humans. In other words, these constraints form a local *association field* that governs how differently-oriented contour elements should interact to form a coherent group (figure 1.2). An association field can be described with two rules: (1) contour elements positioned on a smooth path (figure 1.2a), and (2) contour el-

ements aligned parallel (or collinearly) along the path (figure 1.2b) are more likely to be perceived as belonging to the same contour.

Pettet et al. (1998) further confirmed that lateral interactions between neighboring contour elements follow well-defined constraints similar to those suggested by Field et al. They extended the work by Field et al. and compared the performance of human subjects to their computational model based on fixed lateral interaction constraints similar to the association field. The performance of the model matched psychophysical data very well. In particular, it was consistent with earlier results with closed vs. open-ended contours by Kovacs and Julesz (1993) in that closed contours were easier to detect than open-ended contours. They reasoned that the lack of lateral interaction between the two ends of the open-ended contour degraded the perceptibility of the contour, while reverberating lateral interaction along the closed loop enhanced its perceptibility.

Geisler et al. (1999, 2001) took a different approach in identifying the conditions that govern contour integration (or *local grouping function*, according to the authors). Instead of proposing lateral interaction rules by observing human performance, they extracted the rules from edge co-occurrence statistics measured in natural images (see Chapter 6, figures 6.4 and 6.5b). Edge-detected natural images were decomposed into outline figures, consisting of short oriented edges. The co-occurrence probability of each pair of edges belonging to the same physical contour in natural images was then calculated. A striking discovery was that such edge co-occurrence statistics in fact are very similar to the lateral interaction rules proposed earlier by Field et al. and Pettet et al. Furthermore, Geisler et al. devised a method of extracting contours using these co-occurrence statistics. They defined *grouping criterion*, where two edges were grouped together if the probability of the edges occurring in the given configuration exceeds the binding criterion β (a threshold). They then applied the *transitive grouping rule*, where edges **a** and **c** are grouped together if **a** and **b** can be grouped together and **b** and **c** can be grouped together. Geisler et al. showed that together, these grouping rules accurately predict human performance.

Thus, Geisler et al. showed that the statistical structure in the environment closely corresponds to human perceptual grouping rules. Sigman et al. (2001) also reported similar edge co-occurrence statistics, but they did not apply their results in predicting human performance. However, since the results are similar, the statistics gathered in the work of Sigman et al. can predict human performance if used in the same manner as Geisler et al. did.

The above psychophysical results suggest that lateral interaction between neurons representing contour elements is crucial for contour integration in humans, and that the interaction patterns are highly specific (i.e. constitute a smooth path). If that is true, then what is the neural mechanism of such lateral interactions? In the next section, the possible answers will be briefly reviewed.

1.3 Neural Mechanisms of Perceptual Grouping

One major question in perceptual grouping is how coherent percepts are represented in the cortex. The task consists of two parts: *binding* is the process of grouping together separate constituent representations in the visual scene into a coherent object, and *segmentation* is the process of segregating such coherently bound representations into different objects. von der Malsburg and Schneider (1986) showed that with static activity, it is hard to dynamically represent binding and segmentation in a constantly changing sensory environment. They proposed that *temporal coding* through synchronization, spike timing, phase differences, and other temporal information, could solve the problem. Indeed, experiments with cats have shown that presentation of coherent objects gives rise to synchronized firing of neurons in the visual cortex, and presenting separate object causes no synchronization (Eckhorn et al. 1988; Gray and Singer 1987; Gray et al. 1989; Singer 1993). Such a coherent firing of neurons may be a possible representation for grouping.

Experiments with cats, monkeys, ferrets and tree shrews have shown that the visual cortex has an orderly structure where stimulus dimensions such as position (retinotopy), orientation (Blasdel and Salama 1986; Blasdel 1992a,b; Bosking et al. 1997; Grinvald et al. 1994; Sincich and Blasdel 2001; Ts'o et al. 1990; Weliky et al. 1995), spatial frequency (Issa et al. 2001), ocularity (Blasdel 1992a; Crowley and Katz 2000; Löwel 1994), and direction (Weliky et al. 1996; Shmuel and Grinvald 1996), are represented in a continuously-changing fashion. Another important feature of the visual cortex is that long-range lateral connections and they project to areas with similar sensory tuning, such as similar orientation preference (Bosking et al. 1997; Dalva and Katz 1994; Gilbert 1992; Katz and Callaway 1992; Löwel and Singer 1992; McGuire et al. 1991; Weliky et al. 1995). Such specific patterns of connectivity suggest that the lateral connections are the anatomical basis of lateral interaction for perceptual grouping tasks such as contour integration.

The second important question then is, how do these specific connections and sensory maps come about? Strong evidence suggests that these structures are self-organized, driven by external input (Hirsch and Spinelli 1970; Blakemore and Cooper 1970; Blakemore and van Sluyters 1975; Hubel and Wiesel 1962; Hubel et al. 1977; Hubel and Wiesel 1974; White et al. 2000, 2001). Even though there is a genetic component as well, the input environment plays a critical role: changing the environment changes the final structure that emerges in the cortex.

Computational models have primarily been developed to account for (1) synchronization in a connected network of spiking neurons (reviewed in section 2.1) (2) self-organization of cortical structures (section 2.2), and (3) perceptual grouping in neural networks (section 2.3). However, each of the components were modeled separately, and their interaction was not taken into account.

In this thesis, I intend to bring these important components together into an integrated computational model of the visual cortex. It is expected to account for the neural basis of contour integration in particular, and provide insights into perceptual grouping in general.

1.4 Approach

Motivated by the results reviewed above, I developed a model of the visual cortex called PGLISSOM (Perceptual Grouping Laterally Interconnected Synergetically Self-Organizing Maps; 2000). The model is based on the following two hypotheses: (1) *synchronized neural activity* represents perceptual grouping, and (2) *self-organized lateral connections* mediate the synchrony.

In the PGLISSOM model, three important components of the visual perceptual phenomena are integrated into a coherent computational framework: (1) statistical structure in the visual environment (2) structure of the visual cortex, and (3) functional performance of the visual cortex. Close interrelationships between these components are revealed by a series of computational experiments. The model demonstrates that visual input shapes the structure of the visual cortex through the input-driven self-organization and the structure determines performance in psychophysical tasks measured as the degree of synchrony among neural populations. The model therefore provides a computational account of the possible neural mechanisms of contour integration.

1.5 Outline of the Dissertation

This dissertation is organized into four parts: Introduction and Background (**Chapters 1 and 2**), Model (**Chapter 3**), Results (**Chapters 4 through 8**), and Discussion, Future Work, and Conclusion (**Chapters 9 through 11**).

In **Chapter 2**, I will review the functional and structural organization of the visual cortex that forms the neural basis for perceptual grouping. Previous computational models of contour integration are also surveyed.

In **Chapter 3**, I will formally define the PGLISSOM model and describe the details of its architecture.

In **Chapter 4**, I will analyze the synchronization and desynchronization properties of a network of spiking neurons through computational simulations, and identify crucial parameters that determine the overall behavior of the model.

In **Chapter 5**, the self-organized afferent and lateral connections of the PGLISSOM model are described and the properties of these connections are analyzed. Through self-organization, an orientation map highly similar to those found in experiments emerges.

Focusing on the self-organized structure, in **Chapter 6**, I will test the contour integration and contour segmentation performance on inputs with varying degrees of orientation jitter. The contour integration performance is shown to be consistent with psychophysical data. Lateral connection statistics are gathered to show the relationship between anatomy and perceptual performance.

In **Chapter 7**, I will test the effect of changes in the input distribution on the structure and performance of the model, and find out the possible cause of hemifield differences in visual per-

ceptual performance. The results suggest that differences in the input distribution cause the cortex to self-organize differently, and the structural divergence in turn causes perceptual performance to differ.

In **Chapter 8**, the patchy long-range lateral connections that develop in the PGLISSOM model are shown to assist in contour completion, a basic task that may be the foundation for a number of illusory contour percepts. The conditions under which contour completion occurs are systematically tested.

In **Chapter 9** the role and extent of temporal coding, how to interpret the temporal firing patterns, implications of the PGLISSOM architecture in understanding the laminar layer structure in the visual cortex, and the possible role of higher-level influence on lower level perception are discussed. Finally, predictions made by PGLISSOM on the role of synchronization and self-organization in perceptual grouping and specific experiments to verify these predictions are presented.

In **Chapter 10**, the most promising future directions of this research in psychophysics, neuroscience, computational sciences, and artificial vision are discussed.

Chapter 11 summarizes and evaluates the contributions of the thesis.

Chapter 2

Background

In this chapter, the anatomy and physiology of the visual system that forms the neural basis for perceptual grouping will be reviewed, and previous work on modeling the neural mechanisms of perceptual grouping phenomena will be discussed. First, possible neural representations of perceptual grouping will be described. Next, anatomical structures found in the visual cortex will be overviewed, and experimental results suggesting how these structures come to exist will be outlined. I will also describe how these structures can contribute to perceptual grouping. A review of previous work on perceptual grouping based on these observations concludes the chapter.

2.1 Temporal Coding

Neural representations employing temporal information such as spike timing, synchrony in spikes or population activities, and spiking events locked to background oscillations are generally known as *temporal coding*. An important question in research into the neural mechanisms of binding and segmentation for perceptual grouping is how such grouping relations are represented in the cortex. In this section, the possibility of representing grouping through temporal coding and experimental evidence for it will be described. Next, computational models derived from these observations will be described and compared.

2.1.1 Why Temporal Coding?

Traditional neural network theories hypothesized that the level of activation, or the firing rate of neurons, forms the representation for perceptual events. However, von der Malsburg (1981) pointed out that such static representations suffer from *superposition catastrophe* (figure 2.1). This problem arises when distributed neural representations of two (or more) separate objects are overlapped. In such a case, it is no longer clear which neuron represents which object (figure 2.1*a*). In contrast, if the representations for the individual objects are interleaved in time, binding and segmentation can occur naturally through temporal coding (figure 2.1*b*). von der Malsburg et al.

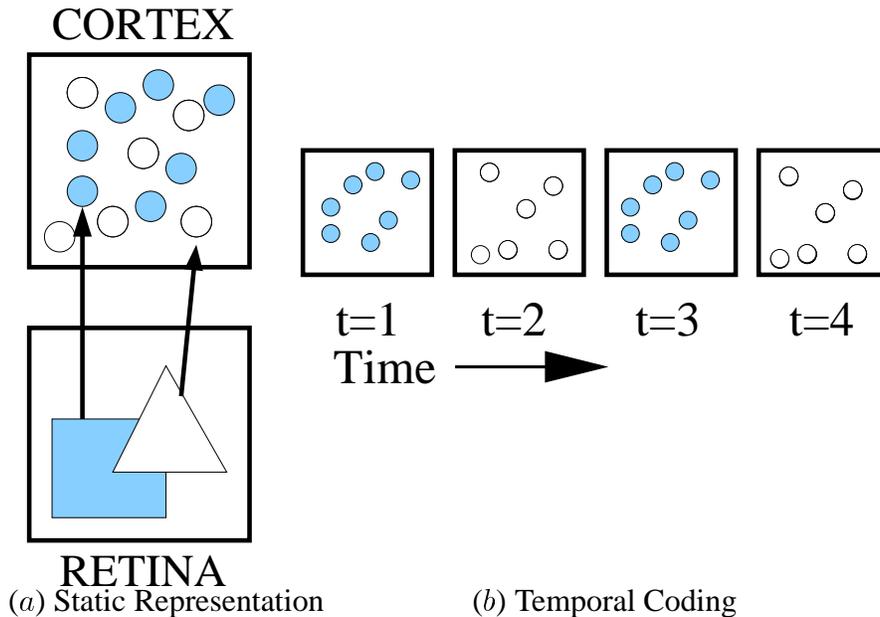


Figure 2.1: **Superposition Catastrophe and its Solution.** If firing rates of neurons alone are used to represent objects, multiple objects in the scene can cause a confusion. In (a), a square and a triangle is presented in the retina, and the neurons representing the two objects are identified and colored the same. When both populations of neurons are active, it is impossible to know which neuron is representing which object. This problem is known as the *superposition catastrophe* (von der Malsburg 1981; von der Malsburg and Schneider 1986). One solution is temporal coding, where temporal information is used to separate the two populations. Neurons representing one object activate at one time step, and neurons representing the other object activate at another time step, as shown in (b).

hypothesized that binding and segmentation are achieved in a similar manner through synchronized and desynchronized firing of neurons through time (von der Malsburg and Schneider 1986; von der Malsburg 1987). In the following, experimental evidence and related work on such temporal representation of grouping will be presented.

2.1.2 Experimental Evidence for Temporal Coding

To test whether such temporal representations are used in the visual system to represent perceptual grouping events, two approaches can be taken. One way is to present inputs to the visual system and measure the temporal properties of neural activation in the cortex. The other is to alter the temporal properties in the input and measure the effect on perceptual performance.

Experimental results showed that coherent oscillations (synchronized high-frequency waves near the 40Hz γ -band) arise within populations of neurons and the activity between two populations with similar properties such as the same orientation preference is synchronized when stimulated with an input (Eckhorn et al. 1988; Gray and Singer 1987; Gray et al. 1989; Singer 1993). The most convincing evidence was found when electrical recordings were made on two sites in

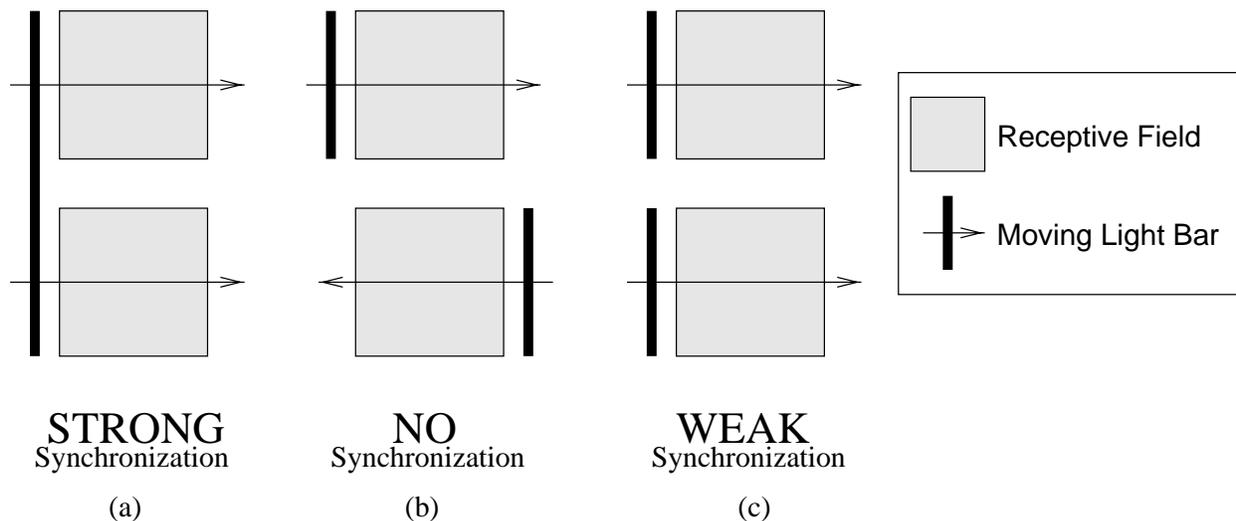


Figure 2.2: **Synchronization Signaling Global Stimulus Configuration.** (a) A single bar moving across two receptive fields results in strong synchronization between the two neuronal populations with these receptive fields. (b) Two separate bars moving in opposite directions results in no synchronization. (c) Two separate bars moving in the same direction results in weak synchronization. (Adapted from Gray et al. 1989).

the cat visual cortex with non-overlapping receptive fields while moving light bar(s) were swept across these receptive fields (figure 2.2). When a single long bar was used as the input, the two populations representing distant sections of the long bar fired synchronously. However, when two short bars were swept in the same location as before in the opposite direction of each other, the firing of the two populations was no longer synchronized. Interestingly, when two separate short bars were swept in the same direction, the two populations showed a weak but synchronized activity (Gray et al. 1989; Engel et al. 1991a,b; Singer 1993). These results suggest that synchronized firing of distant populations of neurons may represent the percept of a single coherent object, and desynchronized firing that of separate objects.

Another piece of evidence for synchronization in perception was obtained by manipulating the temporal properties in the visual input. Usher and Donnelly (1998) hypothesized that if neuronal synchrony plays a major role in binding, altering the timing properties of objects in the visual input stream relative to the background would result in differences in perceived grouping of the objects. They presented inputs where the object to be detected and the background were either flashed in synchrony (both object and background blink at the same time) or flashed asynchronously (object and background blink at different phase) over a period of time. The time-scale of the flashing was made shorter than the integration time of the visual system so that such flashing could not be consciously perceived. Given such input, the subjects were asked to identify where the object appeared among one of four areas in the background. They found that the percentage of correct responses was consistently higher when the object and background were *asynchronously* flashing. They also found that the percentage of correct responses increased as the phase differ-

ence between the flashing of object and background was increased. The explanation was that the timing of the inputs caused the temporal properties of neuronal firing to change and in turn caused the detection performance to differ. Flashing the object and background at different times would cause a slight phase shift between the neurons representing the object and background, and such a shift helped distinguish the object from the background. Increasing the interval between the flashes would cause the phase shift to become larger, therefore allowing better separation of neural representations of the object and the background. Similar results have been reported by Fahle, Leonards, and Singer (Fahle 1993; Leonards et al. 1996; Leonards and Singer 1998) and Lee and Blake (1999).

Motivated by these results, several neural network models with temporal dynamics have been developed. Next, the major categories of such computational models will be overviewed.

2.1.3 Modes of Synchronization for Temporal Coding

There are two ways in which synchrony can occur: (1) synchronized (i.e. simultaneous) firing of individual neurons, and (2) synchronized oscillations of population activity (i.e. number of neurons in the population firing per unit time). Population oscillations are more general and include synchronized firing as a special case. They are also more biologically plausible for several reasons. Due to the stochastic nature of neuronal firing, it seems unlikely that individual neurons could synchronize their actual firing events. However, they could fire within a short time window so that the spikes are approximately aligned, and the whole group could exhibit synchrony (Lisman 1998; Menon 1990). Theoretical results also suggest that the oscillations found in the cortex result from a collective behavior of neurons (Wilson and Cowan 1972). Such population oscillations should be more robust and tolerant of random fluctuations (Menon 1990). For example, Eckhorn et al. (1988) discovered that synchrony in individual neurons is hard to find even when the multi-unit activity (MUA) and local field potential shows coherent oscillation, suggesting that population oscillation is the major mode of operation for binding of percepts. There is also indirect experimental evidence to support this hypothesis. When two almost simultaneous clicks are presented to a subject, they are initially heard as a single click, but as the interval between the two clicks increases, the subject starts hearing two clicks instead of one. Interestingly, this transition from one click to two clicks occurs exactly at the frequency of population oscillations (Joliot et al. 1994), suggesting that neuronal firing events within a single cycle of population oscillation are bound together even though the exact timing does not match, whereas the firings that occur in different cycles are perceived as separate. For these reasons, most of the synchronizing models, including the model in this thesis, adopt the definition of synchrony in terms of population oscillations rather than that of individual neurons.

2.1.4 Computational Models of Temporal Coding

Biological neurons consist of (1) dendrites where contact with the input neurons (or pre-synaptic neurons) are made, (2) the cell body (or soma) where the integration of incoming currents is carried out, (3) the axon hillock where the accumulated charge triggers a spike, and (4) the axon that transmits the generated spike to other neurons (Kandel et al. 1991). A neuron is a cell enclosed in a fatty (or lipid) membrane with various ion channels that conduct electric current into and out of the cell body. The voltage across the membrane (membrane potential) changes as these ions come in and go out of the neuron through the ion channel, and it is this voltage that determines whether the neuron generates a spike or not. Such dynamic change in state over time together with the generation of spikes gives the neuron rich temporal dynamics.

Several computational models have been developed based on such temporal dynamics. There are two major classes distinguishable by how the dynamics are formulated; (1) coupled oscillators and (2) integrate and fire neurons. These *temporal neuron* models have been used to explore the idea that binding and segmentation in the cortex can be expressed in the form of synchronized or desynchronized firing of neuronal populations.

2.1.4.1 Coupled Oscillators

In the first class, the temporal dynamics of each neuron are described by two variables that represent coupled oscillators consisting of an excitatory unit and an inhibitory unit (Horn and Opher 1998; Terman and Wang 1995; von der Malsburg 1987; von der Malsburg and Buhmann 1992; Wang 1996, 1995; Wilson and Cowan 1972): some models use a single complex variable to describe the state of the neuron, resulting in a similar model (Chakravarthy and Ghosh 1996). The two units can represent states of a single neuron, two different neurons, or two separate populations of neurons.

These models originate from the Hodgkin-Huxley (HH) model of excitable membrane (Hodgkin and Huxley 1952), which can be written as:

$$C \frac{dV}{dt} = -I_{ion}(V, W_1, \dots, W_n) + I(t) \quad (2.1)$$

$$\frac{dW_i}{dt} = \Phi \frac{[W_{i,\text{inf}}(V) - W_i]}{\tau_i(V)}, \quad (2.2)$$

where V is the membrane potential, C is the membrane capacity, and I_{ion} is the sum of V and t -dependent currents through each ion channel type, $I(t)$ is the input current, W_i for each ion channel type quantify the fraction of ion channels open, τ_i are the time constants of W_i , Φ is time scale factor for the ion channels, and $W_{i,\text{inf}}(V)$ is the steady-state value of W_i . Such a detailed model allowed a close match with experimental data, but the formulation involves multiple variables which made it hard to analyze. Subsequently, FitzHugh (1961) and Nagumo et al. (1962) reduced the model into a two-variable model, convenient for analysis and visualization in a 2D phase-plane. For a historical account of these models, see Rinzel and Ermentrout (1999). The two units

are coupled with an excitatory connection from the excitatory unit to the inhibitory unit, and with an inhibitory connection from the inhibitory unit to the excitatory unit. A typical behavior in such a dynamical system is limit-cycle oscillation, and such an oscillation can be easily visualized and analyzed in a 2D phase portrait of the activities of the two units.

In a typical implementation, the coupled oscillators are arranged in a 2D topology with local excitatory connections and global inhibition (Terman and Wang 1995; von der Malsburg and Buhmann 1992; Wang 1996, 1995). The excitatory units receive excitatory input from neighboring oscillators and inhibitory influence from a global inhibitor. These local excitatory connections drive the phases of the neighboring oscillators closer to each other. The global inhibitor pools the activity of all excitatory units and gives feedback inhibition to all excitatory units, thus implementing segmentation among distant oscillating populations. The peaks and valleys of the activity are compared to see if they are synchronized or not. These types of networks have been applied to texture segmentation (von der Malsburg and Buhmann 1992), aerial photo and brain scan image segmentation (Wang 1996, 1995), and cluster analysis (Horn and Opher 1998) with impressive results. However, the lateral connections in these models were limited in spatial extent and could not learn input correlations, thus grouping was only based on proximity or similarity.

2.1.4.2 Integrate-and-Fire Neurons

In the other class of temporal neurons, a single variable corresponding to the membrane potential of a neuron is used to describe the state. A common form of the dynamics of this class of neurons involve, (1) accumulation of voltage from incoming signals, (2) comparison of the membrane potential with a threshold to generate a spike, and (3) resetting the membrane potential after spiking. A typical formulation is:

$$C \frac{dV}{dt} = I(t) - \frac{V}{R} \quad (2.3)$$

where C is the capacity, V is the membrane potential, R is the resistance (leak term), and $I(t)$ is the input current (Lapicque 1907; see Gabbini and Koch 1999 for a review). The incoming activity $I(t)$ accumulates over time until the membrane potential reaches a fixed threshold. The neuron is allowed to spike, and is reset to 0. These types of neuron models are generally called integrate-and-fire neurons. Many different subclasses of such models exist, including leaky integrate-and-fire neurons (Campbell et al. 1999; Nischwitz and Glünder 1995), and leaky synapse with dynamic threshold models (Eckhorn et al. 1990; Reitboeck et al. 1993; I will refer to these as LSDT). Gerstner (1998b) showed that these subclasses are equivalent under his general framework of Spike Response Model (SRM).

The LSDT model of Eckhorn et al. (1990) and Reitboeck et al. (1993) is biologically the most detailed, including explicit dendritic integration (leaky synapses). For this reason, this model is extended and used in the current thesis. In LSDT, each synapse is a leaky integrator where pre-synaptic action potentials accumulate in the post-synaptic dendritic membrane and decay over time. The membrane potential, accumulated over space across the dendritic arbor and over time by

the decayed sum of past activity, is compared to a dynamic threshold and the neuron is allowed to fire when the cumulative sum of activity crosses the threshold. The threshold is dynamic in that it is increased acutely after the neuron fires, and the amount of increase is decayed over time. Such an increase models the refractory period of the neuron. Both the leaky synapse and the dynamic threshold are formulated using the same leaky-integration mechanism based on convolution (*):

$$X(t) = Z(t) * I(V, \tau, t), \quad (2.4)$$

where $X(t)$ is the leaked sum at time t , $Z(t)$ is the input, and $I(\cdot)$ is defined as:

$$I(V, \tau, t) = \begin{cases} V \cdot \exp(-t/\tau) & \text{if } t \geq 0 \\ 0 & \text{otherwise,} \end{cases} \quad (2.5)$$

where V is the amplification factor and τ is the time constant of the leak kernel. The convolution is calculated using the digital filter equation:

$$X[t] = X[t - 1] \cdot \exp(-1/\tau) + V \cdot Z[t], \quad (2.6)$$

where t increases in discrete time steps of duration 1.

The integrate-and-fire neurons have also been arranged in a 2D topology with local excitatory connections and global inhibition for static object segmentation (Eckhorn et al. 1990; Campbell et al. 1999), moving object segmentation (Reitboeck et al. 1993), and aerial photo and brain scan image segmentation (Campbell et al. 1999) with good results. However, as in coupled oscillator models, the lateral connections were limited in extent and could not self-organize to map complex input correlations, thus grouping was only based on proximity or similarity.

2.1.5 Conditions of Synchronization

The temporal neurons described above are usually connected with excitatory and inhibitory connections to elicit synchrony and desynchrony. Analytical and computational studies have been conducted to find out the conditions under which such connected neurons synchronize and desynchronize. Synchronization properties have been studied in terms of excitatory vs. inhibitory connections, axonal conduction delay, and noise.

The main results are: (1) excitatory connections with no delay cause synchrony (Campbell et al. 1999; Gerstner and van Hemmen 1992; Han et al. 1998; Horn and Opher 1998; Mirollo and Strogatz 1990; Terman and Wang 1995; Wang 1995, 1996), (2) excitatory connections with delay cause desynchrony (Nischwitz and Glünder 1995), (3) inhibitory connections without delay cause desynchrony (Han et al. 1998; Horn and Opher 1998; Nischwitz and Glünder 1995), and (4) inhibitory connections with delay cause synchrony (Horn and Opher 1998; Kirillov and Woodward 1993; Lytton and Sejnowski 1991; Nischwitz and Glünder 1995; van Vreeswijk and Abbott 1994), and (5) noise helps desynchronization of separate populations (Han et al. 1998; Horn and Opher 1998; Terman and Wang 1995; Wang 1995). However, the effect of varying synaptic decay on synchronization has not been analyzed: this effect will be tested in this thesis.

2.1.6 Summary and Discussion

In this section, I reviewed experimental results suggesting that perceptual grouping in the visual cortex is based on temporal representations, and computational models that describe the temporal dynamics of neurons. I also surveyed studies on the conditions of synchrony in networks of such model neurons. As these results show, lateral connections are essential in mediating synchrony and desynchrony in populations of neurons. However, due to spatially symmetric and/or fixed lateral connection patterns in the previous models, binding was based on proximity or similarity only, and as a result, more complex Gestalt grouping task such as integration could not be performed. For such tasks, the connection patterns need to be more functionally specific and anisotropic, and they need to encode the correlational structure found in the input distribution. In the next section, the neural substrate of such specific lateral interactions will be reviewed and computational accounts of how such a structure emerges during development will be given.

2.2 Self-Organization in the Visual Cortex

In this section, the structure of the visual cortex relevant to perceptual grouping will be reviewed and evidence of input-driven self-organization of these structures will be presented. Then, the computational models of self-organization of such structures will be reviewed.

2.2.1 Structure of the Visual Cortex

The visual cortex, like any other part of the neocortex, is primarily a two-dimensional sheet of neurons and connections. The sheet is arranged into six layers with different types of neurons occupying each layer (Henry 1989). The layers are numbered 1 to 6 from the surface (figure 2.3). At any location on the cortical sheet, the neurons within a vertical cylindrical column have similar feature preferences, and such vertical groups of neurons are termed *columns* (Gilbert and Wiesel 1989). Lateral connections exist between neighboring columns, and they tend to connect columns with similar feature preferences (figure 2.3).

Each column receives input from a limited area in the visual field called a receptive field, and prefers specific values for different feature dimensions such as orientation, spatial frequency (size), and ocular dominance (left, right, or both eyes). Optical imaging and voltage sensitive dyes have been used to measure orientation preference in mammals, and an intricate pattern of orientation maps has been observed (Blasdel and Salama 1986; Blasdel 1992a,b; Grinvald et al. 1994; O'Keefe et al. 1998; Ts'o et al. 1990; Weliky et al. 1995). Figure 2.4 shows such an orientation map. Orientation preference smoothly changes over the map and cycles through all orientations at regular, repeating intervals. Extensive lateral (or horizontal) connections exist between neighboring orientation columns, and generally, these connections are clustered at regular intervals and target columns possessing similar orientation preference (figure 2.4; Bosking et al. 1997). Such

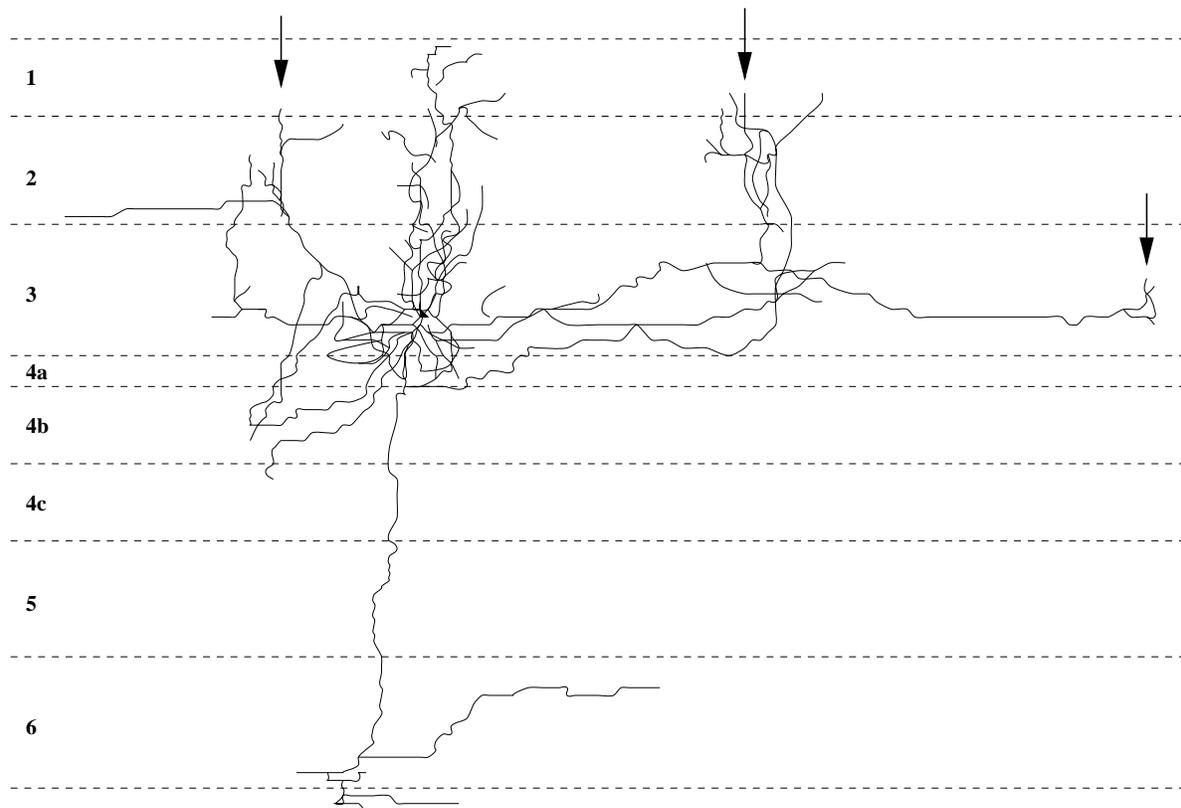


Figure 2.3: Layered Structure and Long-Range Lateral Connections in the Visual Cortex. Lateral connections, also sometimes called intrinsic horizontal or horizontal connections, run tangential to the cortical surface. In the visual cortex they reach distances corresponding to several degrees of the visual field, and sprout synaptic boutons at regular intervals (marked by arrows). The synaptic bouton clusters are highly concentrated on areas where the target neurons have similar feature tuning as the source neuron. Source: Mikkulainen and Sirosh (1996), adapted from (Gilbert and Wiesel 1989).

specific lateral connections can implement, e.g. the Gestalt law of good continuation, and perceptual grouping in general.

2.2.2 Evidence of Input-Driven Self-Organization

How these complicated yet orderly structures come about in the cortex was a big question until the 1960s. At that time, Hubel, Wiesel and their colleagues conducted a number of classic experiments where they showed that altering the visual environment drastically changes the organization of the visual cortex (Hubel and Wiesel 1962; Hubel et al. 1977; Hubel and Wiesel 1974). For example, if a kitten's vision is impaired by suturing the eyes shut, the visual cortex does not develop a normal organization, and orientation maps and ocular dominance columns do not appear. Such an effect is most dramatic during the critical period, typically within a few weeks after birth (Hubel and Wiesel 1962). Keeping the eyes shut until after the critical period caused the animal to become blind. This

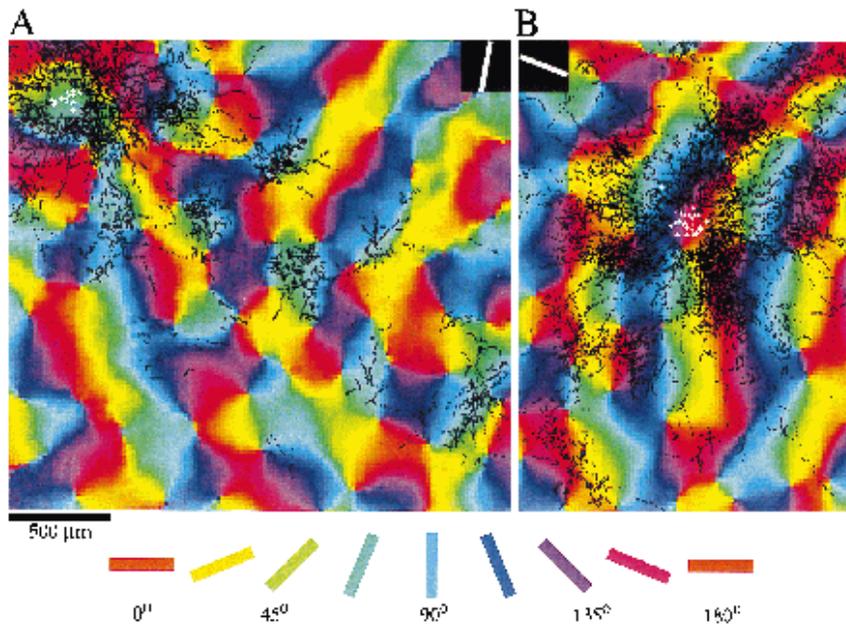


Figure 2.4: **Orientation Map and Lateral Connections in Tree Shrews (color figure)**. The cells are colored according to the orientation of the lines to which they are most sensitive. Two characteristics are immediately noticeable: (1) orientation preference changes smoothly across the cortex, and (2) a full cycle of orientation preferences repeats at regular intervals. The black dots mark the axon terminals projecting out from the area marked with white dots. In **A**, the source neurons are sensitive to the 90 degree orientation (cyan color). Near the source, the connections are found on all orientations, but as the connections reach out further, they are more likely to end up on similarly orientation-tuned areas. This connectivity pattern can also be found in **B**, where the source area is tuned to 0 degree (red color). Such specific lateral connections are believed to implement perceptual grouping rules. Source: Bosking et al. (1997).

result shows how important normal visual stimuli are during the critical period for normal development of the visual system. In another experiment, kittens were raised in an environment with only vertical or horizontal features. As a result, the kittens were unable to respond well to other orientations (Hirsch and Spinelli 1970; Blakemore and Cooper 1970; Blakemore and van Sluyters 1975). Recent experiments reported similar results in dark-reared and eye-sutured animals with stronger disruption found for abnormal visual experience than for deprivation (White et al. 2000, 2001). Similar results were also reported for ocular dominance columns in ferrets (Issa et al. 1999). The auditory cortex was shown to become sensitive to visual inputs when the projections from the retina or visual area of the thalamus (lateral geniculate nucleus or LGN) were surgically connected to the auditory cortex, suggesting that input strongly influences cortical development (Sur et al. 1988; Sharma et al. 2000).

Similar mechanisms are believed to be involved in the development of lateral connections. They were found to form gradually during early development based on visual input (Callaway and Katz 1990, 1991; Löwel and Singer 1992), and around the same time as the orientation maps and ocular dominance (Burkhalter et al. 1993; Dalva and Katz 1994; Katz and Callaway 1992). Lateral

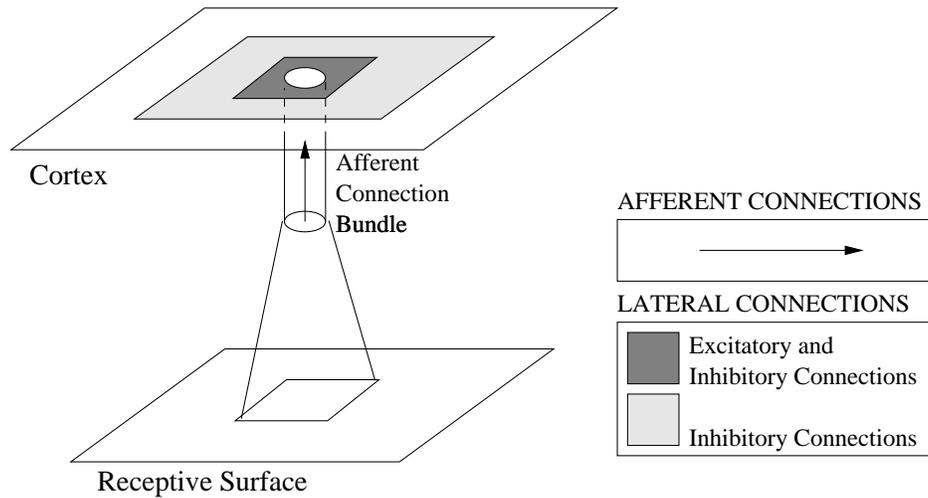


Figure 2.5: **The RF-LISSOM Architecture.** The Receptive-Field Laterally Interconnected Synergetically Self-Organizing Maps(RF-LISSOM) is a model of the visual cortex with self-organizing afferent and lateral connections. Each cortical neuron has a small receptive field on the retina where visual input enters the visual system. Connections are adapted according to the normalized Hebbian learning rule (Hebb 1949).

connections are also believed to play an important role in cortical development and recovery in the mature cortex (Gilbert and Wiesel 1992; Kapadia et al. 1994; Pettet and Gilbert 1992).

