

A mechanistic model of motion processing in the early visual system

by

Aziz Hurzook

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2012

© Aziz Hurzook 2012

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

A prerequisite for the perception of motion in primates is the transformation of varying intensities of light on the retina into an estimation of position, direction and speed of coherent objects. The neuro-computational mechanisms relevant for object feature encoding have been thoroughly explored, with many neurally plausible models able to represent static visual scenes. However, motion estimation requires the comparison of successive scenes through time. Precisely how the necessary neural dynamics arise and how other related neural system components interoperate have yet to be shown in a large-scale, biologically realistic simulation.

The proposed model simulates a spiking neural network computation for representing object velocities in cortical areas V1 and middle temporal area (MT). The essential neural dynamics, hypothesized to reside in networks of V1 simple cells, are implemented through recurrent population connections that generate oscillating spatiotemporal tunings. These oscillators produce a resonance response when stimuli move in an appropriate manner in their receptive fields. The simulation shows close agreement between the predicted and actual impulse responses from V1 simple cells using an ideal stimulus.

By integrating the activities of like V1 simple cells over space, a local measure of visual pattern velocity can be produced. This measure is also the linear weight of an associated velocity in a retinotopic map of optical flow. As a demonstration, the classic motion stimuli of drifting sinusoidal gratings and variably coherent dots are used as test stimuli and optical flow maps are generated. Vector field representations of this structure may serve as inputs for perception and decision making processes in later brain areas.

Acknowledgements

It is because of the support of my family that I have been able to complete this work. My enduring gratitude goes to my wife and daughter for their help in clearing the way for me to pursue this endeavour. I am also grateful to my mother, brother and sister for their continued interest in my often perplexing pursuits. My wife's side of the family has also provided an abundance of encouragement and well-wishes that I greatly appreciate.

I am much obliged to Terry Stewart of our lab group for his selfless assistance in my efforts, and to the rest of the group for their unusual ability to be as clever as they are kind. I wish to have spent more time with them, but I hope our relationships will continue to grow over the stretches of time and distance. There have also been many talented group members over the years who have contributed to the simulation software that makes these large scale models feasible and theoretical results possible. Their contributions are used thousands of times per day.

My thanks must also be extended to Professors Bryan Tripp and Jeff Orchard for their helpful revisions to earlier drafts of this work.

Of course, to my supervisor Chris Eliasmith, who has enabled and inspired both me and this work, I am greatly in debt. I hope I have been able to add to the ideas he has offered, to make the time he has spent on my behalf worthwhile.

Dedication

For my Dad.

Table of Contents

| | |
|--|-------------|
| List of Figures | xv |
| List of Tables | xvii |
| Preface | 1 |
| Structure of the thesis | 1 |
| Motivation and scientific contribution | 1 |
| 1 Physiology and theory of visual motion | 3 |
| 1.1 The process of seeing motion: anatomy and function | 3 |
| 1.1.1 Phototransduction in the retina | 3 |
| 1.1.2 The lateral geniculate nucleus | 5 |
| 1.1.3 Spatiotemporal filtering in V1 | 6 |
| 1.1.4 Motion integration and velocity representation in MT | 8 |
| 1.2 Current theoretical approaches | 11 |
| 1.2.1 Linear filtering as a modelling tool | 11 |
| 1.2.2 Receptive fields are results of prior computation | 13 |
| 1.2.3 The Gabor filter for V1 simple cell spatial responses | 13 |
| 1.2.4 Models of temporal processing in LGN | 15 |
| 1.2.5 Models of V1 direction and speed selectivity | 16 |
| 1.2.6 Models of motion processing in MT | 18 |
| 1.3 Relationship of the proposed model to previous work | 20 |
| 2 Theoretical principles of the proposed model | 23 |
| 2.1 The LIF neuron model | 24 |
| 2.1.1 Suitability of the LIF neuron model | 26 |
| 2.2 Neural network modelling principles | 26 |
| 2.2.1 Vector encoding | 27 |
| 2.2.2 Vector decoding: principles of optimal linear estimation | 28 |
| 2.2.3 Vector decoding: noise and variability | 30 |

| | | |
|----------|---|-----------|
| 2.2.4 | Vector decoding: precision | 31 |
| 2.2.5 | Vector decoding: spiking neurons and time-dependence | 31 |
| 2.3 | Vector transformation | 32 |
| 2.3.1 | Linear transformations | 32 |
| 2.3.2 | Nonlinear transformations | 33 |
| 2.4 | Dynamics | 34 |
| 2.4.1 | The state-space method: an example using a damped oscillator | 35 |
| 2.4.2 | Adapting state-space methods to neural ensembles | 38 |
| 2.4.3 | The oscillator interference (OI) concept | 39 |
| 3 | A mechanistic model of motion processing in the early visual system | 43 |
| 3.1 | Algorithmic description | 43 |
| 3.2 | Implementation details | 45 |
| 3.2.1 | The Nengo simulation environment | 45 |
| 3.2.2 | Stimulus movie generation | 45 |
| 3.2.3 | Basis of representation | 45 |
| 3.2.4 | Centre-surround edge extraction and subsampling | 47 |
| 3.2.5 | Spatiotemporal filtering by LGN/V1 | 48 |
| 3.2.6 | Spatial integration in MT | 49 |
| 3.2.7 | Neural ensemble network and parameters | 51 |
| 3.2.8 | Physiological parameters for neural ensembles | 52 |
| 4 | Experimental results | 53 |
| 4.1 | The generation of velocity selectivity | 53 |
| 4.2 | Predicted impulse response of V1 simple cells | 53 |
| 4.3 | Velocity selectivity in MT | 57 |
| 4.3.1 | Distribution of direction responses | 57 |
| 4.3.2 | Distribution of speed responses | 57 |
| 4.3.3 | Contrast invariance | 61 |
| 4.4 | Optical flow maps in MT | 61 |
| 4.4.1 | Response to sinusoidal gratings | 61 |
| 4.4.2 | Response to variable coherence dot patterns | 65 |
| 5 | Discussion and conclusions | 67 |
| 5.1 | Strengths of the model | 67 |
| 5.2 | Weaknesses of the model | 68 |
| 5.3 | Extensions of the model | 69 |
| 5.3.1 | Decision threshold mechanism by integration of MT responses over time | 69 |
| 5.3.2 | Integration of MT responses over scale space | 69 |
| 5.3.3 | Burst/derivative mechanisms in LGN/V1 | 69 |

List of Figures

| | | |
|------|---|----|
| 1.1 | The motion processing pathway. | 4 |
| 1.2 | Simplified anatomy of the human eye and retina. | 4 |
| 1.3 | Idealized and experimental centre-surround receptive fields. | 5 |
| 1.4 | Episodic spike trains from the LGN. | 7 |
| 1.5 | One-dimensional time courses for nonlagged (left) and lagged (right) LGN neurons. | 7 |
| 1.6 | Diversity of spatial RF structure in V1 simple cells. | 8 |
| 1.7 | Idealized separable and inseparable RF time courses. | 9 |
| 1.8 | Example time courses from V1 simple cells. | 10 |
| 1.9 | Contrast invariant orientation selectivity in V1. | 11 |
| 1.10 | Spiking responses from direction selective neurons in macaque MT. | 12 |
| 1.11 | Distribution of response latencies in macaque MT under direction test. | 12 |
| 1.12 | An example of convolution. | 13 |
| 1.13 | Gabor function variation of parameters. | 14 |
| 1.14 | The LGN temporal impulse response predicted by Dong and Atick. | 15 |
| 1.15 | The motion energy model by Adelson and Bergen. | 17 |
| 1.16 | Intersection of constraints solution to the aperture problem. | 19 |
| 1.17 | The divisive normalization model by Simoncelli and Heeger. | 19 |
| | | |
| 2.1 | Neuron communication. | 23 |
| 2.2 | An RC circuit. | 24 |
| 2.3 | The voltage time course for a LIF neuron. | 24 |
| 2.4 | Tuning curves of a neural ensemble. | 27 |
| 2.5 | Error due to noise and static distortion vs number of neurons. | 31 |
| 2.6 | A linear transformation network. | 33 |
| 2.7 | A nonlinear transformation network. | 34 |
| 2.8 | Impulse response of damped, driven oscillator in one dimension. | 36 |
| 2.9 | A damped, driven oscillator in one dimension with resonant input. | 37 |
| 2.10 | LTI block diagram. | 38 |
| 2.11 | A 2-D damped, driven oscillator. | 39 |

| | | |
|------|--|----|
| 3.1 | Conceptual diagram of the proposed ‘oscillator interference’ (OI) model. | 44 |
| 3.2 | Sample of basis functions used. | 46 |
| 3.3 | Representation of structures in the basis. | 46 |
| 3.4 | Edge extraction, subsampling and V1 array structure. | 47 |
| 3.5 | LGN burst timing. | 49 |
| 3.6 | Damped frequency vs minimum sampling rate. | 50 |
| 3.7 | Decay vs damped frequency. | 50 |
| 3.8 | MT array configuration. | 51 |
| 3.9 | Velocity-selective microcircuit schematic. | 51 |
| 4.1 | The velocity selectivity microcircuit. | 54 |
| 4.2 | Predicted impulse response components of a V1 state phase oscillator. | 55 |
| 4.3 | Predicted vs actual one-dimensional time courses (‘x-t’ maps) for receptive fields of V1 simple cells. | 56 |
| 4.4 | Predicted vs actual tuning curve for an MT neuron. | 57 |
| 4.5 | Direction selectivity of MT responses. | 58 |
| 4.6 | Speed selectivity of MT ensemble responses. | 59 |
| 4.7 | Effect of contrast on MT responses. | 60 |
| 4.8 | MT velocity maps for drifting sinusoidal grating stimulus. | 62 |
| 4.9 | MT velocity maps showing strong transparent retrograde motion. | 63 |
| 4.10 | MT velocity maps for moving dots of varying coherence. | 64 |
| 5.1 | The spatial structure of RFs affects convolution or correlation responses significantly. | 70 |

List of Tables

| | | |
|-----|--|----|
| 3.1 | Neural ensemble parameters used in the simulation. | 52 |
|-----|--|----|

Preface

Structure and content of the thesis

The essential problem for this study is to explain, through neurally plausible simulation, how the early visual system in some mammals is able to compute the velocity of patterns moving in the visual field. The proposed model explains how the observed signal characteristics from the retina and lateral geniculate nuclei (LGN) are transformed in the spatiotemporal domain by recurrent networks of V1 simple cells to produce a retinotopic map of velocity responses in the middle temporal area (MT or V5).

Chapter 1 contains a condensed review of the relevant functional anatomy of the early visual system. This is followed by a brief overview of the prevailing models of visual motion, including a discussion of some important shortcomings of and challenges to those hypotheses.

The theoretical and mathematical principles that apply to the proposed model are detailed in Chapter 2. A model of velocity computation is detailed in Chapter 3 and draws on the material covered earlier. The results of simulations of the early visual system under a variety of stimuli are provided in Chapter 4, with analysis and comparison to electrophysiological data in Chapter 5.

Motivation and scientific contribution

A fundamental characteristic of the velocity-selective response patterns or receptive fields (RFs) of neurons in the primary visual cortex (V1) of primates is the apparent translation of Gabor-like spatial functions over time, called *phase evolution*. According to theoretical models of electrophysiological data, this spatiotemporal characteristic allows the simple cells of V1 to function as a type of velocity detector. However, the mechanism by which phase evolution occurs is not known, and hence, velocity selection and motion processing in the brain are not fully understood.

In this work, a biologically constrained, spiking neural network model achieves velocity selectivity by incorporating known orientation selection methods with a novel speed selection solution. The velocity-selective mechanism correlates an efficiently-coded intensity pattern with the dynamic representational state of a neural network oscillator. The model is implemented in a large-scale simulation with all connection weights and physiological time constants specified. The results explain the generation of transient and sustained

components of separable and inseparable RFs in V1 simple cells, and how characteristic bursts of activity from the LGN interact with V1 networks to produce velocity selectivity.

Some of the same stimuli employed in classic psychometric tests on monkeys and cats are used as test input for the motion processing circuit. The results are in close agreement with the experimentally measured neurometric behaviour of V1 simple cells and MT, supporting the essential claims of the model. It is not the purpose of the model or these results to explain the psychometric, perceptual or behavioural data of test subject animals.

To my knowledge, there is no other biologically plausible, large-scale dynamical neural network simulation giving results that approximate the neural activity data described herein.

Chapter 1

Physiology and theory of visual motion

According to Helmholtz, seeing is an unconscious inference [1]. It is the result of a computational process performed by millions of interconnected cells in the eyes and brain. This system, known as the *early visual system*, uses as its input a noisy, unstable, narrow band of electromagnetic intensities projected over millions of photoreceptive cells in two displaced retinæ. The visual system of humans and many other animals converts these flickers to forms, colours, objects, perceptions, decisions and memories using vast arrays of neurons [2] that communicate with each other every ten to one hundred milliseconds or so using primarily electrochemical means [3].

Research in vision has historically been subdivided into anatomical domains such as the retina, the LGN, V1 (or primary visual cortex, or striate cortex), MT, as well as functional domains including phototransduction, edge, colour and motion detection along with many others. This chapter attempts to synthesize the prevailing understanding of the functional, anatomical and computational structure of the early visual system as it contributes to the appearance of velocity-tuned neurons in MT.

1.1 The process of seeing motion: anatomy and function

The synthesis that follows is a brief account of some prevailing interpretations of anatomical structure and function of the major structures implicated in the perception of motion: the retina, LGN, V1 and MT (Figure 1.1). The complex chemophysical phenomena that underlie the behaviour of these tissue structures are largely beyond the scope of this work but are mentioned when they bear on computational issues. Discussion of quantitative theoretical models related to the aforementioned anatomical structures and velocity computation begins in section 1.2.

1.1.1 Phototransduction in the retina

The process of *phototransduction* occurs in the retina at the cellular level, where the photons of visible light are absorbed selectively by cells called *photoreceptors* (rods and




Figure 1.1. The essential anatomical pathway for motion processing. The diagram depicts a simplified pathway for the processing of light intensity signals on the retina to the retinotopic motion map in area MT. See text for details.

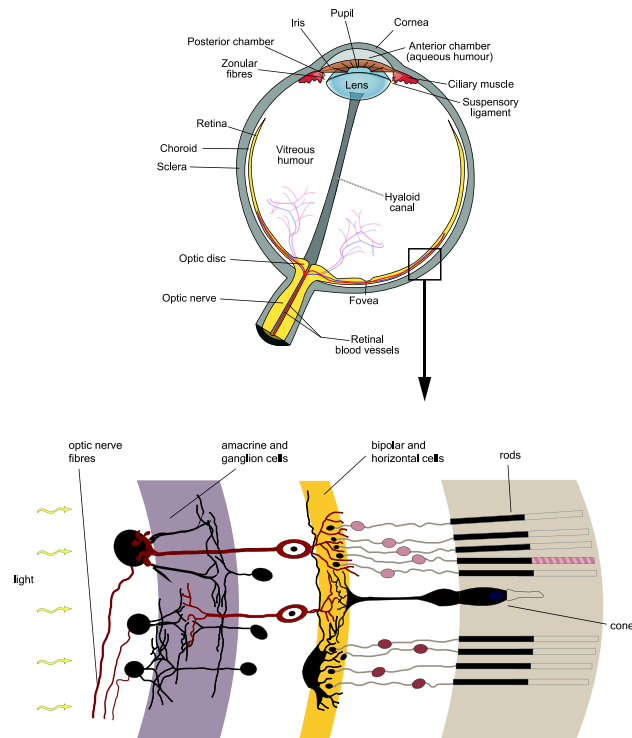


Figure 1.2. Human eye anatomy. *Top:* Simplified anatomy of the human eye. Public domain image. *Bottom:* Axial organization of the retina. Public domain image adapted from *Histologie Du Système Nerveux de l'Homme et Des Vertébrés*, Maloine, Ramón y Cajal, 1911.)




Figure 1.3. The centre-surround concept. *Left, middle:* Idealized centre-surround receptive fields of each type, generated from the second derivative (Laplacian) of a Gaussian function. The positive signed regions indicate areas of visual space in which the presence of light excites the cell. Light cast onto the negative region causes inhibition. *Right:* Contour plot of the spatial RF from the LGN of a cat with visual area of 3x3 degrees, similar to those of the retinal ganglion cells. Adapted, used with permission of DeAngelis [6].

cones) (Figure 1.2) [4]. The energy from this absorption induces a complex sequence of events that result in a change in the cell membrane potential [4]. Photoreceptors relay information chemically to *bipolar cells* (BPCs) through an interface called a *synapse* [5].

BPCs have two classes based on the spatial distribution of sensitivity to light and dark: *on-centre (off-surround)* and *off-centre (on-surround)* [3]. For on-centre cells, light on the centre or darkness on the annulus causes excitation. The opposite pattern of light/dark causes inhibition. For off-centre cells, darkness on the centre or light on the annulus causes excitation. The opposite pattern causes inhibition [3]. These activation patterns are called *receptive fields* (RFs), depicted in Figure 1.3. These types of RFs are essential for contrast, edge and colour detection later in the visual system, as discussed in section 1.1.3.

When the BPCs relay information to the retinal ganglion cells (RGCs), the centre-surround RFs are inherited [3]. The projections of the RGCs comprise the optic nerves and terminate at the LGN structures in the thalamus on the brainstem [3]. Two important types of RGCs are the *parasol* cells, which provide important information for motion and spatial analysis, and *midget* cells, which are necessary for shape processing in later vision [3]. Only at the RGC stage of the transduction process do cells transmit electrical pulses called *action potentials* or *spikes* [3].

1.1.2 The lateral geniculate nucleus

The lateral geniculate nuclei (LGN) are the two nuclei at the end of the optic nerve projections at the base of the thalamus. The left (or right) nucleus receives signals from the retinal ganglia of the right (or left) hemi-field of both eyes [3]. LGN neurons are separated into six layers: the two *magnocellular* (*M*) (called *Y* in non-primates) layers receive input from parasol cells of the retina and make up about 95% of the total cells in the LGN [3]. They respond rapidly to motion, binocular disparity and intensity information from rods [4]. The four *parvocellular* (*P*) (called *X* in non-primates) layers receive input from retinal midget cells and make up about 5% of LGN cells [3]. They respond more slowly to signals from long- and medium-wavelength cones [3]. Between these layers is the

koniocellular (*K*) layer (called *W* in non-primates), which responds to short-wavelength cone signals [7]. The major projections from the LGN are called the *optic radiations* which terminate in V1 [3].

A topic of much current LGN research is motivated by the fact that only 10-15 % of the connection inputs to the LGN are from the retina, while the remaining 85-90% are connections *from* the visual cortex and brainstem [8]. It has been proposed that these cortical inputs may modulate the retinal receptive fields to enhance more relevant areas of the visual field determined by conscious attention [8].

While the spatial response properties of the LGN are known, its functional role with respect to spatial information is not fully understood. This is because each LGN neuron essentially maintains the centre-surround spatial receptive fields of just a few retinal ganglia and preserves relative retinal spatial arrangement (called *retinotopy*) with respect to the contralateral hemi-field [3]. In other words, while LGN appears to function in part as a signal relay between the RGCs and V1, there is no clear transformation performed on the spatial information contained in the input signals from the retina.

Temporally, however, the LGN shows rich dynamics and diversity. Recent studies have shown that more than 50% of spikes emitted by RGCs do *not* lead to a corresponding spiking event in the LGN neurons [9]. Furthermore, the spiking output from the LGN is episodic as it occurs in short bursts of spikes with relatively long interceding periods of inactivity (Figure 1.4) [9].

There are two dynamical types of LGN cells, *lagged* and *nonlagged* (see Figure 1.5) [10]. The lagged neurons show the highest response in the later temporal phase of their time course, while nonlagged neurons peak in the early phase. [6] It has been shown that some LGN outputs to V1 are of both lagged and nonlagged types [11].

Recent work has shown that RGCs have no functional counterpart to the lagged cell behaviour in LGN [12]. Further, these dynamics cannot come from the voltage-to-spike dynamics of RGCs alone and must be a *network* property [10, 13, 14].

1.1.3 Spatiotemporal filtering in V1

The primary visual cortex (V1) is a long-studied region of the visual system in humans, primates and other mammals [15]. Many V1 responses act as spatiotemporal filters, selective for combinations of local contrast, direction, speed, spatial frequency, temporal variation and other visual features [15]. Anatomically, functionally similar assemblies of neurons are organized in columns extending vertically from the surface of the cortex to the white matter [16]. Projections from the magnocellular layer of the LGN terminate in layer 4C α of V1, which in turn projects to layer 4B of V1 [3]. Projections from the parvocellular layer of the LGN terminate in layer 4C β of V1 [3]. Importantly, the geniculocortical connections in V1 are estimated to constitute only 2-28% of V1 afferents, while the majority are intra-striate or *feedback* connections, which are necessarily *non-linear* (c.f. section 1.2.1 below) [17].

In their landmark work of 1959, Hubel and Wiesel described V1 *simple cells* as those neurons in the striate cortex having RFs with distinct, elongated, excitatory and inhibitory regions, sensitive to direction, orientation, speed, size and contrast. These neurons also

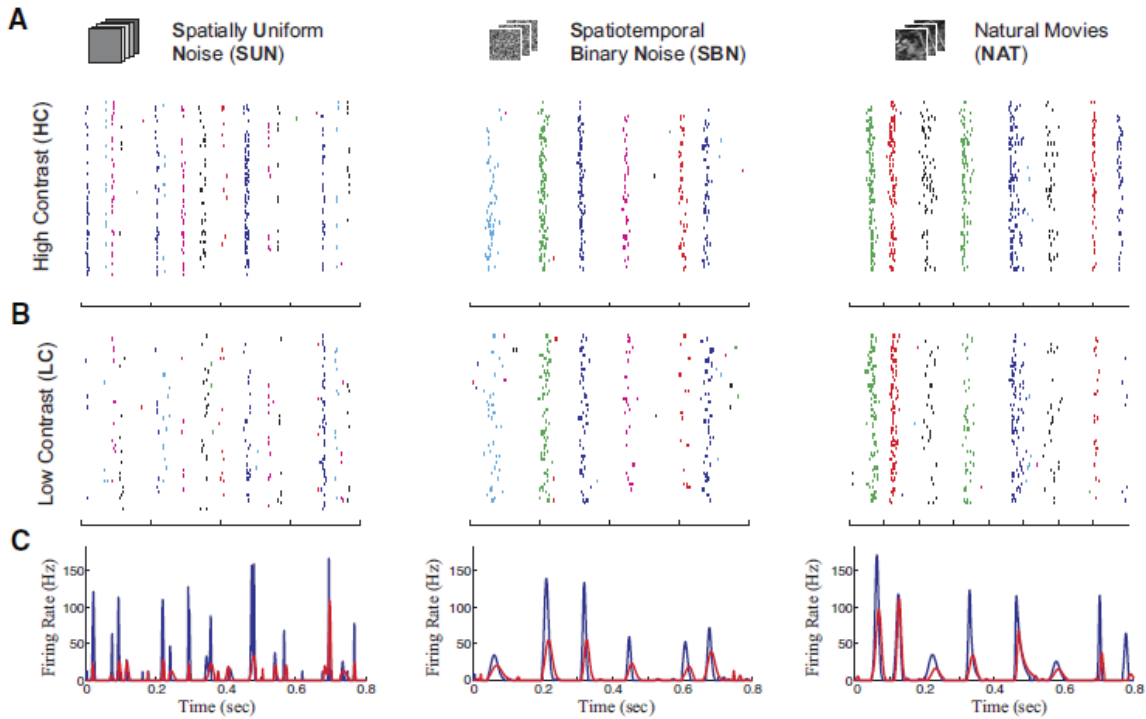


Figure 1.4. Burst signals from LGN. The spike trains from cat LGN X-cell are episodic in response to a variety of time-varying stimuli. Used with permission of Butts [9].

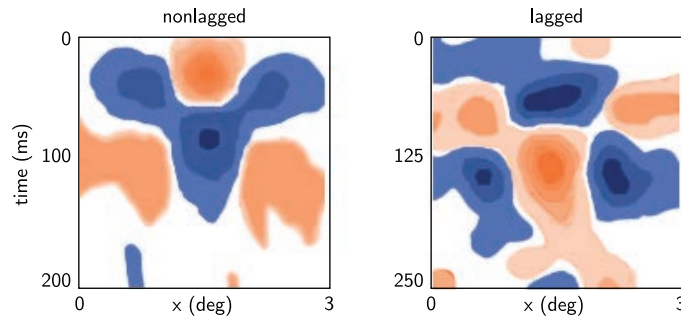


Figure 1.5. Time courses of LGN RFs. One-dimensional (1-D) time courses for nonlagged (left) and lagged (right) LGN neurons in the cat. Each pixel-row is a 1-D slice through the two-dimensional (2-D) RF. Time progresses downward. Adapted, used with permission of DeAngelis [6].

responded in an essentially linear manner to the increase or addition of preferred stimulus in the RF [18]. For this reason they are often modelled as linear filters.