These experimental results convincingly show that the connections in the cortex are shaped by the environmental input, which in turn implies that any statistical regularities found in nature will be reflected in the structure of the visual cortex.

2.2.3 Computational Models of Self-Organization

von der Malsburg (1973) discovered that simple computational rules can drive the development of an ordered map of oriented receptive fields based on visual input. Kohonen (1981, 1982a,b) also showed how a topographically ordered map can arise in the sensory system through unsupervised learning (the Self-Organizing Map, or SOM model). These discoveries sparked the interest in self-organization as a basis of the development of cortical structures. Several computational models have since shown how receptive fields and their global organization in the cortical network can develop through self-organization of afferent synapses (Erwin et al. 1995; Goodhill 1993; Kohonen 1982b; Miller 1994; Miller et al. 1989; Obermayer et al. 1990; Piepenbrock et al. 1997; Shouval and Cooper 1996; Shouval et al. 1997). However, these models have not taken the lateral interactions between cells explicitly into account, or have assumed that they are preset and fixed and have a regular symmetric (or isotropic) profile. Only recently have laterally-connected models started to emerge (Bartsch and van Hemmen 2001; Kalarickal and Marshall 1997; Sirosh 1995). These models can potentially account for a wider set of developmental and functional phenomena than self-organizing map models without explicit lateral connections.

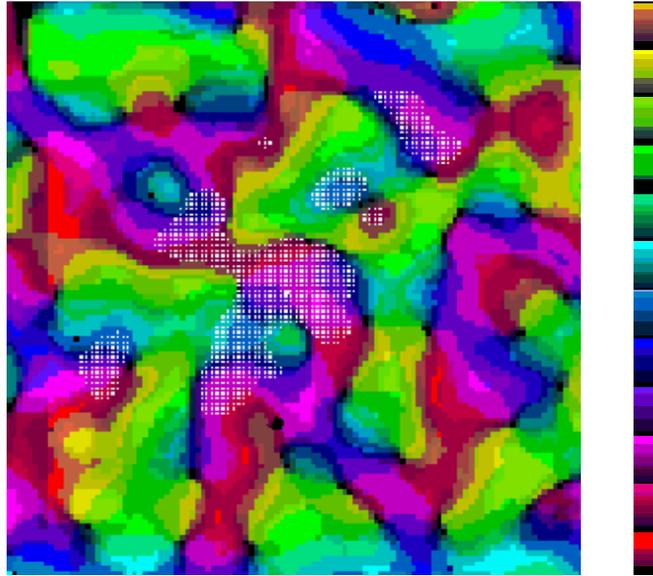


Figure 2.6: **Orientation map and lateral connections in RF-LISSOM (color figure).** The orientation preference (the orientation that gives maximal response) of each neuron is color-coded such that the transitions red \rightarrow magenta \rightarrow blue \rightarrow green \rightarrow orange represents 0 to 180 degrees. Saturation of each color represents orientation selectivity (how well it responds to the optimal orientation compared to other orientations). The white dots along the 60 degree axis represent long-range lateral inhibitory connections from the central neuron. The central neuron is tuned to inputs with 60 degree angle, and the lateral connection profile is oriented roughly along that axis, connecting to neurons with similar orientation preference. Such connection patterns could form the basis for Gestalt principles in perceptual grouping. From Sirosh et al. (1996b).

In the early stages of visual cortex development, lateral connections are believed to self-organize in synergy with the afferent connections to form a topological map of the input space (Burkhalter et al. 1993; Dalva and Katz 1994; Katz and Shatz 1996). This process can be modeled computationally, showing how structures such as oriented receptive fields, orientation columns, and patterned lateral connections form based on input-driven Hebbian learning process, where neurons that activate simultaneously gradually develop a stronger connection (Hebb 1949). The first such a comprehensive model was RF-LISSOM (the Receptive Field - Laterally Interconnected Synergetically Self-Organizing Map; Miikkulainen et al. 1997; Sirosh 1995; Sirosh and Miikkulainen 1996, 1997; Sirosh et al. 1996a; figure 2.5). Figure 2.6 shows the orientation map and patchy lateral connections that emerged in RF-LISSOM.

RF-LISSOM showed that the lateral connections are not only important in the development of delicate structures in the cortex, but that they also have a functionally significant role. The inhibitory lateral connections encode the correlation of activity in the map and perform redundancy reduction. Kurtosis measures of the activities before and after lateral interaction showed that the settled activity after lateral interaction is more sparse than before the lateral interaction. RF-LISSOM was also applied to hand-written digit recognition task, where the lateral connections

were shown to help form a better representation of the input space than a regular SOM, resulting in improved performance, and suggesting that lateral connections provide immediate functional advantage in high-level tasks (Choe 1995; Choe et al. 1996). Lateral connections were also shown to have a major influence in the response properties of the neurons after retinal and cortical lesions, thus playing a significant role in cortical reorganization as well.

Lateral connections may also play a direct role in forming visual representations. They can mediate synchronization and desynchronization of spiking activity, thus facilitating feature binding and segmentation. Such binding and segmentation can in turn serve as a basis for perceptual grouping by Gestalt principles. During self-organization, lateral connections learn correlations between activities in distant areas, and thereby provide a natural way for encoding Gestalt principles. For example, particularly strong lateral connection patterns form between neurons with similar orientation tuning, aligned along the preferred orientation axis, because these areas are often active together (figure 2.6). These connections could then be used for contour integration, in effect implementing the Gestalt law of good continuation.

2.2.4 Summary and Discussion

In this section, the input-driven self-organizing nature of development in the visual cortex was reviewed. The properties of self-organized lateral connections are suitable for complex perceptual grouping tasks, and can provide a physical substrate for synchrony. A computational model of grouping can be constructed by combining self-organizing map models with explicit lateral connections such as RF-LISSOM and the spiking neuron models. By studying the relationship between the lateral connections and the performance in perceptual grouping tasks in such a model, it will be possible to ground psychophysical phenomena to neurophysiology.

2.3 Computational Models of Contour Integration

In the previous section, the developmental and functional role of lateral connections was discussed. Adaptive lateral connections capture the Gestalt principles embedded in the input and thereby could form a basis for perceptual grouping. One concrete example of such perceptual grouping with ample psychophysical data is contour integration, as introduced in Chapter 1. In this section, computational models of contour integration will be reviewed, and the limitations of existing approaches will be discussed.

Several neural network models of contour integration have been developed, showing that specific lateral interactions are sufficient for contour integration (Li 1998, 1999; Yen and Finkel 1997, 1998). The models were able to detect and enhance smooth contours of oriented Gabor patterns embedded in a background of randomly-oriented Gabor patterns (as in figure 1.1*b*).

However, these models used fixed formulas in determining the interactions. For example, Yen and Finkel (1997, 1998) connected the units with long-range lateral excitatory connections.

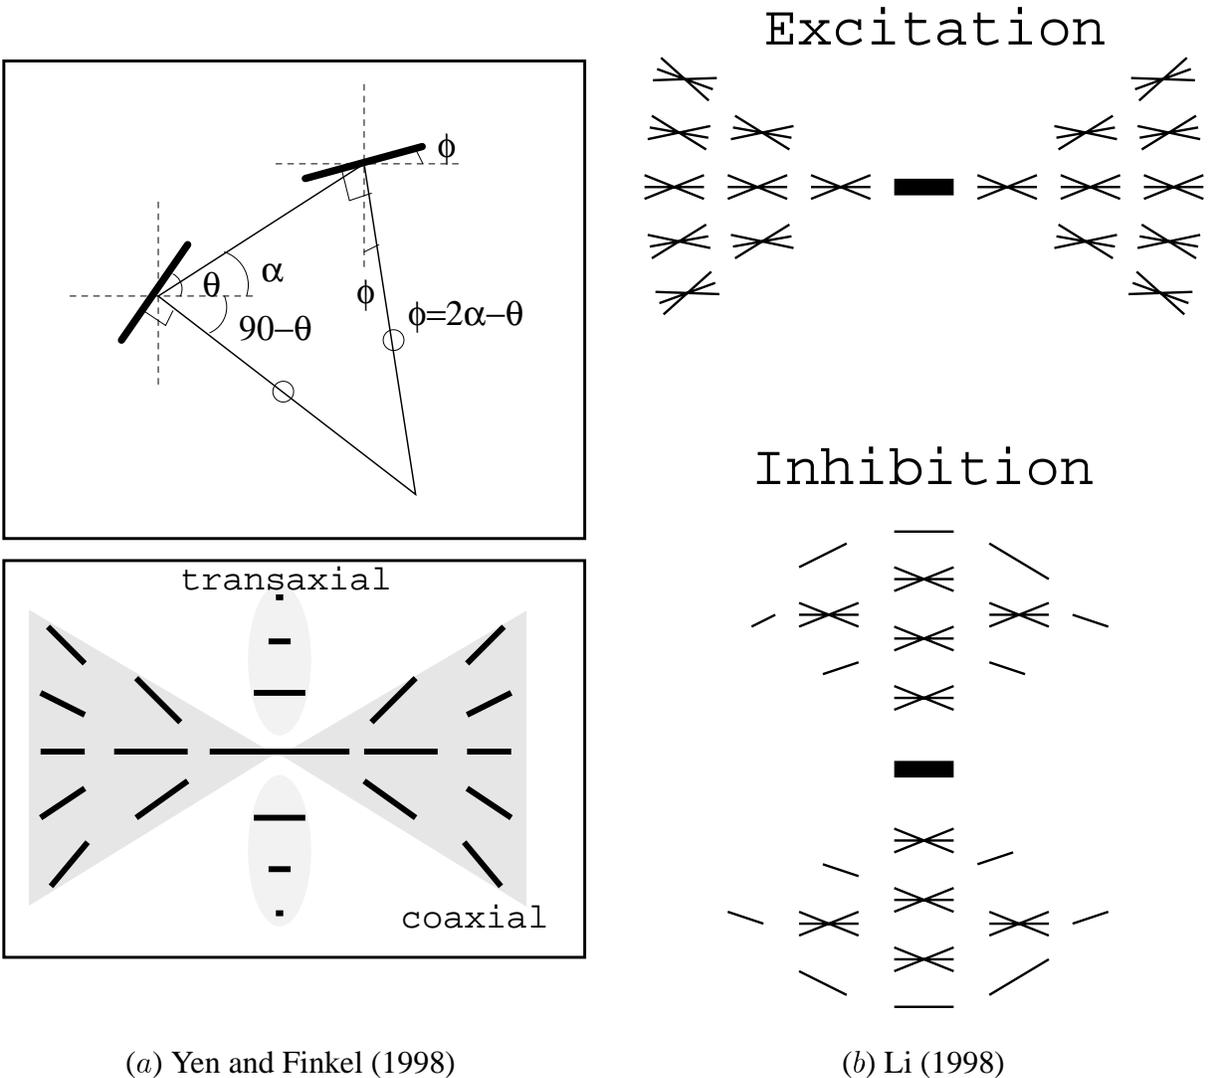


Figure 2.7: Lateral Connection Profiles. The lateral connection profiles of two computational models of contour integration are shown. (a) The model of Yen and Finkel (1998). The measures defining the relationship between edges are shown on the top. In the bottom figure, the co-axial (bow-tie shaped gray region) and the trans-axial (two oval shaped gray regions) excitatory connection schemes are shown. The length of the edges signifies the connection strength. Inhibitory connections were broad and isotropic, and are not shown here. (b) The model of Li (1998). The excitatory (top) and inhibitory (bottom) connections are shown. Excitatory connections go to co-linear and co-circular targets, while inhibitory connections go to areas on the flank on the top and bottom of the edge in the center. The general form of excitatory connections preferring co-circular arrangements can be seen in both connection schemes. These models, however, do not address how such connections could come about during cortical development. Adapted from Yen and Finkel (1998) and Li (1998), respectively.

The magnitude and time course of the synaptic interactions depended upon the position and orientation of the connected units. Excitatory connections were confined within two regions. One extended out to the axis of preferred orientation of the central unit, where co-circular connection scheme was used (co-axial). The other extended out to the flanks orthogonal to the preferred axis (trans-axial), and the area was smaller than in the co-axial case (figure 2.7a, bottom). Given the location of two units and the orientation preference of one unit, the preferred orientation of the second unit was calculated as $2\alpha - \theta$, where α is the slope of the straight line connecting the center of the two units, and θ is the orientation preference of the first unit (figure 2.7a, top). Inhibitory connections went to the rest of the surrounding neurons that did not receive excitatory connections. The connection strengths were gradually decreased by distance to have a Gaussian profile, with the highest at the center. The model was able to predict human contour integration performance, and showed that specific lateral interactions are necessary for contour integration.

Li (1998, 1999) took a different approach. Fixed rules derived from specific constraints defined the excitatory and inhibitory interactions as follows; (1) the system should not generate activity patterns spontaneously, (2) neurons at a region border should give relatively high responses, and (3) the same neural circuit should perform contour enhancement. Coupled oscillators were used to describe the dynamics of the orientation-selective cells, and mean-field techniques and dynamic stability analysis were used to calculate the lateral connection strengths and the connectivity pattern according to these three constraints. The resulting lateral connection strengths were very similar to that of Yen et al., except that there were no trans-axial excitation, and the inhibition had a specific shape instead of a broad profile (figure 2.7b). The model also predicted contour integration performance well, and again, the model showed that specific lateral connections are necessary to accurately predict human contour integration performance.

2.4 Open Issues in Contour Integration

Although the models described in the previous section have been successfully applied to explain experimental data, several questions remain: why do the seemingly different approaches result in similar connectivity patterns, how do these kinds of constraints develop in the brain in the first place, and what is their anatomical basis? It is possible that statistical regularities in the visual environment have properties similar to the connections in these models, and the brain may be encoding such properties. The computational approaches may have discovered the same properties while trying to accurately mimic human performance. Given such statistical properties in the input, an adaptive neural network that detects and encodes such properties may be able to explain why such patterns of connectivity emerge naturally. In such a network, statistical properties are implemented in the connections between neurons, and these could be the anatomical basis of perceptual grouping. Demonstrating these ideas computationally is the main goal of this thesis.

Moreover, the fixed lateral interaction approach cannot account for the differences in human contour integration performance across different parts of the visual field (Hess and Dakin 1997;

Rubin et al. 1996). Such a difference may be caused by difference in lateral interaction rules in the cortical areas, depending on the parts of the visual field they are mapping. Because the models described above are bound on fixed constraints, they result in uniform interaction patterns across the whole network and cannot explain such location-specific performance differences. However, if lateral connections that modulate contour integration could adapt according to the input, such results could be explained. If the input distribution differs in the different parts of the visual field, the lateral connections will develop different patterns to accommodate such differences in input. Such differences in anatomy can lead to differences in perceptual performance. This process can easily be modeled by the adaptive lateral connections in RF-LISSOM, as will be done in this thesis.

2.5 Conclusion

In this chapter, I have reviewed experimental evidence and computational model studies suggesting a possible neural mechanism of perceptual grouping: (1) temporal neurons represent binding and segmentation, and (2) self-organized connections encode activity correlation and input statistics. I also showed that existing models have limitations in explaining important experimental observations, and provided the motivation for an integrative model encompassing the known biological evidence and computational theories. In the next chapter, a model motivated by these previous work will be presented and described in detail.

Chapter 3

The PGLISSOM Model

The PGLISSOM (Perceptual Grouping LISSOM) is a two-layer model of the visual cortex integrating the functions of self-organization and grouping. The first layer is similar to RF-LISSOM, and it performs self-organization. The second layer performs grouping through long-range excitatory lateral connections that adapt to encode input correlations. Each layer is a two-dimensional map of spiking neurons, and temporal synchronization of spiking activity forms the representation for grouping. In this chapter, I will describe the architecture of PGLISSOM and its components in detail, and show how the network is initialized, activated, and trained.

3.1 Motivation

I will (1) review the previous work that led to the development of the PGLISSOM model, (2) discuss the shortcomings of the previous models, (3) propose a new architecture, and (4) provide biological justifications for the concepts and architectural components introduced in PGLISSOM.

PGLISSOM is based on RF-LISSOM, a laterally-connected self-organizing map model of the visual cortex (Sirosh 1995; Sirosh and Miikkulainen 1997; section 2.2.3). RF-LISSOM consists of firing-rate neurons represented as real-numbered units. As we saw in Chapter 2, superposition catastrophe can occur when activities of firing rate neurons representing separate objects are combined.

To overcome this limitation, the RF-SLISSOM model was developed (Choe and Miikkulainen 1997, 1998). The overall structure of RF-SLISSOM is similar to RF-LISSOM, but the firing-rate neurons are replaced with spiking neurons. RF-SLISSOM model showed how self-organization and segmentation of simple objects could be achieved in a unified model of self-organization and temporal coding. The long-range inhibitory lateral interactions play a crucial role in both behaviors: they establish competition that drives self-organization, and they establish desynchronization that drives segmentation. The model did not include any long-range excitatory connections because they were not found necessary to model self-organization and simple segmentation of objects.

However, to perform grouping on more complex features such as long contours, long-range excitatory lateral connections are necessary. Excitation is necessary for synchronization, and connections need to be long to group together several contour elements. Making the short-range excitatory lateral connections longer in a single RF-SLISSOM map does not work, because it results simply in a large active area in the map. The neurons would respond to any input through lateral excitation, resulting in a map where the afferent receptive field of the neurons look almost the same due to averaging over the inputs. The two functions of perceptual grouping and self-organization therefore have conflicting requirements (long vs. short excitatory lateral connections).

In order to model both functions simultaneously, the PGLISSOM model described in this thesis was developed. PGLISSOM includes both short-range and long-range excitatory lateral connections. To prevent the long-range excitatory connections from interfering with the self-organizing process, PGLISSOM includes two layers (or maps). In the first map (MAP1), excitatory lateral connections are short range to allow self-organization and in the second map (MAP2), they are long range and implement perceptual grouping.

The two-layered design of PGLISSOM is biologically motivated. The two layers model the known connectivity in the visual cortex. As shown in figure 2.3, the visual cortex consists of six distinct layers. In the deeper layers, afferent inputs reaching layer 6 in turn innervate layer 4 in an on-center off-surround fashion (Ferster and Lindström 1985; Grieve and Sillito 1995b). MAP1 in PGLISSOM abstracts the center-surround connectivity into short-range excitatory and long-range inhibitory lateral connections in a single map. In layer 2/3 of the visual cortex, two types of long-range lateral connections are found: direct excitatory connections and disynaptic inhibitory connections (Hirsch and Gilbert 1991; McGuire et al. 1991). MAP2 models these long-range connections. In the cortex, connections are also found along the vertical section between layers 2 through 6. These are termed intra-columnar connections because they are concentrated in cortical columns. Such connections are included in PGLISSOM as well, and they allow MAP2 to self-organize properly in the presence of long-range excitation.

In sum, PGLISSOM extends RF-LISSOM in two important ways: spiking neurons for representing grouping through temporal coding, and long-range excitatory lateral connections to coordinate synchronization for temporal coding. The design models the layered architecture found in the visual cortex. In the following sections, I will present the overall organization of the PGLISSOM model and describe the components of the model in detail.

3.2 Network Organization

The overall organization of the PGLISSOM model is shown in figure 3.1. The model consists of two layers (or maps), one overlaid (or stacked) on top of the other. Both maps are based on RF-SLISSOM, but the extent of lateral connections in the two maps differ.

The two maps receive afferent input from the retinal neurons. In the brain, inputs received in the retina first project to the lateral geniculate nucleus (LGN; the vision relay cluster in the

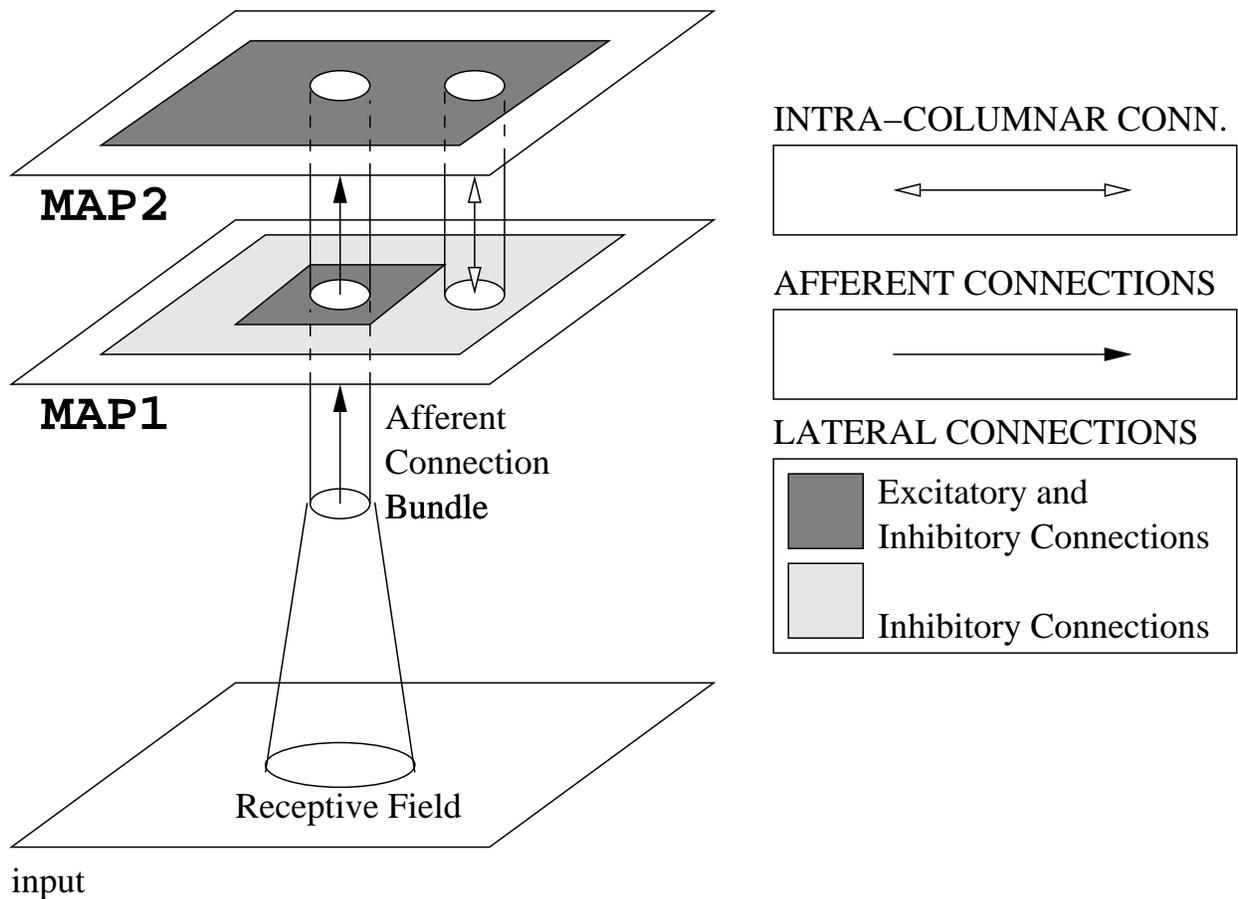


Figure 3.1: **The Overall Organization of the PGLISSOM Network.** The cortical network consists of two layers (or maps): the lower map (MAP1) has short-range excitation and long-range lateral inhibition, and drives the self-organization of the model. In the upper map (MAP2), both excitation and inhibition have a long range, establishing segmentation and binding. The two maps receive input from a model retina, and neurons in the vertically corresponding locations on the two maps are connected via intra-columnar connections, tying such neurons together into a functional unit (i.e. a cortical column).

thalamus), and then to the visual cortex, but in PGLISSOM, LGN is bypassed for simplicity. The lower map, MAP1, has short-range excitatory lateral connections and long-range inhibitory lateral connections. The excitatory connections establish a local neighborhood that enforces local correlation, while the inhibitory connections decorrelate distant activities and thus perform redundancy reduction (Barlow 1985, 1994; Sirosh 1995). Through these connections, MAP1 drives self-organization in the model. The upper map, MAP2, has both long-range excitatory and long-range inhibitory connections. The excitatory connections in this map form the basis for perceptual grouping: through the self-organizing process, they learn to encode correlations in the input distribution, and the strength of the connections controls the degree of synchronization across the neurons. The inhibitory connections are broad and long-range, causing two or more synchronized populations of neurons to desynchronize, thus establishing background inhibition for segmentation.

Neurons in the vertically corresponding locations in the two maps form a functional unit (cortical column), and they are connected to each other through vertically projecting intra-columnar connections. These connections influence the activity on the opposite map so that both self-organization and grouping behaviors are shared by both maps.

Next, the details of the spiking neuron model in PGLISSOM will be presented.

3.3 Neuron Model

A schematic diagram of the spiking neuron is shown in figure 3.2. The model is based on that of Eckhorn et al. (1990) and Reitboeck et al. (1993). Each neuron has three components: leaky synapses, weighted summation, and a spike generator. The synapses continuously calculate the decayed sum of incoming spikes over time. Four different kinds of input connections contribute to the weighted sum: afferent, excitatory lateral, inhibitory lateral, and intra-columnar connections (figure 3.1; both excitatory and inhibitory connections are just shown as *lateral connections*). The activations of the different kinds of inputs are summed and compared to the dynamic threshold in the spike generator. A spike is generated if the activations exceed a dynamic threshold. When the neuron is silent, the threshold remains at the base level. As the neuron fires, the threshold increases, modeling the refractory period in biological neurons. In addition, the neuron is forced to be silent for several time steps immediately following a spike. Details of each component will be discussed next.

3.3.1 Leaky Synapse

Each connection is a leaky integrator that continuously calculates an exponentially decayed sum of incoming spikes:

$$s(t) = \sum_{n=0}^t x(t-n)e^{-\lambda n}, \quad (3.1)$$

where $s(t)$ is the current decayed sum at time step t , $x(t-n)$ is the input spike (either 0 or 1) n time steps in the past, and λ is the decay rate. Different types of connections have separate decay rates: afferent connections (λ_a), excitatory lateral connections (λ_e), inhibitory lateral connections (λ_i), and intra-columnar connections (λ_c). The most recent input has the most influence on the activity, but inputs from past also have some effect. This sum can be defined recursively as:

$$s(t) = x(t) + s(t-1)e^{-\lambda}, \quad (3.2)$$

where $s(t)$ and $s(t-1)$ are the current and previous decayed sums, $x(t)$ is the current input spike and λ is the decay rate ¹. Such a formulation allows for a computationally efficient implementation

¹This equation has the same form as the one derived from convolution equations in Eckhorn et al. (1990; see equation 2.4).

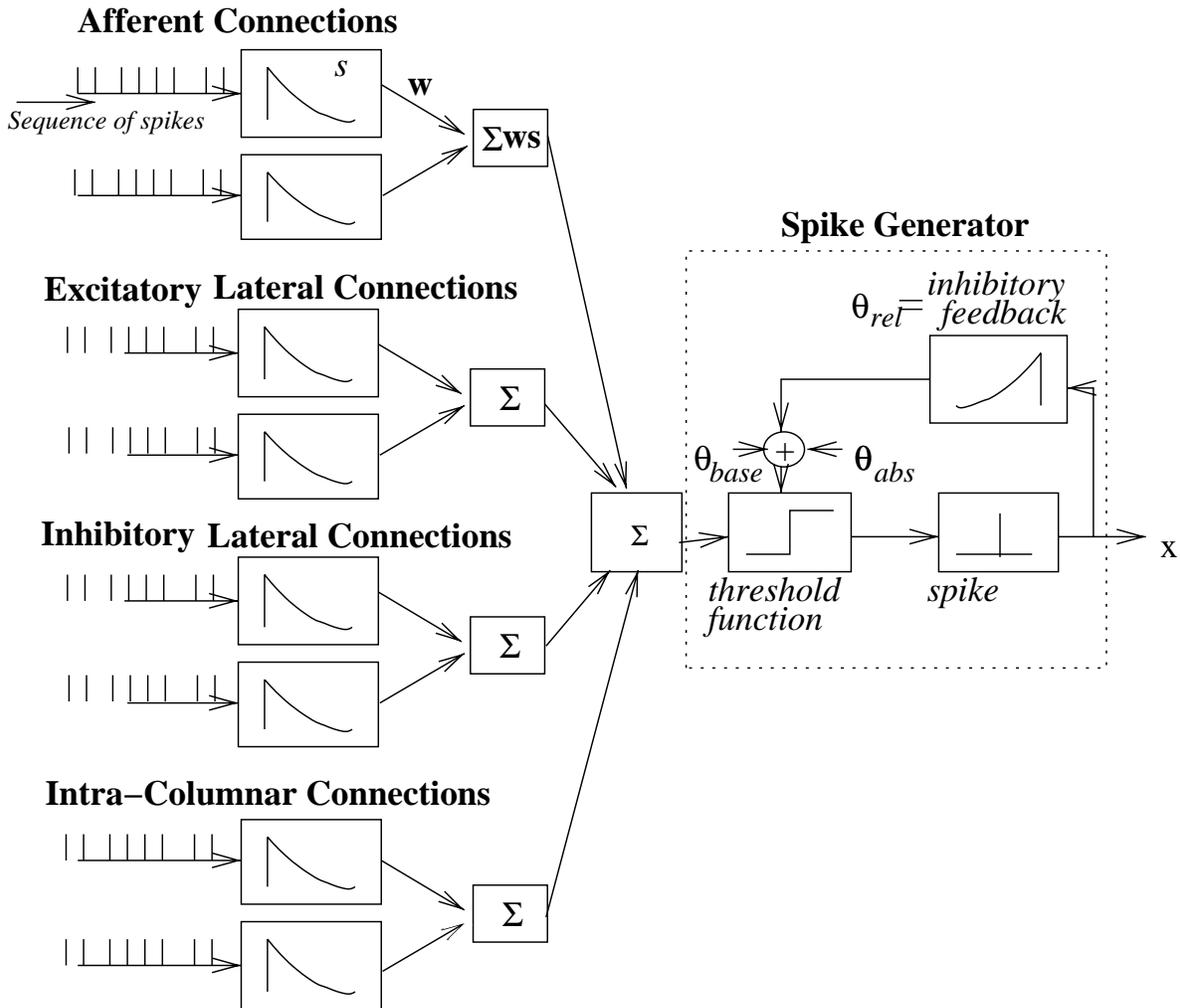


Figure 3.2: **The Neuron Model.** Leaky integrators at each synapse perform decayed summation of incoming spikes, and the outgoing spikes are generated by comparing the weighted sum to the dynamic spiking threshold. Four types of inputs contribute to the activity: afferent, excitatory lateral, inhibitory lateral, and intra-columnar connections. The dynamic threshold consists of the base threshold θ_{base} , the absolute refractory contribution θ_{abs} , and the relative refractory contribution θ_{rel} . The base threshold has a fixed baseline value, and the absolute refractory term has a value of ∞ for a short time period immediately following an output spike. The relative refractory contribution is increased as output spikes are generated, and it decays to 0 if the neuron stays silent.

of the model, since the past spike values $x(\cdot)$ do not need to be stored, nor do they have to be decayed repeatedly.

Such a leaky integrator models the Post-Synaptic Potential (PSP) that decays exponentially over time in biological neurons. The formulation is equivalent to the finite difference approximation of the widely used leaky integrate-and-fire neurons (equation 2.3). By adjusting the decay rate λ , the synapse can function as either a coincidence detector or as a temporal integrator. When the synaptic decay rate is high, the neuron can only activate when there is a sufficient number of inputs coming in from many synapses simultaneously. On the other hand, when the decay rate is low, the neuron accumulates the input. Thus pre-synaptic neurons can have a lingering influence on the post-synaptic neuron. By varying the decay rates for different types of connections, the relative time scales of the different connection types can be controlled to obtain desirable synchronization behavior (see Chapter 4).

The neuron receives incoming spikes through many input connections, and the decayed sums of the synapses are calculated according to equation 3.2. These sums are then multiplied by the connection weights and summed to obtain the input activity. Next, I will explain how the weighted sums of different connection types are combined.

3.3.2 Calculating the Input Activation

The input activation $\sigma_{i,j}(t)$ to the spike generator of the cortical neuron at location (i, j) at time t consists of (1) the input from a fixed-size receptive field in the retina, centered at the location corresponding to the neuron's location in the cortical network (i.e. afferent input), (2) from neurons in the same column in the other map (i.e. intra-columnar input), (3) excitation from neurons around it in the same map (excitatory lateral input), and (4) inhibition from neurons around it in the same map (inhibitory lateral input):

$$\begin{aligned} \sigma_{i,j}(t) = & g\left(\gamma_a \sum_{r_1, r_2} \xi_{r_1, r_2} \mu_{ij, r_1 r_2} + \gamma_c \sum_{p_1, p_2} \zeta_{p_1, p_2} \nu_{ij, p_1 p_2} \right. \\ & \left. + \gamma_e \sum_{k, l} \eta_{kl}(t-1) E_{ij, kl} - \gamma_i \sum_{k, l} \eta_{kl}(t-1) I_{ij, kl}\right), \end{aligned} \quad (3.3)$$

where $\gamma_a, \gamma_c, \gamma_e$, and γ_i are the relative strengths of the afferent, intra-columnar, excitatory, and inhibitory contributions, ξ_{r_1, r_2} is the activity at the retinal neuron (r_1, r_2) , $\mu_{ij, r_1 r_2}$ is the afferent connection weight from retinal neuron (r_1, r_2) to cortical neuron (i, j) , ζ_{p_1, p_2} is the decayed sum of spikes from neuron (p_1, p_2) in the other map in the same vertical column, $\nu_{ij, p_1 p_2}$ is the intra-columnar connection weight from neuron (p_1, p_2) to neuron (i, j) , $\eta_{kl}(t-1)$ is the decayed sum of spikes from the map neuron (k, l) at time $t-1$, $E_{ij, kl}$ is the excitatory and $I_{ij, kl}$ the inhibitory lateral connection weight between neuron (i, j) and (k, l) . The function $g(\cdot)$ is a bounded linear function

$$g(x) = \begin{cases} 0 & \text{if } x < \delta \\ 1 & \text{if } x > \beta \\ \frac{x-\delta}{\beta-\alpha} & \text{otherwise,} \end{cases} \quad (3.4)$$

where δ is the threshold and β is the ceiling. This function is used to keep the input activity to the spike generator between 0.0 and 1.0.

The input activity is then passed on to the spike generator, where comparison with the dynamic threshold is made, and a spike is fired if the input activity exceeds the threshold. Next, I will describe the components of the dynamic threshold mechanism in the spike generator.

3.3.3 Threshold Mechanism

Biological neurons cannot generate a spike during a short period of time immediately after they have spiked. This short interval is called the *refractory period* and it consists of two parts: (1) during the *absolute* refractory period, the neurons cannot fire no matter how large the input is, and (2) during the *relative* refractory period, neurons normally cannot spike, but if sufficient input is received. they can generate spikes

The dynamic threshold in the spike generator implements a refractory period by providing a base threshold and raising the threshold dynamically, depending on the neuron's spike activity. The spike generator compares the input activity to the dynamic threshold and decides whether to fire a spike (figure 3.2). The threshold $\theta(t)$ is a sum of three terms:

$$\theta(t) = \theta_{\text{base}} + \theta_{\text{abs}}(t) + \tau\theta_{\text{rel}}(t), \quad (3.5)$$

where θ_{base} is the base threshold, $\theta_{\text{abs}}(t)$ represents the contribution of the absolute refractory period, $\theta_{\text{rel}}(t)$ is contribution of the current relative refractory period, and τ balances the two contributions.

The relative refractory contribution is an exponentially decayed sum of the output spikes (figure 3.2), which is a leaky integrator similar to the leaky synapses. It can be defined recursively as follows:

$$\theta_{\text{rel}}(t) = y(t) + \theta_{\text{rel}}(t-1)e^{-\lambda_{\text{rel}}}, \quad (3.6)$$

where $\theta_{\text{rel}}(t)$ is the current relative refractory contribution, $y(t)$ is the current output spike (1 if the neuron spikes at this time, 0 otherwise), $\theta_{\text{rel}}(t-1)$ is the relative refractory contribution from the previous time step, and λ_{rel} is the decay rate. The absolute refractory contribution $\theta_{\text{abs}}(t)$ is defined as:

$$\theta_{\text{abs}}(t) = \begin{cases} \infty & \text{if } y(t-i) = 1 \text{ for } i \leq \kappa_{\text{abs}} \\ 0 & \text{otherwise,} \end{cases} \quad (3.7)$$

where κ_{abs} determines the length of the absolute refractory period.

The threshold is compared to the weighted sum of inputs and a spike is generated if the input exceeds the threshold:

$$y(t) = \mathcal{H}(\sigma(t) - \theta(t-1)), \quad (3.8)$$

where $y(t)$ is the output spike at time t , $\sigma(t)$ is the current input activity, $\theta(t-1)$ is the previous

dynamic threshold, and $\mathcal{H}(\cdot)$ is the Heaviside step function:

$$\mathcal{H}(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{otherwise.} \end{cases} \quad (3.9)$$

The spikes generated this way are propagated through the connections and the average firing rates of the neurons in a small time window are gathered. Based on these firing rates, the connection weights are adapted. Next, the details on activating and training the model will be explained.

3.4 Self-Organization

All connection weights in the network are initialized to uniformly random values within 0.0 and 1.0. The network is trained by presenting visual input, and adapting the connection weights according to the Hebbian learning rule. In the experiments reported in this thesis, synthetic visual inputs are used to have tight control over stimulus configurations. It is also possible to train the network with natural images as has been done with RF-LISSOM (Bednar and Miikkulainen 2000; this possibility will be discussed in detail in section 10.1).

The input to the network consists of oriented Gaussian bars, defined as:

$$\xi_{r_1, r_2} = \exp\left(-\frac{((r_1 - x)\cos(\phi) - (r_2 - y)\sin(\phi))^2}{a^2} - \frac{((r_1 - x)\sin(\phi) + (r_2 - y)\cos(\phi))^2}{b^2} \right), \quad (3.10)$$

where ξ_{r_1, r_2} is the desired activity of the retinal neuron at location (r_1, r_2) , a^2 and b^2 specify the length along the major and minor axes of the Gaussian, and ϕ specifies its orientation. To generate a Gaussian bar, the input neurons spike at different frequencies. In the center of the Gaussian, the spike rate is maximal, and the spike rate would decrease gradually for neurons farther away from the center. However, since the inputs are always precisely specified and such a generation of a Gaussian is computationally expensive in the current simulations, the spiking input neurons were replaced with real-valued neurons whose values represent the leaky sums at the afferent input synapses.

An input with a random orientation is placed at a random location in the retina, and the activity levels of the retinal neurons are transmitted through the afferent connections of the network. The cortical neurons of MAP1 and MAP2 then allowed to generate spikes and propagate spikes, according to equations 3.1 through 3.8.

The connection weights in the network are adaptable, modified according to the short-term firing rates of the neurons. In the standard RF-LISSOM model, the input is kept constant while the cortical response settles through the lateral connections, forming a concentrated, redundancy-reduced activation pattern (as measured by the kurtosis of the map activity; Field et al. 1993a;

Sirosh 1995). PGLISSOM goes through a similar settling process. The retinal neurons generate graded activity constantly at each iteration and the cortical neurons generate and propagate spikes. After a while, the neurons reach a stable rate of firing, and this rate is used to modify the weights. The average spiking rate $V(t)$ of neurons was calculated as a running average of spikes as follows:

$$V(t) = \tau_{avg}V(t-1) + (1 - \tau_{avg})y(t), \quad (3.11)$$

where τ_{avg} is the retention rate, $V(t-1)$ is the previous average firing rate, and $y(t)$ is the output spike at time t (either 0 or 1). With this method, a short-term firing rate in a limited time window is calculated. The averaging window size is proportional to the parameter τ_{avg} .