V1 simple cell RFs have diverse spatial structures (Figure 1.6) that change during the response period, enabling time-dependent filtering operations (Figure 1.7) [19]. The selectivity of V1 simple cells to the direction of motion of local intensity gradients (edges) is called *direction selectivity* (DS). It should be stressed that DS is not *velocity* selectivity (VS), which combines DS and *speed* selectivity (SS).

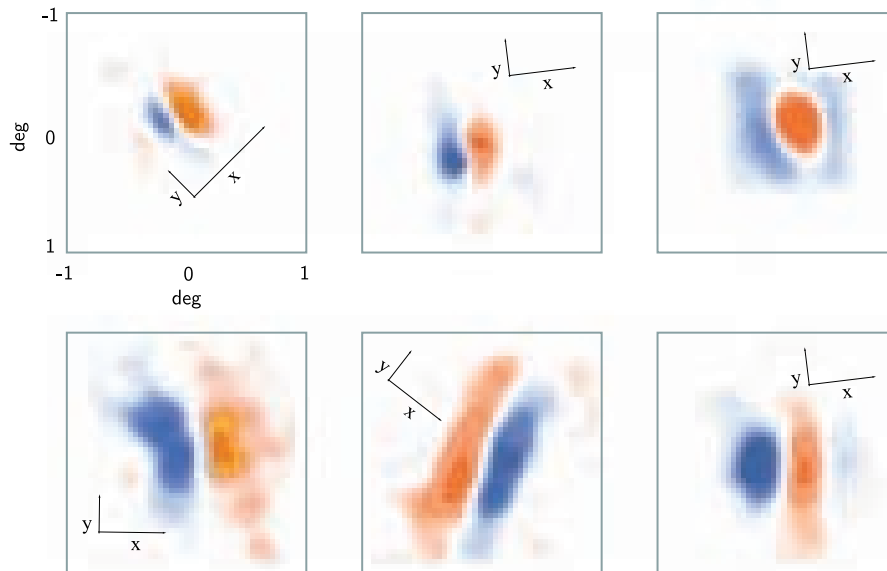


Figure 1.6. Diversity of spatial RF structure in V1 simple cells. A sample of the diversity of two-dimensional spatial RF structure from individual simple cells in macaque V1. Adapted, used with permission of Livingstone [22].

Orientation tuning typically appears after delay of 30-45 ms and persists for 40-85 ms [20]. The structure of V1 RFs may retain on-off patterns of two or three lobes or show more complex behaviour such as sign inversion, drift or secondary peaking (Figure 1.8) [20]. The V1 input layers, $4C\alpha$ and $4C\beta$, which receive input from the M and P layers of LGN, respectively, show a gradual ‘appearance’ and ‘disappearance’ of a constant orientation preference. Neurons in the output layers 2, 3 and 5 show *inseparability* or phase evolution, layer 4B exhibits *separability*, and layer 6 produces a late secondary peak (see Figure 1.8) [20]. While many of these V1 RFs also show a sharpening of orientation tuning not achievable by simple feedforward connectivity [20], the essential behaviour can be modelled by a feedforward process [21]. Orientation tuning has been shown to be contrast invariant in a study of the cat visual cortex (Figure 1.9).

1.1.4 Motion integration and velocity representation in MT

The cortical area MT receives the majority of its inputs from layer 4B (magnocellular) of V1. Like V1, MT shows retinotopic and columnar organization [25]. An important MT study on the macaque by Albright, et al., established that about half of MT neurons respond to moving stimuli regardless of colour, shape, length or orientation. The spatial length of a stimulus affected the response magnitude but not the preferred direction. Cell responses could be divided into those that respond only to one direction of motion (60%), those with a strong preference to one direction but a weaker response in the opposite (24%), orientation selective cells (8%), and those responsive to motion in any direction (8%) [16]. Electrophysiology on behaving monkeys shows that MT neurons are selective for the direction of a moving pattern or component of pattern motion [26], with an average response time (latency) of about 87 ms for a distribution like that shown in

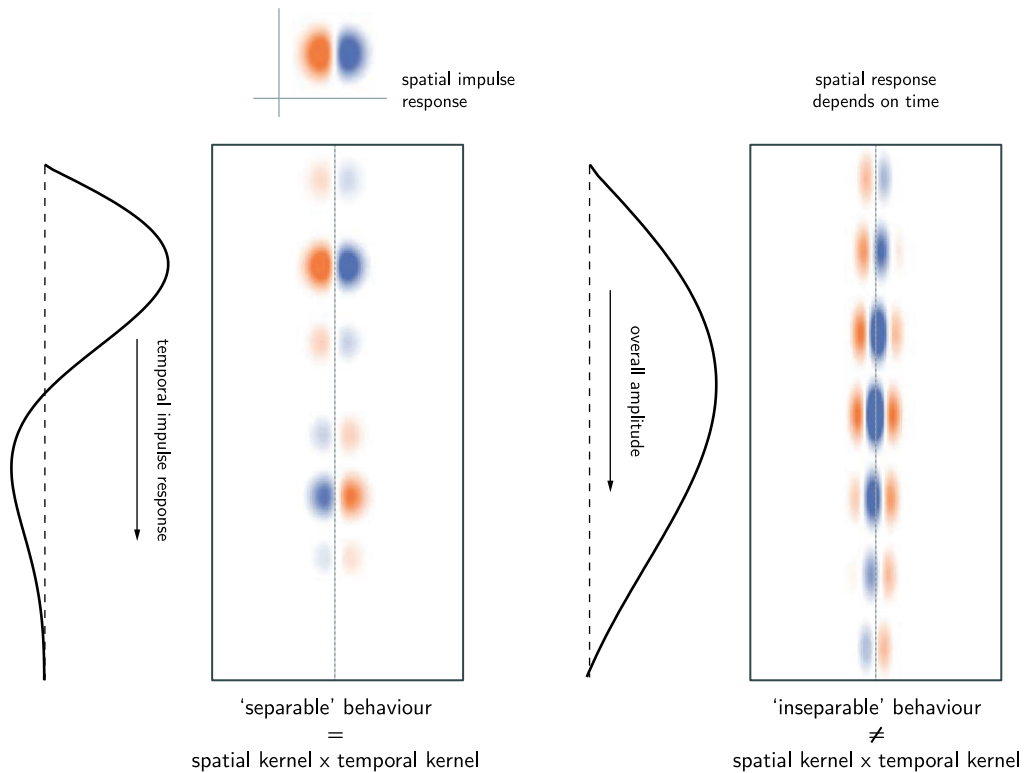


Figure 1.7. The separable and inseparable RF concepts. A idealized depiction of *separable* and *inseparable* spatiotemporal RFs. *Left:* The separable RF with the spatial impulse response along the horizontal and the temporal impulse response along the vertical. Space and time are independent factors of the product (centre). *Right:* The spatial response is a function of time. The spatial response undergoes a phase shift, causing the lobes to translate left. It is assumed in the motion energy model that this type of RF can produce velocity selectivity. Both separable and inseparable types show inversion of amplitude sign late in the time course.

Figure 1.11 [24]. It is believed that MT is involved in the pooling or integration of local motion signals into pattern motions [27], but the precise mechanism of such processing is not fully understood. The functional organization of MT can be depicted as a field or array of velocity sensitive columns, in which each MT receives input from a number of V1 neurons. The RF widths in MT from primate experiments in some studies was shown to be ~ 4 –25 degrees in width, whereas V1 RFs ranged from 0.3–1 degree [28, 29]. A more recent study has shown that the distribution of RF sizes for V1 and MT often overlap [25].

Projections from MT terminate in other vision-related cortical areas that have been implicated in the perception of optic flow, visuospatial decision-making and saccadic eye movements. These include the medial superior temporal area (MST), ventral intraparietal sulcus (VIP), lateral intraparietal sulcus (LIP), the frontal eye fields (FEF), superior colliculus (SC) and dorsolateral pons [25].

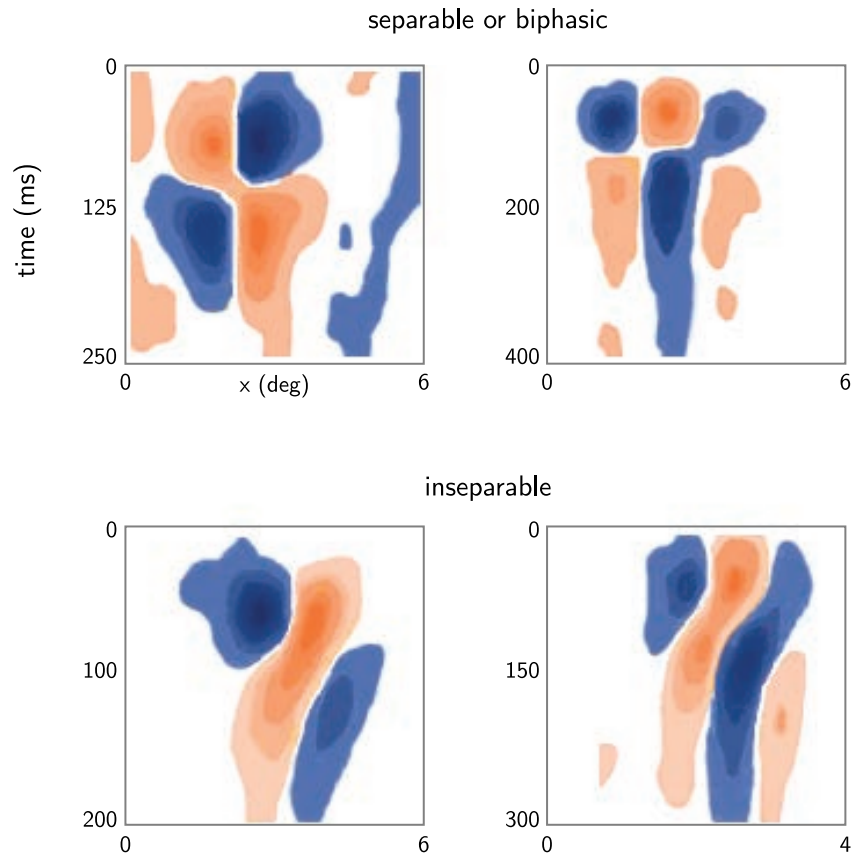


Figure 1.8. Example time courses from V1 simple cells. 1-D time courses ('x-t' plots) from simple cells in cat V1. *Top row:* The separable or biphasic type. *Bottom row:* The inseparable type. Adapted, used with permission of DeAngelis [6].




Figure 1.9. Contrast invariant orientation selectivity in V1. The preferred orientation of this cat V1 cell is invariant under changes of stimulus contrast. Each trace indicates the an averaged tuning curve for a number of neurons selective for 0-degree orientation. The relative contrast of the stimulus is indicated by grey level, where black is strongest contrast. The spiking response is roughly proportional to contrast level. Redrawn from Priebe [21].

1.2 Current theoretical approaches

1.2.1 Linear filtering as a modelling tool

In his seminal papers of 1959 and 1961, H. B. Barlow combined Claude Shannon’s new quantitative theory of information [30] with a previously qualitative notion that a primary function of the sensory nervous system was “recoding to reduce the redundancy of our internal representation of the outer world” [31]. Barlow’s was the *efficient coding hypothesis*, which states that “*sensory relays recode sensory messages so that their redundancy is reduced but comparatively little information is lost*” [31,32]. In the language of modern signal theory, such an operation is called *decorrelation* and, in the linear method, requires an *impulse response function* (or *linear filter*, or *kernel*) to be convolved with an input signal. This idea and its mathematical formulation have become fundamental to the quantitative study of vision, neuroscience and many other disciplines.

In physical terms, a *filter* is a system that responds to a signal in a manner that produces another more useful signal. Filtering is a form of computation since such systems perform calculations on quantifiable signals; they *transform* a given input quantity into an output quantity. *Linearity* refers to the properties of a transformation that acts on the vector quantities x and y , with scalar c , such that the conditions are satisfied:

1. $T(x + y) = T(x) + T(y)$, (superposition) and
2. $T(cx) = cT(x)$ (homogeneity).

Intuitively, these statements mean: 1) combining any two vectors *before* transformation gives the same result as combining them *after* transformation, so there is no interaction

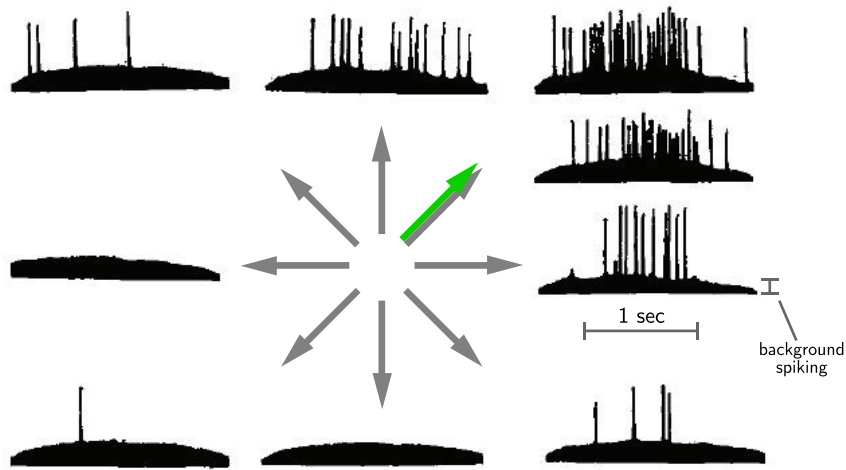


Figure 1.10. Spiking responses from direction selective neurons in macaque MT. The grey arrows indicate the preferred directions of motion for the corresponding MT neurons. The green arrow is the direction of motion of the stimulus, a bright moving bar. Recreated from Dubner and Zeki [23].




Figure 1.11. Distribution of response latencies in macaque MT. The response latency refers to the time for MT neurons to begin responding to directional stimuli. The starred (*) bar indicates the mean latency time, 87 ms. Recreated from Raiguel, et al [24].




Figure 1.12. Convolution, conceptually. An example of the effect of convolution with image intensities (left) and the centre-surround (difference of Gaussians) kernel. The result is a new image consisting primarily of edges.

between any two vectors within the transformational system, and 2) an increase or decrease in the magnitude of the input vector produces a *proportional change* in the output.

In the convolution operation the *impulse response* or *kernel* is denoted $h(\cdot)$. It describes how the system interacts with each unit impulse of input. The linearity conditions mean that the output signal of a linear system is the sum of impulse responses weighted by the input. This linear transformation is called *convolution* and is expressed as the integral

$$y(n) = (x * h)(n) = \int x(n')h(n - n')dn' ,$$

where $x(\cdot)$ is the input signal and the convolution is usually over a domain in time, frequency or space. Thus, for a known input signal, the response of the linear system is completely characterized by the kernel [33]. Determining this function for neurons in the visual system is fundamental to using linear systems as an analytical tool.

In vision research, the kernel is also referred to as the *receptive field map* [34,35], where the signal domain is typically restricted to a small area or ‘patch’ in the visual field. A system of large numbers of overlapping patch functions is a *convolution array*, which is the essential large-scale system description of V1 and other cortical visual areas.

So the primary computational role of the major anatomical structures in the study of vision is that of linear filtering, characterized by the receptive field. These notions will play a central role in the rest of the thesis. A conceptual example is shown in Figure 1.12.

1.2.2 Receptive fields are results of prior computation

It must be stressed that the RF is *not* an intrinsic property of the neuron. Through a complex, time-dependent chemophysical process, neurons release neurotransmitters through spikes in membrane potential as a result of the input of ionic currents that arise from the spike responses of many other preceding neurons (see section 2.1). So the RF is actually the result of *prior computations* that determine the pattern of incoming currents in a given neuron. To understand *how* RFs arise is to understand the computations that determine *why* a neuron is selective in its response.

1.2.3 The Gabor filter for V1 simple cell spatial responses

Gabor functions are used as filter kernels to approximate the spatial responses of V1 simple cells, extracting the locations and directions in space where the greatest changes in intensity (edges) occur. Mathematically, they can be interpreted as a sinusoidal wave

with an exponential decay factor called a *Gaussian window* whose width is σ , that shapes the radial spread of the function. The Gabor is also shaped by the ellipticity, γ , which affects the elongation of the Gabor in the y direction (see Figure 1.13). So for a given orientation θ in the xy -plane, with wavenumber k and phase ϕ , the amplitude at any point in the (x, y) is given by

$$f(x, y) = \exp\left(-\frac{x'^2 + \gamma^2 y'^2}{2\sigma^2}\right) \cos(2\pi kx' + \phi), \quad (1.1)$$

where $x' = x \cos \theta + y \sin \theta$ and $y' = -x \sin \theta + y \cos \theta$. Changing the parameters affects the function as shown in Figure 1.13. The physiological correlates of these parameters, especially the phase angle ϕ , serve an important role in the model of velocity computation proposed in this paper.

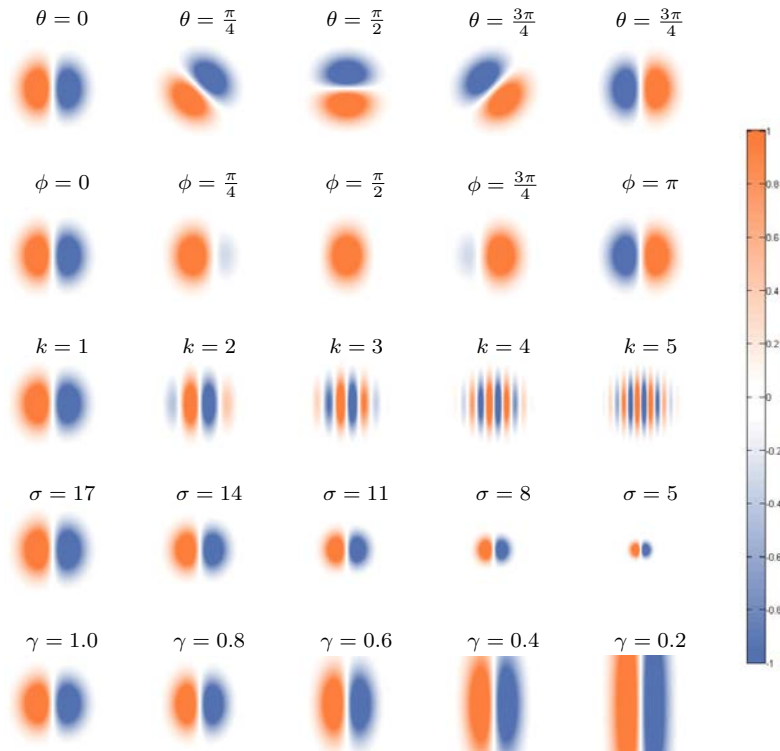


Figure 1.13. Effect of variation of Gabor parameters. V1 simple cell spatial receptive fields approximated by Gabor functions. The effects of changing one parameter while holding others constant: θ , orientation; ϕ , phase angle; k , wavenumber; σ , spread; γ , ellipticity. Values for k , σ and ϕ should be considered relative since the patch area is arbitrary. The Gabor functions used in the proposed model are biologically realistic, using fixed values for γ , λ and σ to give high responses to low spatial frequencies within the receptive field at each resolution.

While the Gabor function is a very common modelling tool in describing V1 RFs, a more general fit has recently been shown by Young, et al, using a Gaussian derivative method that accommodates three dimensional spatiotemporal RFs (2-D space and 1-D time) and approximates a wide variety of primate V1 RFs [36].




Figure 1.14. The predicted temporal (left) and spatial (right) impulse responses for an LGN X cell. Adapted, used with permission of Dong [10].

1.2.4 Models of temporal processing in LGN

Temporal decorrelation

Dong and Atick proposed that a fundamental role of the LGN is the *temporal decorrelation* of the intensity signals passed from the retina. Just as spatial redundancies are removed by efficient coding with on-off RFs in the retina, temporal redundancies may be removed by the LGN by its lagged and nonlagged RFs [10].¹ Starting from the assumption of linear filter behaviour and the efficient coding hypothesis, Dong and Atick used natural image statistics under noise to derive the spatiotemporal kernels that would most efficiently encode a time-dependent intensity signal. The predicted spatial and temporal profiles are shown in Figure 1.14, in excellent agreement with experimental data [10].

The predicted temporal kernel is of particular importance to the proposed model of motion processing as it exhibits the characteristics of a *damped oscillator* (see Figure 2.8). Such dynamics can be generated by a recurrently connected network of V1 simple cells and LGN (see Section 3.2.5). Thus, the proposed model can help explain the predicted and observed LGN dynamics.

The Dong-Atick model does provide a theory of the dynamical responses of LGN based on efficient coding, making a strong argument for the computational role of the structure based on biological evidence and natural image statistics. However, it does not purport to be a network model or even a neural one. Although it provides reasons for LGN behaviour in a larger motion processing system, it does not offer an explanation of *how* the observed RFs arise in LGN.

¹This suggests a distinct specialization in spatial and temporal filtering roles between the retina and the LGN, respectively. This is not strictly true since RGCs show some temporal decorrelation along with spatial responses, but it is a good approximation since the temporal bandpass response of RGCs is essentially flat, and the same holds for the spatial bandpass behaviour of the LGN [10].

1.2.5 Models of V1 direction and speed selectivity

Feedforward models and extensions

Feedforward models are computational algorithms without feedback or recurrent connections. They are typically simpler in construction than feedback models, using a forward-only sequence of linear and/or nonlinear transformations.

A classic and influential feedforward model is the *motion energy* hypothesis of Adelson and Bergen [19]. They cast the problem of velocity computation as one of detection of an orientation vector in space-time. They used a measure called motion energy, a positive magnitude resulting from the convolution of the local intensity signals of a moving stimulus and the impulse response of a spatiotemporally-tuned simple cell in V1 [19].

The first type of linear filter they proposed, called a *separable* impulse response, is the product of independent spatial and temporal impulse responses. A second, *inseparable* response has a time-dependent spatial component. These two types are shown in Figure 1.8.² Since the separable type can select for orientation but not direction of motion, they theorized that velocity selectivity in V1 simple cells would likely come from inseparable-type responses.

To eliminate the problem of negative or null responses from a stimulus feature that trailed a drifting grating by a phase angle of 90 degrees, Adelson and Bergen proposed that the V1 responses were associated in *quadrature pairs* [19], the responses of which were rectified by squaring then summation to give a total motion energy. The model is illustrated in Figure 1.15.

The motion energy model produces a response *distribution* from the collection of RF responses. However, if the stimulus is one-dimensional in structure when viewed through the RF apertures, the extraction of object velocity from the population of V1 responses may not be possible since it requires reconstructing a two-dimensional quantity (velocity components V_x and V_y) from one-dimensional information. This is called the *aperture problem*, and is illustrated along with its geometric solution in Figure 1.16. It is a long-standing problem in vision research. A realistic interpretation of this *decoding* problem and solution is presented in section 2.2.2.

As an extension to this and other feedforward models, Wimbauer, et al, describe the formation of DS RFs as a process whereby the spatiotemporal correlations among LGN afferents (permutations of the on, off, lagged and nonlagged types) are learned using a Hebbian rule [37]. Shon, et al, have also recently specified a learning mechanism for DS in V1 through spike-timing dependent plasticity [38]. A recent review by Priebe and Ferster show that contrast invariance and orientation selectivity require feedforward mechanisms only [21].

Weaknesses of feedforward models

An immediate problem with feedforward approaches is the lack of anatomical justification. As mentioned in section 1.1.3, the majority of connections in V1 is intracortical, suggesting a strong role for non-linearity in the spatiotemporal processing of early visual signals.

²Inseparable spatiotemporal RFs are typically depicted as sinusoidal gratings translating in one direction.




Figure 1.15. The motion energy model of Adelson and Bergen. As described in [19]. The quadrature pair of RFs consists of two phase-progressing Gabor functions, differing only by a 90-degree phase shift, that act as convolution kernels. One dimension of the 2-D filter is shown along each horizontal axis, with time along the vertical. The output of each convolution operation is squared to eliminate negative values. The sum of the two is termed the *motion energy*.

Another major weakness of feedforward schemes is the disparity with electrophysiological data. A feedforward model requires sustained, high input magnitudes from the LGN to drive DS in the preferred direction in the cortex [17, 39]. Also, suppression of neural activity in the null direction requires a high level of inhibitory conductance. Evidence of these physiological requirements has not been observed. In fact, a significant absence of large inhibitory conductance changes has been measured [17, 40, 41].

Crucially, feedforward models offer no account of the *dynamics* that must play a role in DS or VS; that is, they do not explain how the observed dynamic characteristics of V1 RFs arise.

Feedback or recurrence models and extensions

To address the lack of intracortical mechanisms in feedforward models and explain observed changes in neural conductances, Douglas and Martin devised a feedback or recurrently connected microcircuit, in which weak input from LGN is amplified in the preferred direction and suppressed in the null direction by combination of excitatory and inhibitory connections within V1 [39]. Further refinements from biologically plausible computer simulation of this micro-circuit were performed by Suarez, et al [17].

Another biophysically detailed, system-theoretic implementation of a DS network and further elaboration of the feedback microcircuit has been provided by Maex and Orban [42]. A central claim of this work is that the interaction of excitatory and inhibitory feedback connections in V1 provide the mechanism for spatiotemporal correlation (filtering) necessary for DS. In the model, DS (not SS or VS) is derived from excitation and subsequent inhibition by spatially adjacent bipolar Gabor-like RFs in the cortex. Essen-

tially, local excitation of an edge-selective RF inhibits its neighbours in the preferred and null directions. Dynamical properties of RFs are implicit in the model, determined by spacing of the on/off regions of the RF and temporal lags in recurrent connections.

Mineiro and Zipsper have attempted to analytically simplify the Maex-Orban model, and in doing so have shown prediction characteristics in recurrent networks that sustain signal representations under noisy visual conditions [43]. Further, they demonstrate that the integration time constant, τ_{RC} , and not the propagation delay of a recurrent signal, is the most appropriate mechanism for speed preference in single neuron models [43].

Weaknesses of feedback models

These feedback models suffer from several common weaknesses. First, they do not capture the observed characteristics of neural selectivity for both direction and speed, i.e. *velocity* tuning, in which velocity is vector quantity of at least two dimensions. Second, the proposed mechanisms do not explain or predict the known time courses (e.g. separable, inseparable) in V1 simple cells. Finally, they are designed using only a few neurons, lacking any means for distributed representation or large scale simulation.

1.2.6 Models of motion processing in MT

Divisive normalization

An extension of the motion energy model has been proposed by Simoncelli and Heeger in an attempt to resolve the aperture problem. They use a “neural implementation” of the geometric solution, called the *intersection of constraints* [44], which is the main thrust of the divisive normalization model. It is an extension of the motion energy model that uses half-wave rectification and squaring of individual neural firing rates to determine the magnitude of the preferred direction represented by each neuron in a pool. These decoded vectors are then scaled by a proportionality constant, common to all neurons in the pool to recover the original stimulus projection (see Figure 1.17) [44]. This algorithm is used for converting the population firing rates of V1 and MT neurons into coefficients of preferred velocity vectors. However, as with the motion energy model, divisive normalization provides not account of dynamics required to generate the required DS RFs.

The cascade extension

The ‘cascade’ model described by Rust, et al, extends the divisive normalization approach [26,45]. Its primary aim is to show how a composite stimulus, composed of a combination of directional components, drives component-tuned V1 DS cell responses that are pooled as a weighted linear sum in an MT neuron. Essentially, the normalization constant from the Simoncelli-Heeger method becomes a directionally-dependent variable in the cascade model, determined by putative suppression within and between V1 RFs [26].

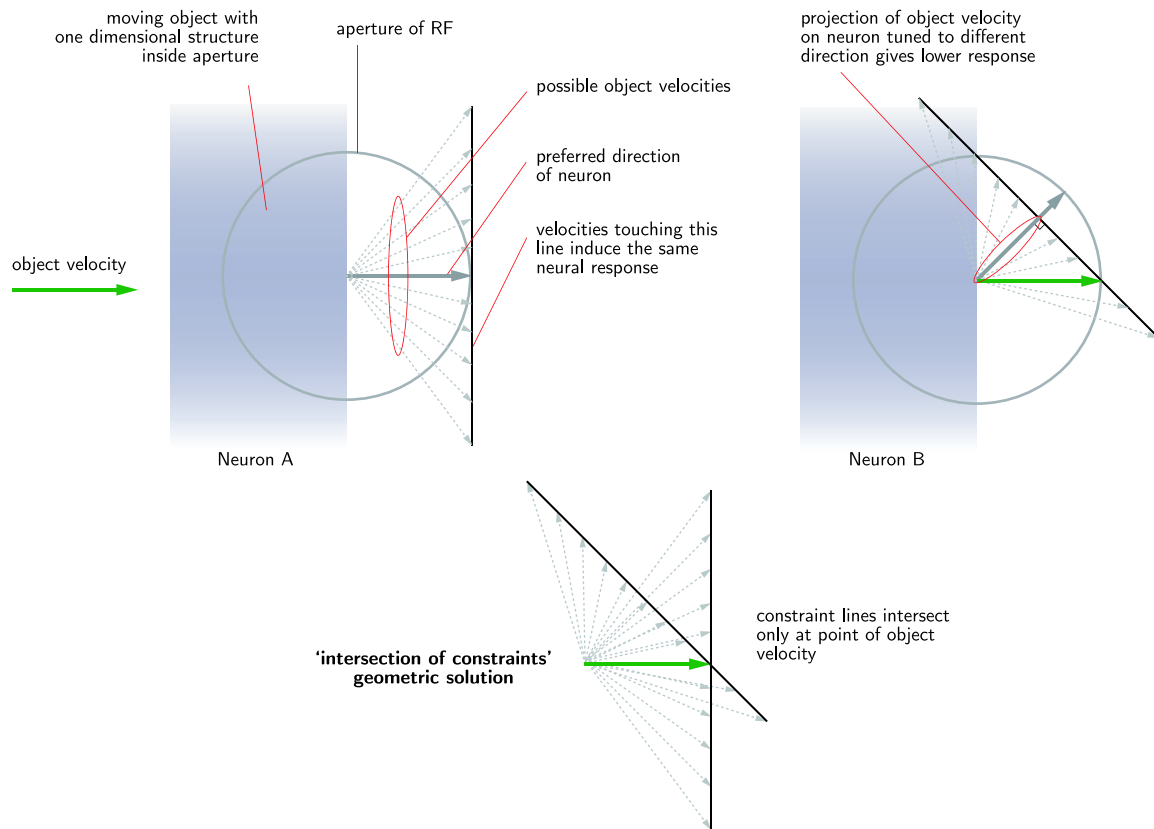


Figure 1.16. The aperture problem and the intersection-of-constraints solution. The velocity of a 1-D structure transiting a local aperture is ambiguous since many possible motions produce the same neural response. *Top left:* A neuron tuned to the object velocity will respond maximally. *Top right:* The same stimulus on a neuron tuned to a different direction produces a lower response. A population of variously tuned neurons would produce a Gaussian-like distribution of responses. *Bottom:* To produce a single vector from the distribution, the ‘intersection of constraints’ construction finds the true vector that intersects all constraint lines [25].

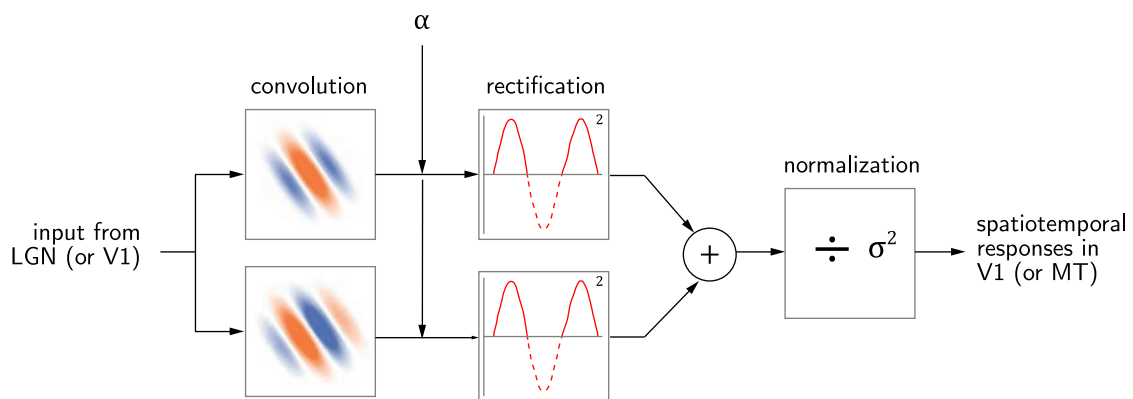


Figure 1.17. The divisive normalization model by Simoncelli and Heeger. Continuous input is convolved with local linear spatiotemporal filters, followed by rectification and summation. The α term is a common background spike rate. A normalization constant, σ^2 , is calculated for the pool of afferents and used to scale the output value of each neuron. The algorithm is the same for LGN or V1 input.

Three-dimensional naturalistic RF extension

In attempt to generalize and simplify the prevailing models of Adelson-Bergen, Simoncelli-Heeger and Rust [19, 26, 44], Nishimoto and Gallant have used naturalistic stimuli on behaving macaques to derive a more accurate MT RF fit in a three dimensional spatiotemporal frequency (STF) domain [46]. One important result of the naturalistic MT RFs is that excitatory fields tend to lie in a common STF plane while inhibitory fields lie off-plane. The studies also showed that MT responses to naturalistic stimuli can be suitably approximated based on MT responses from idealized stimuli such as moving bars, gratings and dots [46].

1.3 Relationship of the proposed model to previous work

The model of motion processing proposed in this thesis shares some similarities to those outlined above. It resembles feedforward models of DS in its use of LGN afferents to drive V1 simple cell activities. It is similar to recurrence models in the use of feedback connectivity in V1 to amplify and sustain thalamic output signals. Like the motion energy scheme, it relies on a measure of frequency similarity between the spatiotemporal structures of the stimulus and V1 simple cell RFs. It is similar to the divisive normalization model of motion integration in MT since it uses weighted pooling of velocity tuned V1 efferents in the construction a retinotopic map of optical flow in MT.

Unlike other models, the proposed model explains *how* the spatiotemporal characteristics of velocity selective RFs (namely, the transient and sustained components of separable/inseparable RFs observed in macaque electrophysiology studies [22, 47]) are generated within V1 simple cell ensembles, and how those populations interact with efferent thalamic signals to produce velocity selectivity in V1. The mechanism for velocity selection is a driven and damped *state phase oscillator*: a spiking neural ensemble in which a time-dependent Gabor function—a vector encoded by the activities of neurons that share a common orientation preference but collectively span the phase angles around a circle—undergoes phase evolution by the rotation of the vector in a two-dimensional space in which the plane axes are any two orthogonal phases. Thus, the rotation of the oscillator state vector in time corresponds to the progression of the phase angle in time. Intuitively, the rotation of the state vector represents the translational ‘motion’ of the Gabor. If the LGN output signal is similar to the intrinsic state of the oscillator at a given time after the onset of oscillation, the interference will be constructive, increasing the state magnitude; if dissimilar, destructive interference will act to decrease the state magnitude. Sensitivity to static signals is eliminated by using V1 efferents that correspond to phase angles late in the oscillation period.

Another important distinction of proposed model is in the biological constraints of the implementation. It uses spiking, noisy, leaky integrate-and-fire (LIF; see section 2.1) neurons in a large scale network with a minimum of 10,000 for each V1-MT microcircuit, and up to 1.29 million neurons when processing the entire stimulus area. It incorporates 49 microcircuits, retinal ganglial output and representational LGN ensembles. Further, the proposal specifies all neural membrane time constants, absolute refractory periods, post-synaptic time constants and distribution of firing rates using the methodology devised by

Eliasmith and Anderson [48]. Finally, the model utilizes and demonstrates the necessity of burst-like signals from LGN to V1.

Chapter 2

Theoretical principles of the proposed model

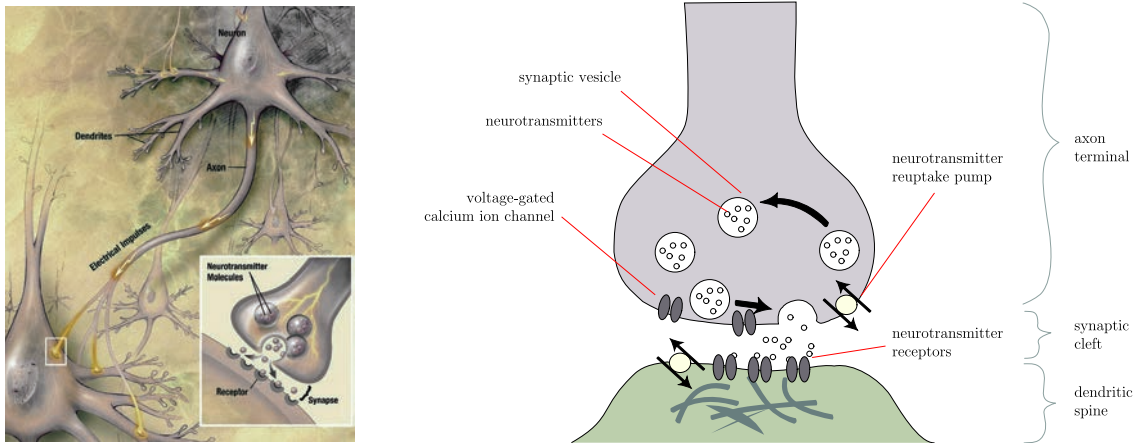


Figure 2.1. The neural connection. *Left:* The essential communication structure between two neurons illustrating a transmitting axon and the many signal-receiving dendrites. Public domain image created by US National Institutes of Health, National Institute on Aging. *Right:* Schematic depiction of the event sequence at the synapse, relevant to signal communication.

Neurons are electrically excitable cells in the nervous systems of vertebrates. In vast networks, they allow animals to sense, represent, process and interact with the environment in order to survive. The left of Figure 2.1 depicts the transmission of an electrical pulse one neuron to another. This pulse, called a *spike* or *action potential* [49], is transmitted along the axon of a neuron in a few milliseconds [50]. The spike is transduced into a chemical signal by a complex process (Figure 2.1, right) that causes an ionic current to flow into a receiving neuron via projections called *dendrites* [50, 51]. The interface between axon and dendrite is called the *synapse* [3]. In the construction of a biologically plausible neural network that can compute the velocity of a visual pattern, we must begin with a quantitative model of the neuron. In this chapter, the single unit approximation of a spiking neuron called the *leaky integrate-and-fire* (LIF) model is presented, followed by a discussion of its suitability in large-scale network simulations.

Next, a general methodology called the *Neural Engineering Framework* (NEF) [48] is

presented to show how vector quantities can be represented and transformed in realistic, time-dependent, noisy networks of neurons. Network dynamics using the state-space method are included in the NEF formalism. An application of the methodology is applied to the damped oscillator, which is an important motivating example for the model proposed in Chapter 3.

2.1 The LIF neuron model




Figure 2.2. An RC circuit. A parallel resistor-capacitor circuit is the electrical analog of the subthreshold range of the LIF neuron. Public domain image.




Figure 2.3. The voltage time course of a LIF neuron. V_{th} : threshold voltage; τ_{RC} : membrane time constant; τ_{ref} : absolute refractory time period. See text for full description. Adapted, used with permission of Eliasmith [48].

The LIF neuron model is a simplification of the successful but more complex Hodgkin-Huxley model [52]. In the subthreshold range the LIF neuron is equivalent to a resistor-capacitor (RC) circuit (or *leaky integrator*) (Figure 2.2) driven by an input current [52]. The change in the quantity of charge $Q(t)$ over time is called the membrane current, $I_M(t)$, where $I_M(t) = \frac{dQ}{dt}$. The membrane capacitance, C_M , results from the differences in ionic concentrations of across the membrane. These terms can be related to the membrane

voltage, $V_M(t)$, by the definition of capacitance:

$$C_M = \frac{Q(t)}{V_M(t)}$$

or $Q(t) = C_M V_M(t)$.

By differentiating with respect to time, we can relate I_M to these quantities:

$$I_M(t) = \frac{dQ}{dt} = C_M \frac{dV_M}{dt}. \quad (2.1)$$

The cell membrane is not a perfect insulator since it leaks ionic current. This gives it resistor properties as well.¹ So a ‘leak’ term is added to (2.1) as the membrane resistance, R_M :

$$I_M(t) - \frac{V_M(t)}{R_M} = C_M \frac{dV_M}{dt}. \quad (2.2)$$

We now drop the subscripts for clarity. Next, we define $\tau_{RC} \equiv RC$, the time required to charge the capacitor to $IR(1 - \frac{1}{e})$ assuming a step input of magnitude 1. Solving (2.2) for $V(t)$ under constant current I gives

$$V(t) = IR \left(1 - e^{-t/\tau_{RC}}\right), \quad (2.3)$$

and is valid for times $0 < t < t_s$, where t_s is the spike time. This is called the sub-threshold regime (Figure 2.3).

At the threshold voltage, V_{th} , the neuron fires. In the LIF approximation, the spiking event is considered stereotypical and is formalized by adding on a delta function at the threshold time (Figure 2.3).

The time of the spike is given by $V_{th} = V(t_s)$, which we can use to solve for t_s in (2.3):

$$t_s = -\tau_{RC} \ln \left(1 - \frac{V_{th}}{IR}\right). \quad (2.4)$$

Note that t_s is just the period of one spike ‘cycle’ assuming a constant current. So, just as the frequency f of a steady-state periodic system with period T is $f = \frac{1}{T}$, the steady-state firing rate of the LIF neuron can be found by the expression

$$a(t_s) = \frac{1}{t_s + \tau_{ref}}, \quad (2.5)$$

where τ_{ref} is the absolute refractory period, a duration in which no current can enter the cell [53]. Substituting (2.4) into (2.5) and using Ohm’s law, $V_{th} = I_{th}R$, gives the steady-state firing rate or *activity* of the neuron under constant current:

$$\begin{aligned} a_i(I) &= \left[\tau_{ref} - \tau_{RC} \ln \left(1 - \frac{V_{th}}{IR}\right) \right]^{-1} \\ &= \left[\tau_{ref} - \tau_{RC} \ln \left(1 - \frac{I_{th}}{I}\right) \right]^{-1}. \end{aligned} \quad (2.6)$$

¹This also allows the neuron to discharge current relatively slowly. Otherwise, it would hold or ‘remember’ past input indefinitely, which is highly unrealistic.

If $I = I(t)$, we can interpret $a_i(I(t))$ as the ‘instantaneous’ firing rate of the neuron. Single neuron models that express activities by rates, as opposed to spiking events, are called *rate neurons*.

2.1.1 Suitability of the LIF neuron model

There is an ample body of literature dealing with the shortcomings of the LIF model as an appropriate simplification of a real spiking neuron [48,54–57]. The main concerns include: the LIF neuron has no volume in space and thus does not account for the variation in potential over the entire cell membrane; ion conductance by the membrane is not taken into account; the leak term in (2.2) is derived from Ohm’s law, where voltage is a linear function of current, yet ionic currents are generally nonlinear functions of membrane potential and time [48,55].

However, a more realistic approximation like the Hodgkin-Huxley model [58] would improve spiking statistics but would not be computationally feasible in a large-scale network simulation [52]. Eliasmith and Anderson have argued that the LIF approximation represents the best trade-off between computability and realism [48].

The importance of a reasonably accurate analytical description of a single neuron is that it allows us to relate an input current to a stimulus signal and the preferred signal of a neuron in order to derive its signal response or *tuning curve*.

In the next section we consider how to connect individual LIF neurons in order to perform information processing in a dynamic network.

2.2 Neural network modelling principles

We approach the task of network modelling using a general neural system design methodology called the *Neural Engineering Framework* (NEF), developed by Chris Eliasmith and Charles H. Anderson [48]. The NEF is a general, biologically constrained, scalable approach that consolidates and extends a variety of classical and modern methods. In particular, the NEF allows us to *solve for* the neural connection weights necessary to represent and transform vector quantities with spiking neurons in noisy, time-dependent networks of arbitrary function.

The central principles of the NEF are:

1. Neural representations are defined by the combination of nonlinear encoding and weighted linear decoding.
2. Transformations of neural representations are functions of variables that are represented by neural populations. Transformations are determined using an alternately weighted linear decoding (transformational decoding).
3. Neural state dynamics are characterized by considering neural representations as state variables in linear time-invariant systems.
4. Neural systems are subject to significant amounts of noise, which must be accounted for in system design and analysis.

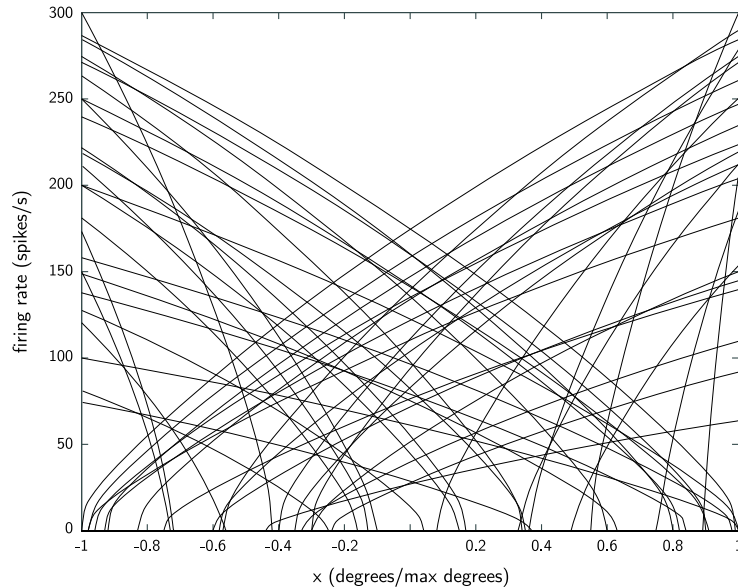


Figure 2.4. Sample tuning curves for a scalar value. The tuning curves of a neural ensemble encoding a scalar value between -1 and 1 with only positive-valued rates. For each input value in the domain, the set of firing rates can be decoded using the principles described to reconstruct the input. Adapted, used with permission of Eliasmith [48].

2.2.1 Vector encoding

The neural ensemble

In many neurophysiological studies of behaving animals, the cortical response induced by a stimulus is distributed over a population of neurons. These functionally related neurons are called an *ensemble* [59]. Each neuron in the ensemble can be characterized as being tuned to a particular stimulus vector, $\mathbf{e}_i = (e_1, e_2, \dots, e_n)^T$, which generates the highest firing rate in that neuron. This means that it receives the highest ionic current, I , when the stimulus $\mathbf{x} = (x_1, x_2, \dots, x_n)^T = \mathbf{e}_i$. We can relate I to \mathbf{x} and \mathbf{e}_i by the inner (dot) product $\langle \mathbf{e}_i, \mathbf{x} \rangle_n$, plus a constant background current, J_i^{bias} . So our rate equation (2.6) becomes

$$\begin{aligned} a_i(\mathbf{x}) &= G_i \left[I(\mathbf{x}) \right] \\ &= G_i \left[\alpha_i \langle \mathbf{e}_i, \mathbf{x} \rangle_n + J_i^{bias} \right], \end{aligned} \quad (2.7)$$

where a_i is the firing rate of neuron i , $G_i[\cdot]$ is the nonlinear function specific to the single neuron model being used (in our case, the LIF neuron), and α_i is the unit conversion and scaling factor. In this way, we can determine the population encoding of \mathbf{x} for any neuron defined by $G_i[\cdot]$, so long as we can specify the encoders and the other parameters in the ensemble. An example of ensemble tuning curves for the scalar range $[-1, 1]$ is shown in Figure 2.4.

2.2.2 Vector decoding: principles of optimal linear estimation

In order to determine how *well* the encoding scheme performs, we require a method of *decoding* the quantities being represented by the neural activities. This is the subject of the remainder of section 2.

In linear algebra, vectors are represented by a unique linear combination of the elements of a *basis*. This combination is determined by the projections (dot products) of the given vector on each basis element. So for vector \mathbf{x} and an n -dimensional orthonormal basis $\{\mathbf{b}_i\}$,

$$\mathbf{x} = \sum_i^n (\mathbf{b}_i \cdot \mathbf{x}) \mathbf{b}_i \quad (2.8)$$

$$= \sum_i^n a_i \mathbf{b}_i. \quad (2.9)$$

If the basis is known, we can write \mathbf{x} more compactly using the set of projection coefficients so that $\mathbf{x} = (a_1, a_2, \dots, a_n)$. In this form, the coefficients *encode* the vector \mathbf{x} . To *decode* or reconstruct the vector in a linear manner, we use the a_i to weight the contribution of each \mathbf{b}_i in the sum over the basis elements, as in (2.9). This system of encoding and decoding defines a *representation*. It is particularly simple since the basis is orthonormal, i.e. consisting of orthogonal unit vectors. By definition, the number of vectors in the basis is the *dimensionality* of the vector space.