For each input presentation, the average spiking rate of each neuron is calculated through several iterations. Then the afferent, lateral, and intra-columnar weights are modified according to the normalized Hebbian learning rule:

$$w_{ij,mn}(t) = \frac{w_{ij,mn}(t-1) + \alpha V_{ij}(t)V_{mn}(t)}{\sum_{mn} [w_{ij,mn}(t-1) + \alpha V_{ij}(t)V_{mn}(t)]}, \quad (3.12)$$

where $w_{ij,mn}(t)$ is the connection weight from neuron (m, n) to (i, j) , $w_{ij,mn}(t-1)$ is the previous weight, α is the learning rate (α_a for afferent, α_c for intra-columnar, α_e for excitatory, and α_i for inhibitory connections), and $V_{ij}(t)$ and $V_{mn}(t)$ are the current average spiking rates of the neurons (i, j) and (m, n) . The adapted weights are normalized (i.e. divided) by the sum of outgoing connection weights. Those connections that become near zero in this process are deleted, modeling death of unused connections during early development in animals (Dalva and Katz 1994; Gilbert 1992; Katz and Callaway 1992; Löwel and Singer 1992). The radius of the lateral excitation in MAP1 is gradually reduced following a preset schedule, resulting in fine tuning of the map (for a theoretical motivation for this process, see Kohonen 1982b, 1989, 1993; Obermayer et al. 1992; Sirosh and Miikkulainen 1997; for neurophysiological evidence, see Dalva and Katz 1994; Hata et al. 1993).

Weight normalization is necessary in Hebbian learning to keep the weights from growing into infinity (Miller and MacKay 1994), but it can also be viewed as a process of redistributing the synaptic resources of each neuron (Choe et al. 2000). Another possible normalization process is neuronal regulation (Horn et al. 1998). In recent experiments, Turrigiano et al. (1998) showed that a change in a single synapse can cause the whole cell's efficacy to change in the opposite direction. This result suggests that local change in the synaptic strength scales the strength of the other synapses of the same neuron. Normalization in PGLISSOM is done this way by keeping the sum of connections weights constant, effectively redistributing limited synaptic resource.

This process of input presentation, activation, and weight adaptation is repeated for a large number of input patterns, and the neurons become gradually sensitive to particular orientations at particular locations. This way, the network forms a global retinotopic orientation map similar to that in the visual cortex (Blasdel 1992b; Blasdel and Salama 1986). The orientation preferences in the map arise as a result of Hebbian adaptation on oriented Gaussian inputs. The self-organized map

with its patchy lateral connections will then synchronize and desynchronize the firing of neurons to indicate binding and segmentation of visual input into different objects.

3.5 Summary

The PGLISSOM model is a biologically inspired model of the visual cortex based on the RF-LISSOM architecture. PGLISSOM has two functionally distinct maps, one for self-organization (MAP1), and another for binding and segmentation (MAP2). The firing rate neurons in RF-LISSOM are replaced with spiking neurons to represent grouping through temporal coding, and long-range excitatory connections are added in MAP2 to coordinate synchrony in populations of neurons. In the following chapters, the results from PGLISSOM experiments will be reported.

Chapter 4

Temporal Coding

Experimental evidence suggests that temporally correlated activity may be the basis for binding and segmentation in perceptual grouping (Eckhorn et al. 1988; Engel et al. 1991b; Gray and Singer 1987; Gray et al. 1989). In PGLISSOM, the neurons are modeled as spiking neurons so that the temporal firing sequence can encode such grouping information. It is important to understand what each component of the model contributes to synchrony and desynchrony so that the model as a whole can be tuned to function properly, and also to gain crucial insights to guide further experimental work. In this chapter, the neuron model of PGLISSOM will be tested to find the conditions under which synchronization and desynchronization occurs under different (1) connection types (excitatory or inhibitory), (2) synaptic decay rates, (3) spatial extent of connections, (4) noise levels, (5) population sizes, and (6) the durations of the absolute refractory period.

4.1 Synchronization

Synchronization is important for binding together populations of neurons that represent input features of the same coherent object. Lateral connections are necessary to coordinate the firing of neurons since without any exchange of information, isolated neurons cannot achieve synchrony. Computational experiments were conducted to find out the conditions under which synchronization occurs. Only one-dimensional networks were considered in this analysis. They are sufficient for testing the various factors governing synchronization, and they are also easier to visualize them in 2-D because the activities can be plotted against time in two dimensions.

4.1.1 Synaptic Decay Rate

Previous models of spiking neurons have either adapted or selected the axonal delays to regulate synchronization behavior (Eurich et al. 2000; Gerstner 1998a; Horn and Opher 1998; Nischwitz and Glünder 1995; Tversky and Miikkulainen 2002). Although delay can be adapted by changing the axonal morphology (length, thickness, and myelination; Eurich et al. 1999), the fine degree of

delay tuning needed in the above models may not be easy to achieve in such a macro structure of the biological neuron.

An alternative to delay adaptation is changing the decay rate of the post-synaptic potential. Decay may be easier to alter in biological neurons since ion channels can be added or removed to tune the leakage of currents through the cell membrane. The number and distribution of ion channels can change through various mechanisms including activity-dependent gene expression and activity-dependent modulation of assembled ion channels (see Abbott and Marder 1995 for a comprehensive review). Nowak and Bullier (1997) studied various mechanisms of decay (or integration time), and further investigations of these mechanisms may well reveal how decay rate can be controlled. Some models already utilize synaptic decay (Eckhorn et al. 1990; Reitboeck et al. 1993), but the influence of different levels of decay on synchronization has not been fully tested.

In PGLISSOM, the neurons have decaying synaptic potentials, and the decay rate was found to strongly influence synchronization. The values of the synaptic decay rate λ (equation 3.1) can differ for different types of connections (excitatory or inhibitory), and the same magnitude of λ results in different behavior depending on whether the connection is excitatory or inhibitory.

For the synaptic decay rate experiments, as well as for the other experiments in this chapter, the following five conditions were assumed unless otherwise noted: (1) one-to-one connectivity between input and output neurons, (2) all input neurons activated to 1.0 at every time step, (3) fixed weights, (4) random initialization of membrane potential, and (5) the parameter values $\delta = 0.0$, $\beta = 3.0$, $\tau = 0.65$, and $\theta_{\text{base}} = 0.1$ (see Chapter 3 for the definitions of the parameters).

Four separate experiments were conducted: A one-dimensional network of 30 neurons with full lateral connections was simulated for 500 iterations. (1) excitatory lateral connections with slow decay ($\lambda_e = 0.1$), (2) inhibitory lateral connections with slow decay ($\lambda_i = 0.1$), (3) excitatory lateral connections with fast decay ($\lambda_e = 1.0$), and (4) inhibitory lateral connections with fast decay ($\lambda_i = 1.0$). Except for the decay rate λ and the connection type, all other parameters were the same in the four experiments: $\gamma_a = 0.8$, $\gamma_i = \gamma_e = 0.01$, and $\lambda_{\text{rel}} = 0.05$.

The results are shown in figure 4.1. Two conditions, excitatory connections with fast decay and inhibitory connections with slow decay, result in synchrony. In contrast, excitatory connections with slow decay and inhibitory connections with fast decay result in desynchrony. This is an interesting result since excitation does not always guarantee synchronization, and inhibition does not always guarantee desynchronization. Nischwitz and Glünder (1995) showed that a similar result is obtained by varying the degree of delay among integrate-and-fire neurons connected via excitatory or inhibitory connections, where short delay with excitatory connections and long delay with inhibitory connections caused the neurons to synchronize, and in the opposite case, to desynchronize. The current result indicates that instead of depending on delay for controlling synchronization behavior, synaptic decay can be utilized. Although synchronization can be achieved through slowly decaying inhibitory connections, excitatory connections are more likely to be responsible for coherent oscillation in the cortex since the faster firing rate in the excitatory case more

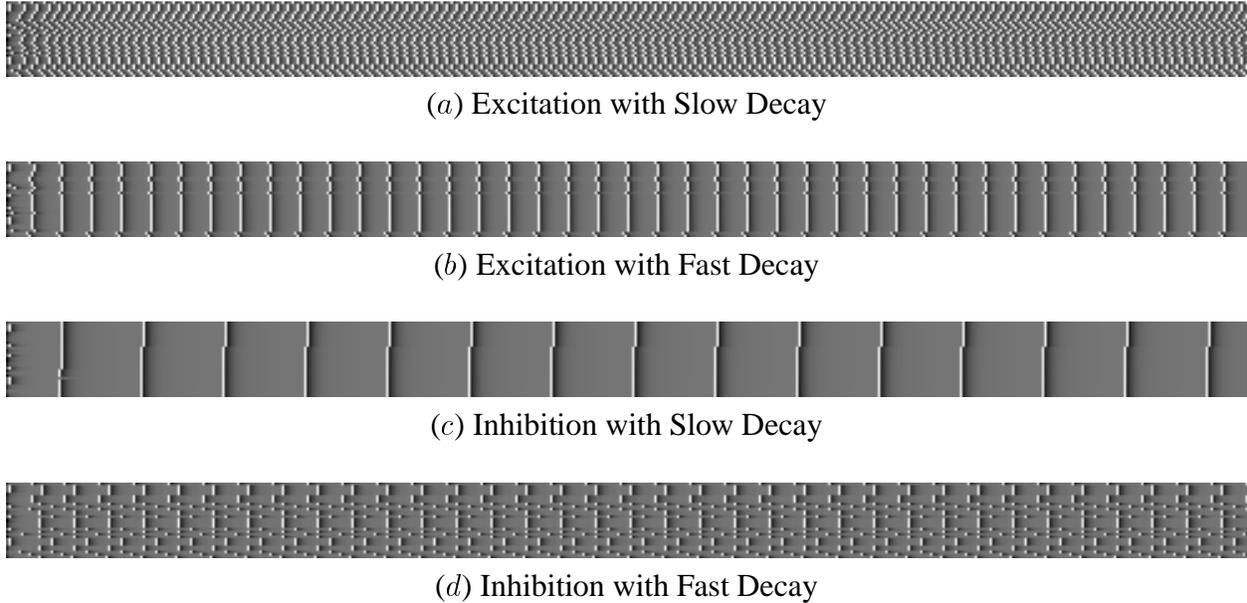


Figure 4.1: Effect of Connection Type and Decay Rate on Synchronization. Thirty neurons with full lateral connections were simulated for 500 iterations. Four experiments were conducted where the type of the lateral connections (excitatory or inhibitory), and the synaptic decay rates (λ) were altered. All other parameters were the same for all four cases. (a) Excitatory connections with slow decay ($\lambda_e = 0.1$) result in desynchronized activity. (b) Excitatory connections with fast decay ($\lambda_e = 1.0$) result in synchronized activity. (c) Inhibitory connections with slow decay ($\lambda_i = 0.1$) result in synchronized activity. (d) Inhibitory connections with fast decay ($\lambda_i = 1.0$) result in synchronized activity. Note that in the two synchronized cases *b* and *c*, the firing rate is higher in *b* because approach to threshold is faster due to higher levels of input through excitation. The results show that synchronization behavior can greatly vary even for the same connection type if the synaptic decay rate differs.

closely matches the fast gamma-frequency oscillation (around 40Hz) observed in the experiments. The rate is faster because the input activity level is higher due to the excitation.

In sum, synaptic decay plays an important role in regulating synchronization, and it is an attractive alternative to models with delay modulation since adapting the decay rate appears to be more biologically plausible.

4.1.2 Effect of the Connection Range

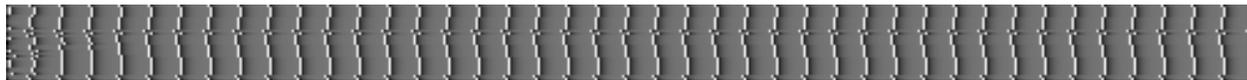
The network was tested to see if local excitatory connections can synchronize a global population. Inhibitory lateral connections were excluded to simplify the experiment. Thirty neurons with varying degrees of excitatory lateral connection radii were simulated for 500 iterations. Five separate experiments were conducted with excitatory connection radii of 30, 10, 5, 2, and 0. Other simulation conditions were the same as in section 4.1.1, except for $\gamma_e = 0.01$ and $\lambda_e = 5.0$, so that the network would synchronize under the smaller radii. As we will see below, this change did not



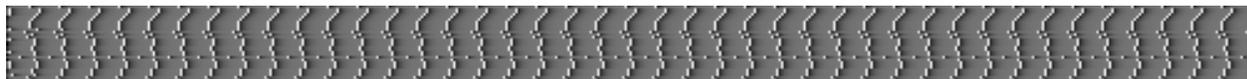
(a) Full connection (connection radius = 30)



(b) Connection radius = 10



(c) Connection radius = 5



(d) Connection radius = 2



(e) No excitatory connections

Figure 4.2: Effect of Connection Extent on Synchronization. A network of 30 neurons with varying extent of lateral excitatory connections was simulated for 500 iterations. Synchronization occurs through the excitatory connections even though the connections did not cover the whole network. From (a) to (e), the lateral excitatory connection radius was reduced from full connectivity (radius = 30) down to no connections at all (radius = 0). All other parameters were the same as before. Synchronization starts to break once the radius reaches 2, but for a fairly local connection radius (down to 5), global synchronization is maintained. As expected, with no connections at all (e), the initial random order of spikes is maintained throughout the simulation.

affect synchrony in the globally connected case.

The results are shown in figure 4.2. Global synchronization is achieved not only in the fully connected network (radius 30), but also in locally connected networks, down to a radius of 5. These results demonstrate that synchronization can propagate through locally connected neurons, which is consistent with other coherent oscillation models with local connections (Terman and Wang 1995; Wang 1995, 1996; Campbell et al. 1999). It shows that synchronization may work as a basis for transitive grouping (Geisler et al. 1999; Geisler and Super 2000; Geisler et al. 2001); if A and B are grouped together and B and C are grouped together, then A and C are perceptually grouped together.

In sum, fully connected networks synchronize well, but it is not necessary to have full connectivity to achieve global synchrony. Global synchronization through local connections in the PGLISSOM model may be a possible mechanism for transitive perceptual grouping.

4.2 Desynchronization

Desynchronization is as important as synchronization since it is the basis for segmentation. While synchronized firing binds together features that belong to a coherent object, desynchronized firing indicates that features belong to different objects. In this section, I will show that inhibitory connections are necessary for segmentation, and that a small amount of noise is necessary for symmetry breaking.

4.2.1 Connection Types

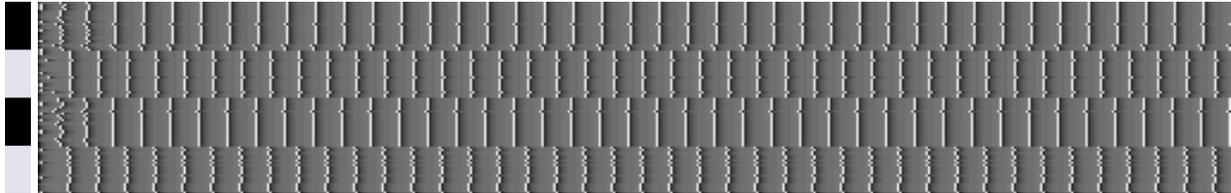
In the previous section, I showed how excitatory connections can facilitate synchronization while inhibitory connections can facilitate desynchronization under high decay rate conditions. For perceptual grouping through temporal coding, both synchronization and desynchronization are necessary, thus a network may require both excitatory and inhibitory lateral connections. In this section, I tested this hypothesis in a one-dimensional network to see if indeed both excitatory and inhibitory connections are necessary, and also to verify that including both types of connections results in a desirable temporal representation for binding and segmentation.

A one-dimensional network of 90 neurons divided into two groups was simulated for 500 iterations. Neurons with indices in the range (1..22) and (43..64) formed the first group, and those indexed (21..42) and (65..90) the second group. Lateral excitatory connections were only allowed to connect neurons within the same group, while inhibitory connections connected the whole population.

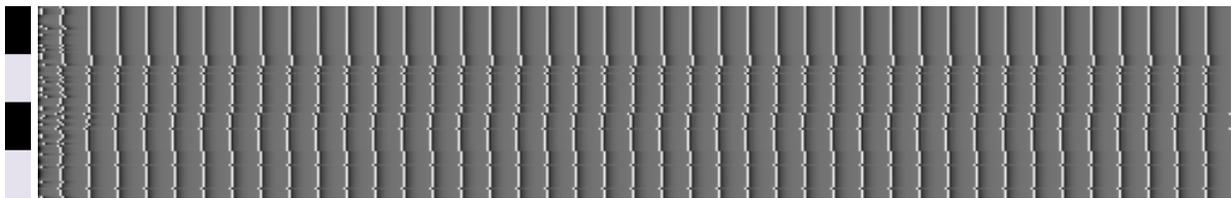
Four separate experiments were conducted: one with both excitatory and inhibitory connections, another with excitatory connections only, the third with inhibitory connections only, and the fourth with no lateral connections at all. Other simulation conditions were the same as in section 4.1.1, except for $\gamma_e = 0.36$ and $\lambda_e = 5.0$, to compensate for the larger size of the network and the addition of inhibitory connections, values $\gamma_i = 0.42$ and $\lambda_i = 5.0$ were used for the inhibitory connections.

The results are shown in figure 4.3. With both excitatory and inhibitory connections, neurons within the same group are synchronized, but across the groups where only inhibitory connections exist, desynchronization occurs (*a*). Such temporal representation is well-suited for perceptual grouping, since binding is signaled by synchrony and segmentation is signaled by desynchrony. Next, I checked if indeed both excitatory and inhibitory connections are necessary to obtain such a behavior. When only the excitatory connections are present, segmentation does not occur (*b*), and when only the inhibitory connections are present, binding does not occur (*c*). Needless to say, without any lateral connections, the neurons fire in random phases (*d*).

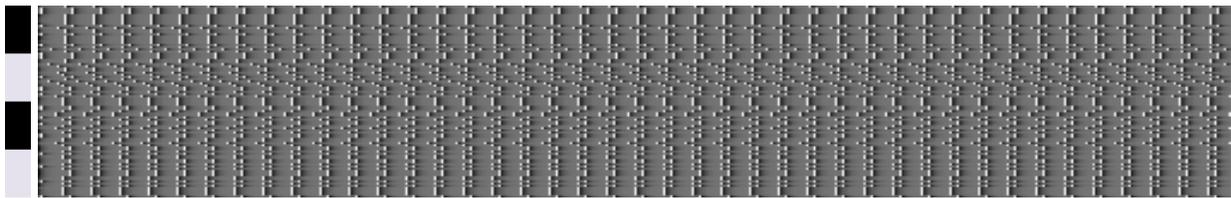
In sum, binding and segmentation can be established in a network with both excitatory and inhibitory lateral connections. Omitting either kind of the lateral connections results in losing the ability to bind, segment, or both.



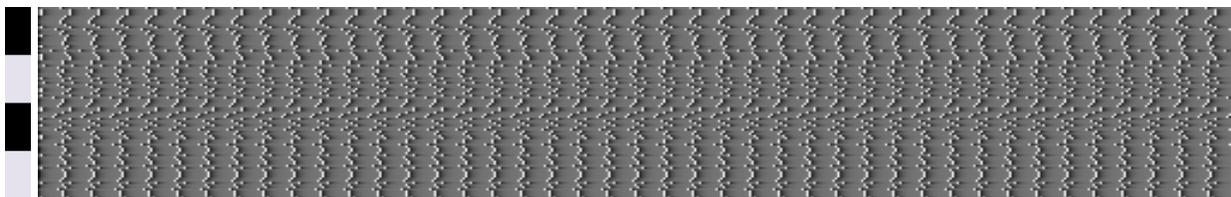
(a) Both excitatory and inhibitory connections.



(b) Excitatory connections only



(c) Inhibitory connections only.



(d) No lateral connection.

Figure 4.3: Binding and Segmentation with Different Connection Types. A one-dimensional network of 90 neurons was divided into two groups and simulated for 500 iterations. The neurons indexed (1..22) and (43..64) formed the first group (rows marked gray on the left), and those indexed (21..42) and (65..90) the second group (rows marked black). Excitatory lateral connections only connected neurons within the same group, and inhibitory connections were global. The neurons are indexed from 1 to 90 from bottom to top. (a) With both excitatory and inhibitory connections, the neurons within the same group are synchronized, while those in different groups are desynchronized. (b) With only the excitatory connections present, the neurons cannot desynchronize. (c) When only the inhibitory connections are present, no coherently synchronized groups emerge. (d) When there are no lateral connections, neurons spike in random phases.

4.2.2 Effect of Noise

In previous sections, the initial membrane potential of each neuron was uniformly randomly initialized. However, whether such an initial perturbation is necessary or not was not explicitly tested. In this section, the role of such initial noise, and also that of continual noise, will be tested. The results are compared to the control case where the simulation was carried out without noise.

A network of 180 neurons with both excitatory and inhibitory lateral connections was simulated for 500 iterations. The network was divided into two groups as in the previous experiment (section 4.3). The neurons indexed (1..22), (43..64), (91..112), and (134..155) formed the first group, and those indexed (23..42), (65..90), (113..133), and (156..180) the second group. Excitatory lateral connections only connected neurons within the same group, and their radius was limited to 90. The inhibitory connections were global.

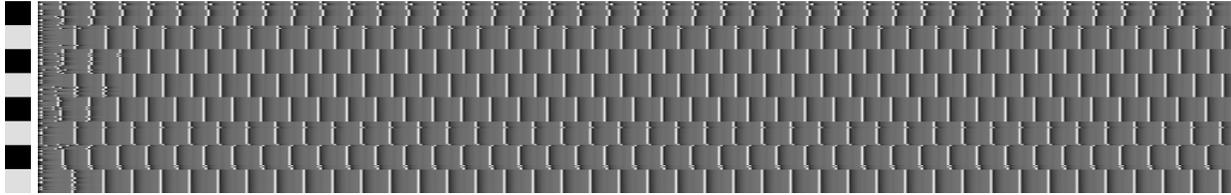
Three separate experiments were conducted: with initial noise only, continual noise only, and without noise. The rest of the parameters were the same in all three experiments: $\gamma_e = 0.48$, $\gamma_i = 0.42$, and $\lambda_e = 5.0$, $\lambda_i = 1.0$. As before, all other simulation conditions were the same as in section 4.1.1.

The results are shown in figure 4.4. With initial noise (i.e. random initial voltage), the neurons within the same group are synchronized while the two groups are desynchronized (*a*). Also, even if the neurons are initialized uniformly, when a small amount of noise (0.1%) is added to the membrane potential at each time step, the two groups will desynchronize (*b*). However, without noise of any kind, symmetry is not broken and the two groups stay synchronized (*c*). So, inhibitory connections alone are not sufficient for desynchronization. Cortical neurons actually operate in a noisy cellular environment, so including such noise in the model is realistic. It also makes the model more robust, as has been shown in other models as well (Horn and Opher 1998; Terman and Wang 1995; Wang 1995). So, a small amount of noise is essential in perceptual grouping, and will be used in the following computational experiments to assist in desynchronization.

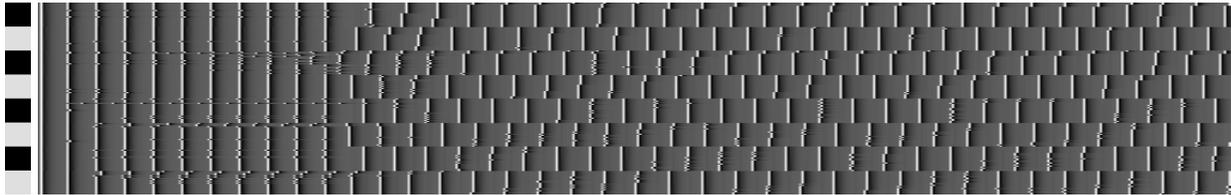
In sum, I showed that noise is needed for desynchronization, by testing two different kinds of noise, (1) initial noise, and (2) continual background noise, and the (3) control case without noise. Next, how to make binding and segmentation robust in the network of spiking neurons will be investigated.

4.3 Robustness of Binding and Segmentation

In previous sections, I have shown how the synaptic decay rate, the type of lateral connections, and the degree of noise can be controlled in the model to achieve synchronization and desynchronization for binding and segmentation. However, there are several factors that can possibly interfere with this process. For example, if the network is presented with different-size inputs simultaneously, the larger input could dominate the smaller input. If the level of noise is raised above a



(a) Random initial voltage.



(b) With 0.1% noise added.



(c) No noise, non-random initial voltage.

Figure 4.4: Effect of Noise on Desynchronization. A network of 180 neurons with both excitatory and inhibitory lateral connections were simulated for 500 iterations. The network was divided into two groups as in the previous experiment (section 4.3). The neurons indexed (1..22), (43..64), (91..112), and (134..155) formed the first group (rows marked by gray), and those indexed (23..42), (65..90), (113..134), and (156..180) the second group (rows marked black). Excitatory lateral connections only connected neurons within the same group, within a radius of 90, and inhibitory connections were global. (a) The membrane potential of each neuron was uniformly randomly initialized, and no noise was added afterward. The symmetry is broken and the two groups are separated as expected. (b) The membrane potentials initially were the same, but perturbed throughout the simulation by adding 0.1% of uniformly random noise. The neurons within the same group are synchronized, while the two groups are desynchronized. (c) When there was no noise (initial or on-going), the symmetry was not broken, and the group remained synchronized.

threshold, noise could dominate and no coherent behavior may be obtained. Some parts of the neuron model can be tuned to overcome such problems, such as increasing the duration of absolute refractory period to overcome high levels of noise. In this section, I will further test the model to see how robust it is against such external factors, and show which components contribute to the robustness.

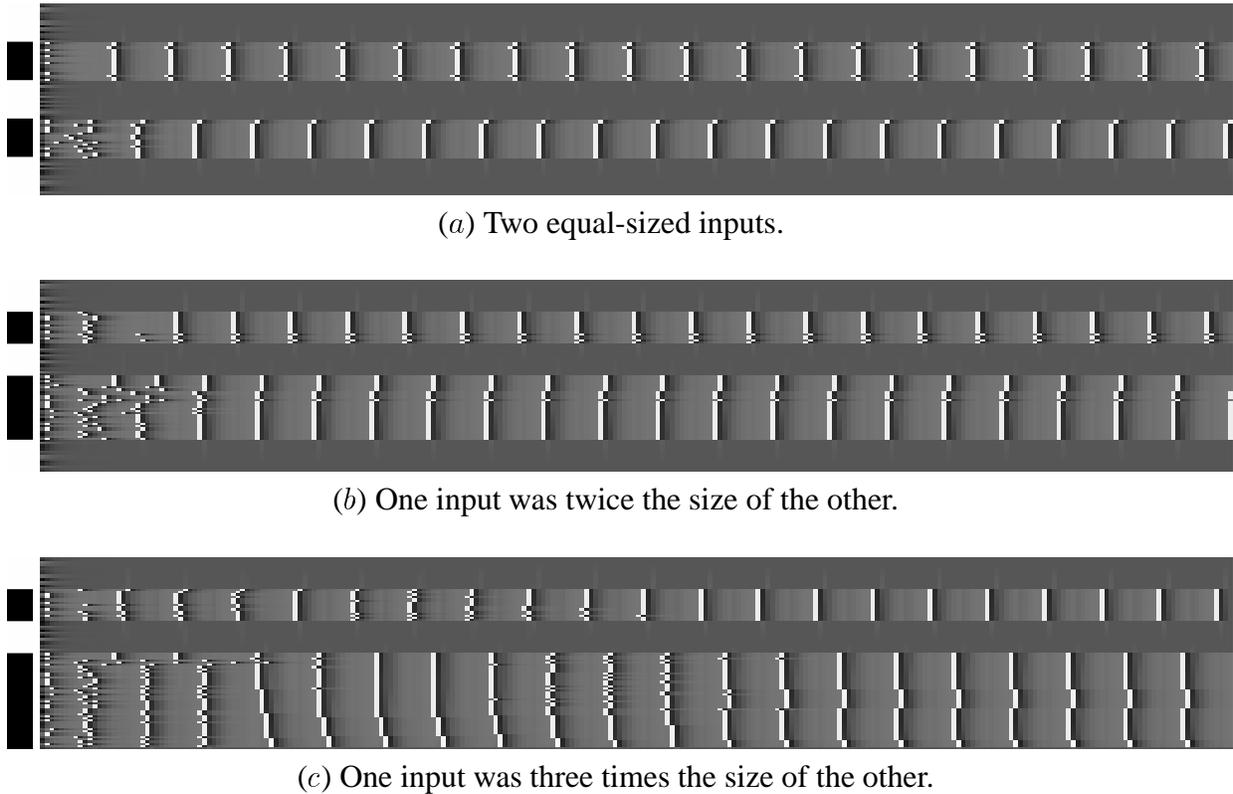


Figure 4.5: Effect of Relative Input Size on Synchronization. A network of 90 neurons with both excitatory and inhibitory lateral connections were simulated for 250 iterations. The excitatory connection radius was 14, and inhibitory connections were global. The network was presented with two spatially separated inputs, and the size of the second input was varied. The rows (neurons) that received input are marked by black solid bars on the left. (a) The two inputs were the same size: neurons indexed (19..36) and (65..82). (b) One input was twice as long as the other input: neurons indexed (16..45) and (61..75). (c) One input was three times as long as the other input: neurons indexed (1..45) and (61..75). In all cases, the inputs are robustly bound and segmented.

4.3.1 Effect of Relative Input Size on Synchronization

One requirement for perceptual grouping is that input features should not be suppressed or promoted on the basis of size only, since smaller but complex input features in the scene can be equally important as large but simple features. Thus, a network of spiking neurons modeling such behavior should tolerate differences in input size.

To test if the PGLISSOM model is robust against such variation, a network of 90 neurons with both excitatory and inhibitory lateral connections was simulated for 250 iterations. The excitatory connection radius was 14 so that neurons representing different inputs were not connected, and inhibitory connections were global. Three separate experiments were conducted by presenting two inputs of relative sizes 1:1, 1:2, and 1:3 (figure 4.5*a,b*, and *c*). The parameter values were the same in the three experiments: $\gamma_e = 0.7$, $\gamma_i = 0.6$, and $\lambda_e = 5.0$, $\lambda_i = 1.0$, to compensate for the

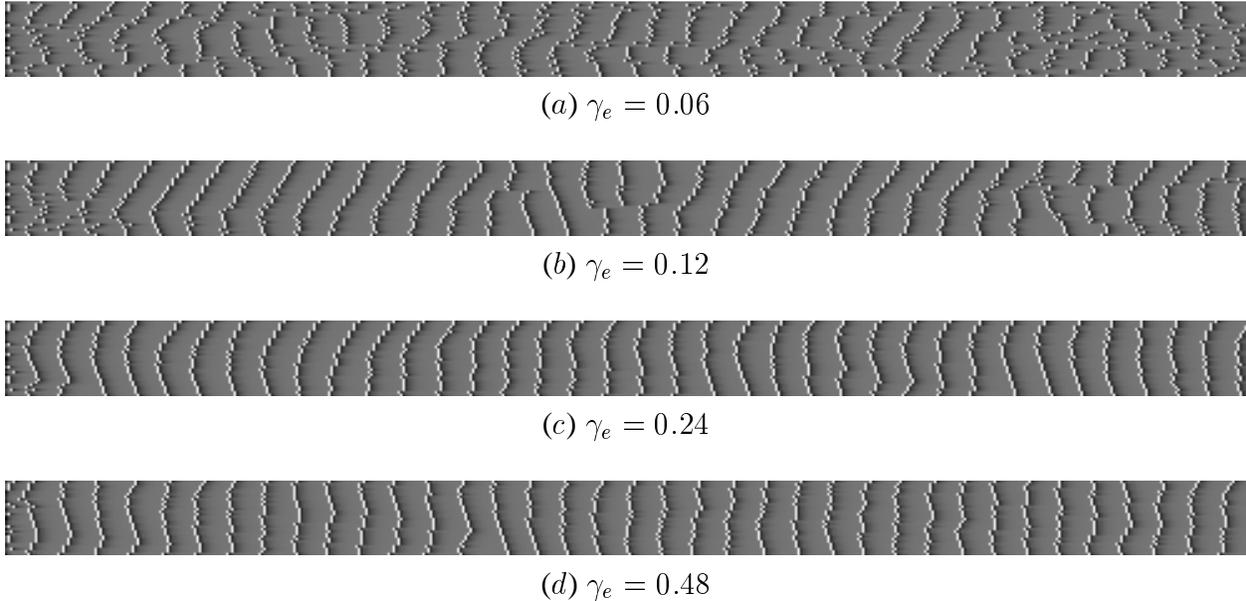


Figure 4.6: **Overcoming Noise with Strong Excitation.** A network of 30 neurons with excitatory lateral connections were simulated for 500 iterations. The excitatory lateral connections were global. A higher-level of noise (1%, i.e. 10 times the noise in section 4.2.2) was added to the membrane potential at each iteration. When the excitatory contribution is low, as in (a) and (b), noise overwhelms and causes the activities to desynchronize. However, if the excitatory contribution gets higher as in (c) and (d), it overcomes noise and achieves synchrony.

larger number of excitatory connections compared to the previous experiments. Other simulation conditions were the same as in section 4.1.1.

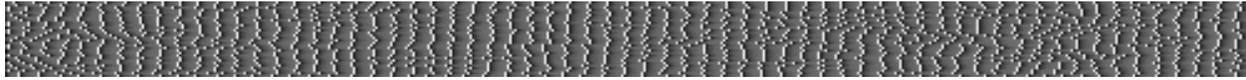
The results are shown in figure 4.5. The two areas of the map representing the two objects are synchronized and desynchronized within and across the group, regardless of the input size. This behavior occurs under identical parameter conditions, and thus demonstrates that the model behavior is not affected by the size of the inputs alone.

4.3.2 Overcoming Noise with Strong Excitation

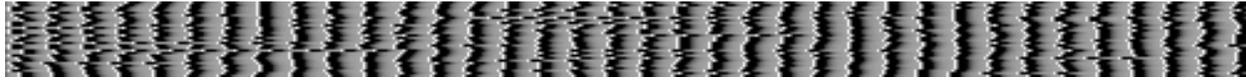
Cortical neurons operate in an inherently noisy environment. Noise can arise due to several causes, such as unreliable synaptic transmission or fluctuations in membrane potential.

In section 4.2.2, I showed that small amount of noise is necessary for desynchronization of group activity to occur. The amount of noise can possibly be much larger, and a model should be robust against such high levels of noise as well.

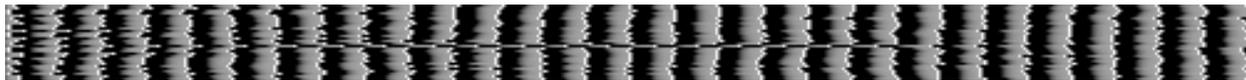
A network of 30 neurons with global excitatory lateral connections was simulated for 500 iterations. Four separate experiments were conducted where the strength of the excitatory contribution was increased in four stages under high-noise (1%, i.e. 10 times the noise in section 4.2.2). The simulation conditions were the same as in section 4.1.1, except for $\lambda_e = 5.0$.



(a) No absolute refractory period.



(b) Absolute refractory period for 3 iterations.



(c) Absolute refractory period for 5 iterations.

Figure 4.7: Overcoming Noise with a Longer Refractory Period. A network of 30 neurons with excitatory lateral connections were simulated for 500 iterations. The excitatory lateral connections were global. A high level of noise (7%, 70 times the noise in section 4.2.2) was added to the membrane potential at each iteration. Such high noise cannot be tolerated by just increasing the lateral excitatory contribution γ . However, enforcing the absolute refractory period can greatly enhance the robustness of the model. (a) When absolute refractory contribution is 0, the activities are random. (b) When the absolute refractory contribution is 3, the activities start to show loose synchronization. (c) When the absolute refractory contribution is 5, the activities show strong synchronization: with longer periods between firing, the noise is washed out.

The results are shown in figure 4.6. With high noise, weak excitatory connections cannot keep the neurons synchronized (a and b), but as the excitatory contribution γ_e is increased, the neurons start to synchronize. This result shows how a network of spiking neurons can robustly synchronize even in high noise conditions. Strong excitation can be used to overcome noise in such cases.

4.3.3 Overcoming Noise with a Longer Refractory Period

Although increasing the excitatory contribution helps, there is a certain threshold where noise cannot be overcome this way. For example, 7% noise would break the synchronization behavior even with extremely strong excitatory connections because noise will dominate the spiking behavior of the network.

A network of 30 neurons with global excitatory lateral connections was simulated for 500 iterations. Three separate experiments were conducted where the length of absolute refractory period was increased. The simulation conditions were the same as in section 4.1.1, except for $\lambda_e = 5.0$, to make synchronization more robust.

The results are shown in figure 4.7. Under very high noise (7%), the excitatory connections cannot keep the neurons synchronized (a), but as the absolute refractory period is increased, the neurons start to synchronize (b and c). This result suggests that absolute refractory periods may have come to exist in biological neurons in part to overcome high levels of noise in the cortical

environment. The time interval during which the neuron can fire is made much smaller than the forced silent period, thus the noise gets washed out.

Together with strong γ_e , the neuron can be made highly robust against noise, suggesting that synchronization can scale up to in real environments.

4.4 Conclusion

In this chapter, I have systematically tested a 1-D network of spiking neurons to find out how the different components of the PGLISSOM model contribute to synchronization and desynchronization of activity. I have shown that (1) synaptic decay rate can induce opposite synchronization behavior in excitatory and inhibitory connections, (2) local excitatory connections can achieve global synchrony, and (3) inhibitory connections and noise are necessary for desynchronization. I also showed that the model is robust to changes in input size, and against very high levels of noise through strong excitation and long absolute refractory periods.

Understanding such qualitative and quantitative factors that affect synchronization is important so that we can predict how a network of spiking neurons behave in a specific configuration and identify why certain mechanisms exist in biological neurons. We will now turn to the second major component of the PGLISSOM model, showing how these connections are self-organized.

Chapter 5

Self-Organization

The visual cortex of mammals has an orderly organization where orientation, spatial frequency, ocular dominance, direction of motion, and other visual features are laid out in a two-dimensional map structure. Horizontally across the map, lateral connections connect areas with similar functionality. Such structures are known to emerge during the early stages of development in animals, but the exact mechanisms are not fully understood. In PGLISSOM, they emerge as a result of input-driven self-organization. In this chapter, I will show how PGLISSOM self-organizes to form (1) orientation maps and (2) functionally specific lateral connection patterns. The properties of the map and the connections will be analyzed and compared to experimental data.

5.1 Simulation Setup

As was described in Chapter 3, the PGLISSOM model consists of two stacked layers receiving input from the retina through afferent connections. The lower layer (MAP1) has short-range excitatory lateral connections and long-range inhibitory lateral connections. The purpose of this layer is to develop a smooth mapping of the input space. In the upper layer (MAP2), the excitatory lateral connections as well as the inhibitory lateral connections have a long range. This layer performs perceptual grouping through the excitatory connections. The intra-columnar connections that connect vertically corresponding locations on the two maps ensure that MAP2 self-organize a smooth feature map similar to MAP1, and MAP1 perform perceptual grouping like MAP2.

In the simulations reported in this chapter, MAP1 consisted of 136×136 neurons, and MAP2 of 54×54 neurons. MAP2 was made smaller to make the simulation run faster and to fit the model within the physical memory limit. However, the intra-columnar connections between MAP1 and MAP2 were proportional to scale, so that the relative locations of connected neurons in the two maps were the same. This connectivity scheme ensures that the global order of the two maps matches as they develop. The excitatory lateral connection radius in MAP1 was initially 7 and was gradually reduced to 3, and inhibitory lateral connections had a fixed radius of 10. Initially, large areas have correlated activity so that global order can be formed, and later on, the

reduced lateral excitatory connections help fine-tune the local order in the map (Kohonen 1982b, 1989, 1993; Sirosh and Miikkulainen 1997). In MAP2, excitatory lateral connections had a radius of 40 and the inhibitory connections were global. Afferent connections from the retina had a radius of 6 in both maps, and intra-columnar connections a radius of 2. The retina consisted of 72×72 receptors. As long as the relative sizes of the map, the retina, and the lateral connection radii are similar to these values, the maps would self-organized well (see Bednar et al. (2002) for a precise equation that allow scaling to maps with different sizes).

The input consisted of oriented Gaussians with major and minor axis lengths of $a^2 = 15.0$ and $b^2 = 1.3$, and gradually elongated to 50.0 and 0.8 to help make the orientation tuning sharper. All weights were initialized uniformly randomly distributed values between 0.0 and 1.0. The relative contribution of afferent (γ_a), lateral excitatory (γ_e), lateral inhibitory (γ_i), and intra-columnar (γ_c) connections were 1.1, 0.8 0.9, and 0.5 for MAP1. The value γ_c was lower than the others so that activity caused by high excitation in MAP2 would not interfere with self-organization in MAP1. For MAP2, these values were 1.1, 0.2, 2.5, and 0.9. The value γ_e was lower and γ_i higher than in MAP1 to prevent the map from becoming too active, and γ_c was higher than in MAP1 to allow the self-organized global order of MAP1 to be transferred to MAP2.

The learning rates of afferent (α_a), lateral excitatory (α_e), lateral inhibitory (α_i), and intra-columnar (α_c) connections were 0.012, 0.008, 0.008, and 0.012 for MAP1 and 0.012, 0.008, 0.0, and 0.012 for MAP2. The inhibitory connections in MAP2 did not adapt ($\alpha_i = 0$) so that the initial broad connectivity remains to provide background inhibition, as explained in Chapter 3. At 5,000 iterations, α_a and α_c in both maps were decreased to 0.008, so that the order in the map can start stabilizing. Initial base threshold θ_{base} for both maps was 0.05. At the beginning of each settling iteration, the θ_{base} value was adjusted to 50% of $\max_{i,j}(\sigma_{i,j}(t))$ so that the network would not become too active or totally silent. Later, the percentile was gradually increased to 57.5% by 15,000 iterations for MAP1, and 65% by 5,000 iteration for MAP2, so that the activities can become sparser. This method has a similar effect as the slope adaptation of sigmoid activation function of Sirosh and Miikkulainen (1994).

The synaptic decay rates were different for the different types of connections. They were 3.0, 0.5, and 1.0 for the lateral excitatory (λ_e), inhibitory (λ_i), and intra-columnar connections (λ_c) in both maps. The decay rate in the spike generator's relative refractory period (λ_{rel}) was 0.5 in both maps. The contribution of the relative refractory period in dynamic threshold calculation (τ) was 0.4 in both maps. The threshold (δ) and ceiling (β) of the activation function $g(\cdot)$ were 0.01 and 1.3 in both maps. For the average spiking rate of neurons, a running average with the rate $\tau_{\text{avg}} = 0.92$ was calculated. These parameter values were found to be effective by running several experiments, and small changes in the parameter values did not affect the global behavior of the model.

The network configured in this way was trained for 40,000 iterations following the procedure described in chapter 3. The training took about 30 wall-clock hours utilizing 178MB of memory on a 1GHz AMD Athlon PC running Linux. The resulting global order was very simi-

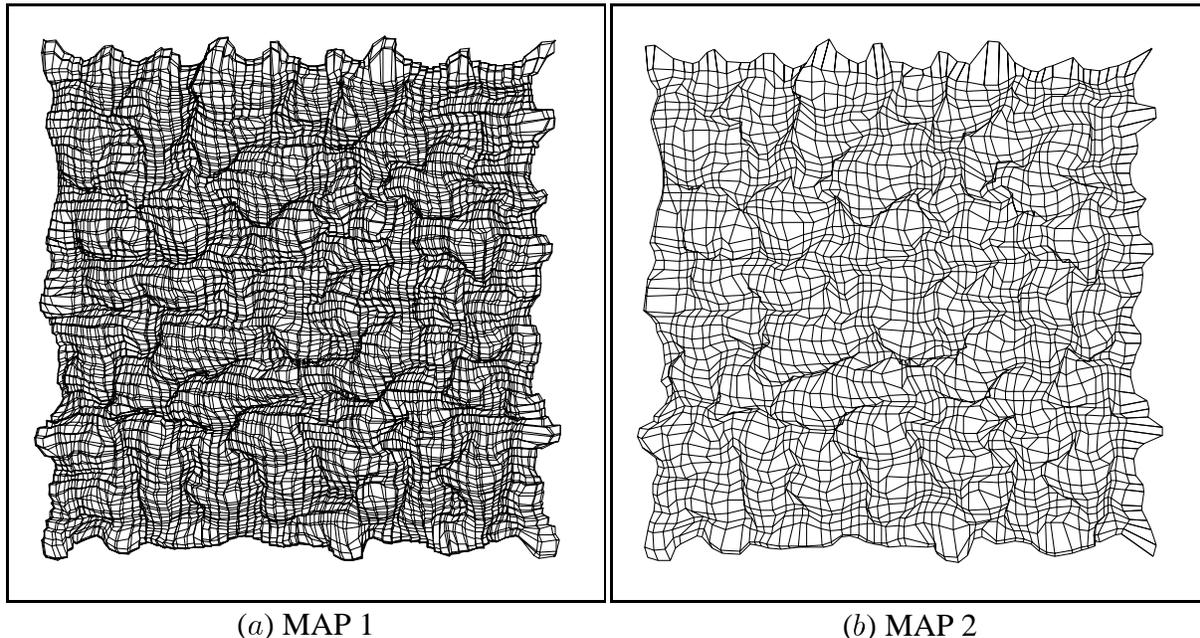


Figure 5.1: Retinotopic Organization. For each cortical neuron, the coordinate of the receptive field center was found by calculating the center of gravity of the afferent weight matrix. These coordinates in retinal space are plotted as nodes. To visualize the location of the receptive field of each cortical neuron relative to neighboring neurons, edges were drawn between the centers of immediately neighboring neurons (a) and (b) show the retinotopic organization of MAP1 and MAP2, respectively. Although the maps are different in size, the overall organization closely matches, and thus neurons in the same cortical column receive input from the same locations in the retinal space. The overall organization of the map does not form a regular, square grid. Instead, there are areas with densely packed receptive fields, while in other areas, receptive fields are sparser, as is the case in experimental data (Das and Gilbert 1997).

lar in the two maps, with the retinotopic organization and the orientation map properties closely matching (figure 5.1–5.4). The long-range excitatory lateral connections in MAP2 formed a functionally specific patchy pattern where similarly orientation-tuned neurons were more likely to be connected (figure 5.11). These results, which closely match experimental data, will be discussed in detail in the following sections.

5.2 Retinotopic Organization

After training, the properties of the resulting maps were measured. The first measure was to plot the receptive field centers, to reveal the retinotopic organization of the PGLISSOM maps. The receptive field centers were found by calculating the center-of-gravity of the afferent weight matrices of each neuron.

The results are shown in figure 5.1. The retinotopic organization in the two maps, MAP1 and MAP2, closely matches because the intra-columnar connections were proportionally scaled.

As seen in the plots, the receptive field centers do not form a square grid. Rather, some areas of the cortex have more densely packed receptive fields, while in other areas they are more sparsely distributed, seen as valleys and hills in the figure. A similar distortion in retinotopic organization was found in cats (Das and Gilbert 1997). Such distortions are believed to be due to differently-distributed lateral connections across the surface of the cortex in these distorted regions. The distortions originate from MAP2, where the long-range excitatory connections have a strong influence. In a separate experiment where the intra-columnar connections between MAP1 and MAP2 were severed, the retinotopic organization of MAP1 showed fewer distortions than in figure 5.1.

In sum, the intra-columnar connections made it possible for the two maps to have similar retinotopic organization. The lateral connections were shown to influence receptive field properties as simple as position (retinotopy), thus providing a possible explanation for small retinotopy distortions found in the brain.

5.3 Orientation Map

In this section, the orientation tuning of the map neurons and the global orientation maps in MAP1 and MAP2 will be analyzed and compared with experimental data.

5.3.1 Receptive Fields, Global Order and Map Features

During self-organization, the input distribution is determined by three variables: the x and y coordinates of the input bar on the retina and its orientation. Since the coordinate information is roughly encoded in the initial receptive field centers already, the feature that most affects the development of the map is orientation.

Circular receptive fields that are randomly initialized in the beginning start to become elongated and smoother as the training proceeds, and this way, the neurons develop orientation preferences. The resulting receptive fields of MAP1 and MAP2 are shown in figure 5.2. A majority of neurons are orientation selective, and the orientation preferences change smoothly across the map. Also, the neurons in MAP1 and MAP2 belonging to the same column have highly similar orientation preferences.

To examine the global order of the orientation maps, the orientation preferences of each neuron were measured using the vector sum method (Bednar 1997; Bosking et al. 1997). For each neuron, the afferent weight matrix was scanned with elongated Gaussian bars at six different orientations. For each orientation, the bar was swiped from one end of the weight matrix to the other and the maximum vector dot product of the weight matrix and the input bar was recorded. Repeating this procedure for six orientations, six (angle,dot-product) pairs of polar vectors were gathered. The vectors were added to obtain a *vector sum*. This sum represents two measures: (1) the angle of the vector represents the *orientation preference* of the neuron (i.e. the orientation to which the neuron responds maximally), and (2) the length of the vector represents the *orientation*

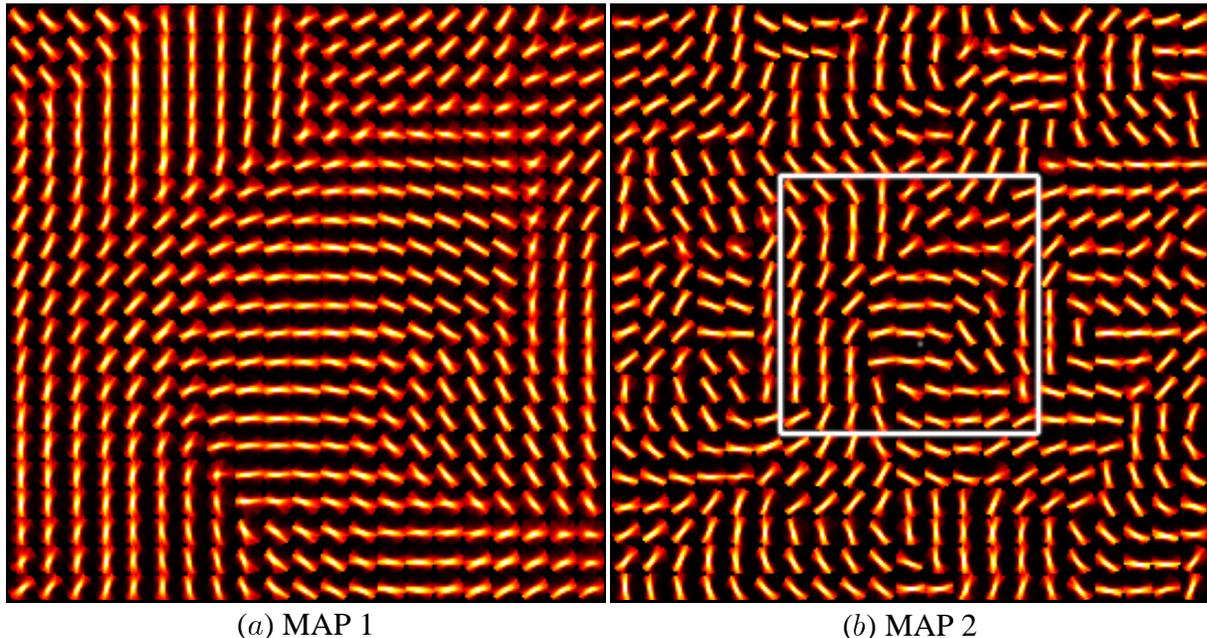


Figure 5.2: **Oriented Receptive Fields.** The receptive fields of the cortical neurons in the central 21×21 area in (a) MAP1 and (b) MAP2 are shown. Each cell in the grid represents the afferent weight matrix of the cortical neuron at that location. Strength of the afferent connection is color coded: from minimum to maximum, black \rightarrow red \rightarrow yellow \rightarrow white. We can see the smoothly changing orientation preferences in both maps. Some neurons have almost circular receptive fields, i.e. they are not tuned to any particular orientation. Since MAP1 and MAP2 differ in size, we have to compare the whole 21×21 area in (a) to the central 9×9 area in (b), marked with the white bounding box. These two areas in the two maps receive projections from exactly the same central area in the retina. Thus, the input distribution they were exposed to during training was identical. As a result, these two areas developed highly similar orientation preferences. As in the retinotopic maps, such a close match in functionality is due to the intra-columnar connections.

selectivity of the neuron (i.e. how well the neuron responds to the preferred orientation compared to other orientations). The orientation preferences resulting from this procedure are shown in figure 5.3. As we saw in figure 5.2, the preferences change smoothly across the map, and MAP1 and MAP2 show highly similar global order. Such order is called an orientation map. The maps have features (highlighted in figure 5.4b) commonly found in mammalian visual cortex such as (1) pinwheel centers (areas with a radial representation of the full 180 degrees; **A–D**), (2) linear zones (a full range of orientations repeating over regular intervals on the cortex; **E**), and (3) fractures (long stretches of cortical areas where orientation preference changes abruptly; **F**)¹. These results show that the architecture and the learning rules in PGLISSOM can develop realistic orientation maps as seen in the visual cortex.

Even though the intertwined organization of the orientation map looks complex, there are

¹Only features of MAP1 are shown in the figure; MAP2 had very similar patterns. MAP2 afferent weight properties closely resemble MAP1 as we saw in the previous figures 5.2-5.3, thus, in the rest of the thesis, afferent weight properties will be shown only for MAP1 unless otherwise necessary.

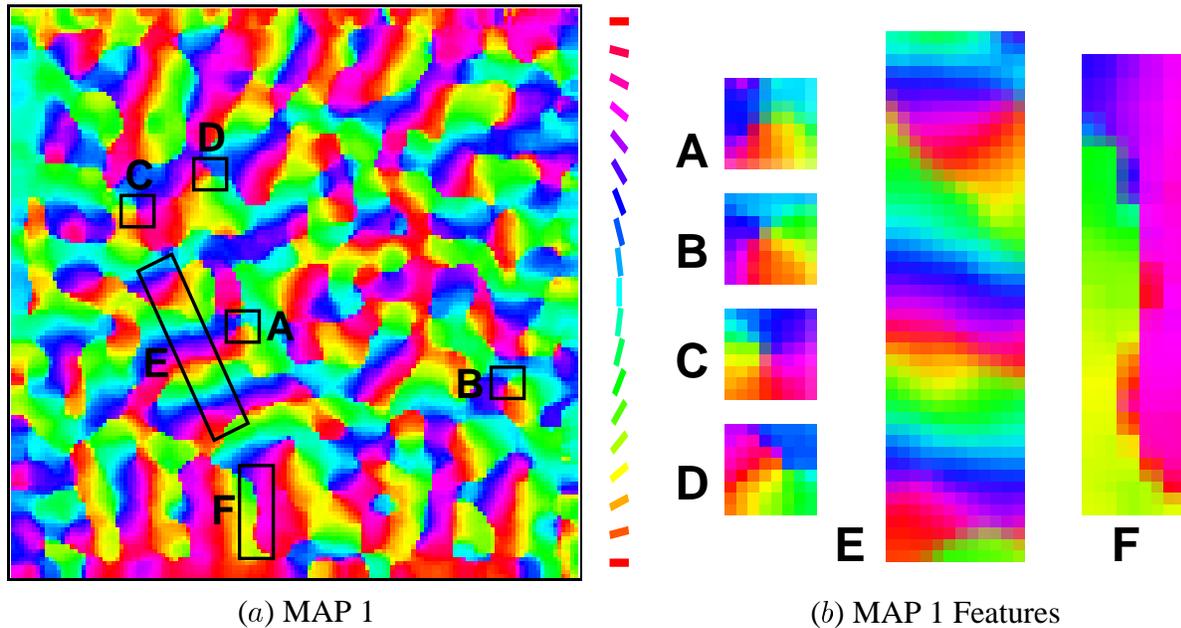


Figure 5.4: **Characteristic Features of the Orientation Map** (*color figure*). The orientation preference at each location on the cortex is coded in color as in figure 5.3. (a) Orientation map in MAP1 is shown. (b) The features in MAP1 are shown: from left to right, pinwheel centers (A–D), a linear zone (E), and a fracture (F). In the pinwheel center, a full 180-degree of orientation preference is represented in a radial cycle. In the linear zone, the full 180 degrees of orientation cycles through along a linear stretch of the cortex. At fractures, the orientation preferences abruptly change from green (50 degrees) to (150 degrees) along a long stretch of the cortex. These features match those found in the visual cortex well.

fields circular.

Since the orientation maps in MAP1 and MAP2 have a close correspondence, even though their sizes differ, their orientation preference histograms are comparable. In MAP2, there are some uneven intervals, which is due to the relatively smaller size of the map, and would be likely to disappear if a larger map were used.

5.3.3 2-D Fourier Power Spectrum

The global organization of the orientation map shows a repeating structure in all angular directions. The 2-D Fourier power spectrum of such a map can reveal the characteristics of the repeating structure quantitatively. This technique is a widely used to verify if orientation maps in computational models have properties similar to the maps in experimental observations. To compare PGLISSOM with biological maps, 2-D Fourier spectra were calculated for both MAP1 and MAP2. The results are plotted in figure 5.7.

A prominent feature of the Fourier spectra in both maps is the ring-shaped high-energy band, which is also observed in experimentally measured orientation maps in mammals (figure 5.6). The regions with same orientation preference in the cortex (and in the model) can be

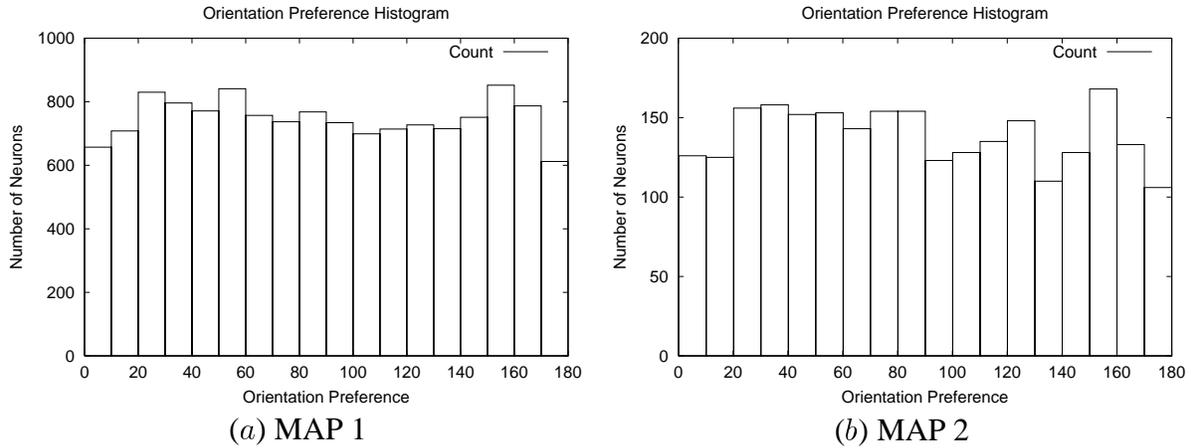


Figure 5.5: Orientation Preference Histograms. A histogram of 18 bins was calculated for (a) MAP1 and (b) MAP2. Each bin covered a 10-degree increment from 0 to 180 degrees. The histogram for MAP1 and MAP2 are both fairly flat, meaning that all orientations are equally well represented.

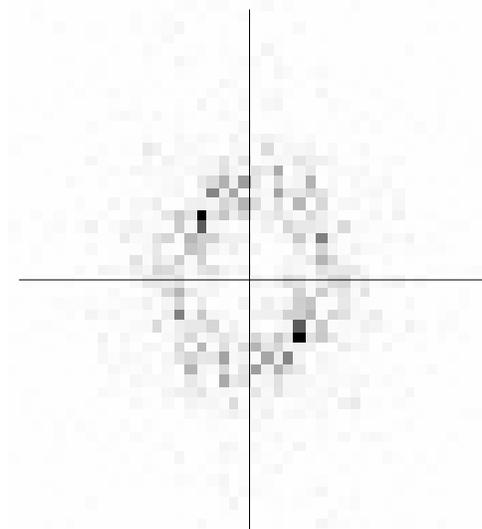


Figure 5.6: Fourier Power Spectrum of a Cortical Orientation Map. A prominent feature of the Fourier spectra in orientation maps in the visual cortex is the ring-shaped high-energy band. Such a shape shows that the areas with same orientation preference can be found around the full 360 degrees, at spatial distance corresponding to the radius of the ring. The image was generated by the author from the monkey data in Blasdel (1992b) using methods described in Erwin et al. (1995).

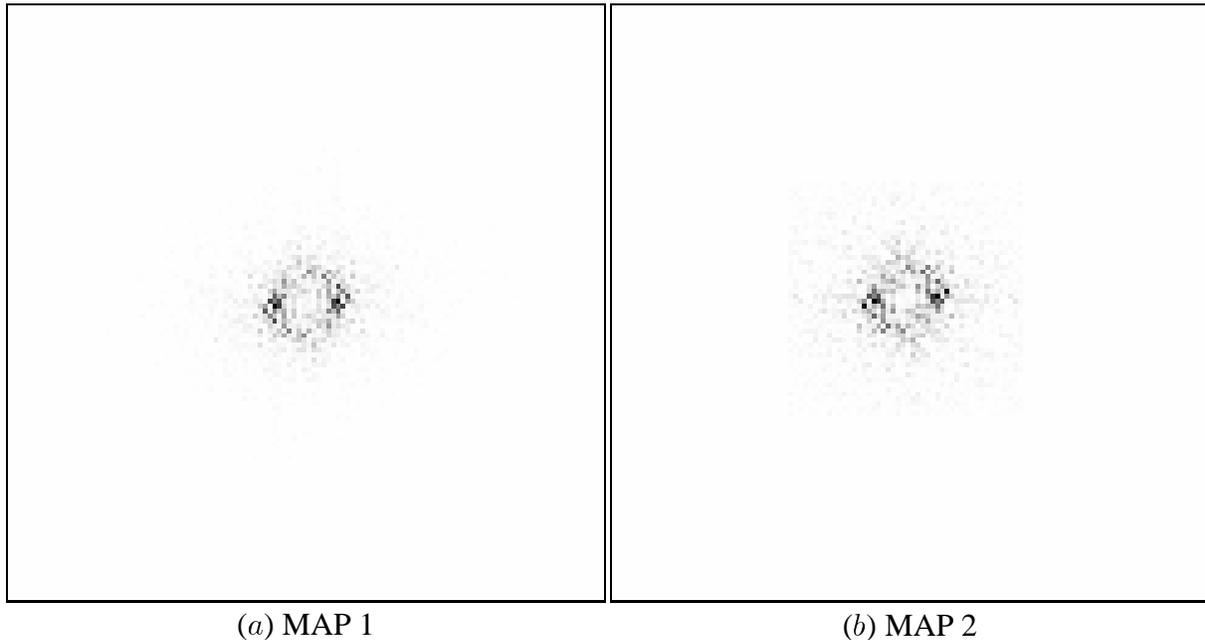


Figure 5.7: **2-D Fourier Power Spectra of PGLISSOM Orientation Maps.** The spectra for MAP1 is shown in (a), and for MAP2 in (b). The power is represented in gray-scale, where white is 0 and black is the maximum. In both maps, the Fourier spectra show characteristic ring-shape seen in figure 5.6.

found in the full 360 degrees at intervals of the length equal to the radius of the ring (Erwin et al. 1995; Sirosh 1995). Such a repetition ensures that at any location in the visual cortex, the full 180 degrees of orientation are fully represented.

5.3.4 Orientation Gradient and Orientation Selectivity

Another quantitative measure of the orientation map is the orientation gradient, which is the local rate of change in orientation preference. Orientation gradient plots are useful for identifying features such as fractures and pinwheel centers, since at these places, the orientation preference abruptly changes. With the orientation gradient plot, it is easier to see and compare the global organization of such features, and compare with experimental results and data from other computational models. Figure 5.8a shows the orientation gradient of MAP1. The gradient is very similar in MAP2, and is not repeated here as mentioned earlier. The plot is comparable to experimental data where such features as fractures and pinwheel centers are highlighted by dark, high-gradient regions (Blasdel 1992b; Koulakov and Chklovskii 2001; Sirosh 1995): figure 5.9.

In optical imaging data, fractures and pinwheels generally have low orientation selectivity, i.e. the degree of tuning for the preferred orientation. High selectivity means that it is hard to elicit response from that neuron when a suboptimally oriented stimulus is presented (section 5.3.1). Thus, the receptive fields with high orientation selectivity tend to be longer and narrower. Figure 5.8b shows the orientation selectivity of neurons in MAP1. As in experimental data, low-

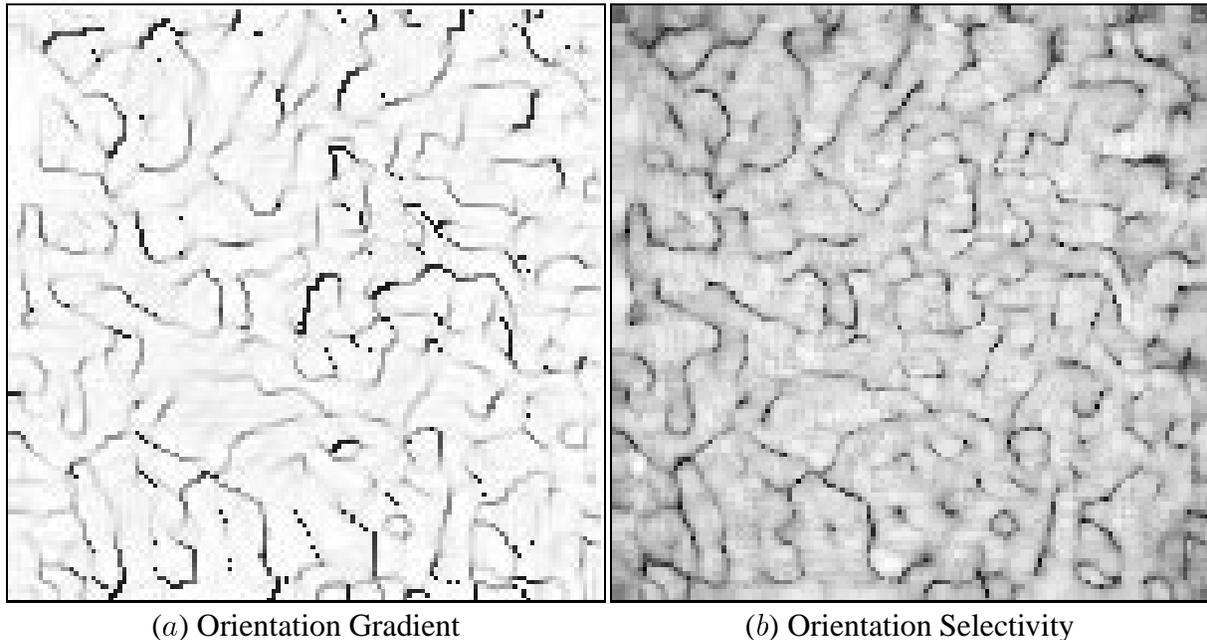


Figure 5.8: Relationship between Orientation Gradient and Orientation Selectivity. (a) The orientation gradient of MAP1 is shown. White areas indicate low gradient, where the change in orientation preference is slow, and black areas indicate high gradient, where the orientation preference of nearby neurons abruptly changes. (b) The orientation selectivity of MAP1. White signifies high selectivity and black indicates low selectivity. Notice that the scale is the reverse of that in (a) for easier comparison. High-gradient areas correspond to low-selectivity areas as can be seen by comparing the black ridges in (a) and (b). Such an orientation gradients plot and selectivity is commonly seen in mammalian visual cortex, and also in computational models (figure 5.9; Blasdel 1992b; Sirosh 1995).

selectivity areas correspond to high-gradient areas. Such a pattern seems to be due to the cortex trying to smoothly map the orientation dimension, suppressing activities arising from abrupt changes in orientation.

In sum, the orientation maps that develop in PGLISSOM have qualitative and quantitative properties that closely approximate experimental data in terms of (1) retinotopy, (2) oriented receptive fields, (3) orientation maps and features, (4) 2-D Fourier power spectrum, (5) orientation gradient, and (6) orientation selectivity. In the next section, the lateral connection patterns will be analyzed and compared to biological data.

5.4 Lateral Connections

As was discussed in Chapter 2, the lateral connections in the visual cortex have two general properties: (1) strong connections exist between neurons with similar orientation preferences (figure 2.4), and (2) the connections extend in the direction along the axis of the source neuron's orientation preference (figure 5.10). Such connections are believed to represent correlational structures in

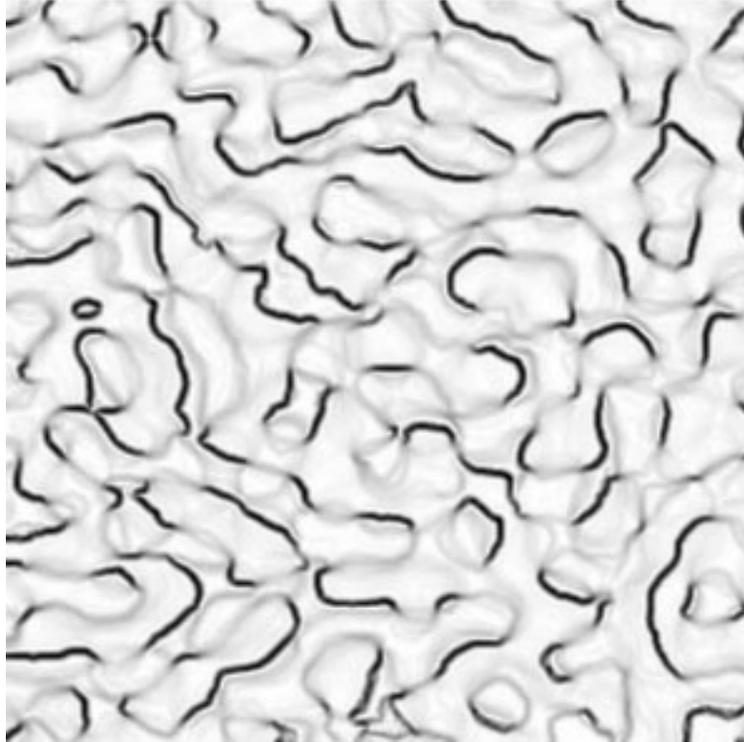


Figure 5.9: **Gradient of a Cortical Orientation Map.** The orientation gradient in the monkey visual cortex is plotted in gray-scale from low to high (gray \rightarrow black). The high-gradient areas (dark ridges) correspond to fractures and at the end of the fractures are pinwheel centers. The gradient map highlights such discontinuities in orientation maps and allows us to compare the global arrangement of these features to those of computational models. Adapted from (Blasdel 1992b).

the input and implement a *local grouping function* (or association field; Chapter 1) for perceptual grouping tasks such as contour integration.

The long-range lateral connections in PGLISSOM also self-organize into intricate patterns that reflect the correlations in the input, and these patterns have the same two properties found in the visual cortex. Examples of these connections in MAP2 are shown in figure 5.11 ² From *b* through *d*, the lateral connections originating from the neurons marked in *a* were plotted. The hue (or color) represents the orientation preferences of the target neurons, and the intensity (how vivid the color is) indicates the strength of the connection. In *b*, the source neuron has an orientation preference of 85 degrees (color between cyan and green), and a majority of its strong connections are found to be of the same color. The same is true for *c*, and *d*, where the preferred orientations were 141 degrees and 174 degrees, and the corresponding color codes were magenta and red. Thus, the figure shows qualitatively that property 1 above holds in the model (cf. figure 2.4). Another feature that is noticeable is the anisotropic projection of connections, i.e. the connections

²Only MAP2 connections are shown because MAP1 only had short-range excitatory connections. For the same reason, synchrony will be measured end illustrated in MAP2 only in the rest of the thesis.

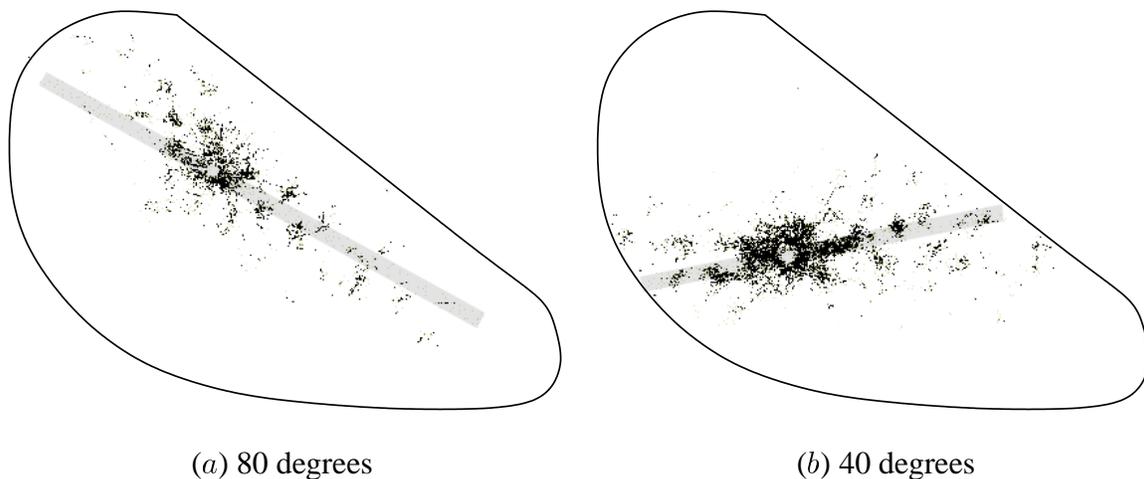


Figure 5.10: Lateral Connections in Three Shrew. The slightly deformed oval area is the entire V1 (primary visual cortex) of the tree shrew (arboreal mammal of Southeast Asia). The synaptic boutons (connection terminals) of the lateral projections are plotted as black dots. The gray bar in the background is the axis of preferred orientation of the neuron from where the projections originate, mapped from the visual space into the cortical space. The connections are dense and circular around the center but farther away from the center they become patchy and align along the axis of preferred orientation of the center neuron. In (a), the neuron in the center has orientation preference of 80 degrees and in (b), 40 degrees. See also figure 2.4 for similar data plotted against the orientation map. Adapted from Bosking et al. (1997).

are stretched along a straight axis of a certain angle, and patches of connections are found along the axis at roughly equal intervals. Such connection patterns are consistent with property 2 exhibited in biological data (figure 5.10).

Such patterns emerge in PGLISSOM because the afferent and lateral connections adapt to encode the statistical structure in the training input. Since the training inputs are long, elongated Gaussian bars, the afferent connections form oriented receptive fields. Neurons with similar orientation preferences whose receptive fields are aligned along a straight axis will be activated simultaneously when a long input happens to fall upon those receptive fields. Due to the Hebbian learning process, such neuron pairs will develop strong lateral connections. Moreover, the connections are not strictly aligned along the axis, but there is also a certain degree of flank in the connections (or an angular deviation from the axis). This flank becomes more pronounced farther from the source neuron in a bow-tie fashion. The same pattern can be seen in biological data as well: see e.g. figure 5.10*b*.

This is an important observation, since such flank allows grouping of not only straight contours, but also co-circular ones. During training, the most prominent feature in the training input was collinearity, since the inputs were straight, elongated Gaussian bars. Neurons with receptive fields aligned exactly on a straight path had the greatest chance of being simultaneously active (coactive), thus they developed strong lateral connections (figure 5.12*a*). However, the orientation-

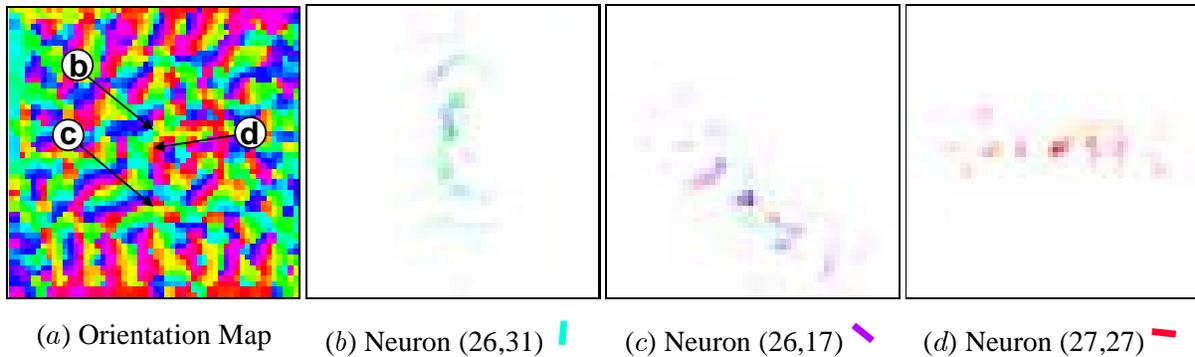


Figure 5.11: MAP2 Excitatory Lateral Connections (color figure). The long-range excitatory lateral connections in MAP2 become patchy and connect neurons with similar orientation preference. (a) The orientation map of MAP2 is shown for comparison. (b)–(d) The excitatory lateral connections of selected neurons (marked in (a) with black arrows) are visualized in color, where the hue represents the orientation preference as before, and the intensity represents the strength of each connection. White areas have very weak or no connections, and vividly colored areas have very strong connections. The neurons are numbered in Cartesian coordinates, where the lower left corner is neuron (1,1) and the upper right corner is neuron (54,54). (b) The excitatory lateral connections of neuron (26,17), which has orientation preference of 85 degrees, are shown. (c) The excitatory lateral connections of neuron (26,31), which has orientation preference of 141 degrees, are shown. (d) The excitatory lateral connections of neuron (27,27), which has orientation preference of 174 degrees, are shown. The lateral connections link similarly orientation tuned neurons (property 1), and the target zones are aligned along the axis extending in the direction of the orientation preference, as is the case in experimental observations (property 2). Specific connections like these are crucial for perceptual grouping such as contour integration. Also see figures 5.10 and 2.4 for a comparison.