In an attempt to use neural activities as a basis, we begin by considering the \mathbf{e}_i as the basis vectors and interpret the firing rates a_i as the coefficients that encode \mathbf{x} . Our temptation, then, is to decode \mathbf{x} as above, where for an ensemble of N neurons

$$\mathbf{x} \stackrel{?}{=} \sum_i^N (\mathbf{e}_i \cdot \mathbf{x}) \mathbf{e}_i \quad (2.10)$$

$$\stackrel{?}{=} \sum_i^N a_i \mathbf{e}_i. \quad (2.11)$$

Unfortunately, this is incorrect. The first problem is that $a_i \geq 0$ always, since firing rates are never negative. But even if we paired each $a_i \mathbf{e}_i$ with some $a_j \mathbf{e}_j = -a_i \mathbf{e}_i$ to resolve negative projections on \mathbf{e}_i in some way, nothing would restrict the number of neural basis elements to equal the number of dimensions in the vector space. Assuming a high degree of representational redundancy, we will have many more neurons in the ensemble than dimensions in the vector space. Thus, we cannot construct a basis by interpreting neural activities in this way; that is, we cannot use the encoders as decoders with any reasonable representational accuracy.

Instead, we can consider the neural activities of an ensemble as an *overcomplete frame*—a mathematical structure well-known in signal engineering and functional analysis. The benefit of overcomplete frames is their representational robustness. Indeed, they perform *better* than orthonormal bases under noise, and if an element is removed, the redundancy of the structure means the encoded vector is not severely degraded as would be the case in an orthonormal basis [48].

How well the neural ensemble can represent any vector within a given vector space depends on how we account for noise, the number and distribution of our encoders, and the impact of linear decoding of vectors encoded in a highly nonlinear manner. The effect of noise is included in section 2.2.3. The effects of the number of ensemble neurons, N , and linear decoding error are addressed in section 2.2.4 but for now we assume N to be suitably high to give a reasonable estimate of \mathbf{x} , denoted $\hat{\mathbf{x}}$.

Following closely the steps provided by Eliasmith and Anderson in [48], to find the decoding vectors in an overcomplete frame we solve for the set of vectors $\{\mathbf{d}_i\}$ such that

$$\hat{\mathbf{x}} = \sum_i^N a_i(\mathbf{x})\mathbf{d}_i, \quad (2.12)$$

where $\hat{\mathbf{x}}$ is of suitably high accuracy. The NEF uses a method of optimal linear estimation using mean square error minimization to solve for the decoders. For neuron j , the decoder \mathbf{d}_j is found by minimizing the the mean square error over a large and diverse sampling of \mathbf{x} , using

$$\begin{aligned} E &= \frac{1}{2} \int [\mathbf{x} - \hat{\mathbf{x}}]^2 dV \\ &= \frac{1}{2} \int \left[\mathbf{x} - \sum_i^N a_i(\mathbf{x})\mathbf{d}_i \right]^2 dV, \end{aligned} \quad (2.13)$$

where dV is the volume element $dx_1 dx_2 \dots dx_n$. To find the minimum, we take the derivative of (2.13) with respect to \mathbf{d}_i , set it to 0 and solve for \mathbf{d}_j :

$$\begin{aligned} 0 &= \frac{d}{d\mathbf{d}_i} \left(\frac{1}{2} \int [\mathbf{x} - \sum_j a_j(\mathbf{x})\mathbf{d}_j]^2 dV \right) \\ \int a_i(\mathbf{x})\mathbf{x} dV &= \sum_j \left(\int a_i(\mathbf{x})a_j(\mathbf{x})dV \right) \mathbf{d}_j. \end{aligned}$$

Writing the left side as vector \mathbf{v} and the right side integral as matrix $\mathbf{\Gamma}$ isolates the decoder we seek:

$$\begin{aligned} \mathbf{v} &= \mathbf{\Gamma}\mathbf{d}_j \\ \text{or} \quad \mathbf{d}_j &= \mathbf{\Gamma}^{-1}\mathbf{v} \end{aligned}$$

where

$$\begin{aligned} \mathbf{v}_i &= \langle \mathbf{x} a_i(\mathbf{x}) \rangle_{\mathbf{x}} \\ \text{and} \quad \mathbf{\Gamma}_{ij} &= \langle a_i(\mathbf{x})a_j(\mathbf{x}) \rangle_{\mathbf{x}}. \end{aligned} \quad (2.14)$$

This technique, along with the encoding scheme from (2.7), provide a method of representing n -dimensional vectors with ensembles of N rate neurons with minimal error.

2.2.3 Vector decoding: noise and variability

Characterizing spike rate variability

Until this point we have assumed that the firing rate of a neuron, $a_i(\mathbf{x})$, is constant if the exciting stimulus, \mathbf{x} , is constant. However, *in vivo* measurements of firing rates of neurons under constant stimulus show a high degree of variability in the inter-spike interval (ISI) times. There are two general views about this variability. The first assumes this is noise, a result of stochastic processes, including the uncertainty of vesicle release in the pre-synaptic neuron; inconsistent timing in the release of neurotransmitters that open ion gated channels in the post-synaptic neuron [48, 60]; variation in the quantity of neurotransmitter components contained in the vesicles [61]; and possible effects of large numbers of interconnected neurons on the ISI for individual neurons [62]. The second view assumes the variability to be the result of precise timing mechanisms from presynaptic processes [62], and thus a carrier of information. The NEF takes the first view, considering the variability in ISI to be the result of random processes, and where the rate of a neuron can depend on $\hat{\mathbf{x}}(t)$ but is independent of other spike times.

To account for the observed variability of real neurons in our encoding and decoding scheme, we must first arrive at a reasonable characterization of it. In their work on the visual cortex of anesthetized cats, Holt, et al, showed that the spike statistics for *in vivo* neurons in Brodmann areas 17, 18 and 19 (areas V1-V5) can be adequately modelled as a Poisson process [63]. Model fitting was applied to experimental data for neurons with an ISI < 50 ms (i.e. spike rates > 20 Hz), with realistic refractory periods, using stimuli duration up to five seconds. The work also showed that *in vitro* spiking under injected current was highly regular, suggesting the major cause of observed variability in behaving cats was due to high background noise and not intracellular triggering mechanisms [63].

The Poisson distribution for many discrete, independent events that combine additively will tend toward a normal (Gaussian) distribution.

Incorporating variability and noise

The NEF assumes that ISI variability, regardless of cause, can be incorporated as a single noise term in the decoding equation (2.12) [48]. A random value per neuron, η_i , is chosen from a Gaussian distribution with a mean of zero and added to signal in to affect the firing rate, a_i :

$$\hat{\mathbf{x}}(t) = \sum_i a_i(\mathbf{x}(t) + \eta_i) \mathbf{d}_i. \quad (2.15)$$

Solving for the decoders is done as before, but to the previous result of (2.14) a noise term is produced:

$$\mathbf{\Gamma}_{ij} = \langle a_i(x) a_j(x) \rangle_x + \sigma^2 \delta_{ij}. \quad (2.16)$$

In this way we can add an appropriate degree of spike variability, assumed to be noise, and thus representational uncertainty to neural networks built on NEF principles. Section 2.2.4 describes the relationship between this noise, the size of the ensemble and the use of linear decoding.

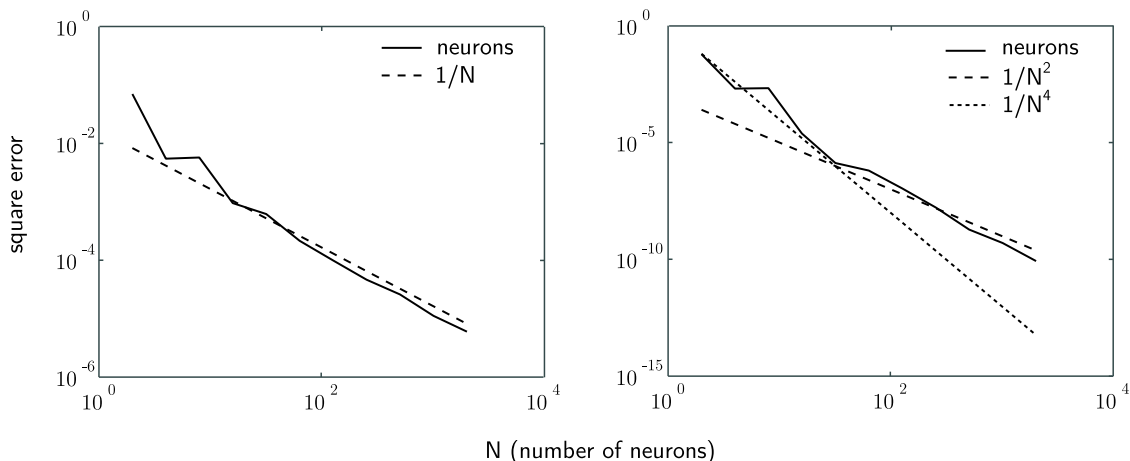


Figure 2.5. The effect of ensemble size on error. Increasing ensemble size rapidly reduces error. *Left:* The effect of increasing N reduces error due to noise by a factor of $\frac{1}{N}$. *Right:* Increasing N reduces error due to static distortion by a factor of $\frac{1}{N^4}$ for low N , and $\frac{1}{N^2}$ for high N . Used with permission of Eliasmith [48].

2.2.4 Vector decoding: precision

In addition to noise, the effects of ensemble size, N , and the use of a linear decoding method must be considered in the assessment of representational precision. Figure 2.5 shows logarithmic plots relating N and the square error of the decoded quantity. This data shows that adding neurons to an ensemble reduces both the effect of noise and the error due to the linear decoding process. This provides confidence in the quantifiable effects of representational error for the large-scale model in this paper and accounts for the impact of noise in neurobiological systems. As well, it shows that errors due to noise are dominant in the assumption of linear decoding.

2.2.5 Vector decoding: spiking neurons and time-dependence

To form a biologically plausible description of vector representation we must progress from rate neurons to time-dependent ones, to account for spiking behaviour. Mathematically, we must adapt our formalism such that the signal and the decoders depend on time:

$$\hat{\mathbf{x}}(t) = \sum_i a_i(\mathbf{x}(t)) \mathbf{d}_i(t) \quad (2.17)$$

$$= \sum_{i,s} \delta(t - t_{is}) \mathbf{d}_i(t), \quad (2.18)$$

where t_{is} is the time of spike s produced by neuron i . From these expressions we can interpret the a_i as being instantaneous firing rates and the \mathbf{d}_i as having infinitesimal lifetimes. Since spikes are, by definition, not continuous, the only way to reconstruct a continuous time signal is through the time-dependence of the decoder. By treating the spikes as impulse inputs into a linear decoding system, the continuous output we seek can

be achieved by finding the *impulse response* in the convolution operation:

$$\hat{\mathbf{x}}(t) = h(t) * \sum_{i,s} \delta(t' - t_{is}) \mathbf{d}_i \quad (2.19)$$

$$= \int_0^T h(t - t') \left(\sum_{i,s} \delta(t' - t_{is}) \mathbf{d}_i \right) dt' \quad (2.20)$$

$$= \sum_{i,s} h(t - t_{is}) \mathbf{d}_i. \quad (2.21)$$

This expression can be interpreted as each spike being decoded by a continuous time signal contribution to $\hat{\mathbf{x}}(t)$, determined by the linear filter $h(\cdot)$. To find the *optimal* linear filter, we minimize the mean square error between $\mathbf{x}(t)$ and $\hat{\mathbf{x}}(t)$ as before. However, in real neurons, the life of a spike ends at the synapse, where it causes an ionic current in the post-synaptic receptor. It is this biologically plausible spike filter—the *post-synaptic current*—that we would like to use so long as it can perform comparably with the optimal decoding filter. The post-synaptic current can be modelled as

$$h_{psc}(t) = \frac{1}{\tau} e^{-t/\tau}, \quad (2.22)$$

where τ is the synaptic time constant. This approximation performs well and is a biologically plausible use of what was initially a mathematical device [48].² So with the decoding filter $h_{psc}(\cdot)$ we have a convenient and, with suitably high N , reasonably accurate means of extracting a time-dependent vector $\hat{\mathbf{x}}(t)$ from the spike trains of an ensemble of neurons.

2.3 Vector transformation

2.3.1 Linear transformations

Suppose we have vector quantities \mathbf{x} and \mathbf{y} , each represented by the activities of a neural ensemble (Figure 2.6). If our decoding is good, i.e. $\hat{\mathbf{x}} \cong \mathbf{x}$ and $\hat{\mathbf{y}} \cong \mathbf{y}$, we can use the formalism we have established to determine the activities of a third neural ensemble that will represent a linear combination of these vectors. So for $\mathbf{z} = c_1 \mathbf{x} + c_2 \mathbf{y}$, where the c_1 and c_2 are scalars, the encoding equation for neurons in ensemble s_k becomes

$$\begin{aligned} s_k(c_1 \mathbf{x} + c_2 \mathbf{y}) &= G_k \left[\alpha_k \langle \mathbf{e}_k, c_1 \hat{\mathbf{x}} + c_2 \hat{\mathbf{y}} \rangle + J_k^{bias} \right] \\ &= G_k \left[\alpha_k \left(\sum_i a_i(\mathbf{x}) c_1 \langle \mathbf{e}_k, \mathbf{d}_i^x \rangle + \sum_j b_j(\mathbf{y}) c_2 \langle \mathbf{e}_k, \mathbf{d}_j^y \rangle \right) + J_k^{bias} \right] \\ &= G_k \left[\sum_i \omega_{ki} a_i(\mathbf{x}) + \sum_j \omega_{kj} b_j(\mathbf{y}) + J_k^{bias} \right], \end{aligned}$$

where $\omega_{ki} \equiv \alpha_k c_1 \langle \mathbf{e}_k, \mathbf{d}_i^x \rangle$ and $\omega_{kj} \equiv \alpha_k c_2 \langle \mathbf{e}_k, \mathbf{d}_j^y \rangle$ are the *connection weights* between ensembles. Networks for other linear combinations can be determined in a similar way.

²The function $h_{psc}(t)$ is also the impulse response for an RC circuit. What makes it biologically plausible is the approximate exponential decay of most post-synaptic currents, and the ability to change tau to reflect different kinds of neurotransmitters.




Figure 2.6. A linear transformation network. The linear transformation $\mathbf{z} = \mathbf{x} + \mathbf{y}$ is computed over three populations of neurons, a_i , b_j and s_k , with connection weights w_{kj} and w_{ki} . Used with permission of Eliasmith [48].

2.3.2 Nonlinear transformations

The method described for the calculation of linear transformations cannot be used for nonlinear transformations. Achieving such transformations (assuming no dendritic mechanisms) requires a *middle* or *hidden layer* that encodes a *higher dimensional* quantity determined by the inputs, which then projects to a population of *lower* representational dimension.

For example, suppose we wish to build a network to multiply two one-dimensional quantities, x and y . Then we compute the activities of the ensemble p_k that will encode the scalar product $z = xy$ (Figure 2.7). The activities of the two-dimensional middle layer are given by

$$\begin{aligned}
 m_l(\mathbf{v}) &= G_l \left[\alpha_l \langle \mathbf{e}_l, \mathbf{v} \rangle + J_l^{bias} \right] \\
 &= G_l \left[\alpha_l (e_l^{(1)} x + e_l^{(2)} y) + J_l^{bias} \right] \\
 &= G_l \left[\sum_i \omega_{li} a_i(x) + \sum_i \omega_{lj} b_i(y) + J_l^{bias} \right],
 \end{aligned}$$

where the superscripts index the vector components, ensembles a_i and b_j represent x and y as before, and where we use a two-dimensional ensemble, m_l , to encode the two-dimensional vector $\mathbf{v} = (x, y)$. The $\omega_{li} \equiv \alpha_l x e_l^{(1)} d_i^x$ and $\omega_{lj} \equiv \alpha_l y e_l^{(2)} d_j^y$ are the connection weights for the middle layer transformation. Next, the activities of the population p_k




Figure 2.7. A nonlinear transformation network. The nonlinear transformation $z = xy$ is computed over four populations of neurons, a_i , b_j , p_k and m_l with connection weights w_{lj} , w_{li} and w_{kl} . The middle layer holds the representation of the two-dimensional quantity $\mathbf{m} = (m_x, m_y)$. Used with permission of Eliasmith [48].

encoding the product are given by

$$\begin{aligned}
 p_k(\mathbf{m}) &= G_k \left[\alpha_k \langle \mathbf{e}_k, \mathbf{m} \rangle + J_k^{bias} \right] \\
 &= G_k \left[\alpha_k \mathbf{e}_k \sum_l m_l(\mathbf{v}) \mathbf{d}_l^{\mathbf{m}} + J_k^{bias} \right] \\
 &= G_k \left[\sum_l \omega_{kl} m_l(\mathbf{v}) + J_k^{bias} \right],
 \end{aligned}$$

where $\omega_{kl} \equiv \alpha_k \langle \mathbf{e}_k, \mathbf{d}_l^{\mathbf{m}} \rangle$. Then p_k can be decoded into the product \hat{z} by

$$\hat{z} = \sum_k p_k(\mathbf{m}) d_k^z, \tag{2.23}$$

where d_k^z is the contribution weight for neuron k in the sum of activities p_k that represent \hat{z} . Networks for other nonlinear functions can be determined in a similar way.

2.4 Dynamics

Biologically relevant stimuli and neural responses generally depend on time. Since we have incorporated time dependence in our account of vector representation, the NEF is able to use the state-space method from control engineering to describe dynamic state transformations. Since the primary contribution of this thesis relates to the design and implementation of a particular kind of dynamical neural network, a discussion of dynamical system modelling and the suitability of the state-space method is provided.

2.4.1 The state-space method: an example using a damped oscillator

Dynamical systems are often described mathematically by differential equations. The *order* of a particular differential equation is the highest derivative of the dependent variable with respect to the independent variable. An *ordinary* differential equation contains one independent variable. The *state-space representation* is the description of an n th order system as a set of n first-order, ordinary differential equations [64]. Unlike the classical frequency domain approach that relates individual system inputs to outputs to determine a transfer function, the state-space method emphasizes the internal workings of the process that converts multiple inputs to multiple outputs. The primary benefit of the state-space approach is the explicit specification of the relations among these independent variables.

The central concept in the analysis is the *system state*, expressed as the vector $\mathbf{x}(t) = (x_1(t), x_2(t), \dots, x_n(t))$. It is defined as the minimum set of variables that provide a complete, but not necessarily unique, quantification of the internal ‘configuration’ of a system at a given time. Knowledge of the state vector for some initial time along with the time evolution properties of the system allow for the prediction of any future state and output for any given input [64]. Systems that can be characterized this way are called *state-determined*. Historically, n refers to the number of independent energy storage elements in the system.

As an illustration of the use of state-space methods and a motivating example of the velocity computation modelled in Chapter 3, we consider a simple dynamical system: a mass on a spring under friction, oscillating in one dimension. The mass m is fixed to a spring with a constant stiffness coefficient, k_1 , moving in one dimension (see Figure 2.8) under a constant friction coefficient, k_2 , but without any driving forces yet. The second-order differential equation for this system is found by equating Hooke’s law, $F_s = -k_1x$, with Newton’s second law of motion, $F_{net} = m\ddot{x}$, then adding the friction (or viscosity) term $-k_2\dot{x}$:

$$\begin{aligned} m\ddot{x} &= -k_1x - k_2\dot{x} \\ \ddot{x} &= -\frac{k_1}{m}x - \frac{k_2}{m}\dot{x} \end{aligned} \tag{2.24}$$

where $x = x(t)$, $\dot{x} \equiv \frac{dx}{dt}$, $\ddot{x} \equiv \frac{d^2x}{dt^2}$. There are two independent energy storage elements: the stiffness of the spring, represented by the position x , and the momentum of the mass, represented by velocity \dot{x} [65]. Thus, the state-space representation of (2.24) will consist of two ordinary, first-order, linear differential equations.

Next we impose a change of variables. Letting $x_1 = x$ and $x_2 = \dot{x}$, we see that \dot{x}_1 and \dot{x}_2 can each be written as linear combinations of x_1 and x_2 . In matrix-vector form this is

$$\begin{aligned} \begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \end{pmatrix} &= \begin{pmatrix} 0 & 1 \\ -\frac{k_1}{m} & -\frac{k_2}{m} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \\ &= \begin{pmatrix} 0 & 1 \\ -\omega_o^2 & -2\gamma \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}, \end{aligned}$$

written $\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}$, (2.25)

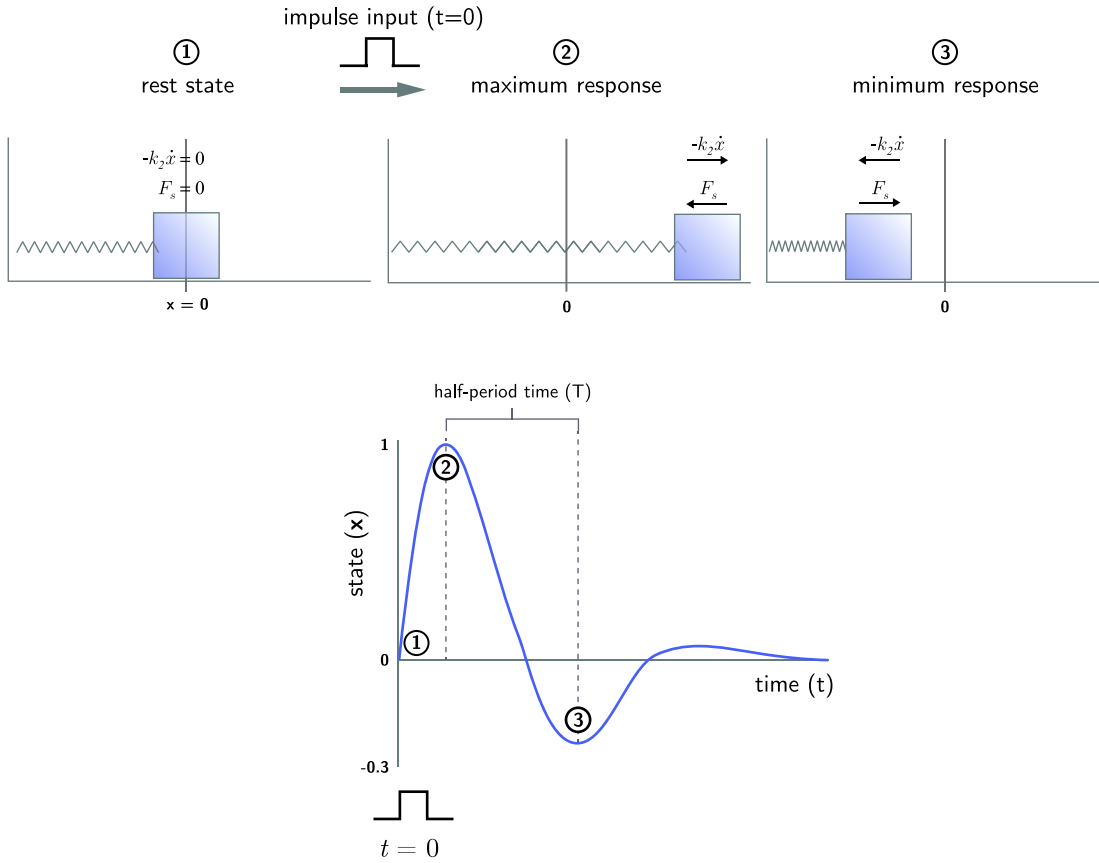


Figure 2.8. Impulse response of damped, driven oscillator in one dimension. A mass on a spring under friction is an example of a damped oscillator. *Top:* The system state, x , is the displacement from the equilibrium point $x = 0$. The friction term acts in the direction opposite the restoring force of the spring. *Bottom:* The time course of the system showing the extrema. The *half-period* time, T , is the interval between the first maximum and following minimum. ① is the system at rest, just before the initial impulse input. ② shows the maximum amplitude after an impulse input. ③ shows the minimum amplitude without further input.

where by convention $\gamma \equiv \frac{k_2}{2m}$ is called the *damping factor* and $\omega_o \equiv \sqrt{\frac{k_1}{m}}$ is called the *natural frequency*³. \mathbf{A} is the *dynamics* matrix instantiated intrinsically by the system⁴. Since the elements of \mathbf{A} do not contain x or any of its derivatives, the system is *linear*. Also, since \mathbf{A} is independent of time, the system will behave the same way whenever the input is applied. Hence, the system is *linear and time-invariant* (LTI). Note that the factors affecting dynamics are contained completely in the dynamics matrix, \mathbf{A} .