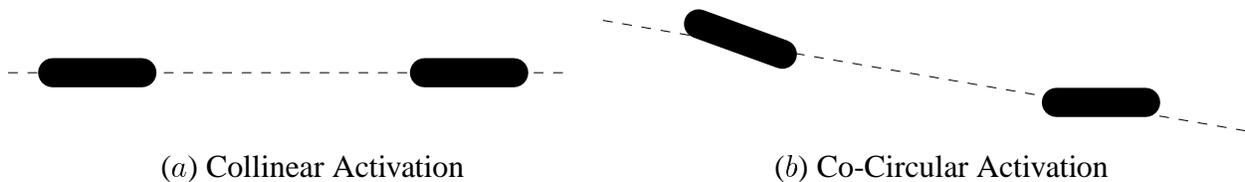


Figure 5.12: Simultaneous Activation of Neurons. The plot shows two representative cases of coactivation, i.e. when two neurons become activated simultaneously, when a long input is presented across the two receptive fields. (a) Collinear arrangement: the two receptive fields (black bars) are precisely aligned on a straight path (dashed line). If a long input is presented on this path, the two neurons will respond maximally, and the connection between them becomes stronger, according to the normalized Hebbian learning rule. (b) Co-circular arrangement: even though the two receptive fields are mis-aligned on the path, they can still weakly activate and the connection will become stronger, although not as strongly as in (a).

tuned neurons not only respond to just the optimal orientation at the optimal position, they also respond to slightly misoriented inputs (figure 5.12*b*). Thus, neurons with receptive fields not exactly aligned on a straight path could also coactivate. Coactivation happens when the two receptive fields can be connected with a straight path, but one or both are slightly misoriented from the axis of the path. Coactivation in such a case gives the lateral connection the cocircular property.

So far, we have seen that the lateral connection patterns in PGLISSOM have similar qualitative properties as those in the visual cortex, and a possible mechanism that gives rise to such patterns was suggested. Next, the connection patterns in the model will be quantitatively measured and the result will be compared to biological data.

To quantitatively measure property 1, the number of MAP2 excitatory lateral connections that connect receptive fields with orientation difference θ was calculated. The result is shown in figure 5.13*a*. The graph peaks at 0 degrees, and rapidly decreases to 0 as the orientation difference increases. It shows that strong excitatory lateral connections in MAP2 are most likely to be found between neurons that have similar orientation preference. Such a result is consistent with experimental measurements (figure 5.13*b*), and conforms with property 1. This measure not only allows us to compare the connectivity in the model and in the biological data, but also suggests a possible functional role for property 1. As we saw in Chapter 1, contours are believed to be grouped through specific lateral interactions between contour elements (representing a local grouping function, or an association field). The measure presented above suggests that lateral connections could be implementing such local grouping functions.

Next, to fully understand the role of property 2 in forming such local grouping functions, a comprehensive quantitative measure of the spatial relationships between pairs of source and target receptive fields were gathered. The results are not reported in this chapter, but instead, they will be presented in Chapter 6 figure 6.5 to highlight the importance of this property in contour integration tasks. The results showed striking similarity to edge co-occurrence statistics in natural images, and the connections properties were well-suited for contour integration tasks.

In this section, we saw that the lateral connections that develop in PGLISSOM have close match to biological data, and the *anatomical* properties of the lateral connections are also consistent with those found in the visual cortex. Quantitative measures were used to precisely compare the model and data a close match was found.

5.5 Conclusion

In this chapter, I explained how the PGLISSOM network self-organizes into orientation maps in MAP1 and MAP2. The two maps form (1) retinotopic mapping of the visual space, (2) orientation maps, and (3) functionally specific lateral connectivity. The self-organized orientation maps in MAP1 and MAP2 correspond closely, forming functional columns, and these properties match experimental data very well. The orientation preferences and the functional lateral connections play an important role in perceptual grouping tasks such as contour integration, which is the subject

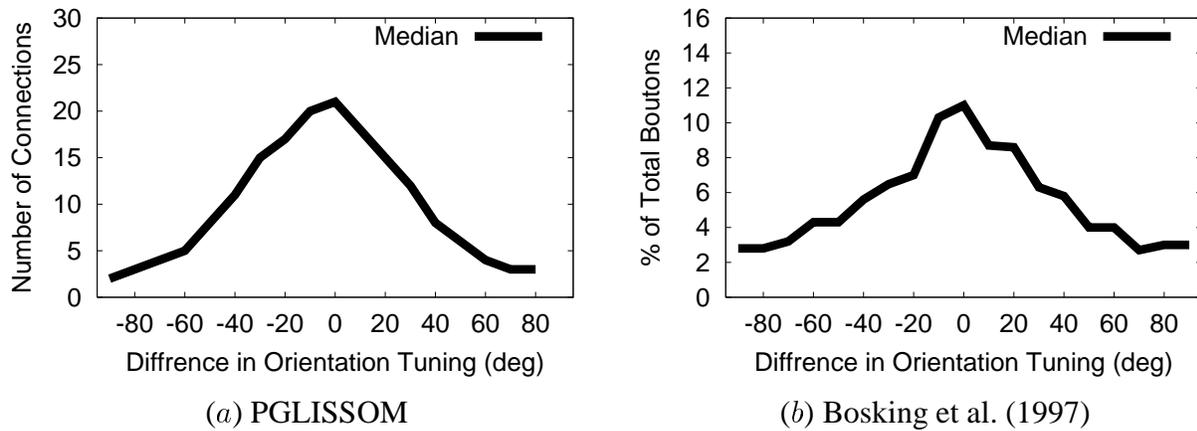


Figure 5.13: Difference in Orientation Preference vs. the Number of MAP2 Excitatory Lateral Connections. The relations between the number of connections and difference in orientation preference in the model and experimental data are shown. (a) The number of excitatory lateral connections in MAP2 that exceed a threshold value of 0.001 were counted, and averaged by the number of neurons. The median of the normalized count is plotted against the angular difference in orientation preference between the two connected neurons. (b) The same measurements in tree shrew visual cortex obtained by staining methods are shown. Adapted from Bosking et al. (1997). The x -axis is the difference in orientation preference of the neurons in degrees. The y -axis is the average number of connections. Both plots peak at 0 degrees, and quickly fall off as the orientation differences become larger. This plot shows that it is more likely to find strong excitatory lateral connections between neurons with similar orientation tuning.

of the next chapter.

Chapter 6

Contour Integration and Segmentation

As discussed in Chapter 1, contour integration means identifying a salient continuous contour consisting of separate edges (or contour elements) embedded in a scene cluttered with randomly-oriented background contour elements. Psychophysical evidence suggests that contour integration requires lateral interaction, and neurophysiological studies have revealed specific lateral connection patterns that might be implementing such functionality. In PGLISSOM, such lateral connections develop through a self-organizing process, as was shown in Chapter 5. In this chapter, I will show how such specific lateral connections combined with synchronized activities representing grouping can account for human contour integration performance, and also show how equally salient contours can be segmented via desynchronized activity.

6.1 Simulation Setup

As was discussed in section 1.2, contour integration accuracy in humans is maximal when orientation jitter in the global contour is 0 degree, and the accuracy decreases as a function of increasing orientation jitter. To verify that the PGLISSOM model exhibits a similar behavior, four contour integration experiments were carried out with varying degrees of orientation jitter.

The PGLISSOM network was trained as described in Chapter 5, with three differences in the setup. The size of the retina was 46×46 receptors so that the cortical magnification factor would be large enough to represent more orientations at the same retinal location. The intra-columnar connection radius from MAP2 to MAP1 was reduced to 1 to make the intra-columnar receptive field size proportional to the map. The size of the training input was scaled down to account for the reduced size of the retina: the length of the Gaussian bar a^2 was set to 45.0, and the width b^2 started out from 0.6, gradually reducing to 0.45. After self-organizing for 40,000 iterations, well ordered orientation maps and patchy lateral connections emerged, similar to those reported in Chapter 5. Lateral excitatory connections in MAP2 with weight ≤ 0.001 were killed afterward to model connection death as discussed in section 3.4.

With the trained network, four contour integration experiments were performed with 0, 30,

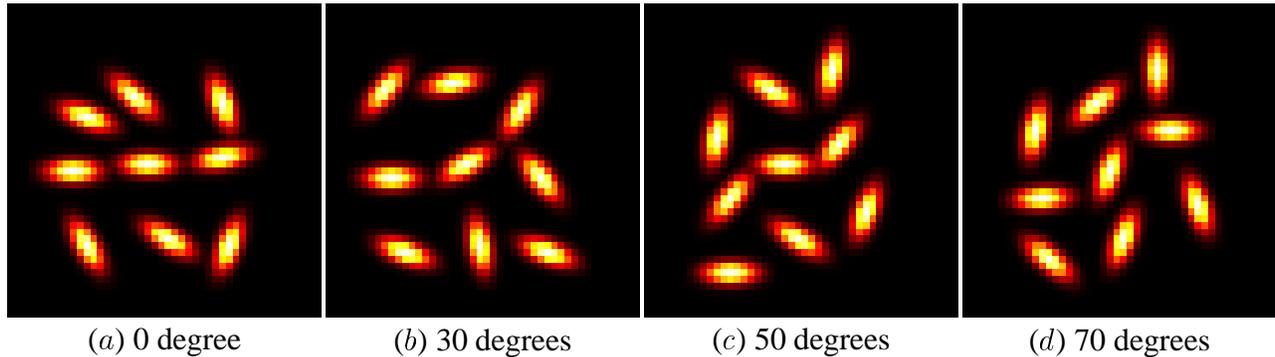


Figure 6.1: Inputs for the Contour Integration Experiments. Each contour element is an oriented Gaussian of length $a^2 = 3.5$ and width $b^2 = 1.5$. The activity levels of the retinal units for the inputs are plotted in color from 0 to maximum (black \rightarrow red \rightarrow yellow \rightarrow white). (a) The contour runs in the middle from left to right (0 degrees of orientation jitter). (b) The contour runs from mid-left to top-right (30 degrees of orientation jitter). (c) The contour runs in the middle from left to right (50 degrees of orientation jitter). (d) The contour runs in the middle from left to right (70 degrees of orientation jitter). It gets harder to identify the contour as the orientation jitter increases.

50 and 70 degrees of orientation jitter (figure 6.1). Each contour element was an oriented Gaussian of length $a^2 = 3.5$ and width $b^2 = 1.5$. For the training, a long Gaussian was necessary, but for the contour integration experiments, they were made short to fit into a single receptive field. The network configuration and parameters were the same as above except for the following changes. The excitatory learning rate α_e in MAP2 was set to 1.0, so that lateral connections could quickly adapt to assist the formation of synchronized populations (von der Malsburg 1981; Wang 1996). Fast adaptation of lateral excitatory connections allow the network to quickly adjust the weights remaining after connection death to a level that allows robust synchronization. As discussed in section 4.3.2, higher excitation results in more robust synchronization, and increasing the excitatory weights had the same effect. To help desynchronization and model synaptic noise, 2% noise was added (section 4.2.2). Previously, for self-organization, the absolute refractory period (κ_{abs}) was set to 0. The firing rates of the neurons were high as a result and the simulation proceeded in a fast time-scale. With a larger κ_{abs} , many more iterations would have been necessary to collect an equal number of spikes to compute an average, causing the simulation to run much longer. However, for the contour integration experiments, a finer degree of temporal resolution was necessary, i.e. we want to stretch time over many simulation iterations so that we can observe the synchronization behavior with more temporal precision. For this purpose, κ_{abs} increased to 1.5. In addition, larger κ_{abs} also helps make synchronization robust, as discussed in subsection 4.3.3. In all experiments in this Chapter, the setup described above was used.

6.2 Contour Integration Performance

Four different sets of inputs (figure 6.1) were presented to the trained network for 600 iterations each, which took about 14 minutes of wall-clock time. The average number of spikes during the 600 iterations was about 60. Assuming that the firing rate is 40Hz (the γ -frequency band; section 2.1.2), the 600 iterations roughly correspond to 1.5 seconds in real-time. The performance of the network was measured in MAP2, since it has the long-range patchy excitatory lateral connections for grouping, and synchronization was most robustly found in MAP2. For each contour element, a distinct area on the map became activated (or spiked). To measure the performance of the model, the degree of synchrony between these active areas needs to be calculated. First, the number of spikes generated in the area of the cortex that responded to each contour element was counted at each time step. As briefly mentioned in section 2.1.3, this quantity is called the Multi-Unit Activity of the response, or MUA, and it corresponds to the collective activity of a population of neurons representing a single feature in the input (Eckhorn et al. 1988, 1990). The MUAs gathered over 600 iterations formed MUA sequences.

Figure 6.2 shows such MUA sequences for each contour element. There are nine rows in each plot, corresponding to nine MUA sequences. The bottom three rows in each figure represent the three contour elements constituting the salient contour. By looking at how these three bottom rows are synchronized, we can determine the performance of the network on that particular set of inputs. For 0-degree orientation jitter, the three bottom MUA sequences are mostly synchronized (figure 6.2a), but as the orientation jitter increases (figure 6.2b through d), the synchronized state does not hold and the phases tend to shift back and forth.

In order to quantitatively measure the degree of synchronization between two areas, the linear correlation coefficient r between their MUA sequences was calculated as follows:

$$r = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2} \sqrt{\sum_i (y_i - \bar{y})^2}}, \quad (6.1)$$

where x_i and y_i , $i = 1, \dots, N$ are the MUA values at time i for the two areas representing the two different objects in the scene, and \bar{x} and \bar{y} are the mean MUA values of each sequence.

Using r as the measure, the contour integration performance of the network in the four different input configurations was calculated. The network was presented with each input for 600 iterations and the MUAs of MAP2 areas corresponding to the nine input contour elements were obtained. Then, the correlation coefficients between all possible pairs of MUAs were calculated. These coefficients represent the perceptual salience of such pairs. The higher these values are, the more synchronized the areas are, thus representing a strong percept of a contour. The average of the within-contour correlations was used as a measure of overall performance of the model (the variation was low). The results are summarized in figure 6.3, together with the human performance data from Geisler et al. (1999, 2001). The plot shows that at low orientation jitter, the model and human performance are both high, but as orientation jitter becomes larger, they both decrease in a similar manner. Correlation coefficients between MUA pairs corresponding to two background

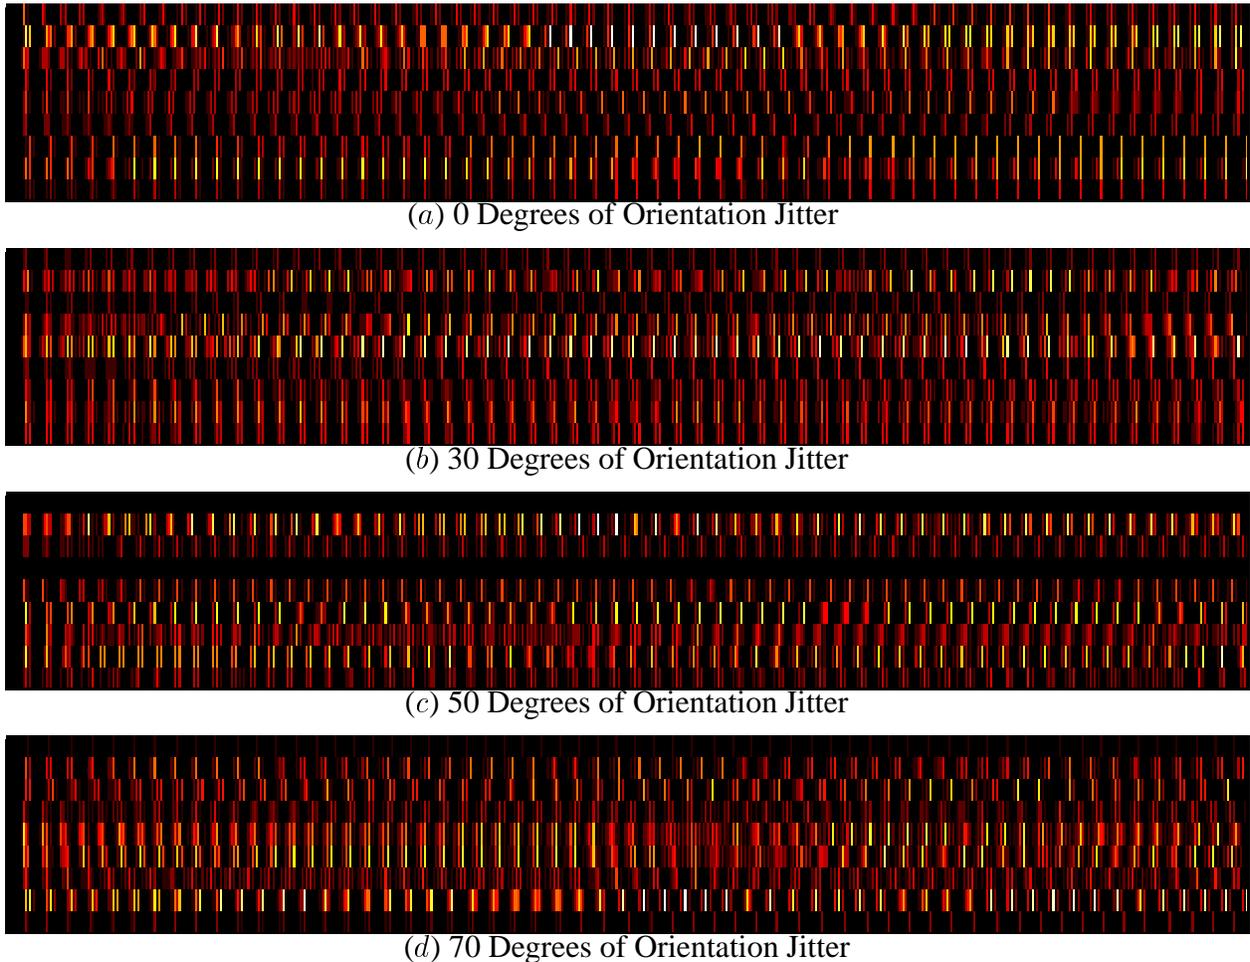


Figure 6.2: Multi-Unit Activities in the Contour Integration Experiments. The Multi-Unit Activities (MUAs) of the active areas are shown in color scale from low to high (black \rightarrow red \rightarrow yellow \rightarrow white). The x -axis is the simulation iteration and the y -axis, consisting of nine rows, is the index of the MUA area for the corresponding input. The bottom three rows represent the MUAs of the salient contour, and the top six rows represent MUAs of the random background contour elements. In all cases (a to d), the background MUAs are unsynchronized. The MUAs corresponding to the continuous contour are highly synchronized for 0 degrees of orientation jitter, but gradually become less synchronized as the orientation jitter increases. A quantitative measure reveals this tendency more clearly (figure 6.3).

contour elements, or pairs between a background and a contour element in the salient contour remained low, usually around 0 (not shown), thus they were not perceptually salient.

The neural mechanism for such a fall-off in perceptual performance with increasing orientation jitter is believed to be linked to the specific lateral connection patterns. Next, the lateral connection statistics of the model will be calculated to examine the connection properties and establish a link to perceptual performance.

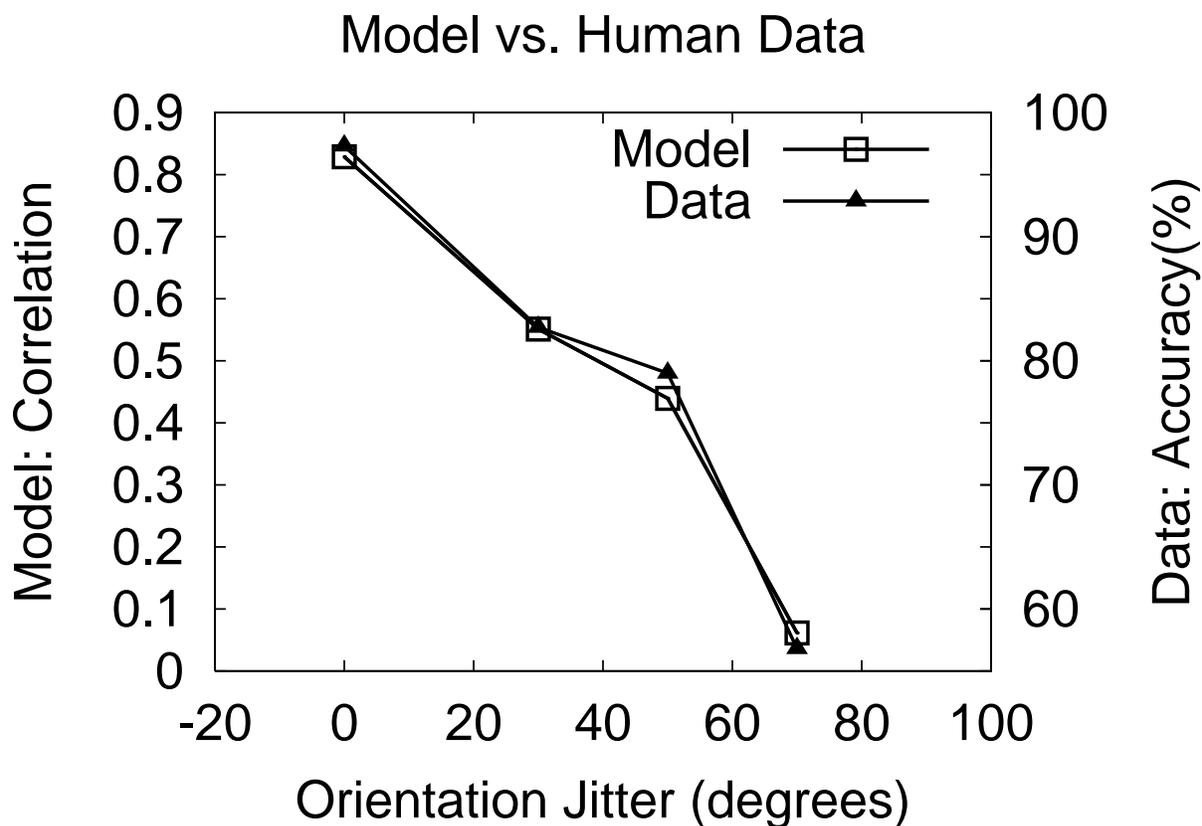


Figure 6.3: **Contour Integration Performance.** The contour integration performance of the model was measured as the average correlation coefficient between MUA sequences constituting the salient contour. The model performance is compared to human contour integration accuracy (Geisler et al. 1999, 2001; RMS amplitude 12.5, fractal exponent 1.5, which is the closest match to their data given our limited input configuration). The x -axis is the orientation jitter in degrees. The y -axis on the left is the correlation coefficient of MUA sequences in the model experiments, and the y -axis on the right is the human contour integration accuracy (percent correct). In both model and human experiments, the performance quickly drops as orientation jitter increases.

6.3 The Role of Lateral Connections

As we saw in Chapter 5, the lateral connections in PGLISSOM have two specific anatomical properties: (1) strong connections exist between neurons with similar orientation preferences, and (2) the connections extend in the direction along the axis of the source neuron's orientation preference. These properties are believed to allow the connections to encode specific local grouping functions (or association fields) for perceptual grouping tasks such as contour integration.

However, to understand the functional role of these connections in visual space (not in cortical space), the relationships between the *receptive fields* of the connected neurons need to be examined. By looking at the receptive fields and the lateral connections of the neurons, we can

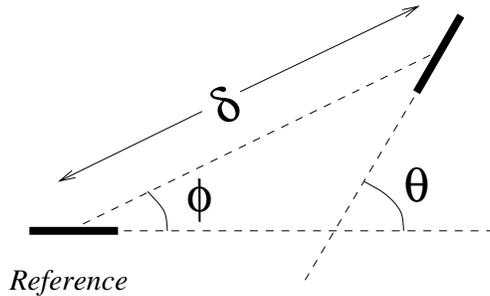


Figure 6.4: **Spatial Relationship between Receptive Fields.** For each excitatory lateral connection in MAP2 remaining after connection death, the spatial relationship between the two receptive fields (thick bars) is defined by three quantities: the difference in relative orientation preference (θ), the direction of target receptive field center relative to the preferred orientation of the reference neuron (ϕ), and the distance between the two receptive field centers (δ). The receptive field center was calculated as the center-of-gravity of the afferent weight matrix. Notice that these values measure the spatial relationship between the two neurons in the retinal (or visual) space, not in the cortical space. These measures were adopted from Geisler et al. (2001) for comparison with experimental data.

see which input features in a scene activate neurons that have strong lateral connections between them. In other words, we can estimate how strongly a pair of input features is bound together in the cortex through lateral connections. In reverse, for each lateral connection, we can see what the input features are that the connection binds together. By comparing the input features that are bound together by lateral connections to the perceptual performance of humans on these features, we can relate perceptual performance to neural structures.

To reveal the precise functional role of the patchy lateral connections, the spatial relationships between receptive fields of connected neurons were gathered from the lateral excitatory connections in MAP2 that remained after connection death. Figure 6.4 defines the quantities that summarize the spatial relationship between a pair of receptive fields. The results are summarized in figure 6.5a.

In figure 6.5a, the reference (or source) receptive field is positioned in the center, aligned along the horizontal axis. The relationship between the source receptive field and the target receptive fields are plotted in two ways: (1) the most probable orientation of the target receptive field (θ) at location (ϕ, δ) is plotted as a black oriented bar, and (2) the relative log-probability (or normalized log-probability) of finding a target receptive field at location (ϕ, δ) in polar coordinates is shown in color scale. Two properties are evident in the plot: (1) the most probable receptive field orientation θ (black oriented bars) of the target neurons are aligned along co-circular paths (not shown) emanating from the center and (2) the locations (ϕ, δ) with high relative probability (the red, yellow, and green areas) form a bow-tie shaped flank along the horizontal axis. In other words, for the neurons connected with strong lateral connections, most target receptive fields are (1) aligned on an arc passing through the source receptive field, (2) they are mostly concentrated on arcs with small curvatures, and (3) they are within a limited spatial extent (i.e. establish local interaction).

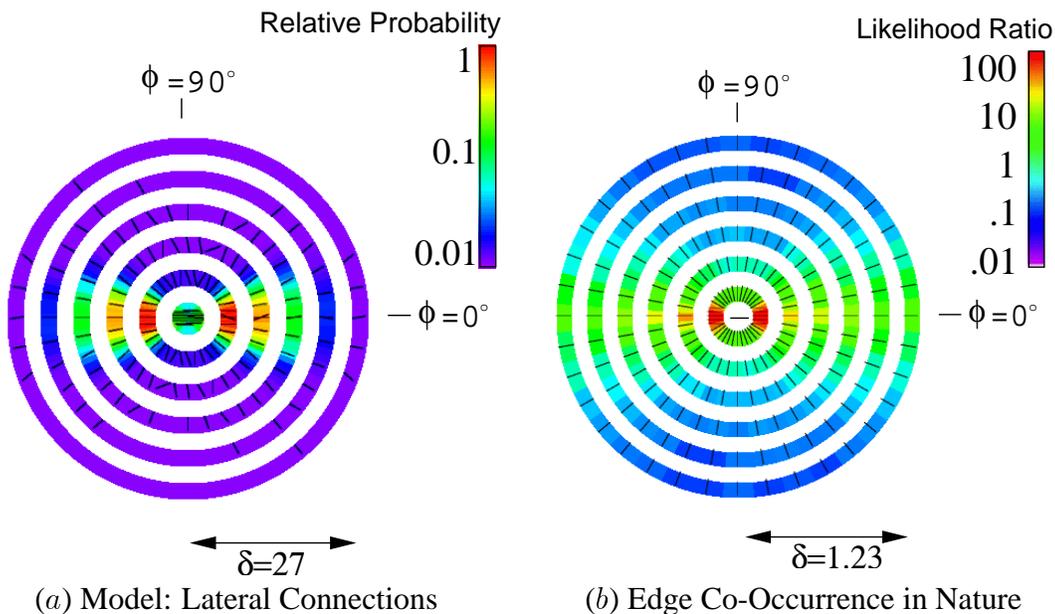


Figure 6.5: **Lateral Connection Statistics and Edge Co-Occurrence Statistics (color figure).** The lateral connection statistics in the model and the edge co-occurrence statistics in nature are compared to see how they relate to perceptual requirements. (a) The θ , ϕ , and δ values for the target receptive fields of lateral excitatory connection in MAP2 are shown (see figure 6.4 for the definitions). The plot at any location (ϕ, δ) consists of two components: (1) the black oriented bars representing the most probable orientation of target receptive field θ , and (2) the color scale polar plot in the background showing the relative log-probability of finding a receptive field at each location (ϕ, δ) . The color scale represents the interval $[0.01, 1.0]$ (purple \rightarrow blue \rightarrow cyan \rightarrow green \rightarrow yellow \rightarrow red). For weak connections, the oriented bar representing θ was not plotted. Two properties are evident in the plot: (1) the most probable target receptive field orientations θ (black oriented bars) are aligned along co-circular paths (not shown) emanating from the center and (2) the positions (ϕ, δ) with high relative probability (the red, yellow, and green areas) form a bow-tie shaped flank along the horizontal axis. Such an arrangement is very similar to the local grouping functions (or association fields) suggested by psychophysical research. (b) The Bayesian edge co-occurrence statistics. Adapted from Geisler et al. (2001). The plot shows the likelihood ratio of the conditional probability of a pair of edge elements in a given configuration (ϕ, θ, δ) belonging to a *same physical contour* vs. *different physical contours* in natural images. A high likelihood ratio indicates that a pair of edges in that configuration is more likely to occur on a *common natural contour* than on *separate contours*. The conditional probabilities were determined through manual labeling of contours in real world images. There is a strong correspondence between this data and the connection statistics in the model, suggesting that the connections encode the statistical regularities in the environment through self-organization.

These statistics show that neurons with receptive fields falling upon a common smooth contour are most likely to be connected with lateral excitatory connections. Such a pattern closely matches the *association field* proposed by Field et al. (1993a; figure 1.2), thus suggesting that perceptual grouping rules can actually be implemented in the intricate patterns of lateral connections in the brain.

In fact, such connection patterns actually predict the contour integration performance shown in section 6.2. Since receptive fields aligned on an arc with smaller curvatures prefer to connect laterally, inputs with smaller orientation jitter would be more strongly bound together than those with large orientation jitter. Such a comparison between anatomy and performance helps us gain insights into how specific neural structures account for perceptual performance.

Furthermore, these functional statistics in the model are similar to the local Bayesian edge co-occurrence statistics in natural images (Geisler et al. 2001; figure 6.5b)¹. The edge co-occurrence statistics shown in the figure summarize the probability of a pair of edges under configuration (ϕ, θ, δ) falling upon a common physical contour, such as the contours of tree trunks, large boulders, etc. The properties found in the PGLISSOM connection statistics are also found in the edge co-occurrence statistics. In fact, Geisler et al. (2001) showed that these edge co-occurrence statistics accurately predict human contour integration performance.

The above comparison suggests that the oriented edge features bound together by lateral connections in PGLISSOM are also most likely to be found on a common natural contour. However, it is surprising that connections with such grouping rules arise in a network that was trained with only straight Gaussian inputs. This result suggests that self-organizing the visual cortex with inputs as simple as elongated Gaussian bars prepares an animal for the visual environment even before birth.

In sum, by comparing the statistics, we can establish a link between the statistical structures in nature and the anatomical structures in the brain. Thus, together with the comparison of perceptual performance and connection statistics, we can bring together (1) the statistical regularities in the natural environment, (2) the anatomical structures in the cortex, and (3) the perceptual performance into an integrated model, where one aspect can be inferred from the other.

6.4 Contour Segmentation Performance

Importantly, the synchronization process that establishes the contour percept can also separate different contours to different percepts. The same self-organized network with the same simulation parameters as listed in section 6.1 were used for the contour segmentation experiment. Three nearly collinear contours in a random background were presented as input (figure 6.6) and the correlations between and across the MUAs representing the input contour elements were calculated.

Figure 6.7 shows the actual MUA sequences of the 9 areas in color code. The bottom three rows correspond to the contour at the bottom, the middle three rows to the contour in the middle,

¹Similar edge co-occurrence statistics were also reported by Sigman et al. (2001).

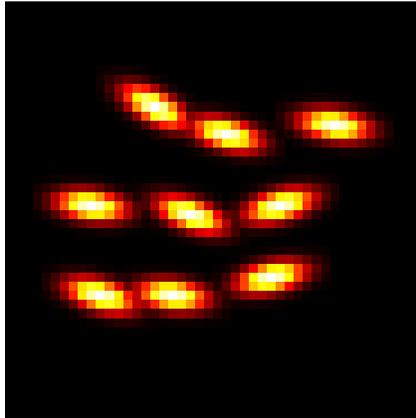


Figure 6.6: **Input for Contour Segmentation.** The input for the contour segmentation experiment consisted of three long horizontal contours, from the top row to the bottom row. The color code and the input sizes are identical to that in figure 6.1.

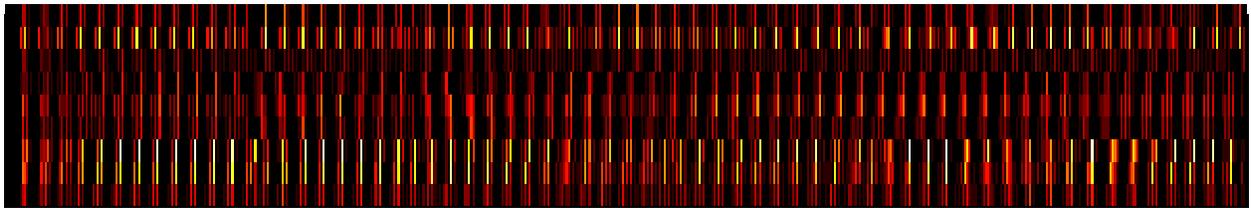


Figure 6.7: **Multi-Unit Activities for the Three Contours.** The MUAs corresponding to the nine contour elements in the input are shown. The plotting conventions are identical to those in figure 6.2. The bottom three rows correspond to the three contour elements in the bottom of the input, and likewise for the middle three rows and top three rows (cf. figure 6.6). The MUA sequences of each group of three are synchronized, and between groups, the activities are desynchronized. In other words, three areas representing the same contour fire together, while other areas remain silent. Such an alternating activation of neuronal groups ensures that a coherent object is represented and is not mixed with representations of other coherent objects.

and the top three rows to the contour at the top in figure 6.6. By comparing the rows in the plot, we can see that in the beginning all areas are mostly synchronized, but as lateral interactions start to take effect, the MUAs start to form three major groups firing in three alternating phases. The correlation coefficients of areas in the same contour are consistently high while those in different contours are very low (figure 6.8), signifying integration within each contour and segmentation across the two contours. This result suggests that the same circuitry responsible for contour integration can also be responsible for segmentation between multiple salient contours.

In the segmentation experiment, PGLISSOM was used to segment three objects, i.e. the three contours. The number of objects that can be segmented in PGLISSOM reaches up to about 6, but above that, representations for some objects will be synchronized instead of being desynchronized. A similar limitation was reported in Horn and Opher (1998) and it was proposed that such limitation may actually account for the limited number of short term memory slots of 7 ± 2 discovered by Miller in 1956 (Horn and Usher 1992). However, in a separate experiment with

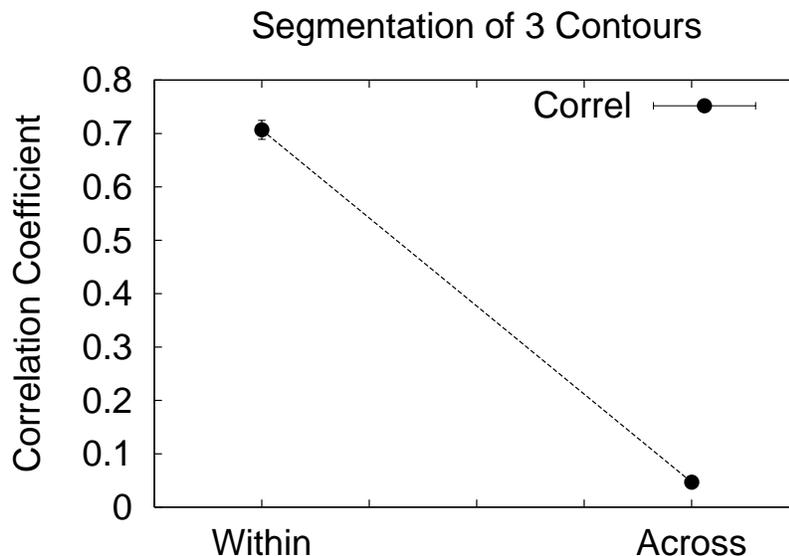


Figure 6.8: **Contour Segmentation Performance.** The average correlation coefficients between pairs of MUA sequences (1) within the same contour and (2) across different contours are plotted. For MUA sequence pairs in the same contour, the correlation is high, but for those belonging to different contours, the correlation is low. This means that neurons within a group are synchronized, while neurons belonging to different groups are desynchronized.

PGLISSOM (not presented in this thesis), it was shown that even when different representations are synchronized, the synchronized state is not permanent. Synchrony eventually breaks, and another pair of representations that was previously desynchronized become synchronized. Therefore, theoretically, even with the small capacity for segmentation, any number of objects can be segmented if the degree of synchrony is measured over a long period of time.

6.5 Conclusion

In this chapter, I have shown how synchronized representations, together with specific, self-organized lateral connections, gives rise to contour integration performance similar to human performance. This result is due to the lateral connectivity pattern that prefers co-circular grouping. The lateral connection statistics were gathered and compared to Bayesian edge co-occurrence statistics in natural images, and were found to be very similar. Correlation coefficients of MUA sequences were used to measure the model performance. Such a measure can be used in neurophysiological experiments to establish a solid link between perceptual experience and neuronal activity.

Chapter 7

Hemifield Differences in Anatomy and Performance

Psychophysical experiments have shown that performance differences exist in illusory contour discrimination and contour integration tasks in different parts of the visual field. PGLISSOM can provide a possible explanation for such differences. As we saw in Chapter 6, the lateral connection patterns in PGLISSOM determine its contour integration performance. Since these patterns are learned through input-driven self-organization, different input distributions in different parts of the visual field will result in different patterns in the corresponding map areas. Such differences in connectivity, in turn, will cause the contour integration performance to differ in different parts of the visual field. In this chapter, I will show how altering the input distributions affects the self-organized structure and contour integration performance of the PGLISSOM model, providing a possible developmental explanation for the observed performance differences in different parts of the visual field.

7.1 Motivation

Perceptual performance has been shown to differ in different parts of the visual field in two ways: (1) lower visual hemifield is more accurate in illusory contour discrimination tasks than upper visual hemifield (Rubin et al. 1996), and (2) contour integration is easier in fovea than periphery (Hess and Dakin 1997). Studying these specific phenomena can provide us with insights into how and why functional divisions occur in the brain in general. In this chapter, I will focus on the lower vs. upper visual hemifield difference, but will also discuss the case of fovea vs. periphery in Chapter 9.

Rubin et al. compared the performance of humans on discriminating the angle made by illusory contours in the lower vs. upper hemifield (see figure 8.1*a* for an example stimulus). The pacman-like disks were rotated by small amounts so that the perceived square in the middle would look either thick (like a barrel) or thin (like an hour glass). Rubin et al. presented the inputs in

either the lower or upper visual hemifields and measured the minimum amount of rotation (i.e. threshold) needed for the subject to reliably tell whether the input was thick or thin. The results showed that the threshold is much higher in the upper hemifield than in the lower hemifield, i.e. the performance was higher in the lower hemifield than in the upper hemifield. Similar results are expected for contour integration tasks, although such experiments have not been carried out so far.

The questions are then, (1) how and (2) why do such differences in performance occur? A possible answer to the first question is that the structure may differ among the cortical areas mapping the different parts of the visual field. The second question then becomes, why do the structures differ? In a self-organizing map model such as PGLISSOM, structural differences are caused by differences in the input distribution, which suggests that the answer may lie in the kinds of inputs that different cortical areas receive.