Next we assume our initial displacement to be $x(t_0) = 0$, so the system does nothing until we add the input $\mathbf{u}(t) = [u_1(t), u_2(t)]$. Then (2.25) becomes the first of the two standard LTI *state equations*,

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}, \quad (2.26)$$

³The *damped frequency* is typically defined as: $\omega \equiv \sqrt{\frac{k_1}{m} - \frac{k_2^2}{4m^2}} = \sqrt{\omega_o^2 - \gamma^2}$.

⁴Notice that due to the redundancy $\dot{x}_1 = x_2 = \dot{x} = \dot{x}_1$, the state-space representation does not add any information or dimensionality.




Figure 2.9. A damped, driven oscillator in one dimension with resonant input. If the frequency of impulse inputs match the natural frequency of the oscillator, the system behaves as if it were undamped (red trace; blue trace from Figure 2.8 for comparison); that is, the input force overcomes the force due to friction. The impulse times are $t \in \{0, T - \epsilon, 2(T - \epsilon)\}$, where ϵ is a short duration that allows the force of the input to be absorbed by the system. The initial phases of the input signal and the oscillator always match since the first input pulse excites the oscillator from rest at ①. The maximum at ② is the same for both scenarios, but the minimum at ③ has a significantly greater magnitude in this case (black arrow). This strong minima is a local measure of the correlation between the intrinsic frequency of the oscillator and the frequency of the input pulses for nearby preceding times. If the input frequency is dissimilar to the oscillator, the interference is destructive, reducing the amplitude. If the impulses are too frequent, the intrinsic dynamics of the system are overwhelmed and the behaviour follows the input function in a scenario called *over-driving*.

where \mathbf{B} is a unit conversion and scaling matrix for the input, but in this example is set to the identity.

If \mathbf{u} is an impulse to the damped oscillator system, the behaviour will resemble that depicted in Figure 2.8. If the input consists of impulses with a frequency equal to the intrinsic frequency of the oscillator, a *resonance* response is produced and the system behaviour will resemble that depicted in Figure 2.9.

The second state equation relates the output of the system, \mathbf{y} , to the system state \mathbf{x} and input \mathbf{u} :

$$\mathbf{y} = \mathbf{C}\mathbf{x} + \mathbf{D}\mathbf{u}$$

where \mathbf{C} and \mathbf{D} are the *output* and *feed-through* matrices, respectively. These matrices are determined by the variables of interest from the output and input. In this example we are only interested in the system state \mathbf{x} without feed-through, so the second state equation is just $\mathbf{y} = \mathbf{x}$. Thus, our discussion is concerned only with the first equation (2.26), which is the state-space formulation for the damped oscillator system.

The state equations are depicted in the standard block diagram in Figure 2.10. In linear systems analyzed in the time domain, the central block or *plant*⁵ integrates the continuous changes in state, which depends on linear combinations of the state $\mathbf{x}(t)$ itself and the input $\mathbf{u}(t)$. The state variables can always be considered as outputs of integrator blocks [64].

⁵Also known as the *transfer function* in frequency-domain analysis.




Figure 2.10. The LTI block diagram. The standard depicting the state equations for a linear time-independent (LTI) system. Elements related to the output equation are greyed as they are not important to the discussion.

2.4.2 Adapting state-space methods to neural ensembles

In order to exploit state-space methods in the design of LTI *neural* systems, the NEF provides a mapping of the standard state equations onto neural ones. With this mapping, we will be able to *define the dynamical behaviour we want* by choosing suitable elements for matrix **A**.

The third NEF principle suggests that the system state vector \mathbf{x} can be understood as a neural representation, just like the decoded signal vector discussed earlier. This means that it is possible to consider the plant is the neural ensemble itself. This does not mean that we take n neurons as n independent energy conserving components, since the activity of any single neuron i does not independently represent any single component x_i in \mathbf{x} . But in the representational ensemble where $\hat{\mathbf{x}} \cong \mathbf{x}$, there exists for each component x_i the linearly independent neural activities that encode x_i .

Eliasmith and Anderson argue that the impulse response of the neural ensemble is dominated by the dynamics of the post-synaptic current, $h_{psc}(t) = \frac{1}{\tau}e^{-t/\tau}$, which acts as a spike (impulse) filter [48]. By the linearity condition, the sum of these responses is the output of the system (weighted by the decoders), $\hat{\mathbf{x}}(t) \cong \mathbf{x}(t)$. So the state equation for the neural ensemble is the convolution

$$\mathbf{x}(t) = h_{psc}(t) * [\mathbf{A}'\mathbf{x}(t) + \mathbf{B}'\mathbf{u}(t)], \quad (2.27)$$

which is essentially (2.19) with state feedback plus an input term. The \mathbf{A}' and \mathbf{B}' are the ‘neural versions’ of **A** and **B**; relating these quantities is the crucial step that will allow us to use standard LTI methods to ascribe the dynamics we want in our system.

For convenience, we now work in the Laplace domain, where $h_{psc}(t) \rightarrow h_{psc}(s) = \frac{1}{1+s\tau}$. So (2.27) becomes

$$\begin{aligned} \mathbf{x}(s) &= \frac{1}{1+s\tau}[\mathbf{A}'\mathbf{x}(s) + \mathbf{B}'\mathbf{u}(s)] \\ &= \frac{\tau^{-1}}{\tau^{-1}+s}[\mathbf{A}'\mathbf{x}(s) + \mathbf{B}'\mathbf{u}(s)]. \\ (\tau^{-1}+s)\mathbf{x}(s) &= \tau^{-1}[\mathbf{A}'\mathbf{x}(s) + \mathbf{B}'\mathbf{u}(s)] \\ s\mathbf{x}(s) &= \tau^{-1}(\mathbf{A}' - \mathbf{I})\mathbf{x}(s) + \tau^{-1}\mathbf{B}'\mathbf{u}(s). \end{aligned} \quad (2.28)$$

Now we can equate the conventional and neural state equations, (2.26) and (2.27) respectively, to express \mathbf{A} and \mathbf{B} in terms of \mathbf{A}' and \mathbf{B}' :

$$\begin{aligned} \mathbf{A}\mathbf{x}(s) + \mathbf{B}\mathbf{u}(s) &= \tau^{-1}(\mathbf{A}' - \mathbf{I})\mathbf{x}(s) + \tau^{-1}\mathbf{B}'\mathbf{u}(s) \\ \Rightarrow \mathbf{A}' &= \tau\mathbf{A} + \mathbf{I} \end{aligned} \tag{2.29}$$

$$\Rightarrow \mathbf{B}' = \tau\mathbf{B} \tag{2.30}$$

With these conversions of the dynamics and input matrices from conventional LTI systems to neural ones, we have a simple way to define the desired dynamics of our neural model.

2.4.3 The oscillator interference (OI) concept

Let us consider Figure 2.11 intuitively. Suppose our physical system is a ball tethered to a post at the origin of a two-dimensional plane with coordinates given by $\mathbf{x} = (x_1, x_2)$. The tether is a spring-like rope, so that the ball is retracted toward the origin unless hit with some outward component of force. The ball has a tendency to orbit the post and is capable of rotation in only one direction. Suppose an initial ‘launching’ input, delivered at the origin, has radial and tangential components of force such that the initial trajectory follows the blue path (left) from the origin. The ball begins to move outward and around the post with an angular frequency dependent on the length of the rope. Without further hits, the ball returns to rest at the origin.

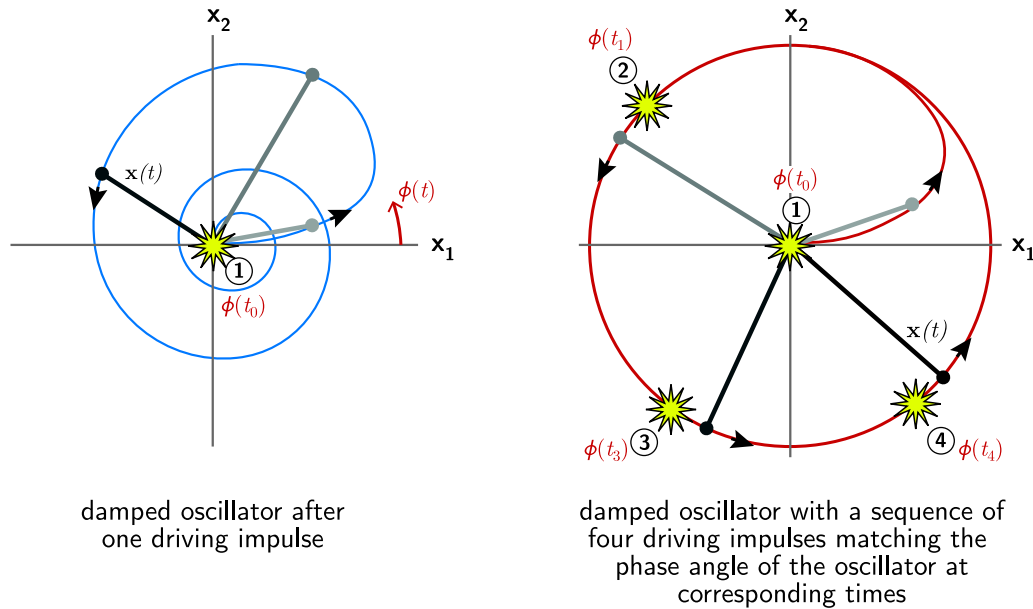


Figure 2.11. A 2-D damped, driven oscillator. Impulse coordinates are indicated by yellow stars, along with time-dependent phase angles denoted $\phi(t_i)$, and order of occurrence (circled numbers). *Left:* Impulse response trajectory of the state vector $\mathbf{x}(t)$. *Right:* The same system subject to resonant input shows sustained oscillatory behaviour. See text for further description.

Now suppose we wish to keep the oscillation going with the help of a few teammates standing around a circle. After the launch of the ball, each player must carefully time her punch (a burst input) of the ball; too early or too late will miss and the ball will begin to retract toward the origin. If the team’s timing is good, the ball will oscillate indefinitely and the radius will be large. In other words, hitting the ball at a position (phase ϕ) and angular velocity (frequency $\dot{\phi} = \omega$) similar to those inherent to the system cause it to *resonate*.

Now if the players hit the ball too weakly, the trajectory will have a smaller radius but the same frequency. If they hit with ample force, the radius will be extended but the frequency preserved (assuming the rope behaves linearly with stretching).

Finally, consider the scenario in which we have too many players, each standing within arms reach of her neighbours. The dynamics of the ball are overwhelmed; the ‘natural’ frequency depends on the input forces alone, without any ‘natural’ behaviour of its own. The system is overdriven and the players can move the ball around as they please; the ball does not have enough time and space to demonstrate its own dynamics.

So, if we observe the game at some time and see that the radius is large, we can infer that there has been a high *correlation* recently between the phase and frequency of the players’ hits and those inherent to the system. In all other cases (notwithstanding overdriving), the radius will be reduced.

We now formalize these notions using state-space methods. The dynamics of the system are determined by the dynamics matrix \mathbf{A} (see section 2.4.1 for details),

$$\mathbf{A} = \begin{pmatrix} 2\gamma & -\omega_o^2 \\ \omega_o^2 & 2\gamma \end{pmatrix}, \quad (2.31)$$

and the system can be described completely by $\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$ (the second equation is just $\mathbf{y} = \mathbf{x}$), along with the initial conditions $\mathbf{x}(0) = \dot{\mathbf{x}}(0) = (0, 0)$. The sign relation of the anti-diagonal terms in the dynamics matrix \mathbf{A} restricts the direction of rotation. The sequence of input pulses is $\{\mathbf{u}(t) \mid t \in [t_0, t_1, t_2, t_3]\}$.

In the resonant case, each input velocity, $\dot{\mathbf{u}}(t) = (\dot{x}_1, \dot{x}_2)$, is similar to the state velocities, $\dot{\mathbf{x}}(t)$, for each time t . Importantly, the magnitude (radius) of the state, $\|\mathbf{x}\|$, is greater than it would be in the non-resonant or uncorrelated case.

Since the oscillator is initially at rest, at least two interactions or *interference events* are required for correlation. The first acts as the initializer, triggering the oscillation. Subsequent events interact (linearly) with the system and extend or reduce $\|\mathbf{x}\|$. Assuming we are considering only recent times (i.e. since the start of the current cycle), we can use $\|\mathbf{x}\|$ during the later phases of its time course as a correlation measure between \mathbf{x} and \mathbf{u} for earlier times or phases.

To achieve this measure, we take the difference between the radius under test stimulus conditions, $\|\mathbf{x}\|_{stim}$, and the radius generated by a single impulse, $\|\mathbf{x}\|_{imp}$. This gives the measure R , where

$$R = \begin{cases} \|\mathbf{x}\|_{stim} - \|\mathbf{x}\|_{imp}, & \text{for } \|\mathbf{x}\|_{stim} \geq \|\mathbf{x}\|_{imp} \\ 0, & \text{otherwise.} \end{cases}$$

Since the proposed model relies on resonant interference between input and oscillator, we refer to it for convenience as the ‘oscillator interference’ (OI) model. The next chapter provides details of the OI concept as applied to the problem of velocity selectivity and its implementation in a large scale neural network simulation.

Chapter 3

A mechanistic model of motion processing in the early visual system

The essential problem for this work is to explain how the observed spatiotemporal filtering properties of V1 simple cells arise and how such behaviour transforms intensity signals on the retina into a map of velocities in area MT. In this chapter, a computational model ('OI', oscillator interference) is presented using as its foundation the physiological data from Chapter 1 and the quantitative principles of Chapter 2. A description of the OI algorithm is followed by a detailed account of the implementation. The results of the simulation are provided in Chapter 4.

3.1 Algorithmic description

The computational sequence of the model, depicted in Figure 3.1, is as follows:

1. **Spatial decorrelation from RGCs**

An input signal of continuous intensity values for a moving stimulus is pre-processed by convolution with an edge detection filter (see Figure 1.12). This simulates the response of the retinal ganglion cells (RGCs) [2].

2. **Temporal decorrelation from LGN**

Changes in the edge signal vector beyond a threshold (super-threshold events) cause a brief (2 ms) burst of spikes representing the moving edge during that interval. This simulates the burst output observed in the LGN [9]. See Figure 3.5.

3. **Directional filtering of V1 input**

The LGN output is projected onto two-dimensional Gabor functions that act as spatial direction input filters for the V1 oscillators. This ensures that the initial phase angle of \mathbf{x} is $\phi = 0$. Each oscillator has an associated direction angle chosen from a set of 8, equally spaced between 0 and 2π .

The oscillator interference (OI) model

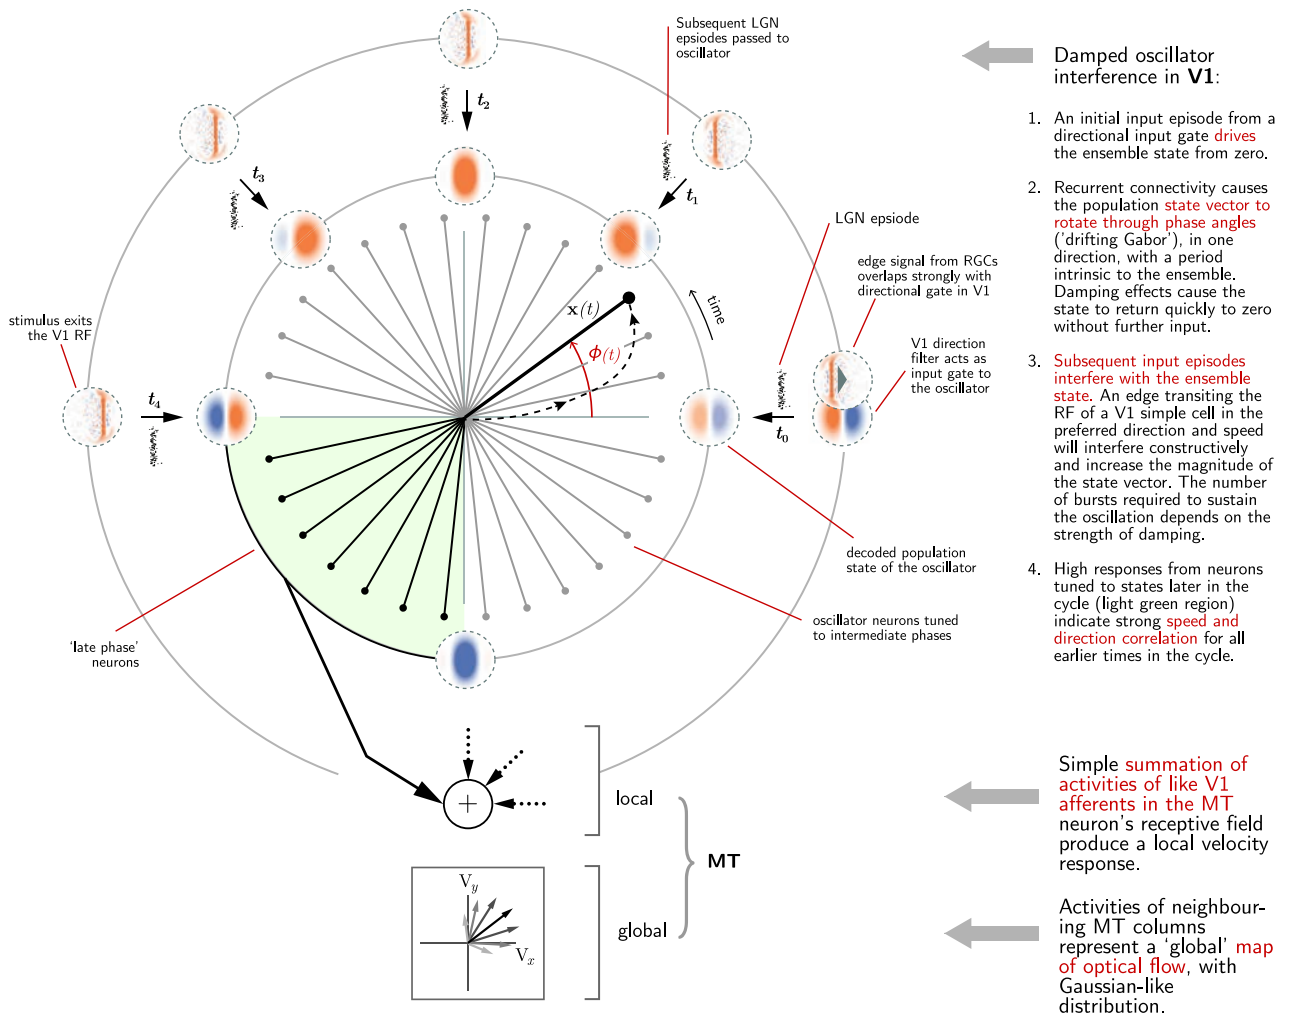


Figure 3.1. Conceptual diagram of the proposed 'oscillator interference' (OI) model. A description of each processing step is included at right. See text for further description.

4. Interference with damped, driven phase oscillators in V1

The V1 simple cell ensembles are driven and damped two-dimensional oscillators able to maintain and evolve their state representations over time. They are composed of several hundred neurons tuned to the same Gabor direction, but with phase tunings distributed from 0 to 2π . Any orthogonal pair of vectors in the phase plane are separated by 90 degrees.¹ The connectivity of the oscillator neurons is specified such that it has a state oscillation frequency ω and damping factor γ .

The rotation through phase space is restricted to one direction by the dynamics matrix detailed in section 3.2.5. Since each oscillator is driven by input and exhibits damping properties, its activity decodes to zero in the absence of stimuli.

¹Any of these pairs constitute the 'quadrature pair' proposed by Adelson and Bergen in [19].

An initial burst drives the state vector, $\mathbf{x}(t)$, to rotate. Subsequent input adds vectorially to $\mathbf{x}(t)$ and thus interferes with the oscillator state. Constructive interference increases the magnitude of $\mathbf{x}(t)$ and furthers its rotation. This occurs only when the direction and phase of the stimulus is similar to $\mathbf{x}(t)$; otherwise, the interference is destructive and decreases the magnitude of the state vector. In other words, the V1 neurons coding the state vector will exhibit greater activity when there is a high correlation between the spatiotemporal frequencies of the stimulus and the oscillator state. This is the essential mechanism for spatiotemporal filtering in the model and is precisely a *resonance* interaction between the input bursts and oscillator dynamics.²

5. Spatial pooling of V1 afferents in MT

The decoded scalar magnitudes of like-tuned, late-phase (where $\phi(t) \in [\pi, 2\pi]$; see Figure 3.1, ‘V1’) oscillator neurons for a number of local patches in the visual field are summed in an MT ensemble. This value is the contribution weight of the preferred velocity of the MT ensemble to the map of optical flow.

3.2 Implementation details

3.2.1 The Nengo simulation environment

Nengo is the name of an open-source, cross-platform Java application developed by the Computational Neuroscience Research Group (CNRG), Centre for Theoretical Neuroscience at the University of Waterloo in Ontario, Canada. Nengo implements all aspects of the Neural Engineering Framework (NEF) as described in this thesis. Using the NEF, Nengo is able to construct spiking neural networks of arbitrary size and function using biologically constrained single neurons, physiological time constants and realistic levels of noise. The program solves for necessary synaptic connection weights to achieve desired linear and nonlinear vector transformations. The Nengo simulator can be downloaded from <http://nengo.ca>.

3.2.2 Stimulus movie generation

Stimulus movies for moving bars and sinusoids were coded in Matlab[®]. Dot motion stimulus movies were generated using the *Psychtoolbox-3* extensions for Matlab[®] [66–68]. The software is freely available at <http://psychtoolbox.org>.

3.2.3 Basis of representation

Other than the burst output from LGN, the computations required by the model are linear transformations. As such, the relevant quantities and operators can be represented in a basis. The difficulty is to find a set of basis functions that can represent small, localized stimuli as well as larger, delocalized structures like Gabor wavelets.

²Resonance requires a signal to match the frequency *and phase* of an oscillator. Since ours are driven from rest by an input already filtered for direction, the initial phase of stimulus and the oscillator is always equal at $\phi = 0$.

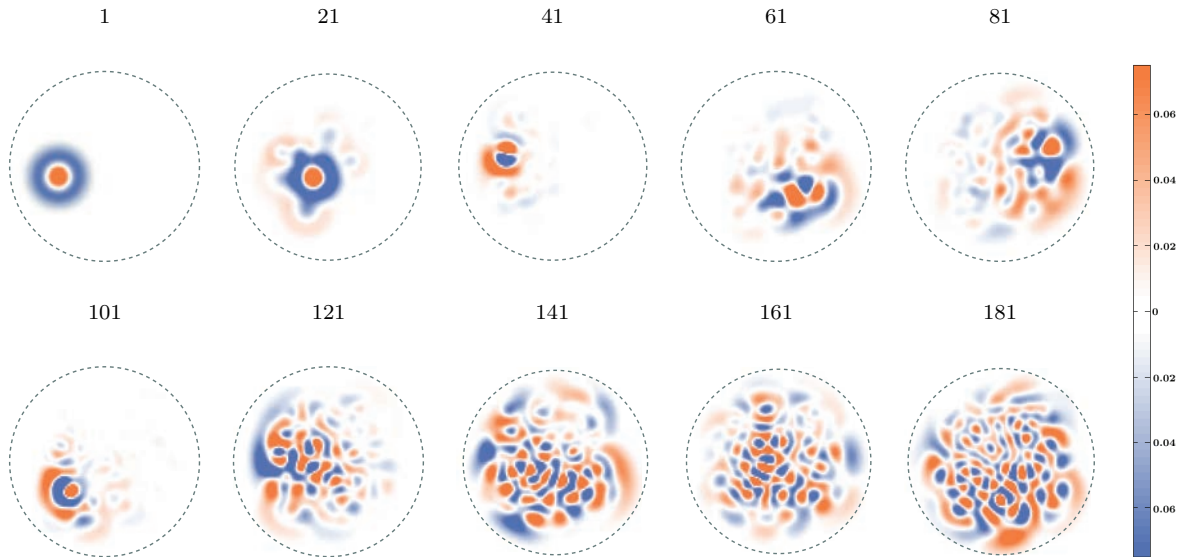


Figure 3.2. Sample of basis functions used. An ordered sample of the orthonormal basis functions (eigenvectors). This ‘wavelet’ basis is used to represent a variety of spatial patterns like localized dots, elongated bars and periodic structures like Gabors. The dashed circle is the RF aperture. When the Gram-Schmidt orthonormalization procedure is applied to pure centre-surround functions like 1 (top left), the results become increasingly delocalized clusters of decaying, alternating peaks.