The inputs can be differed in two ways: (1) due to passive environmental biases and (2) due to active attentional biases. There is evidence for both cases. For example, in animals with high dexterity such as monkeys, the lower hemifield would receive much more detailed input features compared to the upper hemifield because the animal manipulates objects mostly in the lower hemifield, thus causing the input distribution to be different in the lower vs. upper hemifield (Gibson 1950; Nakayama and Shimojo 1992; Previc 1990). Second, there is evidence that human saccades and attention cause input distributions to be different in the fovea vs. periphery. Zador and Pearlmuter (1996) presented natural images to humans and gathered statistics about the locations in the image to which the human attended, by tracking eye movements. They showed that the areas of attention had higher than average luminance contrast. Since the attended areas mostly project to the fovea, the statistical properties will differ in the fovea vs. the periphery. Such evidence suggest that input statistics can differ in different parts of the cortex due to attentional bias, and provides a possible cause for the differential development of areas.

In the following, the validity of the developmental explanation will be tested in computational experiments using PGLISSOM.

7.2 Simulation Setup

To test the effect of altered input distributions on the structure and function of the model, the frequency of input presentation in the lower vs. upper hemifield was made different. For each training iteration, a randomly oriented elongated Gaussian bar was presented to the visual field at a random location. In the lower visual hemifield, the input was kept intact. In the upper visual hemifield, the input was only allowed to appear 50% of the time, at uniformly random iterations. In other words, the pixels in the upper visual field were turned off 50% of the time.

The PGLISSOM network was trained otherwise with the same configuration as in Chapter 5.

To make it easy to interpret the results, the upper visual hemifield was projected to the upper half of the PGLISSOM network, and the lower visual hemifield to the lower half.

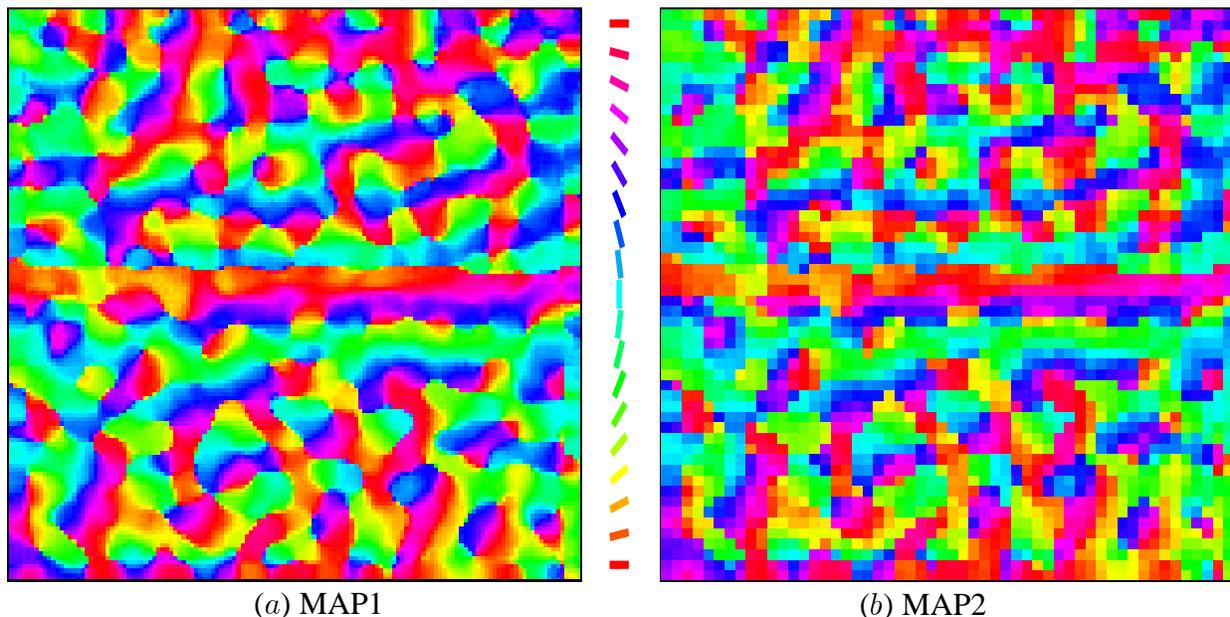


Figure 7.1: **Orientation Maps in Lower vs. Upper Half of PGLISSOM (color figure).** The orientation preference of each cortical neuron was calculated using the vector sum method (section 5.3) for (a) MAP1 and (b) MAP2. The plotting convention is the same as in figure 5.3 where the color represents orientation preference. The orientation maps are comparable to the previous results. The red horizontal band in the middle of each map is an artifact and will be excluded from the analysis in the rest of the chapter. Also, the lower and upper portions of the map are very similar, thus the difference in input distribution affect the orientation preference.

7.3 Differences in Afferent and Lateral Connections

After training the network for 40,000 iterations, orientation maps comparable to those in the previous self-organization experiments emerged (figure 7.1). However, there is a noticeable artifact. In the middle of each map, there is a boundary between the lower and upper half. These neurons have receptive fields that span both sides of the boundary. They receive more input in the lower half, and as a result, their receptive fields become half-circles. Such receptive fields prefer horizontal inputs, resulting in horizontal preference bands in the middle of each map. Because this band is just an artifact, the area surrounding the band will not be used in the analysis in the rest of this chapter.

As expected, two kinds of differences emerged in the lower vs. upper half of the cortex: difference in (1) orientation selectivity and (2) lateral connection statistics.

Selectivities of neurons in MAP1 was calculated using the vector sum method described in section 5.3, and the results are plotted in figure 7.2a. Orientation selectivity in MAP2 are similar (not shown). We can see that the lower half is brighter than the upper half of the cortex, suggesting that orientation selectivity is generally higher in the lower half of the map. The orientation selectivity histogram demonstrates this tendency more clearly (figure 7.2b). The histogram for the

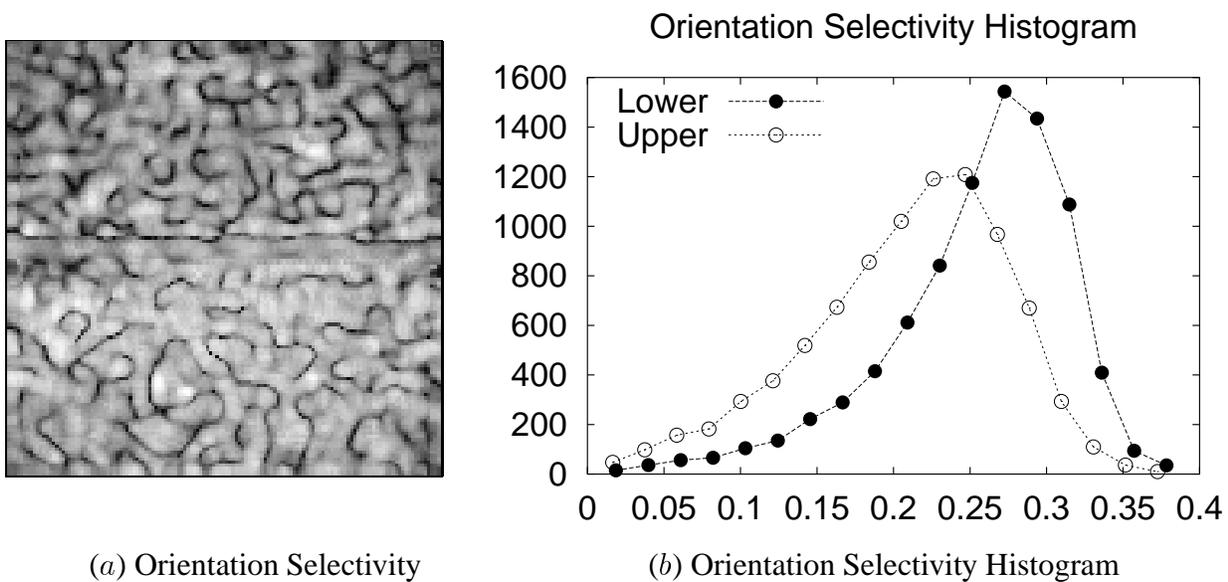


Figure 7.2: Orientation Selectivity in Lower vs. Upper Half of MAP1. The orientation selectivity of each MAP1 neuron was calculated using the vector sum method (section 5.3). (a) Orientation selectivity of each neuron is plotted in gray-scale from low to high (black \rightarrow white). The upper half is darker, suggesting less selectivity for the upper visual hemifield. (b) Orientation selectivity histograms for lower (filled circles) and upper (open circles) half of the cortex are plotted. The 136×10 horizontal area in the middle was excluded from the histogram calculation to avoid the band artifact from skewing the histogram. For neurons in the upper half, the histogram peaks at around 0.27, while for the lower half neurons, the peak is around 0.24. The histograms show that the neurons in the lower half, exposed to more frequent inputs, develop higher orientation selectivity than the upper half neurons. Thus, sharp edges in the visual field will draw higher response from the lower half of the map than from the upper half.

lower half (filled circles) is skewed toward higher orientation selectivity, compared to the upper half (open circles), confirming the above observation. Thus, sharply orientation-tuned edges will generate higher response if placed in the lower visual hemifield than in the upper visual hemifield.

The difference in orientation selectivity can be explained in terms of Hebbian learning. In the beginning of the training, the circular receptive fields are initialized to have random weights, and these receptive fields have very low orientation selectivity. As they are trained with elongated Gaussian bars, they start to take on an oriented shape. Early on, the receptive fields look like thick Gaussians, but as they get trained by more thin Gaussians, they will become thinner. The upper half of the map trained less, and the receptive fields tend to remain thicker than those in the lower half of the map.

We have seen that the frequency of input presentation alone can cause the afferent connection properties to differ in the two areas of the cortex. The lateral connections in PGLISSOM also self-organize according to the input, thus the difference in input distribution causes the final structure of lateral connections to differ as well. To compare the properties of the lateral connection

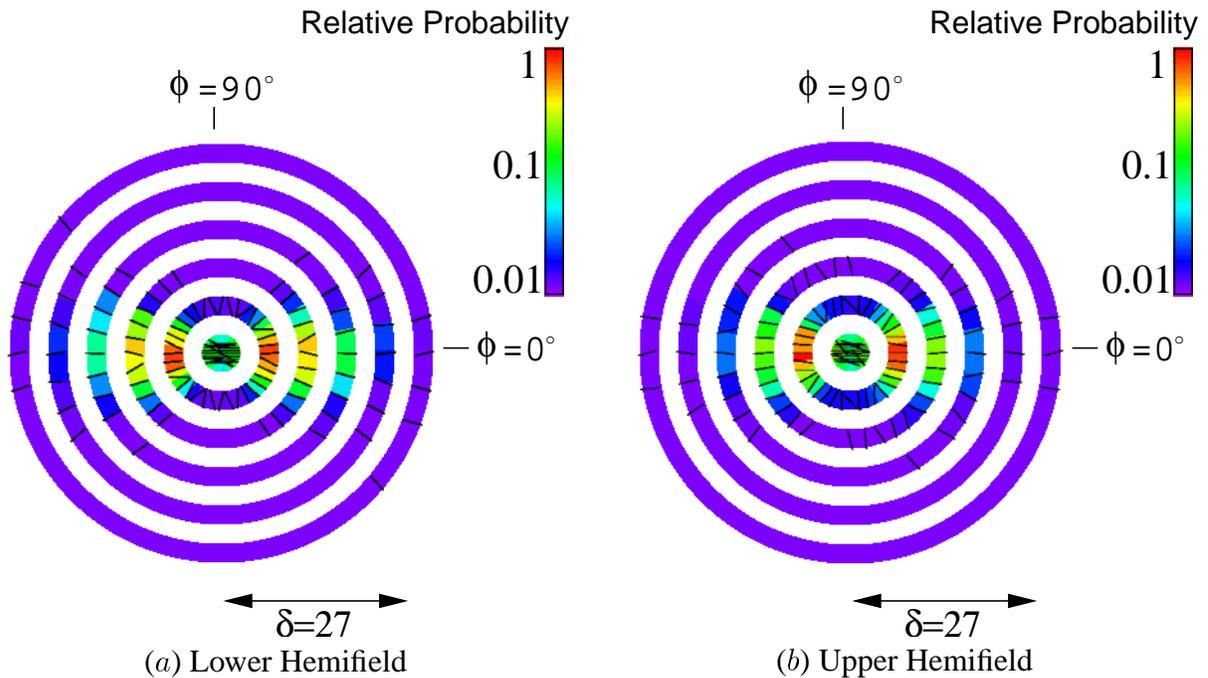


Figure 7.3: Excitatory Lateral Connection Statistics in Upper vs. Lower Half of MAP2 (color figure). The excitatory lateral connection statistics in MAP2 were gathered and plotted as in Chapter 6, for the (a) lower and (b) upper half of the map. Again, the 136×10 horizontal area in the middle was excluded. There are two noticeable differences: (1) high probability areas (green, yellow, and red) along the horizontal axis in (a) are longer than in (b), and (2) the most probable θ (black oriented bars) are co-circular in the lower half (a), but in the upper half (b) they are mostly collinear or aligned on separate parallel paths. Such results predict that curved contours may be easier to detect in the lower visual hemifield than in the upper hemifield.

patterns in the lower vs. upper half of the map, the (ϕ, θ, δ) statistics was calculated as in Chapter 6. The results are shown in figure 7.3. There are two major differences in these statistics. First, the high probability areas extend out longer on the horizontal axis for the lower half of the map than the upper half. Thus, the lower half can group more distant inputs than the upper half. Second, the most probable θ for a given (ϕ, δ) location shows a more co-circular tendency in the lower half of the map, while in the upper half, it is more collinear (the black edges are more parallel). Thus, the difference suggests that in the lower half of the map (lower visual hemifield), (1) lateral connections exert influence over longer distance in the visual field, and (2) curved contours are easier to detect than in the upper half (upper visual hemifield). The only difference during training was the frequency of input presentation, thus such a difference could be the only cause for the differing lateral connections.

As we saw in section 5.4, collinearity is the most prominent feature learned by lateral connections, but co-circularity can also be learned due to the co-activation of neurons with misaligned receptive fields (figure 5.11). However, the upper half of the map could only learn the primary property (collinearity) because of the scarce input presentations. Co-circularity develops slower

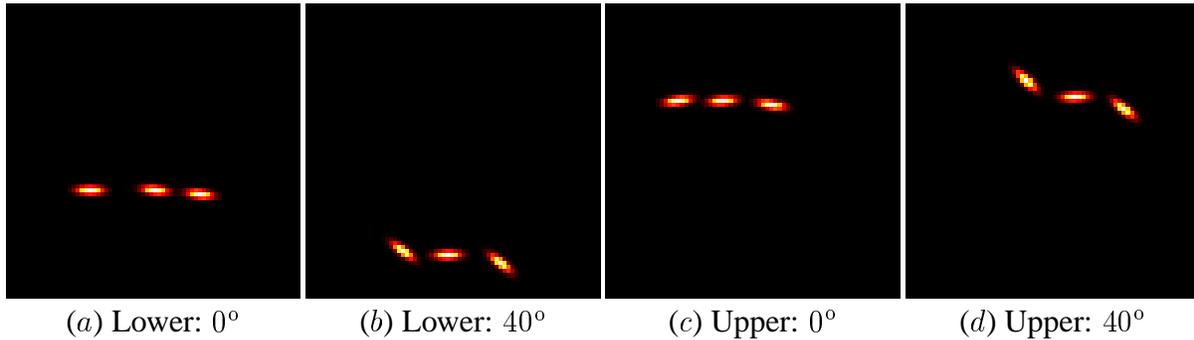


Figure 7.4: **Inputs in the Lower and Upper Hemifields.** The inputs for the four contour integration experiments are shown, consisting of oriented Gaussians of size $\sigma_x = 3.5$ and $\sigma_y = 1.5$. The intensity of the input is coded in color scale, from low to high (black \rightarrow red \rightarrow yellow \rightarrow white). The location of the contour (lower vs. upper hemifield) and the orientation jitter (0-degree and 40-degree) were varied: (a) Lower hemifield, 0 degrees of orientation jitter. (b) Lower hemifield, 40 degrees of orientation jitter. (c) Upper hemifield, 0 degrees of orientation jitter. (d) Upper hemifield, 40 degrees orientation jitter.

than collinearity because of the activities are lower in the co-circular arrangement. The lower half of the map had enough input presentations and was able to learn the secondary (co-circularity) property as well.

In sum, the areas of the map that receive less frequent input have lower orientation selectivity and the lateral connection statistics show more collinear than co-circular characteristics. Such difference in structure predicts that contour integration performance will differ in the two areas, which will be tested in the next section.

7.4 Differences in Contour Integration Performance

Four contour integration experiments were performed to see if there is a difference in performance in the lower vs. upper half of the map trained with different input presentation frequencies. The network trained in section 7.3 was used for these experiments, with the same setup as described in section 6.1. The only difference was the input layout (figure 7.4).

For each experiment, the network was activated for 600 iterations and the MUA sequences corresponding to the three contour elements were gathered. Figure 7.5 shows the complete MUA sequences. The MUAs are more synchronized for the lower-hemifield inputs than for the upper-hemifield inputs even when the orientation jitter is the same (compare figure 7.5a vs. c, and b vs. d). The correlation coefficients of these MUA sequences confirm this observation (figure 7.6). From this plot, we can see that the MUAs in the lower half of the map, are more synchronized (i.e. the MUA values are vertically aligned in the plot) than the MUAs in the upper half of the map. The results show that in both 0 degree and 40 degrees of orientation jitter, the lower half of the map performed better than the upper half. Another noticeable feature in the plot is that the performance gap is more pronounced in the 40-degree case than the 0-degree case.

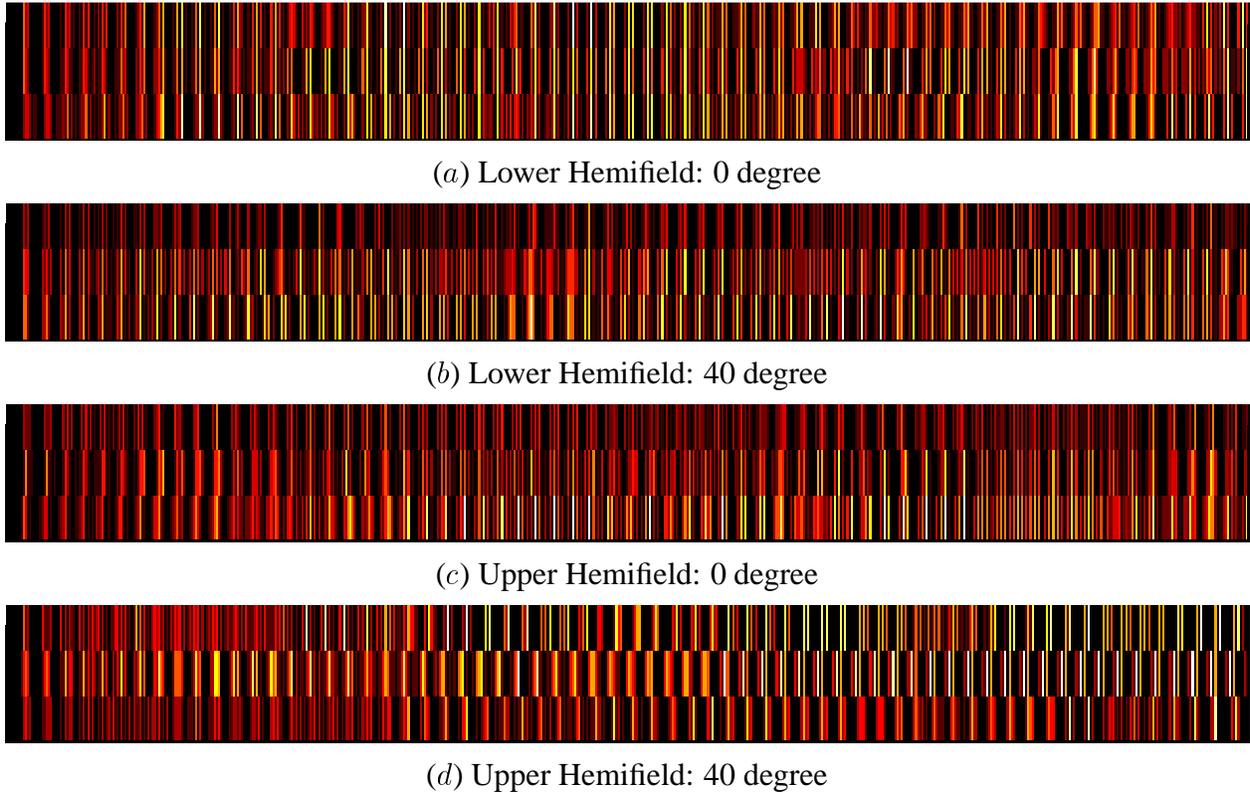


Figure 7.5: Multi-Unit Activities in the Lower vs. Upper Half of MAP2. The MUAs for the four input conditions in figure 7.4 are shown. The magnitude of the MUA is coded in color scale, from low to high (black \rightarrow red \rightarrow yellow \rightarrow white). In each plot, the three rows correspond to the MUA sequences for the three contour elements in the input. For the same orientation jitter range (0 or 40-degree), we can see that the lower hemifield shows a higher degree of synchrony; for example, compare (a) vs. (c), and (b) vs. (d). The correlation coefficients between the MUA sequences in each experiment confirm this observation (figure 7.6).

Such a performance difference is predicted by the afferent and lateral connection patterns presented in 7.3. For 0-degree orientation jitter, the difference is mostly due to the difference in afferent connection properties, i.e. orientation selectivity. In both lower and upper half of the map, lateral connections can group collinear contours. With 40-degrees of orientation jitter, the performance gap is larger than in the 0-degree case. Since each neighboring pair of contour elements can be aligned on a co-circular path and lateral connections in the lower half of the map can group co-circular contours, contours of even high orientation jitter can be detected. In contrast, lateral connections in the upper half can only group collinear contours. Thus, the larger performance gap is due to the different lateral connection patterns.

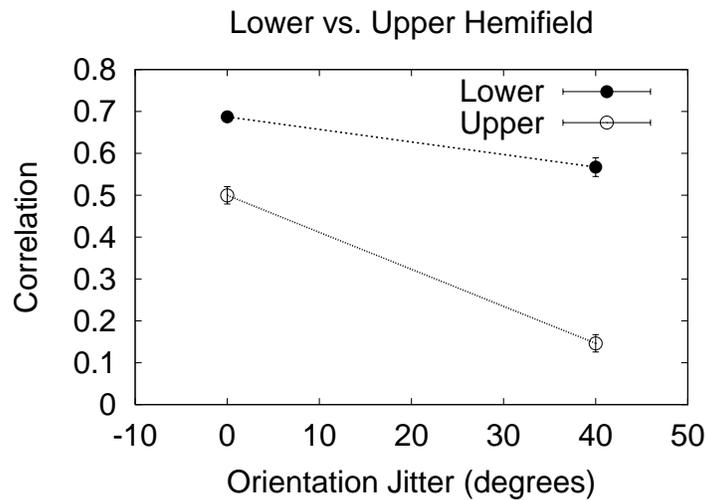


Figure 7.6: Contour Integration Performance in Lower vs. Upper Hemifield. For each input consisting of 3 contour elements, the correlation coefficients between the MUA sequences in figure 6.3 were calculated, and the average was used as the measure of performance, as in Chapter 6. For both 0-degree and 40-degree orientation jitter, lower hemifield (filled circles) had higher correlation than the upper hemifield (open circles). The performance gap in the lower vs. upper hemifield is more pronounced in the 40-degree case, as predicted by the lateral connection patterns in figure 7.3.

7.5 Conclusion

In this chapter, I have shown that altering the input distribution as simple as frequency of presentation can change the structural organization of the maps, and this in turn affects contour integration performance. Such differences in structure and function are due to the input-driven nature of self-organization, and provide a possible developmental explanation for the hemifield differences in contour integration found in psychophysical experiments.

Chapter 8

Illusory Contours and Contour Completion

Illusory contours are clear borders that are perceived in images that do not have luminance contrast corresponding to the borders. There is experimental evidence that such phenomena occur early on in the visual system (V1 and V2, i.e. the primary and secondary visual cortex), the same level where contour integration is believed to take place. Lateral interactions may be necessary in illusory contour detection tasks as well as in contour integration tasks. In this chapter, I will show how the same self-organized lateral connections in PGLISSOM used in contour integration can account for certain illusory contour phenomena.

8.1 Motivation

Examples of illusory contours are shown in figure 8.1, where (a) the triangle, and (b) the circle at the center are immediately visible, but there are no luminance borders around these objects. Following the initial discovery by Schumann (1904), illusory contours were made popular by Ehrenstein (1941) and Kanizsa (1955), and they have become an important subject in visual perception because of their relation to figure-ground separation, occluded object recognition, and perceptual grouping in general (Leshner and Mingolla 1995).

Early on, those were the main theory about the cause of illusory contour perception: (1) bottom-up brightness theory and (2) top-down cognitive factor theory. Brightness theory maintained that illusory contours arise from a low-level mechanism that gives illusory brightness to areas enclosed by illusory borders. On the other hand, cognitive theorist argued that illusory contours are purely a high-level cognitive phenomenon. However, evidence has started to appear suggesting that neither of these theories can account for the full range of illusory contour phenomena. Kanizsa (1976), Parks (1980), and Prazdny (1983) discovered that illusory contours can arise from image configurations without subjective brightness, providing a counterexample to the brightness theory (see Hoffman 1998 for an overview). On the other hand, Peterhans et al. (1986) and von der Heydt and Peterhans (1989) showed that V2 cells respond to illusory contours, and Redies et al. (1986) and Sheth et al. (1996) showed that so do V1 cells. Such experimental results showed that illusory

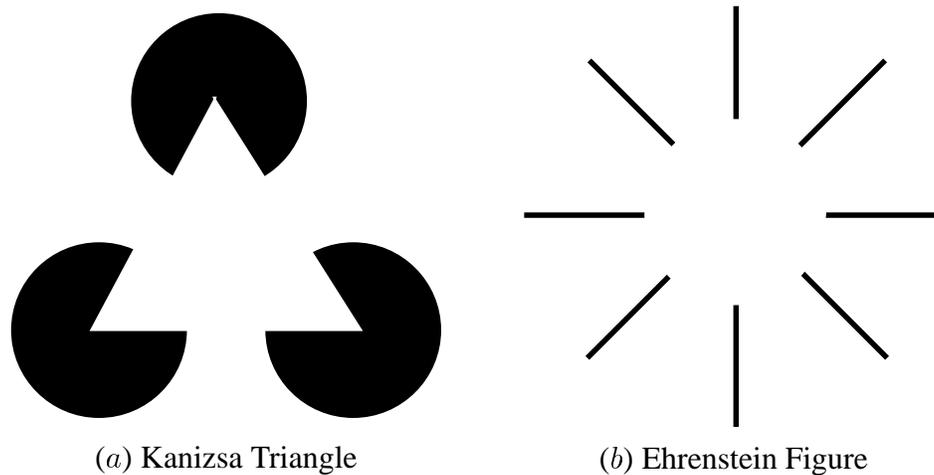


Figure 8.1: **Two Types of Illusory Contours.** Illusory contours formed along two different stimulus configurations are shown: (a) empty space between separate, but collinearly aligned edges (the three sides of the bright triangle in the middle), or (b) line-ends (the border of the bright circle in the middle). The edges belonging to the mouths of the three black pacmans in (a), and the line-ends around the illusory circle in (b) are thought to be responsible for giving rise to (or *inducing*) the illusory percept. Such edges and line-ends are commonly known as *edge inducers* and *line-end inducers* of illusory contours (Leshner and Mingolla 1995).

contour can be processed early on, unlike what the cognitive theories suggested.

Meanwhile, neural network models based on low-level neural architectures were introduced. These models were based on the observation that illusory contours can be triggered by two types of stimulus configurations: *edge inducers* and *line-end inducers* (see figure 8.1 for a definition of these terms). The Kanizsa triangle in figure 8.1 is a representative of edge-induced illusory contours, where the contour forms parallel (or collinear) along the inducing edges of the pacmans. The Ehrenstein figure in figure 8.1 is a good example of line-end induced illusory contours, where the border of the circle is orthogonal to the line ends near the center.

The first neural network model of illusory contour was based on edge inducers, and illusory contours were processed through contour completion (Ullman 1976). Figure 8.2 shows an example of how such an edge-induced contour can be formed by contour completion. However, subsequent models were based on line-end inducers: they modeled illusory contours defined by edge inducers as a special case. In these models, the corners where the edges meet (e.g. inside the mouths of the pacmans in figure 8.1a) and the tip of convex angles (e.g. the tip of the mouths of the pacmans) were processed as line-ends with good results (Peterhans et al. 1986; Finkel and Edelman 1989). Shipley and Kellman (1992) later found that models strictly based on line-end inducers cannot account for psychophysical results where increasing the length of the inducing edge causes the illusory contour to become clearer. The line-end inducer models predicted that such an increase in clarity would not happen. Thus, no single model could account for both inducer types.

In the model developed by Grossberg and Mingolla (1985), the two inducer types were

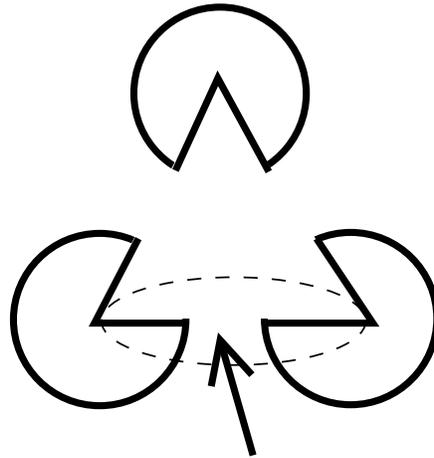


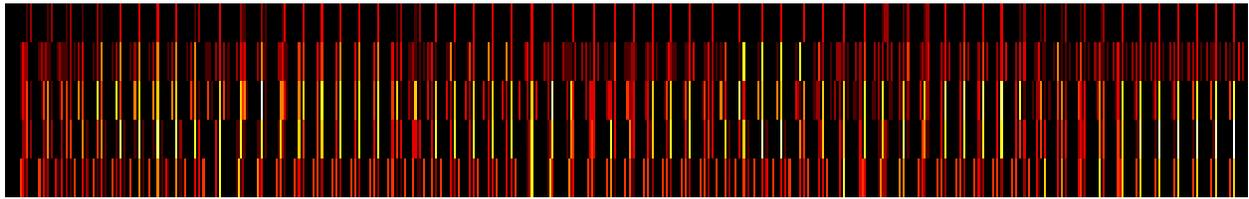
Figure 8.2: **Contour Completion across Edge-Inducers.** Contour completion is a possible mechanism for illusory contours arising from edge inducers. After the retinal and thalamic preprocessing, the inputs received by the primary visual cortex cells become similar to edge-detected versions of the original images. An example of the edge-detected image of figure 8.1a is shown above in solid lines. The three sides (one side is marked with a dashed oval) of the edge-detected illusory triangle in the center have gaps (such as the one pointed by an arrow) in the middle of each side. A triangle boundary can be generated by completing the discontinuous contours on the three sides.

incorporated. In the first stage of their model, borders formed orthogonal to the line-end inducers, and in the second stage, borders formed parallel (or collinear) to both line-end and edge inducers. The model was based on known neurophysiological data and was able to predict and confirm various illusory contour phenomena. In the second stage, neurons called *cooperative bipole cells*, with bow-tie shaped receptive fields were used (figure 8.3). These receptive fields were formed by combining the responses of two first stage neurons. Such a neuron activated only when stimuli were present on both lobes of receptive fields. This property is called the *bipole property* and the neurons with such a property can be used in contour completion. Neurons with such properties are found in V2 (von der Heydt and Peterhans 1989), but in V1, a conclusive evidence yet need to be discovered. With such neurons, the model and its successors were able to successfully model illusory contours where contour completion had an important role (Gove et al. 1993; Grossberg et al. 1997; Grossberg and Williamson 2001; Grossberg 1999; Ross et al. 2000).

It is possible that the self-organized lateral connections in PGLISSOM shown in Chapter 6 can mediate contour completion as well. In this chapter, the PGLISSOM will be tested with contour completion tasks, and various conditions under which completion occurs will be analyzed.

8.2 Contour Completion Performance

In PGLISSOM, the self-organized excitatory lateral connections have the appropriate structure to perform contour completion. To test if the existing circuitry in PGLISSOM is able to fill in gaps in



(a) Contour Completion



(b) Single Edge

Figure 8.5: Multi-Unit Activities for Contour Completion Experiments. (a) The MUA sequences for the four contour elements (bottom four rows) and the gap (top row) are shown. Even though there were no inputs in the middle, the cortical area representing the gap is activated, and the activations are synchronized with the other four MUA sequences. This behavior indicates that contour completion occurred and the gap is perceived as an illusory edge. (b) The MUA sequences for the two contour elements (bottom two rows) and the gap (top row) are shown. The MUA sequence for the gap shows no activity, indicating that contour completion did not occur. Thus, both sides of the gap need to be stimulated for the gap to be perceived as an edge. This behavior is similar to the bipole property (figure 8.3).

sequence representing the gap (top row) indeed shows activity, and this activity is synchronized with the rest of the MUA sequences (bottom four rows). This way, the gap is perceived as part of a contour. However, for the single edge input (figure 8.4b), the MUA sequence representing the gap (top row) shows no activity, while the rest of the MUA sequences (bottom two rows) are active and synchronized. Thus, the same self-organized circuitry in PGLISSOM that is responsible for contour integration can also account for contour completion, and can be the neural mechanism for illusory contours arising from edge inducers.

To further test the ability of the model to support a more realistic illusory contour tasks, a simplified illusory triangle was presented to the network (figure 8.6b). The triangle with gaps in each of the three sides approximates the edge-detected Kanizsa triangle in figure 8.2. The network was also tested with one vertex of the triangle removed (figure 8.6c) to see if both sides of the gaps are necessary for contour completion. Figure 8.13 shows the actual Kanizsa triangles corresponding to these inputs.

The simulation setup and data gathering methods were the same as described above in this section. The inputs were numbered from 1 to 9 (figure 8.6a), where 2, 5, and 8 are the gaps. The corresponding MUAs for the contour elements and the gaps were gathered and the results are shown in figure 8.7.

The results show that for the full triangle (figure 8.6b), the gaps 2, 5, and 8 are all activated and synchronized with the neighboring contour elements, indicating that the gaps are perceived as part of the contours. However, for the broken triangle, the gaps 2 and 8 were no longer filled in,

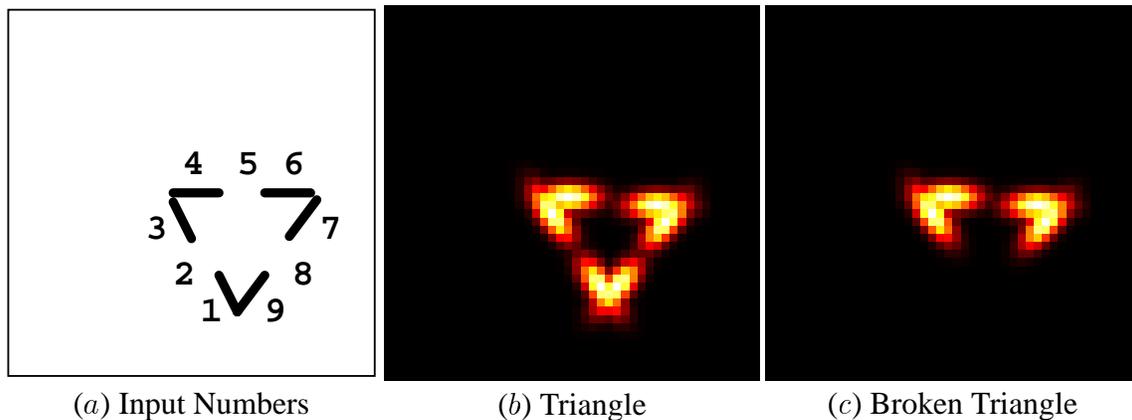
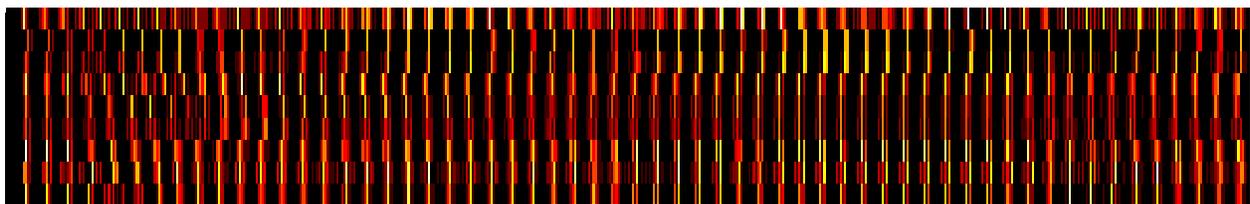
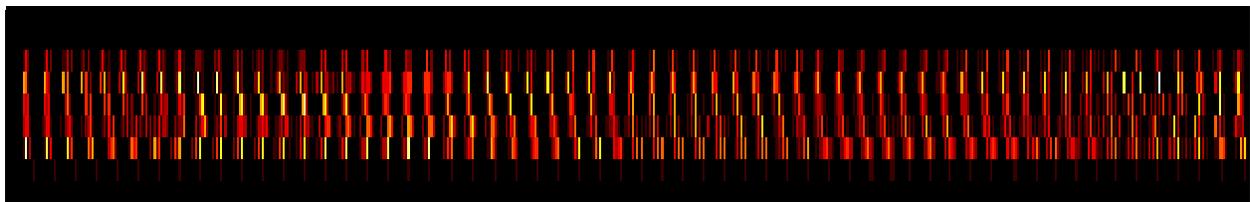


Figure 8.6: **Illusory Triangle Inputs.** The inputs for the illusory triangle experiments and the input numbers are shown. The color code and the contour element sizes are the same as those in figure 8.4. (a) The contour elements and the gaps are numbered from 1 to 9, with the gaps numbered 2, 5, and 8. (b) A full triangle with gaps in the middle of each side is shown. This input is the central triangular part of the edge-detected Kanizsa triangle in figure 8.2. (c) A broken triangle where one vertex (inputs 1 and 9) was removed. This input was used to test if gaps 2 and 8 will be filled in when the two inducing edges are removed. See also figure 8.13 for the corresponding Kanizsa triangles.



(a) Triangle



(b) Broken Triangle

Figure 8.7: **Multi-Unit Activities for the Illusory Triangle Inputs.** The MUA sequences for the nine input regions are shown. The rows are numbered 1 to 9 from bottom to top, and they correspond to the input numbers in figure 8.6a. (a) The MUA sequences for the triangle input in figure 8.6b are shown. The rows corresponding to gaps 2, 5, and 8 are all active and synchronized with the neighboring inputs. Overall, the three MUA sequences representing one side of the triangle are also synchronized with the MUAs representing other sides of the triangle (see also figure 8.8). (b) The MUA sequences for the broken triangle in figure 8.6c are shown. Without the inducing inputs 1 and 9, the gaps 2 and 8 are no longer active.

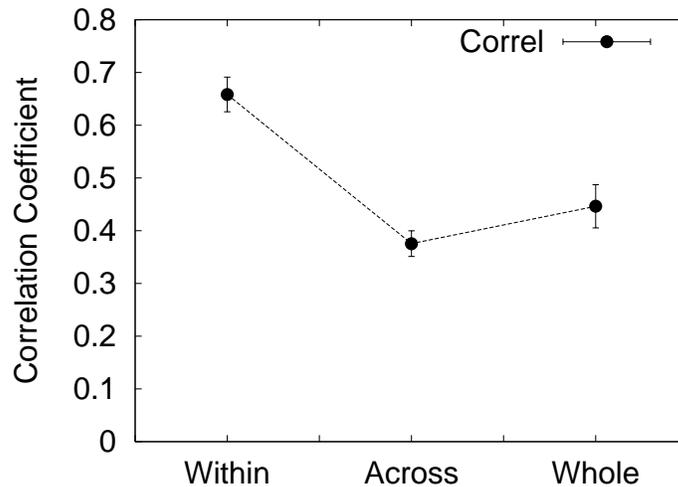


Figure 8.8: Within and Across Contour Correlation. For the complete triangle input (figure 8.6b), each side of the triangle is represented by a group of three MUA sequences, and constitutes a separate contour. There is synchrony within each group, but there is also global synchrony across the groups. The average correlation coefficients were calculated for (1) within the same group, (2) across different groups, and (3) between all pairs of MUAs from the whole object. For the “within” contour case, the correlation is high (0.66), but even in the “across” case, it is quite strong (0.38) This behavior indicates that the three sides are not regarded as separate contours, and the triangle as a whole has a high degree of salience.

whereas gap 5 was filled in. These results are consistent with the experiments done earlier in this section with a single contour.

The result is particularly interesting because the three sides of the triangle, which actually constitute three independent contours, are synchronized. Figure 8.8 shows that the correlation coefficients for the MUA sequence pairs from within each group, across the groups, and among the whole triangle have high values, suggesting that the MUAs are synchronized across the whole triangle. Based on the analysis in section 6.4 however, the three separate contours would be expected to be desynchronized.

There is an interesting and important reason why a global synchrony emerges in this case, and not in section 6.4. The cause can be found at the three vertices of the triangle. At the vertices, two contour elements with different orientation preference are overlapped to form an angle. Since the afferent receptive fields in PGLISSOM are topologically organized, the two cortical areas responding to the two edges at the vertex are close by on the map. As shown in figure 5.11, the excitatory connections not only connect to neurons with similar orientation preferences, but at a close range also to those with fairly different orientation preferences. Thus, proximity of the inputs, as well as the good continuation of contours, determines the degree of synchronization. At the vertices, the two abutting inputs cause the corresponding cortical areas to synchronize (due to proximity), and this in turn causes the three sides of the triangle to synchronize. As a result, the network represents the whole triangle as a coherent object.

These results show that PGLISSOM can do contour completion, and also form representations for whole objects such as a triangle through synchronization of close by inputs and inputs aligned on a smooth path. Such mechanisms may form a basis for complex illusory contour tasks in general.

8.3 Effects of Afferent and Excitatory Lateral Connections

Such filling-in of gaps in the PGLISSOM model is to be expected, given that specific excitatory lateral connections project from the neighboring areas into the area representing the gap. However, in principle, it is also possible that a small amount of afferent input may be causing the completion. In animals and in the PGLISSOM model, receptive fields in neighboring areas in the cortex overlap. Thus, it is possible that the cortical area representing the gap receives enough afferent input from the two tips of the contour elements around the gap, and is thereby activated the rest of the contour element representations.

To check the amount of afferent input received by the gap, the net afferent contribution in MAP2 was measured (figure 8.9). A 2D intensity plot (*a*) and a cross-section plot (*b*) show that the central area receives some amount of afferent input (the central area in *a* and the small peak in the middle in *b*, both numbered 3). Such spurious afferent input may be responsible for the activation of the area representing the gap.

The question is whether the afferent contribution alone, or whether the lateral excitatory contribution alone, can cause the filling-in effect, or whether the phenomenon requires both kinds of contributions. To answer this question, two additional experiments were performed with the same setup as in section 8.3: (1) no afferent connections to the gap area, and (2) no excitatory lateral connections to the gap area. The MUA sequences for the two experiments are shown in figure 8.10. In both cases, the MUA sequence representing the gap in the contour (top row of MUA plot) shows no activity at all, suggesting that contour completion did not occur in either case.

To contrast with the contour completion experiment in the previous section, the correlation coefficients between each MUA sequence from the four contour elements and the MUA sequence from the gap area were calculated (figure 8.11). The figure summarizes the result that contour completion in the PGLISSOM model occurs only when there are both a small amount of sub-optimal afferent contribution and lateral excitatory contribution. Such condition can only occur when the input contour elements are aligned along a smooth path, since only in such a case, the central receptive field can look at the tip of input in neighboring areas of the visual field, and the co-circular projection of lateral connections can help bring the sub-optimal afferent activation above threshold.

The results show that contour completion in PGLISSOM is not simply driven by overlapping receptive fields in the gap region, thus the specific excitatory lateral connections are necessary to fill in the gap. The results also indicate that the model needs small amounts of afferent input to

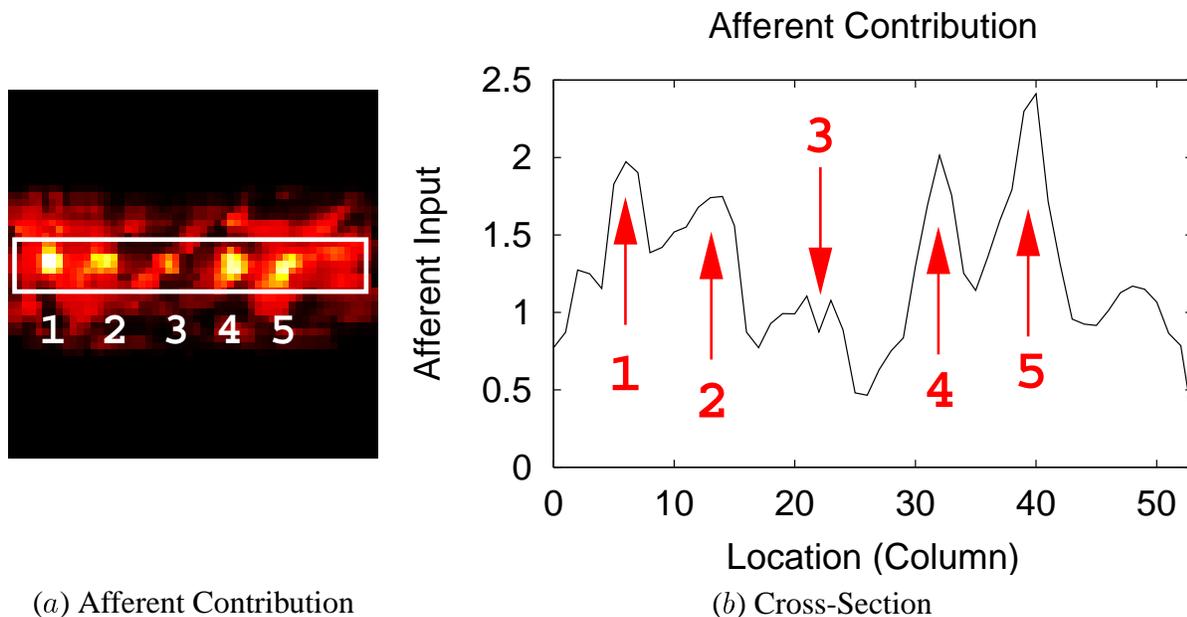


Figure 8.9: **Afferent Contributions in MAP2.** The afferent contributions to MAP2 neurons are shown. (a) The contributions to the 54×54 map are plotted in color. The four areas (numbered 1, 2, 4, and 5) corresponding to the four contour elements in the input (figure 8.1a) show high level of afferent input. The center area corresponding to the gap (numbered 3) also receive some afferent input due to slight overlap with neighboring regions in the retina. The white box surrounding the middle area is where the calculations for the plot in (b) are done. (b) The 8×54 map area in the middle of (a) marked by the white box was extracted into a matrix, and column by column, the average of each column vector of size 8×1 was calculated and plotted. As observed in (a), there are four peaks corresponding to the four input segments (numbered 1, 2, 4, and 5), and a small peak in the middle where there is the gap in the input (numbered 3).

fill in the gap so that arbitrarily large gaps will not be filled in. Such conditions allow us to predict the limits of contour completion performance in humans.

8.4 Saliency of Closed vs. Open Contours

Further interpretation of the illusory triangle experiment in section 8.2 provides a possible explanation for another contour detection result reported by Kovacs and Julesz (1993) and Pettet et al. (1998): closed contours were found to be easier to detect than open contours.

The complete and broken triangles in figure 8.6 form closed and open contours. To measure the saliency of the two objects, the average correlation coefficients between the nine elements of the full triangle (inputs 1 through 9) and five elements of the broken triangle (inputs 3 through 7) were calculated. The results are shown in figure 8.12. The results indicate that the activities in the network for the closed contour are significantly more synchronized than those of the open contour, indicating that closed contours are more salient.

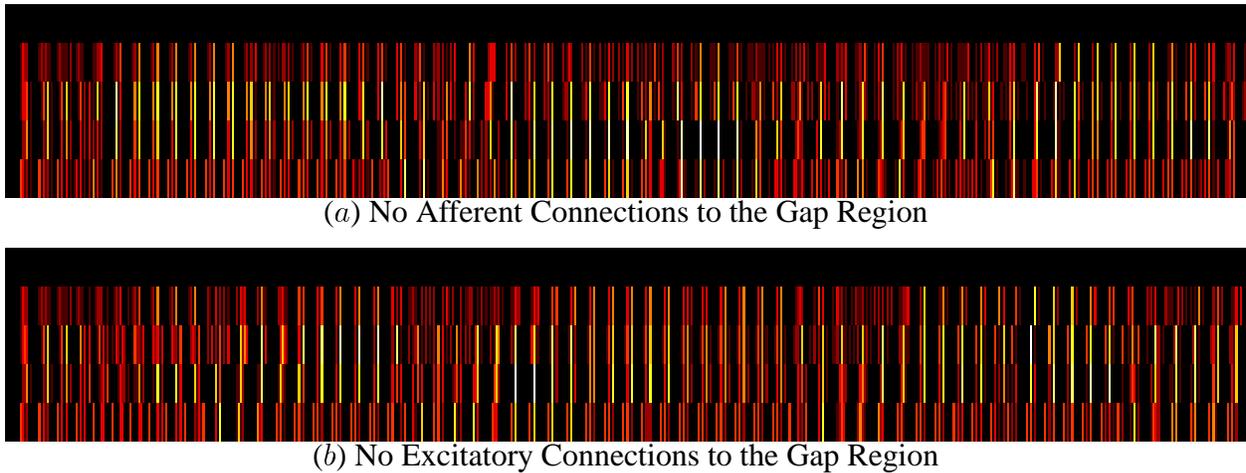


Figure 8.10: Multi-Unit Activities for Contour Completion with Removed Connections. The MUA sequences for the (a) afferent-deprived vs. (b) lateral excitation-deprived networks are shown. The MUA sequences for the four input contour elements (bottom four rows) and for the gap (top row) are shown in each plot. In both cases, the MUA sequence for the gap in the contour shows no activity, suggesting that filling-in did not occur.

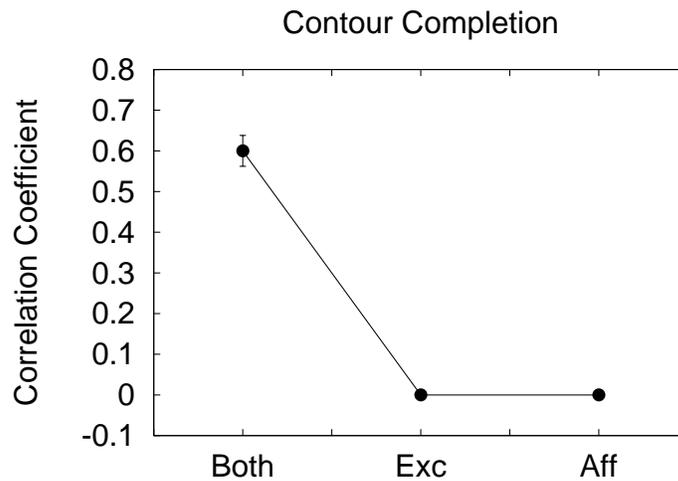


Figure 8.11: Contour Completion Performance with Removed Connections. The average correlation coefficients for the four MUA sequences representing the four input contour elements vs. the MUA sequence representing the gap are shown. From left to right, (1) both afferent and excitatory connections exist (Both), (2) afferent connections are removed from the center and excitatory connections remain (Exc), and (3) the excitatory connections are removed from the center and only afferent connections remain (Aff). The plot shows that both afferent and excitatory contributions are necessary for contour completion.

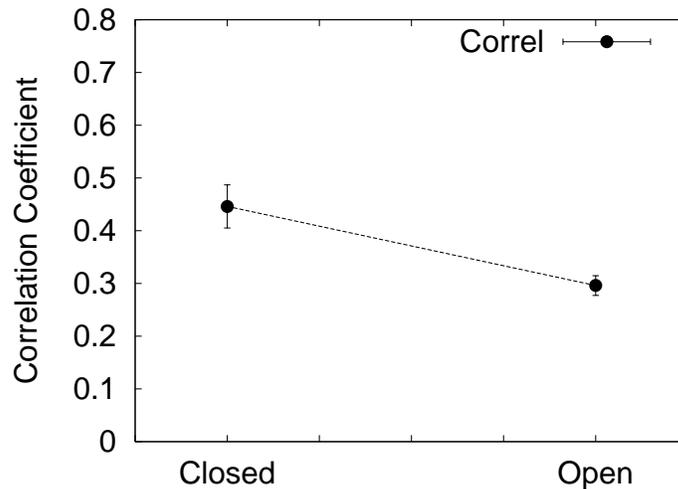


Figure 8.12: **Saliency of Closed vs. Open Contours.** The average correlation coefficients between the MUA sequences in the complete (Closed) and broken (Open) triangles (figure 8.6*b* and *c*) are shown. The correlation is higher in the closed configuration, consistent with the results reported by Kovacs and Julesz (1993) and Pettet et al. (1998), where closed contours were easier to detect than open contours.

This behavior is consistent with the psychophysical results. Such a behavior is believed to occur through a recursive lateral facilitation around the contour when the contour is closed to form a loop (Pettet et al. 1998). A similar argument applies in PGLISSOM. In a closed loop, every part of the contour receives excitatory lateral contribution from *both* neighboring areas, thus the overall synchronization is reinforced along the whole loop. However, for open contours, at the two ends of the contour, neurons only receive lateral excitation from one neighboring area instead of two, thus the synchrony does not reach the same level of saliency. Such mechanism may explain why the perception of an illusory triangle breaks when one pacman is removed from the Kanizsa triangle (figure 8.13).

These results suggest that the contour integration mechanism in PGLISSOM can also account for the enhanced saliency of closed contours.

8.5 Conclusion

In this chapter, I have shown that the circuitry in PGLISSOM responsible for contour integration and segmentation can perform contour completion as well. Such a behavior was due to two factors, (1) small amount of afferent contribution and (2) lateral excitatory reinforcement. The results show that lateral excitatory connections are necessary for contour completion, but not sufficient. I have also shown that lateral excitation together with the topological organization of PGLISSOM can explain the higher saliency of closed contours over open contours. Understanding these conditions allow us to know what to expect from low-level neural circuitry, and provides us with a foundation

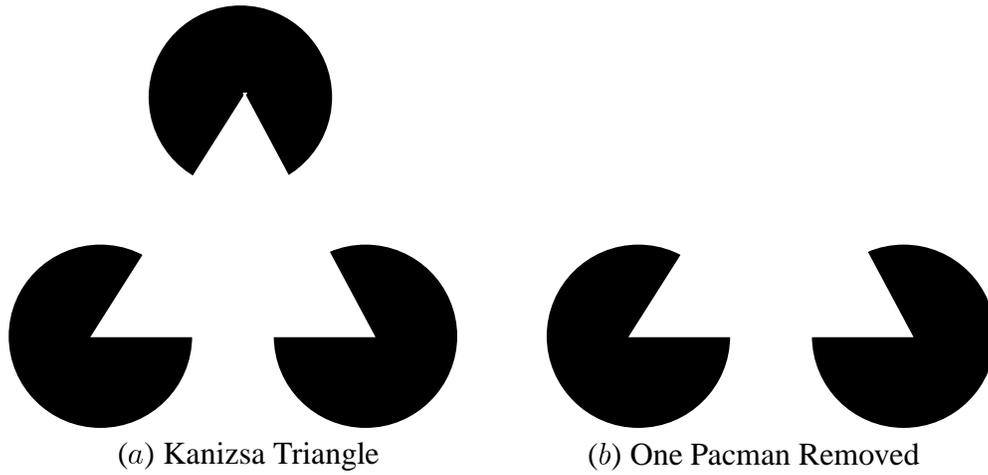


Figure 8.13: **Salience of Kanizsa Triangles.** The perception of an illusory object is quite strong in the original Kanizsa triangle (a). However, when a pacman is removed, the vivid perception suddenly disappears (b).

for performing more complex visual tasks.

Chapter 9

Discussion

In this thesis, I have shown that the degree of synchrony in neural populations can predict various perceptual grouping phenomena. I have also shown that self-organized lateral connections that capture the statistical regularities in the environment are essential for mediating synchrony and predicting task performance. However, several open issues remain: (1) how can the temporal coding hypothesis be further verified, and how can we interpret information carried by temporal events, (2) how and why does functional divisions occur in the layered architecture and across different areas in the cortex, and (3) do cognitive (or higher-level) processes influence perceptual performance? In this chapter, these open issues will be described and possible solutions outlined.

9.1 Temporal Coding in Neural Systems

In this section, I will discuss how the link between perceptual performance and neural synchrony can be verified, and how the temporal code can be interpreted by the later stages of visual processing.

9.1.1 Synchrony as a Perceptual Representation

The hypothesis that synchronous neural activity predicts perceptual performance is based on experiments where either (1) input properties and synchrony in the animal's neural activity are compared (Eckhorn et al. 1988; Gray and Singer 1987; Gray et al. 1989; Singer 1993), or (2) psychophysical performance in humans and timing between input features are compared (Fahle 1993; Leonards et al. 1996; Leonards and Singer 1998; Usher and Donnelly 1998). In case 1, we do not know whether the animal really perceived the input as coherent or not, and in case 2, we do not know whether the neurons actually fired in synchrony. Thus, to explicitly verify that synchronized firing gives rise to perceptual experience, the degree of synchrony in neural activity and perceptual performance need to be measured simultaneously and compared. In the following, I will discuss how such a comparison can be established.

9.1.1.1 Neural Activity and Behavior

The assumption in such experiments is that the perceptual experience of the animal is faithfully represented in its behavior. For over a decade, Celebrini and Newsome (1994) and their colleagues have been measuring neural activity and animal behavior simultaneously in the Middle Temporal (MT) area of monkeys, i.e. the motion detection center in the brain (Bair et al. 2001; Britten et al. 1992; Salzman et al. 1990). Early studies showed that microstimulation of neurons in MT causes a significant change in motion detection tasks (Salzman et al. 1990). Subsequent results showed that the spike count of a single neuron in MT for a short duration accurately predicts the behavior of the monkey in motion detection tasks (Britten et al. 1992; Celebrini and Newsome 1994).

Although Bair et al. (2001) *did not* find periodically modulated synchronization (or coherent oscillations) in MT neurons contrary to what was observed earlier in V1 and V2 of cats (Eckhorn et al. 1988; Gray and Singer 1987; Gray et al. 1989; Singer 1993), the strong relation they established between neural activity and psychophysical performance in animals greatly advanced our understanding in neural representations of perception. The methodologies used in MT can be applied to areas where coherent oscillations have been found (e.g. areas V1 and V2 in cats). Instead of just comparing the coherence in input and the neural activity as in the previous experiments in cats, the animal can be trained to respond to the stimuli and act out the decision it made about the stimuli. Then, all three measures, (1) coherence in input features, (2) synchrony in neural firing, and (3) perceptual experience (manifested as behavioral performance), can be measured and compared to explicitly verify that correlated firing of neurons represents perceptually salient events. Thus, interpretation of temporal sequences and the role of different types of temporal codes in perceptual grouping are important open questions that should be further investigated.

9.1.1.2 Degree of Neural Synchrony and Perception

According to the temporal correlation hypothesis, to explain the graded response in behavior, a graded degree of synchrony in the neural activity must be possible. How to measure the degree of synchrony is a difficult problem (see also section 9.1.2). In PGLISSOM, the graded response in perception can be measured by the correlation coefficients between two MUA sequences. Thus, the two sequences can come into and go out of sync, but in the end, the overall correlation of the two sequences is what determines the degree of synchrony. Such an assumption can be tested in the experiments outlined above for comparing neural activity and behavior. If a different kind of graded measure of synchrony is developed later, that measure can be performed on the model's behavior as well.

When degrees of synchrony are used as graded perceptual grouping representations, a computational problem arises in relation to the transitive grouping rule proposed by Geisler et al. (1999; 2001; Geisler and Super 2000). If a grouping rule determines whether the two representations group together or not (i.e. a binary decision), transitive grouping rule works well. The question is, how would such a transitive relation apply to graded responses (or graded grouping rules). For

example, if representations (A, B) are related by amount x , and (B, C) are related by amount y , how related are (A, C) ? Should it be $\max(x, y)$, $\min(x, y)$, or some other quantity? In neural terms, given the degree of synchrony in (A, B) and (B, C) , how synchronized should (A, C) be? Since connections within a local area in the cortex (excluding the ones that go through white matter) are limited in spatial extent, such a problem can arise when two populations within a local area without a direct connection try to participate in forming a coherent perceptual representation through synchrony. As was seen in this thesis (and also Terman and Wang 1995), synchrony can be achieved in a locally connected network of temporal neurons, but how to represent the *degree* of synchrony in such a locally connected network is still an open issue. Computational studies focusing on possible encoding schemes for such information will provide us with deeper insights into the role of synchronization in perceptual grouping at a larger cortical scale.

9.1.2 Interpretation of Temporal Codes

How can the coherence of synchronized activity be measured in the cortex, and how can we apply those understandings in designing models and also analyzing the behavior of a model? In PGLISSOM, coherence is measured as the correlation coefficient between MUA sequences, but are there other measures? In a broader sense, such questions fall under the general question of how the information in the spike sequence (or spike train) can be interpreted?

Rieke et al. (1997) presented an extensive review of this topic. They discussed a wide range of techniques including Bayesian inference, information theory, and statistical methods to read out the information embedded in spike trains. These techniques can be equally valuable in interpreting the activities in PGLISSOM. Applying such techniques on computational models can help verify the robustness and accuracy of the model and the efficiency of the techniques. On the other hand, techniques developed on the model can be applied to experimental settings.

So far, I have discussed synchrony as one of the most likely temporal codes employed by the brain. However, there are other forms of temporal information that neural systems may utilize as well. In the comprehensive review of temporal coding by Gerstner (1998b), the candidate temporal codes were identified as: (1) time to first spike (Maass 1998), (2) phase difference relative to background oscillation (O'Keefe and Reece 1993), and (3) reverse correlation from spike trains (i.e. stimulus reconstruction; Rieke et al. 1997). Whether such different types of temporal codes already exists in PGLISSOM, and whether certain parameter values can give rise to behavior relevant to such temporal codes are important questions. In computational models such as PGLISSOM, the possible causes of such temporal codes, and their possible roles can be relatively easily assessed because we can closely monitor and alter the state of the model.

Thus, existing methods of interpreting temporal sequences and the role of different types of temporal codes in perceptual grouping should be further investigated.

9.2 Functional Divisions in the Visual Cortex

There are functional differences across the vertical layers in the cortex, and also across different areas in the cortical sheet. In this section, the question of how and why such functional divisions in the visual cortex occur will be discussed. I will outline how the two layers of PGLISSOM correspond to the layers in the visual cortex and suggest a possible reason for such a division. Further, the asymmetry in contour integration performance reported in the fovea and periphery will be discussed, and possible ways to explain it will be outlined.

9.2.1 Layers of the Visual Cortex

The two-layer organization of PGLISSOM was developed because long-range inhibitory connections were found necessary for ordered self-organization (MAP1) and long-range excitatory connections for grouping (MAP2). However, such an organization was not an arbitrary design choice: a similar arrangement is found in the layered architecture of the visual cortex (Chapter 2, figure 2.3). The short-range excitatory lateral connections and long-range inhibitory connections in MAP1 are abstractions of the on-center off-surround projection from layer 6 to layer 4 in the visual cortex (Ahmed et al. 1997; Grieve and Sillito 1995a). On the other hand, the long-range excitatory and inhibitory lateral connections in MAP2 correspond to the long-range axonal projections in layer 2/3 (Gilbert and Wiesel 1989). Their possible role in grouping was made more plausible by the discovery of a group of fast-spiking cells known as *chattering cells* in precisely the layers 2 and 3 (Gray and McCormick 1996). These cells were postulated to contribute to coherent oscillations and perceptual grouping, and hits at the role of layer 2/3 in perceptual grouping. The third important architectural component in PGLISSOM consists of the intra-columnar connections, modeling such known connections between layers 2/3, 4, 5 and 6 in the cortex (Gilbert and Wiesel 1979).

Thus, the layered structure of the visual cortex motivates the architectural design of PGLISSOM. Also, in reverse, the functionality of PGLISSOM suggests that the specific anatomical arrangements of the visual cortex may be due to different functional requirements. A related computational model of the six layer architecture proposed by Grossberg (1999) and Grossberg and Williamson (2001) addressed the same issue. In their model, layers 4 and 6 help stabilize development, and lateral connections in layer 2/3 perform perceptual grouping. Thus, both PGLISSOM and the model of Grossberg and Williamson predict that different layers in the visual cortex perform organization and grouping. The question is, does the layered architecture originate from different functional requirements? Is there a way to experimentally verify it? These questions will be addressed in Chapter 10.

Computational models such as PGLISSOM can be used to model such different functions found in the layered architecture of the visual cortex, and help gain insight into how the functional divisions may occur and how they interact.

9.2.2 Fovea vs. Periphery

Functional differences exist not only in the layered architecture of the cortex, but also across a fairly uniform area of the cortex such as the visual cortex. I briefly mentioned in Chapter 7 the results by Hess and Dakin (1997), showing that contour integration performance in the periphery is significantly worse than in the fovea. The same developmental argument as advanced in Chapter 7 may apply in the case of fovea vs. periphery. Due to attentional bias, the input statistics may differ in fovea vs. periphery, resulting in structural and functional division in the two areas. The difference in the distribution may drive the two areas of the cortex to develop differently, resulting in a difference in performance similar to that observed by Hess and Dakin (1997). To firmly verify if there exists such a difference in the input distribution, we can collect statistics from different parts of the visual field using eye-tracking devices, while human subjects are freely browsing the environment. Such statistics will account for both environmental and attentional biases, thus giving us an accurate portrayal of the input distributions in the different parts of the visual field. This information will allow us to predict perceptual performance of the different cortical areas.

However, the comparison is more complicated in the fovea vs. periphery case than in the lower vs. upper hemifield case because we have to account for the difference in optics (more blurring in the periphery) and photo-receptor density (far smaller number of photo-receptors in the periphery). Details of a small input that can easily be seen in the fovea may not be as visible when presented to the periphery. It also means that larger details of larger objects *may* be visible in the periphery. Therefore, it is possible that when inputs are larger, contour integration in the periphery can approximate that of the fovea¹. Thus, these factors should also be fully considered when gathering input statistics and reasoning about the possible causes of functional divisions in areas such as the fovea vs. periphery.

9.3 High-Level Influence on Perceptual Grouping

This thesis has focused on the representation of perceptual events in the primary visual cortex. An interesting question is whether synchronized activities exist in higher-level visual or cognitive areas of the brain, and whether the higher levels have influence on lower-level perception and behavior.

In fact, precisely correlated spike events were found in the frontal cortex of awake monkeys (Abeles 1991; Abeles et al. 1993). These correlated spike events, which he named *synfire chains*, are highly specific spatio-temporal patterns where a synchronized population of firing neurons activates another population in a successive, feed-forward manner. Later, Abeles et al. (1995) showed that the occurrence of synfire chains in monkey frontal cortex is highly correlated to the behavior of the animal. Such synfire chains can have an effect on the degree of synchrony in the lower-level areas. If synfire chains in the higher areas project back to the lower-level areas, it is possible that the higher-level can influence perception through modulating the synchrony at lower

¹This idea is due to personal communication with W. Geisler.

levels. Such feedback modulation seems to be happening at least in the lower-level hierarchies of the visual pathway. For example, Sillito et al. (1994) showed that to achieve synchrony, the LGN needs feedback from V1. These results show that synchronized events indeed exist in the higher-level cortical areas, and they can have influence on lower-level processing.

Illusory contour phenomena constitute further evidence of high-level influence on low-level perceptual grouping. The approach to illusory contour detection in this thesis was fundamentally bottom-up, based on experimental observations of illusory contour sensitive cells in V1 and V2 (Peterhans et al. 1986; Sheth et al. 1996; von der Heydt and Peterhans 1989). However, there exist illusory contour stimulus configurations that cannot be explained by lower-level mechanisms (Hoffman 1998). In fact, connections between lower and higher visual areas are not strictly feedforward, they are reciprocal (Nelson 1995). Therefore, activity response to illusory contours in V1 and V2 may not be purely due to low-level mechanisms, although mostly driven by afferent inputs. For this reason, even though the most basic illusory contours can be processed by the low-level mechanisms in V1 and V2 as shown in this thesis, higher-level influences should also be included for a complete explanation of the illusory contour phenomena.

9.4 Predictions of the PGLISSOM Model

The PGLISSOM model makes several concrete predictions about the perceptual grouping and structural organization in the visual cortex, and about how statistical regularities in the input shape this structure. In this section, the predictions based on the results in PGLISSOM will be summarized, and pointers to sections where these predictions are discussed in detail are provided.

1. Synchronized Activity and Representation of Perceptual Grouping.

- Degree of synchronization in populations of neurons (MUA) correlates to perception. If MUA correlation is disturbed, perception will change.

Section 9.1

- Decay rate at synapses acts like delays. Adapting the decay rate may be easier than adapting delay, and it is possible that such mechanisms exist in real neurons.

Section 4.1.1

- V1 mechanisms can account for edge-induced illusory contours through contour completion. Specific excitatory lateral connections that are responsible for contour integration are also crucial for such processes. Binding through proximity, together with contour completion, can help form synchronized representations of coherent objects.

Chapter 8

2. Self-Organization and Structural Development in the Visual Cortex.

- Layered architecture in the visual cortex may exist because of the different functional requirements in the cortex: (1) self-organization (MAP1; layer 4 and 6 in V1) and (2) perceptual grouping through synchronization (MAP2; layer 2/3 in V1).

Section 9.2

- Differences in structure and functional performance in different hemifields are due to differences in input distributions and the input-driven nature of self-organization. Environmental and attentional bias can alter input distributions. Detailed image statistics of the environment that take into account attentional biases can predict inter-area differences in the visual cortex.

Chapter 7 and Section 9.2.2

- Oriented Gaussian inputs give rise to a rich correlational structure in lateral connections, approximating co-circularity of edge co-occurrence in natural images. Hebbian learning, based on suboptimal responses, is responsible for this effect.

Section 5.4

9.5 Summary

In this chapter, I discussed the issues of temporal coding, self-organization, and high-level influence on perception that were raised by observations of the structure and behavior of the PGLIS-SOM model. Concrete predictions on perception, anatomy, and development made by the PGLIS-SOM model were presented. How to experimentally verify the predictions and how to answer the open issues presented in this chapter by extending PGLISSOM are the main theme of the next chapter.

Chapter 10

Future Work

The open issues discussed in the previous chapter can lead to further research in several areas. In this chapter, such future directions are outlined by discipline: (1) psychophysics, (2) neuroscience, (3) computational science, and (4) artificial vision.

10.1 Psychophysics

As the psychophysical experiments on contour integration tasks suggest, there are several stimulus dimensions other than orientation that affect perceptual performance (Field et al. 1993a; Geisler et al. 1999, 2001; McIlhagga and Mullen 1996; Pettet et al. 1998). These are, to name a few, inter-element distance, background randomness, fractal exponent of the global contour (how jagged the path is), spatial frequency (size), relative phase of successive Gabor contour elements on the contour, color, and contrast. PGLISSOM can be trained with more complicated input patterns to include such stimulus dimensions, and be used to predict human perceptual performance on such stimuli. Trained with such inputs, PGLISSOM can be used to identify the input statistics that are important for each stimulus dimension, by analysing the lateral connection patterns and receptive field properties.

Instead of such artificial inputs, the network could be trained with natural images. In fact, RF-LISSOM was trained with natural images and orientation maps similar to those reported in this thesis were obtained (personal communication with J. Bednar; see also section 10.4). In PGLISSOM, natural images cannot be used because of higher computational requirements of PGLISSOM than RF-LISSOM. However, with scaling-up techniques suggested in section 10.3, it will be able to train a larger network to handle natural images. Another possibility is to train with multiple edges that are generated from known natural image statistics (suggested by W. Geisler). The images need to be appropriately filtered based on known biological processes in the retina and LGN, because PGLISSOM is a model of V1 and assumes that such processes have already taken place. The network can then be trained as described in this thesis. Since stimulus dimensions are mixed in natural images, PGLISSOM will attempt to encode all dimensions at once and

unexpected structures could arise. For example, cells could become responsive to combinations of stimulus dimensions or only to a single dimension and form segregated clusters in the map. Such results can be verified against known biological data, such as the combined orientation and spatial frequency map reported by Issa et al. (2001), and propose new structures and properties to look for in the cortex. By analyzing the functional role of the novel structures it might even be possible to predict yet unknown psychophysical phenomena involving multiple stimulus dimensions in tasks such as contour integration.

10.2 Neuroscience

PGLISSOM predicts that different areas of the visual cortex can have different patterns of connections, and such differences cause perceptual performance to differ in corresponding areas of the visual field. Such predictions can be tested experimentally by measuring lateral connection patterns. Finding such evidence is an important step in understanding the cause of functional differences within the visual cortex.

Several specific predictions of the model can be verified this way. First, lateral connectivity can be measured in upper vs. lower hemifield and fovea vs. periphery to see if there are any differences between these areas as the model suggests. Next, if such differences are found, the cause of the differences can be experimentally tested. Animals can be reared in controlled visual environments where the input distributions as simple as frequency of input presentations in the lower and upper visual hemifields are different. The model predicts that, in this case, the hemisphere with more frequent input will have more co-circular lateral connections, while the other hemisphere will show more collinear patterns.

The layered architecture in the visual cortex may arise due to multiple functional requirements as discussed in section 9.2. Such functional divisions can be tested in the cortical layers experimentally. Each layer can be selectively disabled, or the intra-columnar connections between the layers can be disrupted. Then, progress of development can be measured in the deep layers (layers 6 to 4) and compared to the shallow layers (layers 2/3). The predictions in section 9.2 will be verified if shallow layers do not properly self-organize. However, this experiment would require new imaging techniques because the existing ones may not be able to selectively record activity in the individual layers. New lesion techniques also need to be developed to selectively disable a particular layer.

Another promising direction in neuroscience is investigating the possible role of synchrony in multi-modal sensory integration. There are coherent oscillations in other sensory areas such as the olfactory bulb and the auditory system. The question is whether synchronization can bind these different sensory modalities together. There is evidence of multi-modal interaction, such as auditory influence on visual perception (Churchland et al. 1994). The role of synchronization in such interactions should be investigated in the areas of the cortex where these sensory modalities converge.

PGLISSOM also predicts that post-synaptic decay rate can change the synchronization behavior and possibly counteract the effects of various delays in the nervous system. Such predictions may be experimentally verified. The work of Nowak and Bullier (1997) on measuring various sources of delay in the nervous system can aid in identifying the locus in a neuron where decay rate may be playing a crucial role in controlling synchrony. Experimental techniques should be developed to measure and alter the decay property of the cell membrane. Such an investigation can lead to a deeper understanding of how neural synchrony occurs and how a fine tuning of temporal behavior is possible when there are various degrees of delay in the system.

10.3 Computational Sciences

Developing techniques to simulate large networks accurately is always a big issue in computational neuroscience research. In PGLISSOM, lateral connections cause most of the overhead. With an $n \times n$ map, there are $O(n^4)$ lateral connections; as the map grows, the number of lateral connections grows rapidly. For example, a 127×127 MAP2 can fit in 2 GB of memory when standard floating point numbers (4 bytes) are used. A network of that size is slightly too small to process realistic input, since smaller map means smaller retina. Thus, techniques for efficiently representing and storing the connections, especially the lateral connections, need to be developed. Also, ways of computationally approximating the lateral interactions in the model should be devised.

In the current model, initial lateral connections are long-reaching and dense, causing a huge memory and computational overhead. However, after self-organization, most of these connections become very weak and are pruned away. Only a small number of patchy connections remain in the final map. Thus, the memory requirement for a fairly large network can be significantly reduced. There are several ways to make use of this observation. The most promising direction is to grow the map gradually (Bednar et al. 2002). In their extension of the RF-LISSOM model called GLISSOM, precise scaling equations determine how the size of the network can be smoothly increased while preserving global order and lateral connectivity. With the addition of spiking neurons and intra-columnar connections, GLISSOM can be extended to run PGLISSOM faster and with much lower memory requirements.

Another way of reducing memory overhead caused by lateral connections is to initialize the connection weights based on known biological connectivity or connectivity derived from image statistics. Since such connectivity is usually sparse, a larger map can be constructed. Although learning cannot be modeled in such networks, large-scale simulations can be run to (1) test specific functional effects of biological connectivity patterns, and (2) test the components of the network in various realistic psychophysical tasks.

A third way to deal with large input space is by scanning the visual space with a small PGLISSOM, one area at a time. How to collect and interpret the responses of the network at a large number of locations in the visual is currently an open question. If a model has only afferent connections, we can simply partition the input space into discrete grids and feed input from each

grid into the network. The responses at each location can then be combined together to form a global output of the network. However, in a model with lateral connections, the problem is how to define lateral interactions between the outputs from successive scans. For example, if the network generates output p at input location x and output q at location y , how should the two outputs p and q interact? Such a problem arises because lateral connections at the borders do not extend outside of the map. To solve this problem, the network could be connected as a torus, i.e. by joining the top and bottom sides of the 2-D map, as well as the left and right sides. However, the question of how to partition and scan the input space still remains. Since lateral interactions occur through recurrent feedback, the time course of lateral interaction in relation to each scanned output needs to be precisely determined. Although lateral interaction is possible in such an architecture, the order of visiting the grids can affect the final output. Such specific algorithms should be developed to approximate a 2-D map.

A question that arises with any computational study like the one presented in this thesis is “what is the computational goal of the architecture?” There are three ways of approaching this question: (1) devise a computational principle first, and then show that architecture satisfies it (Lee et al. 2000b; Linsker 1986; Oja 1992; Rao and Ballard 1999; Wachtler et al. 2001) (2) test the input and output relations of the model to infer the computational principles that describe it (Sirosh 1995), and (3) analyze the properties of the input directly to infer what kind of representations are needed to efficiently encode them (Barlow 1985, 1994; Field et al. 1993b).

For example, Oja (1992) and Linsker (1986) showed that learning rules similar to Hebbian can be derived using principal components analysis and information maximization. Lee et al. (2000b) and Wachtler et al. (2001), and Rao and Ballard (1999) showed that realistic receptive fields arise from independent components analysis and predictive coding framework. Barlow (1985, 1994), Field et al. (1993b), and Olshausen and Field (1996) showed that redundancy exists in the environmental input, and these inputs can be efficiently represented by sparse coding of neural activity. Further investigation should be done to find out if these computational theories apply to PGLISSOM. We can study if the computational principle describes the computations in PGLISSOM, and if there are other computational goals implicit in it.