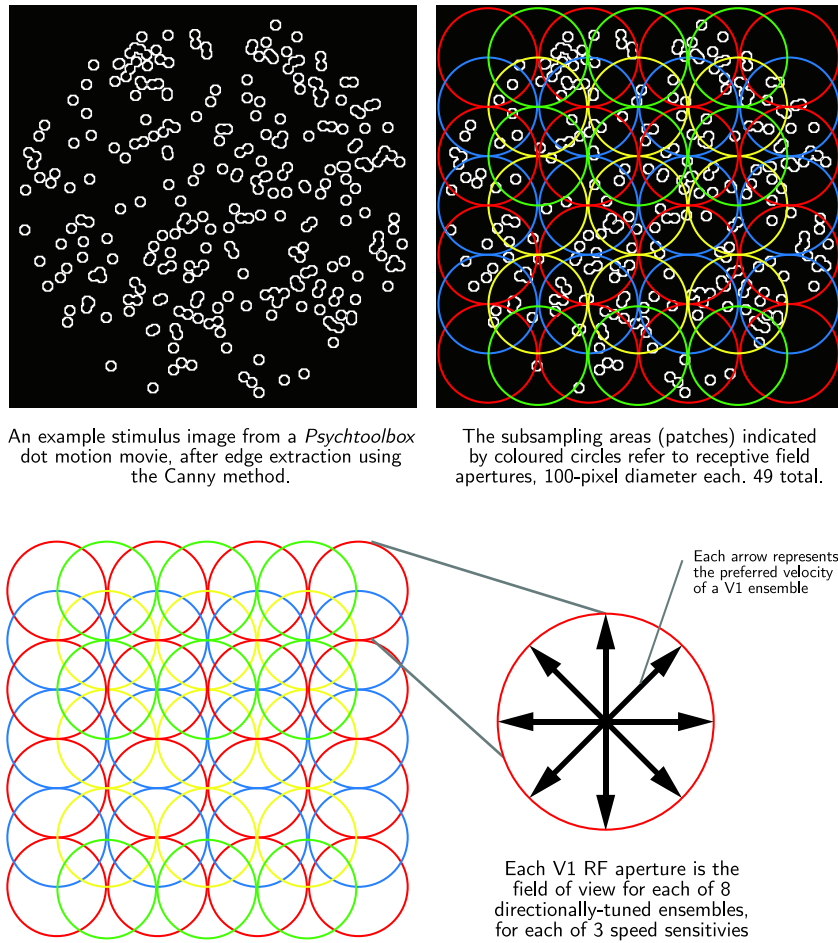


Figure 3.3. Representation of structures in the basis. The structures on the left of each pair are represented as coordinate vectors in the basis and reconstructed on the right. Structures near the edge of the RF aperture do not reconstruct well due to edge effects.

It is widely believed that the Gabor-like spatial RFs of V1 neurons are determined by combinations of centre-surround responses (Figure 1.3) from the LGN [18]. While it is possible to construct a basis from such functions, the number required to adequately represent small, localized structures as well as large periodic ones used in this work was found to be about 1000. To ease the implementation requirements of computer memory and processing time for loading and running the model, a reduced basis of 200 functions was constructed to suit the stimuli used in these experiments. This was achieved by using an idealized centre-surround function to seed the Gram-Schmidt orthonormalization algorithm [69]. This procedure is a recursive linear algorithm in which each successive candidate basis function is replaced by an orthogonal linear combination of all preceding functions. A sample of the reduced basis functions is illustrated in Figure 3.2. A sample of vectors reconstructed from this basis is shown in Figure 3.3.



An example stimulus image from a *Psychtoolbox* dot motion movie, after edge extraction using the Canny method.

The subsampling areas (patches) indicated by coloured circles refer to receptive field apertures, 100-pixel diameter each. 49 total.

Figure 3.4. Stimulus edge extraction, field subsampling and V1 array structure. *Top left:* An example image from a dot motion stimulus movie after edge extraction. *Top right:* Overlapping RF apertures (‘subsamples’ or ‘patches’) cover the stimulus area. *Bottom:* Each patch contains the field of view for 24 V1 oscillators (8 preferred directions \times 3 preferred speeds each).

3.2.4 Centre-surround edge extraction and subsampling

The first transformation in the model is the extraction of edge information from an intensity image. This is accomplished as follows (see Figure 3.4, top left):

1. The stimulus is a greyscale movie with 1 ms time steps, 400×400 pixels.
2. The movie is filtered with a centre-surround edge (Canny) detection filter in Matlab[®] to approximate the output of RGCs.
3. The movie is partitioned into 49 patches with 50% overlap (see Figure 3.4, top right).
4. The pixel values from each patch are projected onto the high-dimensional basis described above to generate the 200 vector coordinates per patch, per time step.

3.2.5 Spatiotemporal filtering by LGN/V1

Oscillation period as a speed measure

Each oscillator i has a damped frequency, ω_i , and an associated period, $2T_i$. The *half-period* time, T_i , is equal to the time interval between adjacent positive and negative peaks, and is chosen from five equally spaced intervals: $T_i \in [20, 35, 50, 65, 80]$ ms. Each T_i is associated with only one ω_i , which in turn describes the rate of change of represented Gabor phase (drift speed) and thus the preferred speed of the direction-selective oscillator. While units of time cannot alone measure speed, we can, for convenience, refer to the preferred speeds of the oscillators by these T_i without ambiguity.

Pre-processing: LGN burst timing

Dong and Atick (see section 1.2.4 and [10]) argue that a primary function of the LGN is the temporal decorrelation of an edge signal. The temporal filter kernel they have proposed (see Figure 1.14, left) can produce a reduced signal with minimal loss of information, but a biologically plausible mechanism to generate such behaviour in LGN is not provided. Such a mechanism is also outside the scope of this thesis.

It is an assumption of the OI model that the timing of an initial LGN burst corresponds to the positive overlap between the stimulus edge signal and a V1 directional filter. Figure 3.5 shows an stimulus moving across a Gabor directional filter. Shortly after the threshold is crossed ($t = t_0$), a 2 ms sample of the stimulus signal is passed to the associated oscillator.³ For each oscillator i , a second sample is passed from LGN after duration T_i , irrespective of the overlap between the stimulus and the directional filter. Thus, the initial peak overlap between the stimulus and the V1 directional filter triggers the oscillator.

Oscillator design

We take as our starting point a damped oscillator (section 2.4.1), and impose the essential characteristics of the Dong-Atick temporal filter. These include:

1. The oscillation undergoes approximately one period, suggesting strong but sub-critical damping.
2. The negative amplitude is between $\frac{1}{2}$ and $\frac{1}{4}$ of the maximum positive amplitude.

As discussed in section 2.4.1, a linear time-invariant system can be realized by differential equations that depend on the system state at each time t , where the dynamical behaviour is determined completely by the matrix \mathbf{A} in the state equation $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$. Since we have already adapted this equation to the representational state \mathbf{x} of a neural ensemble in which $\mathbf{A}' = \tau\mathbf{A} + \mathbf{I}$ and $\mathbf{B}' = \tau\mathbf{B}$, all that remains is to specify the linear

³This duration is kept short, as the longer the duration, the more the intrinsic dynamics of the oscillator are suppressed. The work of Butts [9] suggests that a duration closer to ~ 20 ms. This poses no problem for the OI model since any burst duration greater than 2 ms is suitable.




Figure 3.5. LGN burst timing. A stimulus moves across a Gabor directional filter. Shortly after the threshold is crossed at $t = t_0$, a 2 ms sample of the signal is passed from the LGN ensemble to the associated V1 oscillator. A second sample is passed after duration T_i , irrespective of the overlap between the stimulus and the directional filter.

transform \mathbf{A} that gives us the dynamics we want. The matrix for the incremental rotation of a damped two-dimensional vector counter-clockwise about the origin is:

$$\mathbf{A} = \begin{pmatrix} 2\gamma & -\omega_o^2 \\ \omega_o^2 & 2\gamma \end{pmatrix}.$$

The opposing signs of the off-diagonal natural frequency terms, $\pm\omega_o^2$, guarantee that the rotation of the state vector occurs in one direction only. Thus, this is the recurrent connection matrix we want for each V1 oscillator. With the damped frequencies known, we can determine the decay constant for the desired temporal filter by the logistic function, shown in Figure 3.7. The input signal is projected onto the oscillators using the associated directional Gabor filters.

3.2.6 Spatial integration in MT

The two-dimensional outputs from the late-phase neurons (see Figure 3.1) in each oscillator within a pooling neighbourhood of 1-9 V1 RFs are transformed to a scalar value and summed in an associated MT ensemble.

Construction of the retinotopic MT velocity field

The construction of the MT velocity field array was achieved by mapping a family of eight directional MT ensembles of common speed selectivity to their associated spatial




Figure 3.6. Damped frequency vs half-period time. The half-period times, T , are chosen from equally spaced intervals.




Figure 3.7. Decay vs damped frequency. Using the damped frequencies, the decay constants can be found along a logistic function.




Figure 3.8. MT velocity field configuration. A 7×7 retinotopic array of MT RF centres (dots). Each arrow represents the maximum velocity associated with an ensemble of MT neurons. The positive scalar weight of each arrow is determined by the V1 oscillator output. This weight is proportional to the length of the arrow illustrated. Each of these ensembles pools a variable number of smaller-area V1 ensembles.

location on a grid (Figure 3.8). Each patch overlapped its neighbours by 50%, except around the edges. The activities of each MT ensemble encode a positive scalar weight for an associated velocity and location in the stimulus field. This weight is proportional to the length of the arrow illustrated. Each of these ensembles pools a variable number of smaller-area V1 ensembles, which in this work ranges from 1 to 9.

3.2.7 Neural ensemble network and parameters

Velocity-selective microcircuit schematic




Figure 3.9. Velocity-selective microcircuit schematic. Each microcircuit or sub-network receives input from one patch of the visual field. N , dimensionality of the ensemble. d the number of neurons in the ensemble. See text for details.

Figure 3.9 shows a microcircuit or sub-network of neural ensembles (circle clusters) associated with each patch i of the stimulus area, for each Gabor direction θ_i and oscillator speed T_i . A pre-processed burst signal is passed to an LGN ensemble for representation only; the neural network here does not implement the temporal filtering behaviour of LGN.

The total number of ensembles and neurons in the simulations are:

- V1 ensembles: 49 patches x 8 directions x 3 speeds = 1176
- V1 neurons: 1176 ensembles x 400 neurons each = 4.70×10^5
- MT ensembles: (1 to 9) x V1 ensembles = 130 to 1176
- MT neurons: (130 to 1176 ensembles) x 400 neurons each = 5.2×10^4 to 4.70×10^5
- LGN ensembles: = V1 ensembles = 1176
- LGN neurons: = V1 neurons = 4.70×10^5

Total ensembles (maximum) = 3528

Total neurons (maximum) = $9.41 \times 10^5 + 3.53 \times 10^5 = 1.29 \times 10^6$

3.2.8 Physiological parameters for neural ensembles

Table 3.1 lists the physiological parameters chosen for the proposed model. NOISE?

| Ensemble parameters | Model Value | Biological Value | Reference |
|---|-------------|------------------|------------------|
| LGN | | | |
| RC constant (τ_{RC})(ms) | 10 | 10-20 | [70, 71] |
| Post-synaptic constant (τ_{psc})(ms) | 5.0 | ~ 6.6 | [72] |
| Abs refractory period (τ_{ref})(ms) | 1 | < 0.85 | [73] |
| Max firing rate (Hz) | 100-250 | 100-200+ | [13, 74, 75] |
| V1 simple cells | | | |
| RC constant (τ_{RC}) | 20 | 10-20 | [70, 71] |
| Post-synaptic constant (τ_{psc}) | 5.0 | ~ 6.6 | [72] |
| Abs refractory period (τ_{ref}) | 2 | 1 [†] | [76] |
| Max firing rate | 100-250 | $\sim 70 - 100$ | [77-79] |
| MT | | | |
| RC constant (τ_{RC}) | 20 | 10-20 | [70, 71] |
| Post-synaptic constant (τ_{psc}) | 5.0 | ~ 6.6 | [72] |
| Abs refractory period (τ_{ref}) | 5 | 1 ^{††} | - |
| Max firing rate | 50-150 | 100-200 | [25, 28, 80, 81] |

Table 3.1. Neural ensemble parameters used in the simulation. Reported biological values are included where available. [†] = value is based on a model estimate. ^{††} = using V1 value. (-) = not available.

Noise allowance

The statistical assumptions of noise and effect of this term on the decoders is detailed in section 2.2.3. In brief, the noise added to the neural activities represents the aggregate from all sources, distributed normally about a mean of 0 with a standard deviation of 10% of the maximum firing rate of each neuron.

Chapter 4

Experimental results

4.1 The generation of velocity selectivity

The simulation results of Figure 4.1 illustrate how velocity selectivity arises in a velocity microcircuit; a subsystem that responds preferentially to spatiotemporal frequencies close to its directional Gabor input filter and intrinsic oscillator dynamics. It consists of connected ensembles in a feedforward pathway from LGN to V1 simple cells to MT. The stimulus in this test is a moving bar of maximum intensity with preferred direction and speed for the given V1 oscillator.

The initial burst from LGN drives the damped oscillation of the V1 ensemble. Subsequent input events interfere with the spatiotemporal time course of the V1 population state by adding vectorially. If the interference is constructive, a resonance effect is produced, as seen in the centre plot of Figure 4.1. The burst input frequency can be higher than two events per cycle, but the oscillator will increasingly become overdriven, losing selectivity. If the input frequency is any lower (a single burst), speed correlation cannot occur in this algorithm. Thus, two bursts constitute the minimum number of input bursts from LGN to the associated V1 oscillator capable of velocity detection. The magnitudes of the oscillator state vector in each of the two phase-orthogonal axes are summed and fed forward to MT.

4.2 Predicted impulse response of V1 simple cells

Figure 4.2 shows the decoded component responses of the two-dimensional state vector over time, from a V1 oscillator under 10 ms of ideal input. The orthogonal directions represent two Gabor functions with the same direction but 90-degrees out of phase. A polar plot of the state trajectory around the plane is shown at inset. The time course of the x_1 component strongly resembles those of observed biphasic (separable) V1 RFs [9], while x_2 shows the characteristics of monophasic V1 RFs [9]. Added together, the $x_1 + x_2$ time course shows the observed inseparable behaviour of V1 RFs. Figure 4.3 shows these same results in one-dimensional form.

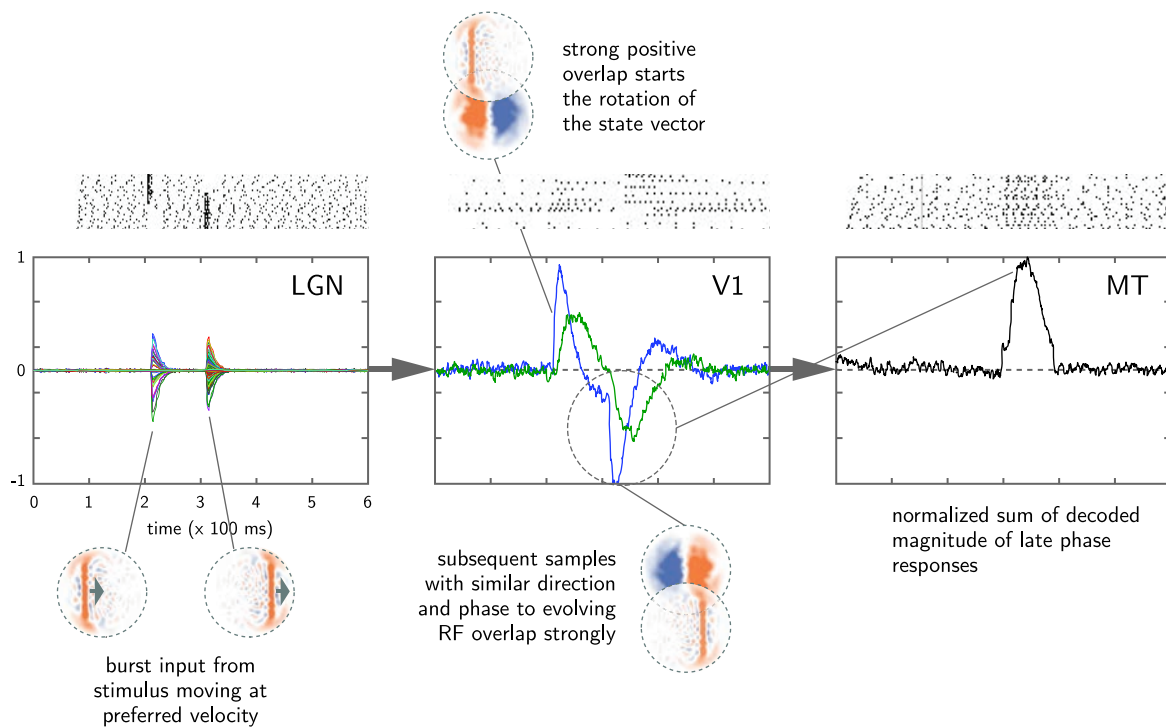


Figure 4.1. The velocity selectivity microcircuit. Each trace shows a decoded vector component value for signals as labelled. The LGN signal is high-dimensional and thus has that number of superimposed traces. Sample spike rasters are shown above. *Left:* A moving bar stimulus produces burst signals in LGN. *Middle:* The initial input drives the oscillation of a V1 simple cell ensemble. Subsequent input can interfere constructively, as shown here, or destructively, with the spatiotemporal frequency of the oscillator state. Fewer than two bursts cannot produce speed correlation in the OI model. *Right:* The magnitudes of V1 efferents that correspond to phase angles late in the oscillation period are summed linearly in MT.

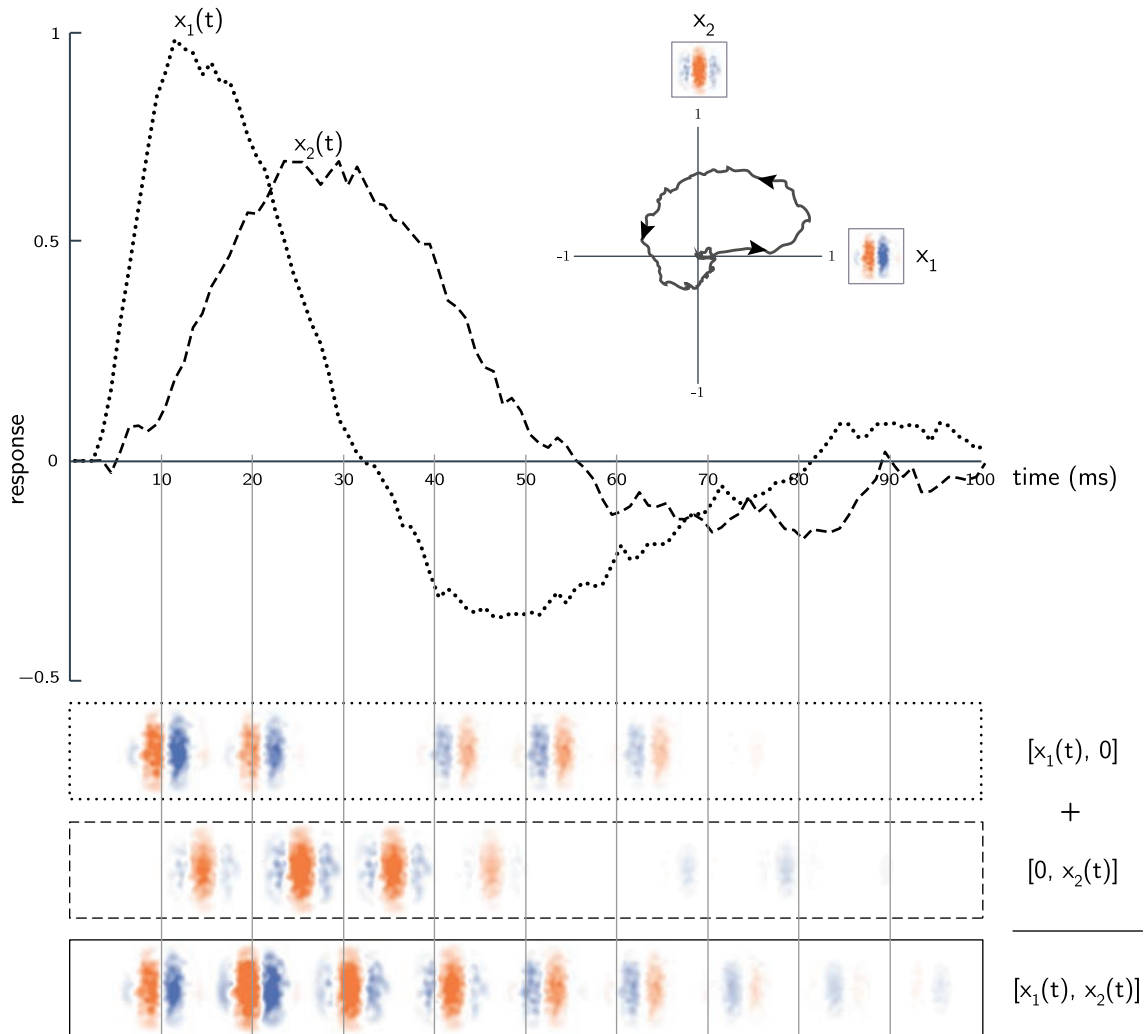


Figure 4.2. Predicted impulse response components of a V1 state phase oscillator. The figure shows an example of the time evolution of the state vector from a sample V1 oscillator, with $T_i = 35$ (or ' T_{35} '). *Inset top right:* The two-dimensional trajectory of the state vector around the phase plane, rotating counterclockwise. The x_1 and x_2 axes represent Gabor functions of the same direction but orthogonal phases. *Top:* Dotted trace showing the time course of the projection of the state vector on x_1 ; dashed trace showing the time course of the projection of the state vector on x_2 . *Below:* The 2-D time courses for x_1 , x_2 and $x_1 + x_2$.




Figure 4.3. Predicted vs actual one-dimensional time courses ('x-t' maps) for V1 simple cell RFs. Each pixel-row is a 1-D slice through the time. *Top row:* Evolving 2-D RFs for the T_{35} oscillator. Since the OI model is scale-invariant, the units of space are relative. *Top left:* The separable x-t map, also known as the biphasic component. *Top middle:* The monophasic component. *Top right:* The inseparable RF x-t map. *Bottom row:* x-t maps from electrophysiological measurements from direction-selective RFs of the separable (left) and inseparable (right) types from V1 neurons in the cat, rescaled and recoloured from Priebe [47].

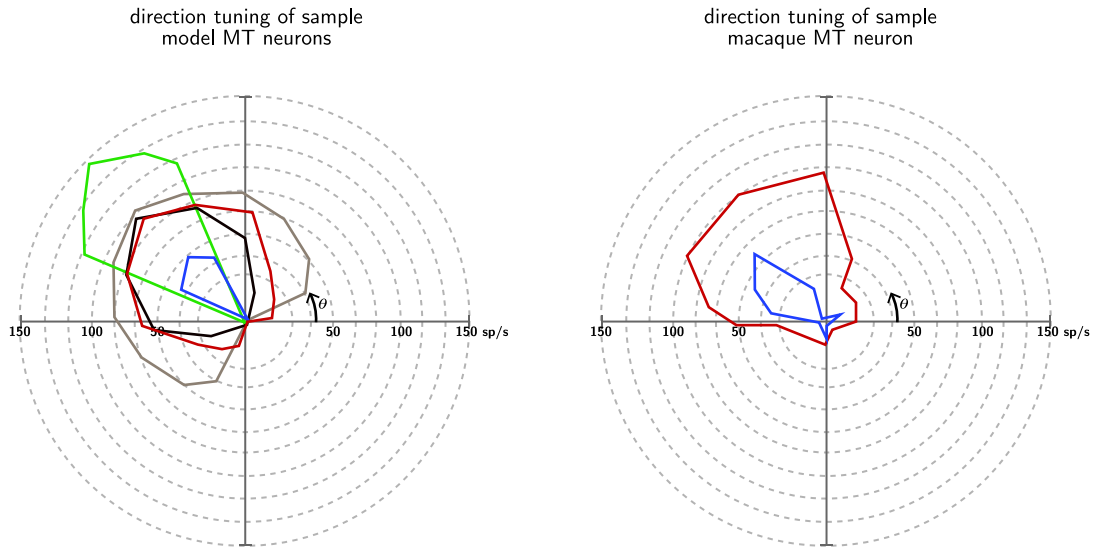


Figure 4.4. Predicted vs actual direction tuning curves for single MT neurons. Each trace is a polar plot of the firing rate of one neuron for a constant speed stimulus moving in direction θ , where $\theta = 0$ is the rightward direction. *Left:* A sample of 5 neurons from an MT ensemble show the approximate range of diversity of direction tuning curves used in the model. *Right:* A sample of one ‘large’ and one ‘small’ direction-sensitive tunings curves for two neurons in macaque MT. Right figure redrawn from, Priebe [82].