10.4 Artificial Vision

One reason for trying to understand the mechanisms of human perceptual grouping is to design a vision system that can automatically process visual information the same way humans do. Such an application would have many practical uses. The questions are then, (1) can PGLISSOM scale up to process real-world images, and (2) what are the extra processing steps necessary for the model to perform complex vision tasks? The first question was partially answered in section 10.3 (scaling up). To answer the second question, the system needs to be extended and pre and postprocessing components need to be added.

As suggested in section 10.1, natural images need to be preprocessed before presenting

them to the network. One way is to use On-Off cells or filters such as oriented Gabor wavelets or difference-of-Gaussians to perform preliminary edge-detection and spatial frequency normalization (where the low-powered higher frequency bands are adjusted to have higher power).

Several postprocessing steps are also necessary. A most important one is to automatically identify active areas in the maps and assign labels to neural populations that respond to a particular stimulus features. This problem is highly related to the problem of interpreting the neural code, and they should probably be investigated together (as discussed in section 9.1.2). Post-processing modules employing such interpretation techniques can be implemented for robust contour detection performance.

Once we detect coherent contours, the next problem is to identify objects. Since PGLISSOM covers a small visual area, and the lateral interaction range is limited, the current model cannot interpret entire scenes. However, a hierarchy of PGLISSOM networks with increasing receptive field size at each level might be able to do it. At the lowest level, preliminary features such as contours can be detected, and at each successively higher layer, the receptive fields will cover more area in the visual space. Thus, continuous contours can be grouped into object representations, and so on. At each level, synchrony will effectively represent coherent objects in the scene and desynchrony will segment different objects. It may then be possible to select between these layers to get grouping information at different levels (Geisler and Super 2000).

Including pre and postprocessing for PGLISSOM and making the maps hierarchical will enable us to build a robust visual perceptual grouping system that operates at multiple levels of complexity. The resulting system would be likely to lead to insights into how high-level grouping mechanisms are implemented in the brain.

10.5 Summary

In this chapter, I have discussed some of the most promising future directions of the research done in this dissertation. Although they span several disciplines, it is possible to see that many of these directions interact. PGLISSOM will make it possible to bring together various disciplines to gain a deeper understanding about the mechanisms of perceptual phenomena.

Chapter 11

Conclusion

The goal of this thesis was to understand the neural mechanisms of perceptual grouping through computational modeling. I developed a model called PGLISSOM with carefully selected computational primitives inspired from biological neural networks: (1) spiking neurons for the representation of perceptual grouping events through temporal coding, (2) afferent connections for input mapping, (3) lateral connections to encode activity correlations and mediate synchrony, (4) layered architecture and intra-columnar connections for functional specialization into self-organization and perceptual grouping, and (5) Hebbian learning mechanism for adapting the connection weights based on visual input.

I showed that in the model, (1) synchronization can be robustly controlled in a network of spiking neurons, (2) realistic orientation maps develop through input-driven self-organization, (3) intricate lateral connection patterns that are similar to patterns found in experimental observations emerge, (4) contour integration performance in the model predicts human performance, (5) contour segmentation is also achieved in the same network, (6) the network also performs contour completion, and (7) differences in input distribution cause the network to develop different structure and functionality.

With these experiments, I demonstrated that the degree of synchrony measured in the neural populations accurately predicts psychophysical performance in contour integration tasks. Thus, synchronized activity may form the neural representation of perceptual grouping. I also showed that lateral connections are crucial in mediating synchrony, and that they learn to implement the grouping rules in the cortex. Such lateral connections can self-organize to encode the correlational structure in the input distribution, and altering the input distribution can cause the structure and performance to change. Thus, the model provides computational evidence for the idea that self-organized lateral connections form the structural foundation for perceptual grouping.

In the future, computational models like PGLISSOM will play a crucial role in understanding perceptual phenomena, by providing a computational framework where ideas from multiple disciplines can be integrated.

Bibliography

- Abbott, L. F., and Marder, E. (1995). Activity-dependent regulation of neuronal conductances. In (Arbib 1995), 63–65.
- Abeles, M. (1991). *Corticonics: Neuronal Circuits of the Cerebral Cortex*. Cambridge, England: Cambridge University Press. First edition.
- Abeles, M., Bergman, H., Gat, I., Meilijson, I., Seidemann, E., Tishby, N., and Vaadia, E. (1995). Cortical activity flips among quasi stationary states. *Proceedings of the National Academy of Sciences, USA*, 92:8616–8620.
- Abeles, M., Bergman, H., Margalit, E., and Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of Neurophysiology*, 70:1629–38.
- Ahmed, R., Anderson, J. C., Martin, K. A. C., and Charmaine, N. J. (1997). Map of the synapses onto layer 4 basket cells of the primary visual cortex of the cat. *Journal of Computational Neuroscience*, 380:230–242.
- Anderson, J. A., and Rosenfeld, E., editors (1988). *Neurocomputing: Foundations of Research*. Cambridge, MA: MIT Press.
- Arbib, M. A., editor (1995). *The Handbook of Brain Theory and Neural Networks*. Cambridge, MA: MIT Press.
- Bair, W., Zohary, E., and Newsome, W. T. (2001). Correlated firing in macaque visual area MT: Time scales and relationship to behavior. *Journal of Neuroscience*, 21:1676–1697.
- Barlow, H. (1994). What is the computational goal of the neocortex? In (Koch and Davis 1994), 1–22.
- Barlow, H. B. (1985). The twelfth Bartlett memorial lecture: The role of single neurons in the psychology of perception. *Quarterly Journal of Experimental Psychology*, 37A:121–145.
- Bartsch, A. P., and van Hemmen, J. L. (2001). Combined hebbian development of geniculocortical and lateral connectivity in a model of primary visual cortex. *Biological Cybernetics*, 84:41–55.

- Bednar, J. A. (1997). *Tilt Aftereffects in a Self-Organizing Model of the Primary Visual Cortex*. Master's thesis, Department of Computer Sciences, The University of Texas at Austin. Technical Report AI97-259.
- Bednar, J. A., Kelkar, A., and Miikkulainen, R. (2002). Modeling large cortical networks with growing self-organizing maps. In (Bower 2002). To appear.
- Bednar, J. A., and Miikkulainen, R. (2000). Self-organization of innate face preferences: Could genetics be expressed through learning?. In *Proceedings of the 17th National Conference on Artificial Intelligence*, 117–122. Cambridge, MA: MIT Press.
- Blakemore, C., and Cooper, G. F. (1970). Development of the brain depends on the visual environment. *Nature*, 228:477–478.
- Blakemore, C., and van Sluyters, R. C. (1975). Innate and environmental factors in the development of the kitten's visual cortex. *Journal of Physiology (London)*, 248:663–716.
- Blasdel, G. G. (1992a). Differential imaging of ocular dominance columns and orientation selectivity in monkey striate cortex. *Journal of Neuroscience*, 12:3115–3138.
- Blasdel, G. G. (1992b). Orientation selectivity, preference, and continuity in monkey striate cortex. *Journal of Neuroscience*, 12:3139–3161.
- Blasdel, G. G., and Salama, G. (1986). Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, 321:579–585.
- Bosking, W. H., Zhang, Y., Schofield, B., and Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in tree shrew striate cortex. *Journal of Neuroscience*, 17(6):2112–2127.
- Bower, J. M., editor (2002). *Computational Neuroscience: Trends in Research, 2002*. New York: Elsevier. To appear.
- Britten, K. H., Shalden, M. N., Newsome, W. T., and Movshon, J. A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, 12:4745–4765.
- Burkhalter, A., Bernardo, K. L., and Charles, V. (1993). Development of local circuits in human visual cortex. *Journal of Neuroscience*, 13:1916–1931.
- Callaway, E. M., and Katz, L. C. (1990). Emergence and refinement of clustered horizontal connections in cat striate cortex. *Journal of Neuroscience*, 10:1134–1153.

- Callaway, E. M., and Katz, L. C. (1991). Effects of binocular deprivation on the development of clustered horizontal connections in cat striate cortex. *Proceedings of the National Academy of Sciences, USA*, 88:745–749.
- Campbell, S. R., Wang, D. L., and Jayaprakash, C. (1999). Synchrony and desynchrony in integrate-and-fire oscillators. *Neural Computation*, 11:1595–1619.
- Celebrini, S., and Newsome, W. T. (1994). Neuronal and psychophysical sensitivity to motion signals in extrastriate mst of the macaque monkey. *Journal of Neuroscience*, 14:4109–4124.
- Chakravarthy, S. V., and Ghosh, J. (1996). A complex-valued associative memory for storing patterns as oscillatory states. *Biological Cybernetics*, 75:229–238.
- Choe, Y. (1995). *Laterally Interconnected Self-Organizing Feature Map in Handwritten Digit Recognition*. Master's thesis, Department of Computer Sciences, The University of Texas at Austin. Technical Report AI95-236.
- Choe, Y., and Miikkulainen, R. (1997). Self-organization and segmentation with laterally connected spiking neurons. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, 1120–1125. San Francisco, CA: Morgan Kaufmann.
- Choe, Y., and Miikkulainen, R. (1998). Self-organization and segmentation in a laterally connected orientation map of spiking neurons. *Neurocomputing*, 21:139–157.
- Choe, Y., and Miikkulainen, R. (2000). A self-organizing neural network for contour integration through synchronized firing. In *Proceedings of the 17th National Conference on Artificial Intelligence*, 123–128. Cambridge, MA: MIT Press.
- Choe, Y., Miikkulainen, R., and Cormack, L. K. (2000). Effects of presynaptic and postsynaptic resource redistribution in Hebbian weight adaptation. *Neurocomputing*, 32–33:77–82.
- Choe, Y., Sirosh, J., and Miikkulainen, R. (1996). Laterally interconnected self-organizing maps in hand-written digit recognition. In Touretzky, D. S., Mozer, M. C., and Hasselmo, M. E., editors, *Advances in Neural Information Processing Systems 8*, 736–742. Cambridge, MA: MIT Press.
- Churchland, P. S., Ramachandran, V. S., and Sejnowski, T. J. (1994). A critique of pure vision. In Koch, C., and Davis, J. L., editors, *Large Scale Neuronal Theories of the Brain*. Cambridge, MA: MIT Press.
- Crowley, J. C., and Katz, L. C. (2000). Early development of ocular dominance columns. *Science*, 290:1321–1324.
- Dalva, M. B., and Katz, L. C. (1994). Rearrangements of synaptic connections in visual cortex revealed by laser photostimulation. *Science*, 265:255–258.

- Das, A., and Gilbert, C. (1997). Distortions of visuotopic map match orientation singularities in primary visual cortex. *Nature*, 387:594–598.
- Eckhorn, R., Bauer, R., Jordan, W., Kruse, M., Munk, W., and Reitboeck, H. J. (1988). Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological Cybernetics*, 60:121–130.
- Eckhorn, R., Reitboeck, H. J., Arndt, M., and Dicke, P. (1990). Feature linking via synchronization among distributed assemblies: Simulations of results from cat visual cortex. *Neural Computation*, 2:293–307.
- Ehrenstein, W. (1941). Über abwandlungen der l. hermannschen helligkeitserscheinung. *Zeitschrift für Psychologie*, 150:83–91. Modifications of Brightness Phenomenon of L. Hermann; translated by A. Hogg, in Petry and Meyer (1987), pp.35–39.
- Engel, A. K., König, P., Kreiter, A. K., and Singer, W. (1991a). Interhemispheric synchronization of oscillatory neuronal responses in cat visual cortex. *Science*, 252:1177–1179.
- Engel, A. K., Kreiter, A. K., König, P., and Singer, W. (1991b). Synchronization of oscillatory neuronal responses between striate and extrastriate visual cortical areas of the cat. *Proceedings of the National Academy of Sciences, USA*, 88:6048–6052.
- Erwin, E., Obermayer, K., and Schulten, K. (1995). Models of orientation and ocular dominance columns in the visual cortex: A critical comparison. *Neural Computation*, 7(3):425–468.
- Eurich, C. W., Pawelzik, K., Ernst, U., Cowan, J. D., and Milton, J. G. (1999). Dynamics of self-organized delay adaptation. *Physical Review Letters*, 82:1594–1597.
- Eurich, C. W., Pawelzik, K., Ernst, U., Thiel, A., Cowan, J. D., and Milton, J. G. (2000). Delay adaptation in the nervous system. *Neurocomputing*, 32–33:741–748.
- Fahle, M. (1993). Figure ground discrimination for temporal information. *Proceedings of the Royal Society of London B.*, 254:199–203.
- Ferster, D., and Lindström, S. (1985). Synaptic excitation of neurons in area 17 of the cat by intracortical axon collaterals of cortico-geniculate cells. *Journal of Physiology*, 367:233–252.
- Field, D. J., Hayes, A., and Hess, R. F. (1993a). Contour integration by the human visual system: Evidence for a local association field. *Vision Research*, 33:173–193.
- Field, D. J., Hayes, A., and Hess, R. F. (1993b). Contour integration by the human visual system: Evidence for a local association field. *Vision Research*, 33:173–193.

- Finkel, L. H., and Edelman, G. M. (1989). Integration of distributed cortical systems by reentry: A computer simulation of interactive functionally segregated visual areas. *Journal of Neuroscience*, 9:3188–3208.
- FitzHugh, R. (1961). Impulses and physiological states in models of nerve membrane. *Biophysical Journal*, 1:445–466.
- Gabbini, F., and Koch, C. (1999). Principles of spike train analysis. In Koch, C., and Segev, I., editors, *Methods in Neural Modeling*, chapter 7, 313–360. MIT Press.
- Geisler, W. S., Perry, J. S., Super, B. J., and Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*. 711–724.
- Geisler, W. S., and Super, B. (2000). Perceptual organization of two-dimensional patterns. *Psychological Review*, 107:677–708.
- Geisler, W. S., Thornton, T., Gallogly, D. P., and Perry, J. S. (1999). Image structure models of texture and contour visibility. In *Proceeding of the NATO Workshop on Search and Target Acquisition*.
- Gerstner, W. (1998a). Hebbian learning of pulse timing in the barn owl auditory system. In Maass, W., and Bishop, C. M., editors, *Pulsed Neural Networks*, chapter 14, 353–377. MIT Press.
- Gerstner, W. (1998b). Spiking neurons. In Maass, W., and Bishop, C. M., editors, *Pulsed Neural Networks*, chapter 1, 3–54. MIT Press.
- Gerstner, W., and van Hemmen, J. L. (1992). Associative memory in a network of spiking neurons. *Network*, 3:139–164.
- Gibson, J. J. (1950). *The Perception of the Visual World*. Boston: Houghton Mifflin.
- Gilbert, C. D. (1992). Horizontal integration and cortical dynamics. *Neuron*, 9:1–13.
- Gilbert, C. D., and Wiesel, T. N. (1979). Morphology and intracortical projections of functionally identified neurons in cat visual cortex. *Nature*, 280:120–125.
- Gilbert, C. D., and Wiesel, T. N. (1989). Columnar specificity of intrinsic horizontal and cortico-cortical connections in cat visual cortex. *Journal of Neuroscience*, 9:2432–2442.
- Gilbert, C. D., and Wiesel, T. N. (1992). Receptive field dynamics in adult primary visual cortex. *Nature*, 356:150–152.
- Goodhill, G. (1993). Topography and ocular dominance: A model exploring positive correlations. *Biological Cybernetics*, 69:109–118.

- Gove, A., Grossberg, S., and Mingolla, E. (1993). Brightness perception, illusory contours and corticogeniculate feedback. In *Proceedings of the World Congress on Neural Network*, vol. 1, 25–28. Erlbaum.
- Gray, C. M., Konig, P., Engel, A., and Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338:334–337.
- Gray, C. M., and McCormick, D. A. (1996). Chattering cells: Superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science*, 274:109–113.
- Gray, C. M., and Singer, W. (1987). Stimulus specific neuronal oscillations in the cat visual cortex: A cortical functional unit. In *Society for Neuroscience Abstracts*, vol. 13, 404.3.
- Grieve, K. L., and Sillito, A. M. (1995a). Non-length-tuned cells in layer II/III and IV of the visual cortex: the effect of blockade of layer VI on responses to stimuli of different lengths. *Experimental Brain Research*, 104:12–20.
- Grieve, K. L., and Sillito, A. M. (1995b). A re-appraisal of the role of layer VI of the visual cortex in the generation of cortical end inhibition. *Experimental Brain Research*, 104:12–20.
- Grinvald, A., Lieke, E. E., Frostig, R. D., and Hildesheim, R. (1994). Cortical point-spread function and long-range lateral interactions revealed by real-time optical imaging of macaque monkey primary visual cortex. *Journal of Neuroscience*, 14:2545–2568.
- Grossberg, S. (1999). How does the cerebral cortex work? Learning, attention, and grouping by the laminar circuits of visual cortex. *Spatial Vision*, 12:125–254.
- Grossberg, S., and Mingolla, E. (1985). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review*, 92:173–211.
- Grossberg, S., Mingolla, E., and Ross, W. D. (1997). Visual brain and visual perception: How does the cortex do perceptual grouping? *Trends in Neuroscience*, 20:106–111.
- Grossberg, S., and Williamson, J. R. (2001). A neural model of how horizontal and interlaminar connections of visual cortex develop into adult circuits that carry out perceptual grouping and learning. *Cerebral Cortex*, 9:878–895.
- Han, S. K., Kim, W. S., and Kook, H. (1998). Temporal segmentation of the stochastic oscillator neural network. *Physical Review E*, 58:2325–2334.
- Hata, Y., Tsumoto, T., Sato, H., Hagihara, K., and Tamura, H. (1993). Development of local horizontal interactions in cat visual cortex studied by cross-correlation analysis. *Journal of Neurophysiology*, 69:40–56.

- Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory*. New York: Wiley.
- Henry, G. H. (1989). Afferent inputs, receptive field properties and morphological cell types in different laminae of the striate cortex. In Leventhal, A. G., editor, *The Neural Basis of Visual Function*, vol. 4 of *Vision and Visual Dysfunction*, 223–245. Boca Raton, Florida: CRC Press.
- Hess, R. F., and Dakin, S. C. (1997). Absence of contour linking in peripheral vision. *Nature*, 390:602–604.
- Hirsch, H. V. B., and Spinelli, D. (1970). Visual experience modifies distribution of horizontally and vertically oriented receptive fields in cats. *Science*, 168:869–871.
- Hirsch, J. A., and Gilbert, C. D. (1991). Synaptic physiology of horizontal connections in the cat's visual cortex. *Journal of Neuroscience*, 11:1800–1809.
- Hodgkin, A. L., and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117:500–544.
- Hoffman, D. D. (1998). *Visual Intelligence: How We Create What We See*. Norton. First edition.
- Horn, D., Levy, N., and Ruppin, E. (1998). Memory maintenance via neuronal regulation. *Neural Computation*, 10:1–18.
- Horn, D., and Opher, I. (1998). Collective excitation phenomenon and their applications. In Maass, W., and Bishop, C. M., editors, *Pulsed Neural Networks*, chapter 11, 297–320. MIT Press.
- Horn, D., and Usher, M. (1992). Oscillatory model of short term memory. In *Advances in Neural Information Processing Systems*, 4, 125–132.
- Hubel, D. H., and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology (London)*, 160:106–154.
- Hubel, D. H., and Wiesel, T. N. (1974). Sequence regularity and geometry of orientation columns in the monkey striate cortex. *Journal of Comparative Neurology*, 158:267–294.
- 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*, 278:377–409.
- Issa, N. P., Trachtenberg, J. T., Chapman, B., Zahs, K. R., and Stryker, M. P. (1999). The critical period for ocular dominance plasticity in the ferret's visual cortex. *Journal of Neuroscience*, 19(16):6965–6978.

- Issa, N. P., Trepel, C., and Stryker, M. P. (2001). Spatial frequency maps in cat visual cortex. *Journal of Neuroscience*, 20:8504–8514.
- Joliot, M., Ribary, U., and Linás, R. (1994). Human oscillatory brain activity near 40 Hz coexists with cognitive temporal binding. *Proceedings of the National Academy of Sciences, USA*, 91:11748–11751.
- Kalarickal, G. J., and Marshall, J. A. (1997). Visual classical rearing and synaptic plasticity: Comparison of exin and bcm learning rules. In *Proceedings of Vision, Recognition, Action: Neural Models of Mind and Machine*, 28.
- Kandel, E. R., Schwartz, J. H., and Jessell, T. M. (1991). *Principles of Neural Science*. New York: Elsevier. Third edition.
- Kanizsa, G. (1955). Margini quasi-percettivi in campi con stimolazione omogenea. *Rivista di Psicologia*, 49:7–30. Quasiperceptual Margins in Homogeneously Stimulated Fields; translated by W. Gerbino, in Petry and Meyer (1987), pp.40–49.
- Kanizsa, G. (1976). Subjective contours. *Scientific American*, 234:48–52.
- Kapadia, M. K., Gilbert, C. D., and Westheimer, G. (1994). A quantitative measure for short-term cortical plasticity in human vision. *Journal of Neuroscience*, 14:451–457.
- Katz, L. C., and Callaway, E. M. (1992). Development of local circuits in mammalian visual cortex. *Annual Review of Neuroscience*, 15:31–56.
- Katz, L. C., and Shatz, C. J. (1996). Synaptic activity and the construction of cortical circuits. *Science*, 274:1133–1138.
- Kirillov, A. B., and Woodward, D. J. (1993). Synchronization of spiking neurons: Transmission delays, noise and nmda receptors. In *Proceedings of the World Congress on Neural Networks*, 594–597.
- Koch, C., and Davis, J. L., editors (1994). *Large Scale Neuronal Theories of the Brain*. Cambridge, MA: MIT Press.
- Koch, C., and Segev, I., editors (1998). *Methods in Neuronal Modeling*. Cambridge, MA: MIT Press. Second edition.
- Kohonen, T. (1981). Automatic formation of topological maps of patterns in a self-organizing system. In *Proceedings of the 2nd Scandinavian Conference on Image Analysis*, 214–220. Espoo, Finland: Pattern Recognition Society of Finland.
- Kohonen, T. (1982a). Analysis of a simple self-organizing process. *Biological Cybernetics*, 44:135–140.

- Kohonen, T. (1982b). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43:59–69.
- Kohonen, T. (1989). *Self-Organization and Associative Memory*. Berlin; New York: Springer. Third edition.
- Kohonen, T. (1993). Physiological interpretation of the self-organizing map algorithm. *Neural Networks*, 6:895–905.
- Koulakov, A. A., and Chklovskii, D. B. (2001). Orientation preference patterns in mammalian visual cortex: A wire length minimization approach. *Neuron*, 29:519–527.
- Kovacs, I., and Julesz, B. (1993). A closed curve is much more than an incomplete one. *Proceedings of the National Academy of Sciences, USA*, 90:7495–7497.
- Lapicque, L. (1907). Recherches quantitatives sur l’excitation électrique des nerfs traitée comme une polarisation. *Physiol. Pathol. Gen.*, 9:620–635.
- Lee, S.-H., and Blake, R. (1999). Visual form created solely from temporal structure. *Science*, 284:1165–1168.
- Lee, S.-W., Bülthoff, H. H., and Poggio, T., editors (2000a). *First IEEE International Workshop, Biologically Motivated Computer Vision 2000*. Springer.
- Lee, T.-W., Wachtler, T., and Sejnowski, T. J. (2000b). The spectral independent components of natural scenes. In (Lee et al. 2000a), 527–534.
- Leonards, U., and Singer, W. (1998). Two segmentation mechanisms with differential sensitivity for colour and luminance contrast. *Vision Research*, 38:101–109.
- Leonards, U., Singer, W., and Fahle, M. (1996). The influence of temporal phase difference on texture segmentation. *Vision Research*, 36:2689–2697.
- Leshner, G. W., and Mingolla, E. (1995). Illusory contour formation. In (Arbib 1995), 481–483.
- Li, Z. (1998). A neural model of contour integration in the primary visual cortex. *Neural Computation*, 10:903–940.
- Li, Z. (1999). Visual segmentation by contextual influences via inter-cortical interactions in the primary visual cortex. *Network: Computational Neural Systems*, 10:187–212.
- Linsker, R. (1986). From basic network principles to neural architecture: Emergence of spatial-opponent cells. *Proceedings of the National Academy of Sciences, USA*, 83:7508–7512.
- Lisman, J. (1998). What makes the brain’s ticker tock. *Nature*, 394:132–133.

- Löwel, S. (1994). Ocular dominance column development: Strabismus changes the spacing of adjacent columns in cat visual cortex. *Journal of Neuroscience*, 14(12):7451–7468.
- Löwel, S., and Singer, W. (1992). Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science*, 255:209–212.
- Lytton, W. W., and Sejnowski, T. J. (1991). Simulations of cortical pyramidal neurons synchronized by inhibitory interneurons. *Journal of Neurophysiology*, 66:1059–1079.
- Maass, W. (1998). Computing with spiking neurons. In Maass, W., and Bishop, C. M., editors, *Pulsed Neural Networks*, chapter 14, 55–85. MIT Press.
- McGuire, B. A., Gilbert, C. D., Rivlin, P. K., and Wiesel, T. N. (1991). Targets of horizontal connections in macaque primary visual cortex. *Journal of Comparative Neurology*, 305:370–392.
- McIlhagga, W. H., and Mullen, K. T. (1996). Contour integration with colour and luminance contrast. *Vision Research*, 36:1265–1279.
- Menon, V. (1990). *Dynamic Aspects of Signaling in Distributed Neural Systems*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin. Technical Report TR-90-36.
- Miikkulainen, R., Bednar, J. A., Choe, Y., and Sirosh, J. (1997). Self-organization, plasticity, and low-level visual phenomena in a laterally connected map model of the primary visual cortex. In Goldstone, R. L., Schyns, P. G., and Medin, D. L., editors, *Perceptual Learning*, vol. 36 of *Psychology of Learning and Motivation*, 257–308. San Diego, CA: Academic Press.
- Miikkulainen, R., and Sirosh, J. (1996). Introduction: The emerging understanding of lateral interactions in the cortex. In Sirosh, J., Miikkulainen, R., and Choe, Y., editors, *Lateral Interactions in the Cortex: Structure and Function*. Austin, TX: The UTCS Neural Networks Research Group. Electronic book, ISBN 0-9647060-0-8, <http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96>.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity of processing information. *Psychological Review*, 63:81–97.
- Miller, K. D. (1994). A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between ON- and OFF-center inputs. *Journal of Neuroscience*, 14:409–441.
- Miller, K. D., Keller, J. B., and Stryker, M. P. (1989). Ocular dominance column development: Analysis and simulation. *Science*, 245:605–615.

- Miller, K. D., and MacKay, D. J. C. (1994). The role of constraints in Hebbian learning. *Neural Computation*, 6:100–126.
- Mirollo, R. E., and Strogatz, S. H. (1990). Synchronization of pulse-coupled biological oscillators. *SIAM Journal of Applied Mathematics*, 50:1645–1662.
- Nagumo, J. S., Arimoto, S., and Yoshizawa, S. (1962). An active pulse transmission line simulating a nerve axon. *Proceedings of IRE*, 50:2061–2070.
- Nakayama, K., and Shimojo, S. (1992). Experiencing and perceiving visual surfaces. *Science*, 257:1357–1363.
- Nelson, J. I. (1995). Visual scene perception: Neurophysiology. In (Arbib 1995), 1024–1028.
- Nischwitz, A., and Glünder, H. (1995). Local lateral inhibition: A key to spike synchronization? *Biological Cybernetics*, 73:389–400.
- Nowak, L. G., and Bullier, J. (1997). The timing of information transfer in the visual system. *Cerebral Cortex*, 12:205–239.
- Obermayer, K., Blasdel, G. G., and Schulten, K. J. (1992). Statistical-mechanical analysis of self-organization and pattern formation during the development of visual maps. *Physical Review A*, 45:7568–7589.
- Obermayer, K., Ritter, H. J., and Schulten, K. J. (1990). A principle for the formation of the spatial structure of cortical feature maps. *Proceedings of the National Academy of Sciences, USA*, 87:8345–8349.
- Oja, E. (1992). Principal components, minor components, and linear neural networks. *Neural Networks*, 5:927–935.
- O’Keefe, J., and Reece, M. (1993). Phase relationship between hippocampal place units and the hippocampal theta rhythm. *Hippocampus*, 3:317–330.
- O’Keefe, L. P., Levitt, J. B., Kiper, D. C., Shapley, R. M., and Movshon, J. A. (1998). Functional organization of owl monkey lateral geniculate nucleus and visual cortex. *Journal of Neurophysiology*, 80:594–609.
- Olshausen, B. A., and Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607–609.
- Parks, T. E. (1980). Letter to the editor. *Perception*, 9:723.

- Peterhans, E., von der Heydt, R., and Baumgartner, G. (1986). Neuronal responses to illusory contour stimuli reveal stages of visual cortical processing. In Pettigrew, J. D., Sanderson, K. J., and Levick, W. R., editors, *Visual Neuroscience*, 343–351. Cambridge, England: Cambridge University Press.
- Petry, S., and Meyer, G. E., editors (1987). *The Perception of Illusory Contours*. New York: Springer.
- Pettet, M. W., and Gilbert, C. D. (1992). Dynamic changes in receptive-field size in cat primary visual cortex. *Proceedings of the National Academy of Sciences, USA*, 89:8366–8370.
- Pettet, M. W., McKee, S. P., and Grzywacz, N. M. (1998). Constraints on long range interactions mediating contour detection. *Vision Research*, 38:865–879.
- Piepenbrock, C., Ritter, H., and Obermayer, K. (1997). The joint development of orientation and ocular dominance: Role of constraints. *Neural Computation*, 9:959–970.
- Prazdny, K. (1983). Illusory contours are not caused by simultaneous brightness contrast. *Perception and Psychophysics*, 34:403–404.
- Previc, F. H. (1990). Functional specialization in the lower and upper visual fields in humans: Its ecological origins and neurophysiological implications. *Behavioral and Brain Sciences*, 13:519–575.
- Rao, R. P. N., and Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2:79–87.
- Redies, C., Crook, J. M., and Cruetzfeldt, O. D. (1986). Neuronal responses to borders with and without luminance gradients in cat visual cortex and dorsal lateral geniculate nucleus. *Experimental Brain Research*, 61:569–481.
- Reitboeck, H., Stoecker, M., and Hahn, C. (1993). Object separation in dynamic neural networks. In *Proceedings of the IEEE International Conference on Neural Networks* (San Francisco, CA), vol. 2, 638–641. Piscataway, NJ: IEEE.
- Rieke, F., Warland, D., de Ruter van Steveninck, R., and Bialek, W. (1997). *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT Press. First edition.
- Rinzel, J., and Ermentrout, B. (1999). Analysis of neural excitability and oscillations. In Koch, C., and Segev, I., editors, *Methods in Neural Modeling*, chapter 7, 251–291. MIT Press.
- Ross, W. D., Grossberg, S., and Mingolla, E. (2000). Visual cortical mechanisms of perceptual grouping: Interacting layers, networks, columns, and maps. *Neural Networks*, 13:571–588.

- Rubin, N., Nakayama, K., and Shapley, R. (1996). Enhanced perception of illusory contours in the lower versus upper visual hemifields. *Science*, 271:651–653.
- Salzman, C. D., Britten, K. H., and Newsome, W. T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, 346:174–177, Erratum 346, 589.
- Schumann, F. (1904). Einige böbachtungen über die zusammenfassung von gesichtseindrucken zu einheiten. *Psychologische Studien*, 1:1–32.
- Sharma, J., Angelucci, A., and Sur, M. (2000). Induction of visual orientation modules in auditory cortex. *Nature*, 404:841–847.
- Sheth, B. R., Sharma, J., Rao, S. C., and Sur, M. (1996). Orientation maps of subjective contours in visual cortex. *Science*, 274:2110–2115.
- Shipley, T. F., and Kellman, P. J. (1992). Strength of visual interpolation depends on the ratio of physically specified to total edge length. *Perception and Psychophysics*, 52:97–106.
- Shmuel, A., and Grinvald, A. (1996). Functional organization for direction of motion and its relationship to orientation maps in cat area 18. *The Journal of Neuroscience*, 16:6945–6964.
- Shouval, H., and Cooper, L. N. (1996). Organization of receptive fields for networks of neurons with hebbian type learning: The connection between synaptic and phenomenological. *Biological Cybernetics*, 73:439–447.
- Shouval, H., Intrator, N., and Cooper, L. (1997). BCM network develops orientation selectivity and ocular dominance in natural scene environment. *Vision Research*, 37:3339–3342.
- Sigman, M., Cecchi, G. A., Gilbert, C. D., and Magnasco, M. O. (2001). On a common circle: Natural scenes and gestalt rules. *Proceedings of the National Academy of Sciences, USA*, 98:1935–1940.
- Sillito, A. M., Jones, H. E., Gerstein, G. L., and West, D. C. (1994). Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. *Nature*, 369:479–482.
- Sincich, L. C., and Blasdel, G. G. (2001). Oriented axon projections in primary visual cortex of the monkey. *Journal of Neuroscience*, 21:4416–4426.
- Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. *Annual Review of Physiology*, 55:349–374.
- Sirosh, J. (1995). *A Self-Organizing Neural Network Model of the Primary Visual Cortex*. PhD thesis, Department of Computer Sciences, The University of Texas at Austin, Austin, TX. Technical Report AI95-237.

- Sirosh, J., and Miikkulainen, R. (1994). Cooperative self-organization of afferent and lateral connections in cortical maps. *Biological Cybernetics*, 71:66–78.
- Sirosh, J., and Miikkulainen, R. (1996). Self-organization and functional role of lateral connections and multisize receptive fields in the primary visual cortex. *Neural Processing Letters*, 3:39–48.
- Sirosh, J., and Miikkulainen, R. (1997). Topographic receptive fields and patterned lateral interaction in a self-organizing model of the primary visual cortex. *Neural Computation*, 9:577–594.
- Sirosh, J., Miikkulainen, R., and Bednar, J. A. (1996a). 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*. Austin, TX: The UTCS Neural Networks Research Group. Electronic book, ISBN 0-9647060-0-8, <http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96>.
- Sirosh, J., Miikkulainen, R., and Choe, Y., editors (1996b). *Lateral Interactions in the Cortex: Structure and Function*. Austin, TX: The UTCS Neural Networks Research Group. Electronic book, ISBN 0-9647060-0-8, <http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96>.
- Sur, M., Garraghty, P. E., and Roe, A. W. (1988). Experimentally induced visual projections in auditory thalamus and cortex. *Science*, 242:1437–1441.
- Terman, D., and Wang, D. (1995). Global competition and local cooperation in a network of neural oscillators. *Physica D*, 81:148–176.
- Ts'o, D. Y., Frostig, R. D., Lieke, E. E., and Grinvald, A. (1990). Functional organization of primate visual cortex revealed by high resolution optical imaging. *Science*, 249:417–420.
- Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature*, 391:845–846.
- Tversky, T., and Miikkulainen, R. (2002). Modeling directional selectivity using self-organized delay-adaptation maps. In (Bower 2002). To appear.
- Ullman, S. (1976). Filling-in the gaps: The shape of subjective contours and a model for their generation. *Biological Cybernetics*, 25:1–6.
- Usher, M., and Donnelly, N. (1998). Visual synchrony affects binding and segmentation in perception. *Nature*, 394:179–182.
- van Vreeswijk, C., and Abbott, L. F. (1994). When inhibition not excitation synchronizes neural firing. *Journal of Computational Neuroscience*, 1:313–321.

- von der Heydt, R., and Peterhans, E. (1989). Mechanisms of contour perception in monkey visual cortex. I. Lines of pattern discontinuity. *Journal of Neuroscience*, 9:1731–1748.
- von der Malsburg, C. (1973). Self-organization of orientation-sensitive cells in the striate cortex. *Kybernetik*, 15:85–100. Reprinted in Anderson and Rosenfeld 1988.
- von der Malsburg, C. (1981). The correlation theory of brain function. Internal Report 81-2, Department of Neurobiology, Max-Planck-Institute for Biophysical Chemistry, Göttingen, Germany.
- von der Malsburg, C. (1987). Synaptic plasticity as basis of brain organization. In Changeux, J.-P., and Konishi, M., editors, *The Neural and Molecular Bases of Learning*, 411–432. New York: Wiley.
- von der Malsburg, C., and Buhmann, J. (1992). Sensory segmentation with coupled neural oscillators. *Biological Cybernetics*, 67:233–242.
- von der Malsburg, C., and Schneider, W. (1986). A neural cocktail-party processor. *Biological Cybernetics*, 54:29–40.
- Wachtler, T., Lee, T.-W., and Sejnowski, T. J. (2001). The chromatic structure of natural scenes. *Journal of the Optical Society of America*, 18:65–77.
- Wang, D. (1996). Synchronous oscillations based on lateral connections. In Sirosh, J., Miikkulainen, R., and Choe, Y., editors, *Lateral Interactions in the Cortex: Structure and Function*. Austin, TX: The UTCS Neural Networks Research Group. Electronic book, ISBN 0-9647060-0-8, <http://www.cs.utexas.edu/users/nn/web-pubs/htmlbook96>.
- Wang, D. L. (1995). Emergent synchrony in locally coupled neural oscillators. *IEEE Transactions on Neural Networks*, 6:941–948.
- Weliky, M., Bosking, W. H., and Fitzpatrick, D. (1996). A systematic map of direction preference in primary visual cortex. *Nature*, 379:725–728.
- Weliky, M., Kandler, K., Fitzpatrick, D., and Katz, L. C. (1995). Patterns of excitation and inhibition evoked by horizontal connections in visual cortex share a common relationship to orientation columns. *Neuron*, 15:541–552.
- White, L. E., Coppola, D. M., and Fitzpatrick, D. (2001). The contribution of sensory experience to the maturation of orientation selectivity in ferret visual cortex. *Nature*, 411:1049–1052.
- White, L. E., Van Loon, K. J., Coppola, D. M., and Fitzpatrick, D. (2000). Effects of visual deprivation on functional maps and horizontal connections in developing ferret visual cortex. In *Society for Neuroscience Abstracts*.

- Wilson, H. R., and Cowan, J. D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. *Biophysical Journal*, 12:1–24.
- Yen, S.-C., and Finkel, L. (1998). Extraction of perceptually salient contours by striate cortical networks. *Vision Research*, 38:719–741.
- Yen, S.-C., and Finkel, L. H. (1997). Identification of salient contours in cluttered images. In *Computer Vision and Pattern Recognition*, 273–279.
- Zador, A. M., and Pearlmutter, B. A. (1996). VC dimension of an integrate-and-fire neuron model. *Neural Computation*, 8:611–624.

Vita

Yoonsuck Choe was born in Freetown, Sierra Leone on January 3rd, 1970, the son of Woojin Choe and Mija Kim. After returning back to Korea with his family, all of his education took place in Seoul, Korea. After receiving his Bachelor of Science degree, majoring in Computer Science, at Yonsei University in Seoul, Korea in the Summer of 1993, he entered the University of Texas at Austin as a graduate student in the Fall of 1993. In December 1995, he obtained his Master of Arts degree in Computer Sciences.

Permanent Address: 3203 Von Trapp Ln.
College Station, TX 77845

This dissertation was typeset with L^AT_εX 2_ε¹ by the author.

¹L^AT_εX 2_ε is an extension of L^AT_EX. L^AT_EX is a collection of macros for T_EX. T_EX is a trademark of the American Mathematical Society. The macros used in formatting this dissertation were written by Dinesh Das, Department of Computer Sciences, The University of Texas at Austin.