4.3 Velocity selectivity in MT

The transformation of the signal from V1 to MT is the sum of the decoded component magnitudes of responses in V1 from phases angles π to $\frac{3\pi}{2}$. Since speed and direction selectivity are spatiotemporal characteristics of the oscillator, each MT ensemble inherits its velocity selectivity from V1. Each MT ensemble pools output from variable number of 1 to 9 V1 ensembles.

4.3.1 Distribution of direction responses

A moving bar stimulus is used to demonstrate the directional selectivity of MT. In this test, eight MT ensembles are tuned to a common speed and one of eight equally spaced directions around a circle. Figure 4.4 provides a qualitative comparison of individual, direction-selective neuron tuning curves to those measured in a macaque electrophysiology study [82]. The results in Figure 4.5 show the highest response in the direction of the stimulus, lower responses for the adjacent directions, and essentially zero response for all other directions.

4.3.2 Distribution of speed responses

In this test, five MT ensembles are tuned to a common direction and one of five equally spaced speeds. The results in Figure 4.6 show the highest response in the ensemble tuned to the speed of the stimulus, and progressively lower responses for the neighbouring speeds.

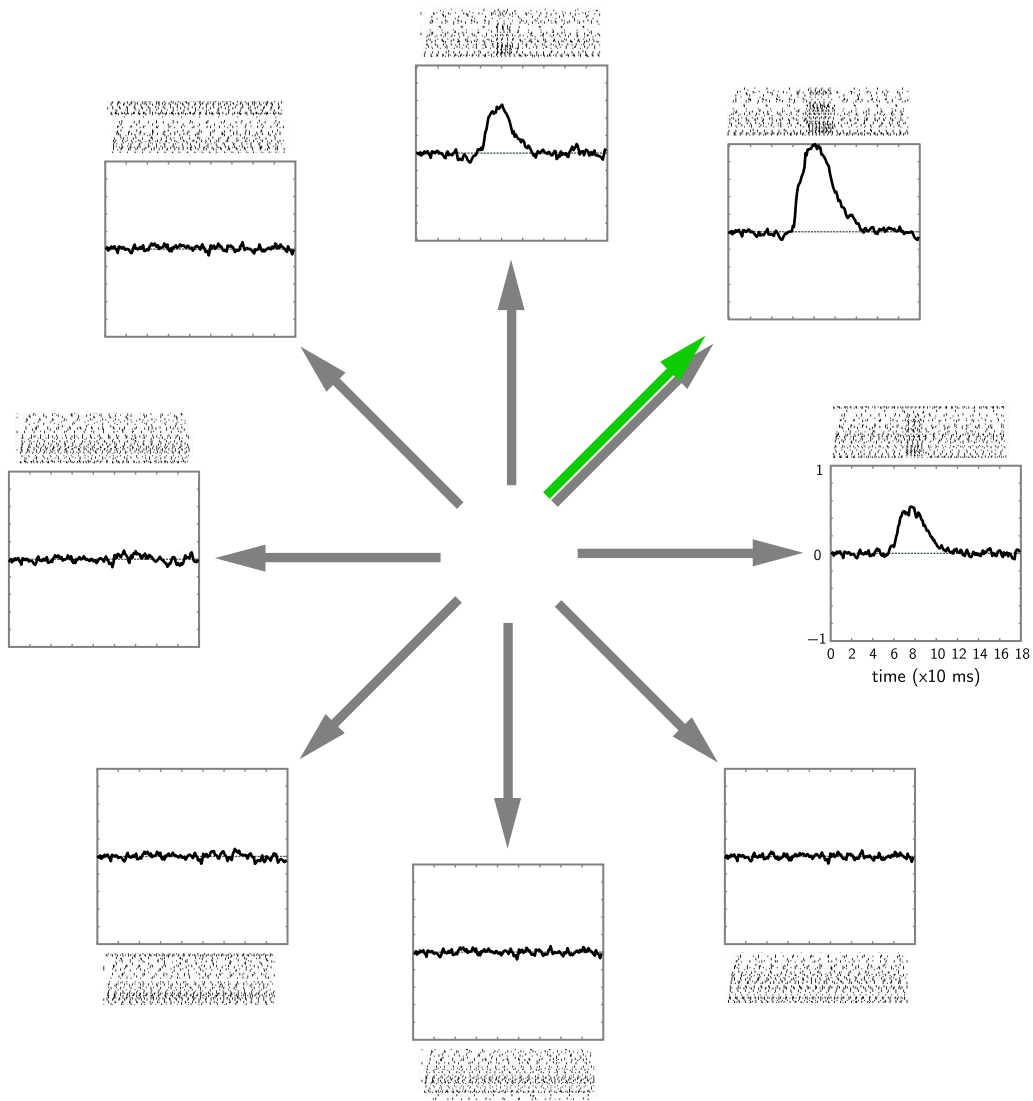


Figure 4.5. Direction selective responses in simulated MT. Each trace shows the decoded scalar response magnitude of the associated MT ensemble, normalized to the maximum of the direction selectivity trials. Spike rasters are above or below the traces. The grey arrows indicate the preferred directions of the associated MT ensemble neurons. The green arrow is the velocity of the moving bar stimulus. The stimulus onset time is $t = 0$. The character of the change in response ('drop-of') from high to low is related to the shape of the V1 RFs be pooled, as discussed in Chapter 5.2. Refer to Figure 1.10 for qualitative comparison.

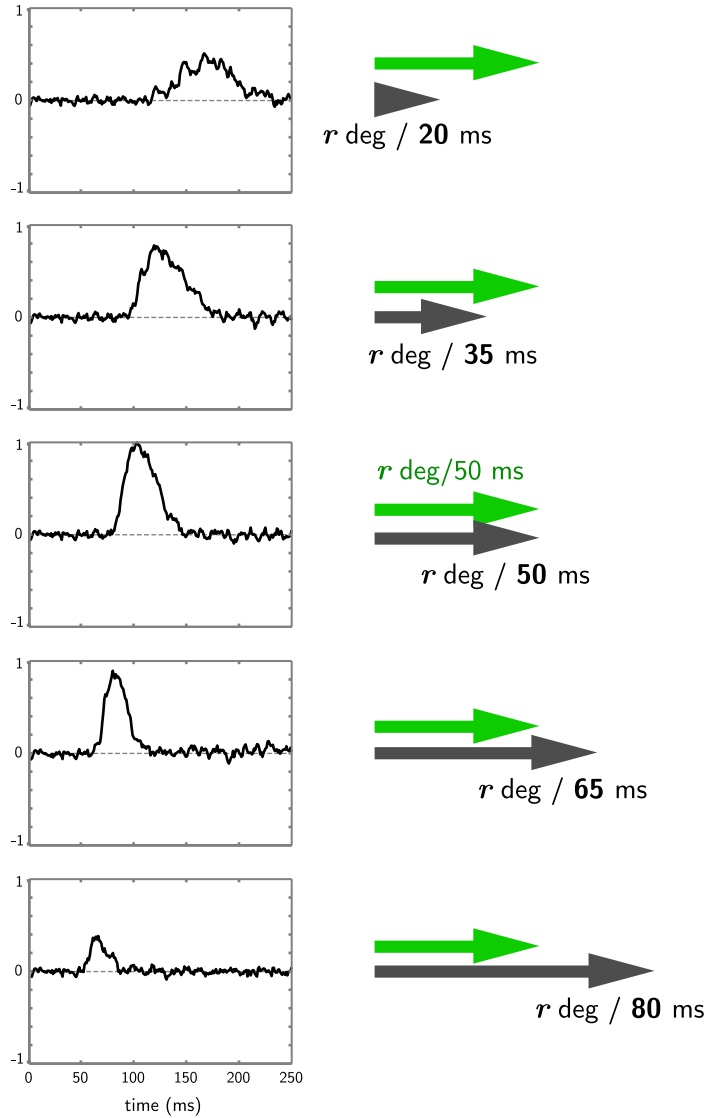


Figure 4.6. Speed selectivity of MT ensemble responses. The grey arrows indicate the preferred velocity of the associated MT ensemble, while the labels below indicate the preferred speed. Since the OI model is scale invariant, the speed value is given in relative spatial units of the visual field, r deg/ms, where r is the chosen RF aperture size. The green arrow indicates the stimulus velocity, the same for each MT ensemble. In these trials, each MT ensemble receives input from one V1 oscillator. Both V1 and MT RF aperture sizes are equal. Each trace shows the decoded value, normalized to the maximum across all speed trials. The MT ensemble tuned to the stimulus velocity (middle) shows the highest response, while the higher and lower speed-tuned ensembles respond less strongly. The stimulus onset time is $t = 0$.




Figure 4.7. Predicted contrast invariance in MT responses. Each trace shows the decoded MT response for the rightward-selective ensemble using a moving bar stimulus of variable contrast. Velocities are shown in green with contrast values noted. The stimulus contrast was varied from 100% to 25%, in quartiles. The rightward-selective ensemble produced the highest response at each contrast level compared to other direction-selective ensembles.

4.3.3 Contrast invariance

The purpose of this test was to show the invariance of direction selectivity of MT ensembles under changes in stimulus contrast. The contrast value of the bar stimulus, moving to the right, was varied from 100% to 25%, in quartiles. The results in Figure 4.7 show the response for the rightward-selective ensemble only, which produced the highest response at each contrast level compared to other direction-selective ensembles. The MT ensembles used for these tests received signal from one V1 ensemble only; thus, the contrast selectivity seen in the model MT is inherited from the spatial properties of the V1 RFs.

4.4 Optical flow maps in MT

The activities of each MT ensemble represent a positive scalar weight for an associated velocity and location in a 400×400 -pixel visual field. This weight is proportional to the length of the arrow illustrated. Further details are provided with Figure 3.8.

The latency times between stimulus onset and MT response varied from 50-65 ms. The time values for the velocity field plots shown in Figures 4.8 to 4.10 show times from 80-140 ms to demonstrate the most characteristic response interval for the given stimuli.

4.4.1 Response to sinusoidal gratings

A drifting sinusoidal grating was used as stimulus input for the simulation, with a speed slightly faster than the preferred speed of the $T = 50$ ms array (T_{50}). The velocity maps for times $t = 80, 100, 120, 140$ ms are shown in Figure 4.8. Overall, the strongest responses are seen in T_{50} , while weaker responses are seen in the ‘faster’ T_{35} array and almost zero response in the ‘slower’ T_{65} array. As seen at $t = 100$ ms in the T_{50} column, there is a Gaussian-like distribution of responses; i.e. the neighbouring velocities appear but with lower magnitude than that matching the stimulus velocity.

Importantly, the T_{35} responses show significant responses in the direction *opposite* the stimulus. In this model, this occurs when the spatiotemporal frequencies of a stimulus are close those those of the V1 oscillators. If this effect were not suppressed by a higher order process (outside the scope of this model), this result would predict the perceptual phenomenon of *transparent retrograde motion*.¹

In attempt to enhance this effect, another experiment was performed using a drifting grating that exactly matched the spatiotemporal frequencies at the $\theta = 0$ components of the T_{50} array. The results are shown in Figure 4.9. As expected, the false detection of simultaneous retrograde motion is magnified. These results are addressed further in Chapter 5.

¹A familiar example is watching a fan or helicopter blades that at certain rotational speeds appear to be rotating opposite the true direction. The same phenomenon is seen in some patterns in translational motion. This phenomenon depends on the spatial frequency in relation to the V1 RFs and the speed of motion.

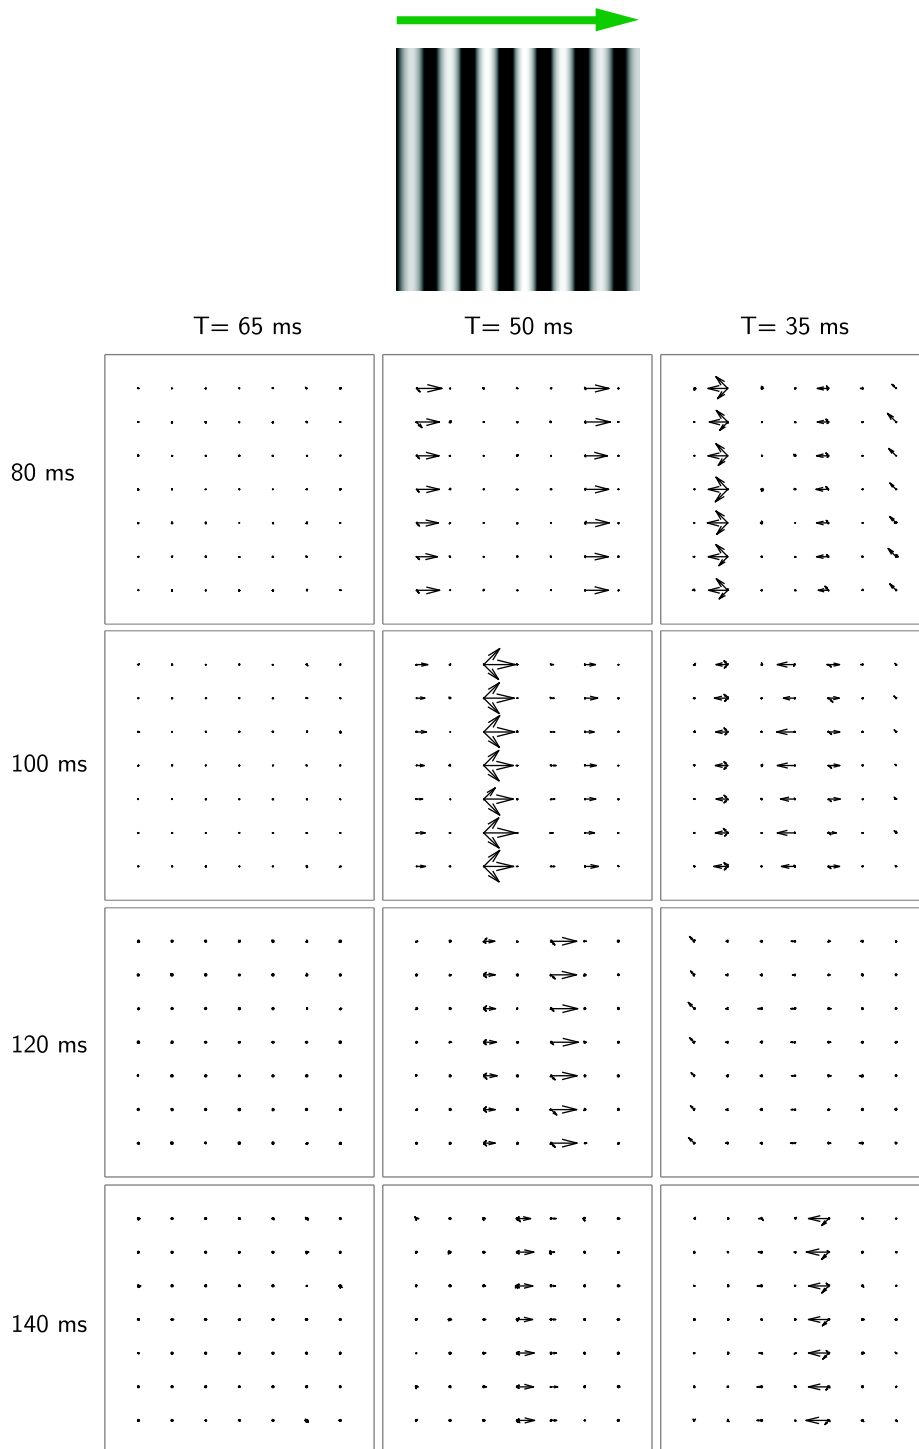


Figure 4.8. Retinotopic velocity maps in MT. Each box contains a 7×7 array of decoded velocities in MT. The three different speed selectivities are organized by column and labelled by the half-period time, T , of the associated oscillator. The stimulus (top) is a sinusoidal grating moving to the right with a spatial frequency slightly lower than, and a temporal frequency slightly higher than, the ‘medium’ T_{50} array (centre column). The strongest response is in T_{50} as expected, with lower responses in the higher speed array T_{35} . Note that the T_{35} array shows some responses contrary to the stimulus direction. The perception of false contrary velocity is known as *transparent retrograde motion* [83]. For a description of how the velocity field is constructed, see Figure 3.8.

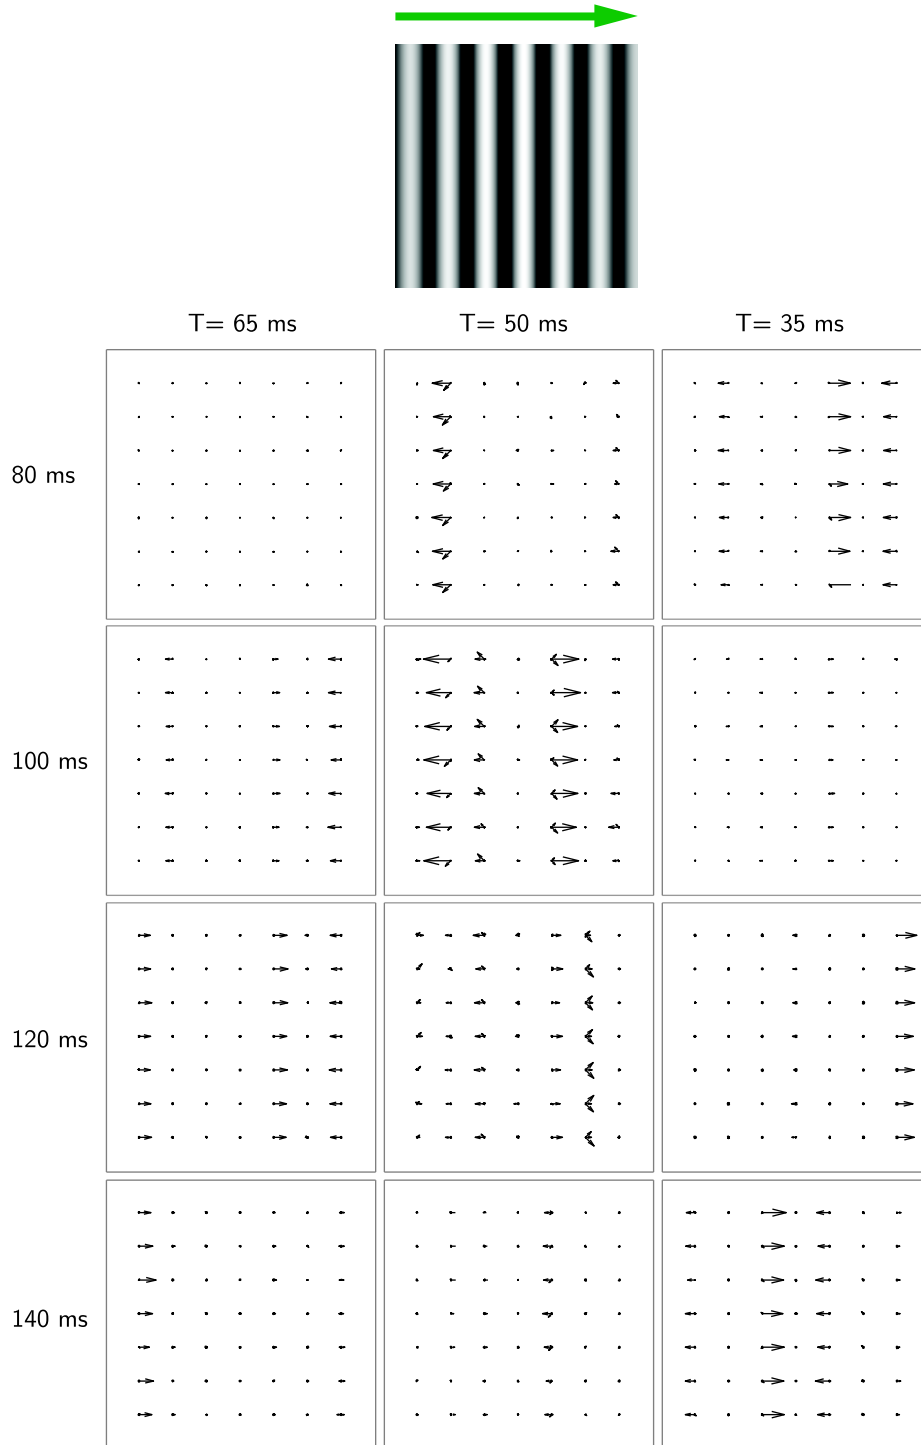


Figure 4.9. Velocity maps showing strong transparent retrograde motion. The stimulus was changed to match the spatiotemporal frequencies of the T_{50} array. At various times, especially 100 ms, the T_{50} array shows strong retrograde motion (to the left) simultaneously with actual rightward velocity.

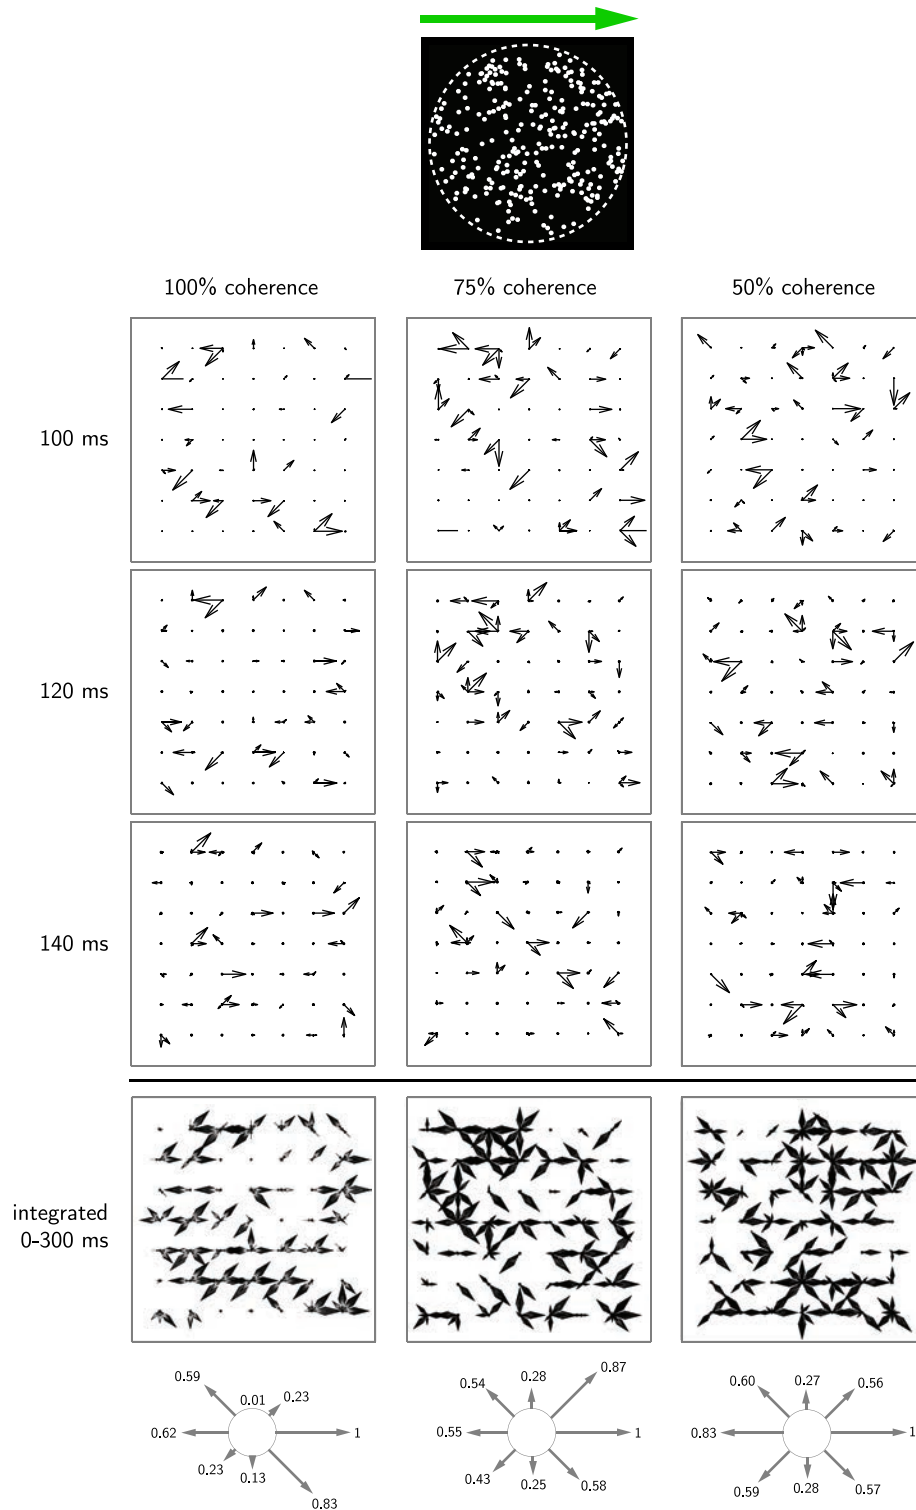


Figure 4.10. MT velocity maps for moving dots of varying coherence. Stimulus coherence levels are categorized by column. The speed selectivity of each array is $T = 50$ ms. In all columns, the stimulus produces a distribution of motion responses from MT. Below each integrated array is a distribution diagram showing the relative weight of each vector in the integration. *Left:* At 100% coherence, the response is greatest in the stimulus direction, but with significant responses in the contrary direction, likely due to retrograde motion effect (see Figure 4.9). The minima occur in the directions orthogonal to the stimulus. *Middle, right:* Decreasing the stimulus coherence causes the response distribution to spread out, as expected.

4.4.2 Response to variable coherence dot patterns

The structure of the moving dot stimulus used in these tests is such that there is little similarity between the spatial frequencies of the stimulus and those of the RFs of the motion processing circuit. Thus, it is a suitable complement to the sinusoidal grating test. In this experiment, the dots move at the same speed with coherence levels of 100% (all moving together, same direction), 75%, and 50% (half moving together, half moving independently in any direction).

Figure 4.10 shows the results for moving dots translating at the preferred speed of the T_{50} array. The velocity maps at $t = 100, 120, 140$ ms do not show any strong indication of the stimulus direction. Moreover, the velocity maps at the same times for coherence levels of 75% and 50% show similar results.

It is only when the MT velocity responses are integrated over the entire stimulus duration (0–300 ms) that any useful information is gained. Figure 4.10 (bottom) shows the integration of all vector responses for each coherence level. At 100%, the greatest responses are in the direction of the stimulus, along with a strong response in the $\theta = -\frac{\pi}{4}, -\pi, \frac{3\pi}{4}$ directions. As the coherence level decreases, the distribution of integrated directions spreads out as expected.

Chapter 5

Discussion and conclusions

The model presented in this thesis has shown through large scale, biologically plausible simulation a possible mechanism for the V1 simple cells RFs required for velocity selectivity, and how such signals can be arrayed in MT to produce a retinotopic velocity map in accord with electrophysiological studies. It has also shown the importance of the temporal structure of burst-like signals passed from the LGN, while not providing a full account for their generation.

The OI model is not meant to predict choice probability of behaving animals in psychometric tests. In fact, drifting grating or moving dots experiments of this work, there was no single ‘winner’ or vector in the velocity map that would strongly indicate the true direction of stimulus motion. Integration of the V1 signals in MT over time improved the outcome, but not convincingly so (see Figure 4.10). A recent study has shown that pooling of neural population activity in MT is a better predictor of choice probability than single or multi-neuron measurements [84], but the prediction value is still weak. Purushothaman and Bradley have demonstrated that fine discrimination motion decisions based depend on the activities of *most informative* neurons, not just pools [85]. Thus, the use of the OI model as a perceptual one is limited, but may serve as the foundation or a component of subsequent processing more directly related to perception and behaviour.

Some important predictions and strengths of the OI model are discussed below, along with some immediate weaknesses and possible extensions.

5.1 Strengths of the model

1. **A mechanism for the generation of observed spatiotemporal receptive fields for V1 simple cells.** The dynamical properties of V1 simple cell RFs are a prediction of the OI model rather than an assumption, as is the case with the other models described in Chapter 1. The OI model shows the separable/inseparable time courses to be one- and two-dimensional aspects of a state vector rotating in phase space. The assumed summation of monophasic and biphasic components in the inseparable time course is shown to be the summation of orthogonal projections of the two-dimensional oscillator state vector.

It is worth stressing that this mechanism is at once both spatial *and* temporal. While it is not intended to fully address the generation of Gabor-like spatial RFs in V1, it does explain how the velocity-selective characteristics of simple cell RFs are generated.

2. **LGN/V1 interference.** The model explains how LGN signals interfere dynamically with V1 oscillators and how resonance conditions correspond to preferred velocities in the stimulus.
3. **Generation of optical flow/velocity maps in MT.** Using a simple sum of the magnitudes of late-phase responses from a pool of V1 oscillators, the production of retinotopic velocity maps is achieved.
4. **Reasonable latency times for MT responses.** The latencies for MT responses in the simulations occurred between 50-65 ms. Studies of macaque MT indicate that >69% of MT latencies are between 50-100 ms for directional stimuli (see Figure 1.11) [24].
5. **Fundamental limit to speed selectivity by lower bound on half-period time.** The shortest half-period time of any oscillator that still produced velocity selectivity was 20 ms. If the neural processes that generate speed perception are limited by the speed sensitivity of V1 simple cells, then the OI predicts that no perception of translational speeds faster than 20 ms (for any unit of length, since the OI model is scale-invariant) is possible.
6. **Prediction of contrast invariance.** This result comes directly from the inner product of convolution of stimulus and direction-selective V1 simple cell RFs and is not a constant of proportionality determined by local population parameters, as in the Simoncelli-Heeger model [44].
7. **Fully specified.** All neural connection weights, time constants and spatiotemporal frequencies required by the model are biologically plausible and/or determined by physiological experiments.

5.2 Weaknesses of the model

1. **No LGN/V1 reciprocal connections.** Since 85-90% of the connections in LGN are *from* the visual cortex and brainstem [8], the possibility that the oscillators proposed in this work are recurrent networks of LGN+V1 is important. If true, the anatomical correspondence of the model would change, but the algorithm would not.
2. **No account of role of V1 complex cells.** The model does not address the function of complex RFs in V1 or the possible consequences of such to the motion processing circuit.
3. **No integration over scales.** While the intent of the model has been to show a scale-invariant mechanism for velocity selection, integration over different levels

in a scale pyramid could reduce or cancel the velocity maps for some stimuli. For example, the retrograde motion that appears in Figure 4.9 may integrate to 0 and thus provide no contribution to the motion percept in higher order visual processing and decision structures, such as the lateral intraparietal area (LIP) [86].

4. **No account of effect of spatial variability in V1 RFs.** As shown in Figure 1.6, V1 RFs have diverse spatial characteristics. In convolution, the spatial structure of the RF has a significant impact on the transformed signal. A conceptual demonstration of this is included in Figure 5.1. The V1 RFs used in this study are relatively simple and idealized; a more diverse and thus more realistic set of RFs would benefit the implementation.
5. **No account of binocular disparity.** The effects of displaced visual signals and potential problems, such as image registration, has not been addressed [25, 82].

5.3 Extensions of the model

5.3.1 Decision threshold mechanism by integration of MT responses over time

An immediate and fruitful extension to the model would be the integration of MT signals in LIP to provide a cumulative measure of motion ‘evidence’ over time. This hypothesis has strong experimental support [87] and is known as the ‘integrate-to-threshold’ mechanism of visual decision. As shown in results of Figure 4.10, integration provides additional information not available in an instantaneous reading of the velocity maps in MT.

5.3.2 Integration of MT responses over scale space

As mentioned, the construction of a scale pyramid containing several resolutions of MT velocity maps may provide important insights into the effects of cancellation of vector representations.

5.3.3 Burst/derivative mechanisms in LGN/V1

The model uses a pre-processed burst signal to simulate output from LGN. As mentioned in section 5.2, the strong projections from V1 to LGN may play an important role in determining the value of edge signal derivatives. The calculation of *when* an edge signal most strongly overlaps a particular V1 RF is not likely achieved in LGN alone and would require feedback connections from V1 simple cells. The inclusion of some type of derivative pathway would provide a more complete picture and support a crucial assumption of the model.

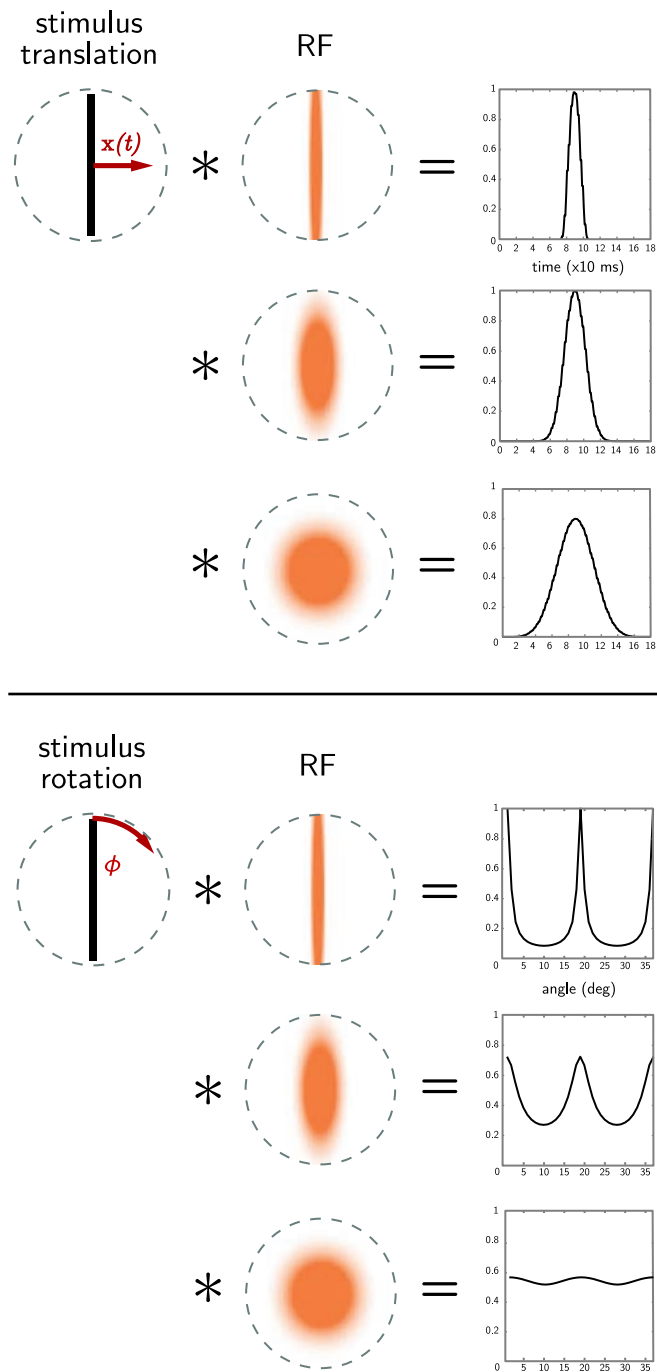


Figure 5.1. The spatial structure of RFs affects convolution or correlation responses significantly. The shape of the RF plays a significant role in convolution with an edge signal. *Top half:* Translational motion of an edge stimulus over time across three different RFs. The plots at right show the results of convolution. The top RF is similar in shape to the stimulus, resulting in a sharp peak response at the time of peak spatial overlap. As the RF becomes less eccentric (more round), the response becomes less peaked for the given stimulus. *Bottom half:* Rotational motion of the stimulus around an axial centre point shows convolution responses that are quite different compared to translational motion, but the eccentricity of the RF affects the sharpness of the response peaks.

References

- [1] H. von Helmholtz. *Treatise on physiological optics*, volume III, chapter “Concerning the perceptions in general”. Dover, 3 edition, 1962.
- [2] David Marr. *Vision: A computational investigation into the human representation and processing of visual information*. W. H. Freeman & Co., MIT Press, 1982, 2010.
- [3] H. Blumenfeld. *Neuroanatomy Through Clinical Cases*. Sinauer Associates, Inc., 2 edition, 2010.
- [4] Yingbin Fu. Phototransduction in rods and cones. In Helga Kolb, Ralph Nelson, Eduardo Fernandez, and Bryan Jones, editors, *Webvision: The Organization of the Retina and Visual System*. 2011.
- [5] Ralph Nelson and Victoria Connaughton. Bipolar cell pathways in the vertebrate retina. In Helga Kolb, Ralph Nelson, Eduardo Fernandez, and Bryan Jones, editors, *Webvision: The Organization of the Retina and Visual System*. 2011.
- [6] G.C. DeAngelis, I. Ohzawa, and R.D. Freeman. Receptive-field dynamics in the central visual pathways. *Trends in Neuroscience*, 18(10), 1995.
- [7] Stephen D. Van Hooser, J. Alexander, F. Heimel, and Sacha B. Nelson. Receptive field properties and laminar organization of lateral geniculate nucleus in the gray squirrel (*sciurus carolinensis*). *Journal of Neurophysiology*, 90:3398–3418, 2003.
- [8] Daniel H. O’Connor, Miki M. Fukui, Mark A. Pinsk, and Sabine Kastner. Attention modulates responses in the human lateral geniculate nucleus. *Nature Neuroscience*, 5(11), 2002.
- [9] D.A. Butts, G. Desbordes, C. Weng, J. Jin, J.-M. Alonso, and G.B. Stanley. The episodic nature of spike trains in the early visual pathway. *Journal of Neurophysiology*, 104:3371–3387, 2010.
- [10] Dawei W. Dong and Joseph J. Atick. Temporal decorrelation: A theory of lagged and nonlagged responses in the lateral geniculate nucleus. *Network: Computation in neural systems*, 6(2), 1995.
- [11] A.B. Saul and A.L. Humphrey. Evidence of input from lagged cells in the lateral geniculate nucleus to simple cells in cortical area 17 of the cat. *Journal of Neurophysiology*, 68(4):1190–1208, 1992.

- [12] A.B. Saul and A.L. Humphrey. Spatial and temporal response properties of lagged and nonlagged cells in cat lateral geniculate nucleus. *Journal of Neurophysiology*, 64(1):206–224, July 1990.
- [13] Jonathan D. Victor. The dynamics of the cat retinal X cell centre. *Journal of Physiology*, 386:219–246, 1987.
- [14] Ehud Kaplan and Ethan Benardete. The dynamics of primate retinal ganglion cells. In C. Casanova and M. Ptito, editors, *Progress in Brain Research - Vision: From Neurons to Cognition*, chapter 2, pages 17–34. Elsevier, 2001.
- [15] Matthew Schmolesky. The primary visual cortex. In Helga Kolb, Ralph Nelson, Eduardo Fernandez, and Bryan Jones, editors, *Webvision: The Organization of the Retina and Visual System*. 2011.
- [16] T.D. Albright, R. Desimone, and C.G. Gross. Columnar organization of directionally selective cells in visual area MT of the macaque. *Journal of Neurophysiology*, 51(1):16–31, 1984.
- [17] Humbert Suarez, Christof Koch, and Rodney Douglas. Modeling direction selectivity of simple cells in striate visual cortex within the framework of the canonical microcircuit. *Journal of Neuroscience*, 15(10):6700–6719, 1995.
- [18] D.H. Hubel and T.N. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *Journal of Physiology*, 160:106–154, 1962.
- [19] E.H. Adelson and J.R. Bergen. Spatiotemporal energy models for the perception of motion. *Journal of the Optical Society of America A*, 2(2), 1985.
- [20] Dario L. Ringach, Michael J. Hawken, and Robert Shapley. Dynamics of orientation tuning in macaque primary visual cortex. *Nature*, 387:281–284, 1997.
- [21] Nicholas J. Priebe and David Ferster. Mechanisms of neuronal computation in mammalian visual cortex. *Neuron*, 75:194–208, 2012.
- [22] Margaret S. Livingstone and Bevil R. Conway. Substructure of direction-selective receptive fields in macaque V1. *Journal of Neurophysiology*, 89:2743–2759, 2002.
- [23] R. Dubner and S.M. Zeki. Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Research*, 35(2):528–532, 1971.
- [24] S.E. Raiguel, D.-K. Xiao, V.L. Marcar, and G.A. Orban. Response latency of macaque area MT/V5 neurons and its relationship to stimulus parameters. *Journal of Neurophysiology*, 82:1944–1956, 1999.
- [25] R.T. Born and D.C. Bradley. Structure and function of visual area MT. *Annual Review of Neuroscience*, 28:157189, 2005.
- [26] Nicole C. Rust, Valerio Mante, Eero P. Simoncelli, and J. Anthony Movshon. How MT cells analyze the motion of visual patterns. *Nature Neuroscience*, 9(11), 2006.

- [27] J.A. Movshon, E.H. Adelson, M.S. Gizzi, and W.T. Newsome. *The analysis of moving visual patterns*, volume 54, 117-151. Vatican Press, 1983.
- [28] D.J. Felleman and J.H. Kaas. Receptive-field properties of neurons in middle temporal visual area (MT) of owl monkeys. *Journal of Neurophysiology*, 52:488–513, 1984.
- [29] A.T. Smith, K.D. Singh, A.L. Williams, and M.W. Greenlee. Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. *Cerebral Cortex*, 11:1182–1190, 2001.
- [30] C.E. Shannon and W. Weaver. *The mathematical theory of communication*. Urbana: University of Illinois Press, 1949.
- [31] H.B. Barlow. *Sensory Communication*. MIT Press, 1961.
- [32] H.B. Barlow. Sensory mechanisms, the reduction of redundancy, and intelligence. In *NPL Symposium on the Mechanization of Thought Process*, number 10, pages 535–539. HM Stationery Office, London, 1959.
- [33] Stephen W. Smith. *Digital Signal Processing: A Practical Guide for Engineers and Scientists*. Newnes, 2002.
- [34] Panos Z. Marmarelis and Vasilis Z. Marmarelis. *Analysis of physiological systems: the white-noise approach*. Plenum Press, 1978.
- [35] Ben Willmore and Darragh Smyth. Methods for first-order kernel estimation: simple-cell receptive fields from responses to natural scenes. *Network: Computation in neural systems*, 14:553–577, 2003.
- [36] Richard A. Young and Ronald M. Lesperance. The Gaussian derivative model for spatial-temporal vision: II. Cortical data. *Spatial Vision*, 14(3,4):321–389, 2001.
- [37] S. Wimbauer, O. G. Wensich, K. D. Miller, and J. L. van Hemmen. Development of spatiotemporal receptive fields of simple cells: I. Model formulation. *Biological Cybernetics*, 77:453–461, 1997.
- [38] A.P. Shon, R.P.N. Rao, and T.J. Sejnowski. Motion detection and prediction through spike-timing dependent plasticity. *Network: Computation in Neural Systems*, 15:179–198, 2004.
- [39] R.J. Douglas and K.A.C. Martin. A functional microcircuit for cat visual cortex. *Journal of Physiology*, 440:735–769, 1991.
- [40] X. Pei, M. Volgushev, T.R. Vidyasagar, and O.D. Creutzfeldt. Whole cell recording and conductance measurements in cat visual cortex *in vivo*. *Neuroreport*, 2:485–488, 1991.
- [41] D. Ferster and B. Jagdeesh. EPSP-IPSP interactions in cat visual cortex studied with *in vivo* whole-cell patch recording. *Journal of Neuroscience*, 12:1262–1274, 1992.

- [42] R. Maex and G.A. Orban. Model circuit of spiking neurons generating directional selectivity in simple cells. *Journal of Neurophysiology*, 75(4), 1996.
- [43] Paul Mineiro and David Zipser. Analysis of direction selectivity arising from recurrent cortical interactions. *Neural Computation*, 10(353-371), 1998.
- [44] Eero P. Simoncelli and David J. Heeger. A model of neuronal responses in area MT. *Vision Research*, 38(5), 1998.
- [45] David J. Heeger. Optical flow using spatiotemporal filters. *International Journal of Computer Vision*, pages 279–302, 1988.
- [46] Shinji Nishimoto and Jack L. Gallant. A three-dimensional spatiotemporal receptive field model explains responses of area MT neurons to naturalistic movies. *Journal of Neuroscience*, 31(41):14551–14564, 2011.
- [47] Nicholas J. Priebe and David Ferster. Direction selectivity of excitation and inhibition in simple cells of the cat primary visual cortex. *Neuron*, 45:133–145, 2005.
- [48] Chris Eliasmith and Charles H. Anderson. *Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems*. MIT Press, 2003.
- [49] Fred Rieke, David Warland, Rob de Ruyter van Steveninck, and William Bialek. *Spikes: exploring the neural code*. MIT Press, 1997.
- [50] Bruce Alberts, Alexander Johnson, Julian Lewis, Martin Raff, Keith Roberts, and Peter Walter. *Molecular biology of the cell*, chapter 16, pages 1227–1228. Garland Science, 2002.
- [51] Peter Dayan and L.F. Abbott. *Theoretical Neuroscience: Computational and mathematical modeling of neural systems*, chapter 7.2, pages 232–233. MIT Press, 2001.
- [52] Helène Paugam-Moisy and Sander Bohte. Computing with spiking neuron networks. In Joost N. Kok Grzegorz Rozenberg, Thomas Bäck, editor, *Handbook of Natural Computing*, volume 1. Springer, 2012.
- [53] Peter Dayan and L.F. Abbott. *Theoretical Neuroscience: Computational and mathematical modeling of neural systems*, chapter 1.1, page 4. MIT Press, 2001.
- [54] C. Koch. *Biophysics of computation: Information processing in single neurons*. Oxford University Press, new edition, 2005.
- [55] J.J.B. Jack, D. Noble, and R.W. Tsien. *Electric current flow in excitable cells*. Oxford: Clarendon Press, 1975.
- [56] W. Softky and C. Koch. *The handbook of brain theory and neural networks*. MIT Press, 1995.
- [57] S. Shamma. *Methods in neuronal modeling*, chapter Spatial and temporal processing in central auditory networks. MIT Press, 1989.

- [58] A.L. Hodgkin and A.F. Huxley. A quantitative description of ion currents and its applications. *Journal of Physiology*, 117:500–544, 1952.
- [59] D.O. Hebb. *The Organization of Behavior: A Neuropsychological Theory*. Psychology Press, new edition, 2002.
- [60] C.F. Stevens and Y. Wang. Changes in reliability of synaptic function as a mechanism for plasticity. *Nature*, 371:704–707, 1994.
- [61] E. Henneman and L.M. Mendell. *Functional organization of motoneuron pool and its inputs*, pages 423–507. Comprehensive Physiology. Wiley, 2011.
- [62] David Heeger. *Poisson Model of Spike Generation*, September 5, 2005.
- [63] Gary R. Holt, William R. Softky, Christof Koch, and Rodney J. Douglas. Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. *Journal of Neurophysiology*, 75(5), 1996.
- [64] Derek Rowell. *MIT Department of Mechanical Engineering, 2.14 Analysis and Design of Feedback Control Systems, Fall Term, 2004*.
- [65] Clarence W. de Silva. *Mechatronics: An Integrated Approach*. CRC Press, 2004.
- [66] M. Kleiner, D. Brainard, and D. Pelli. What’s new in Psychtoolbox-3? *Perception*, 36, 2007. ECVF Abstract Supplement.
- [67] D.G. Pelli. The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10:437–442, 1997.
- [68] D.H. Brainard. The Psychophysics Toolbox. *Spatial Vision*, 10:433–436, 1997.
- [69] Gilbert Strang. *Introduction to Linear Algebra*, chapter 4.4: Orthogonal bases and Gram-Schmidt. Wellesley-Cambridge Press, 3 edition, 2003.
- [70] David A. McCormick, James W. Lighthall Barry W. Connors, and David A. Prince. Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *Journal of Neurophysiology*, 54(4), 1985.
- [71] Michael N. Shadlen and William T. Newsome. Noise, neural codes and cortical organization. *Current Opinion in Neurobiology*, 4:569–579, 1994.
- [72] Donald S. Faber and Henri Korn. Single-shot channel activation accounts for duration of inhibitory postsynaptic potentials in a central neuron. *Science*, 208(4444):612–615, 1980.
- [73] P.O. Bishop and W.A. Evans. The refractory period of the sensory synapses of the lateral geniculate nucleus. *Journal of Physiology*, 134:538–557, 1956.
- [74] Pratik Mukherjee and Ehud Kaplan. Dynamics of neurons in the cat lateral geniculate nucleus: In vivo electrophysiology and computational modeling. *Journal of Neurophysiology*, 74(3), 1995.

- [75] Prakash Kara, Pamela Reinagel, and R. Clay Reid. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron*, 27:635–646, 2000.
- [76] Stacia Friedman-Hill, Pedro E. Maldonado, and Charles M. Gray. Dynamics of striate cortical activity in the alert macaque: I. Incidence and stimulus-dependence of gamma-band neuronal oscillations. *Cerebral Cortex*, 10:1105–1116, 2000.
- [77] Matteo Carandini and David Ferster. Membrane potential and firing rate in cat primary visual cortex. *Journal of Neuroscience*, 20(1):470–484, 2000.
- [78] Jennifer A. O'Brien Peter Schwindt and Wayne Crill. Quantitative analysis of firing properties of pyramidal neurons from layer 5 of rat sensorimotor cortex. *Journal of Neurophysiology*, 77:2484–2498, 1997.
- [79] Duane G. Albrecht, Wilson S. Geisler, and Alison M. Crane. Nonlinear properties of visual cortex neurons: Temporal dynamics, stimulus selectivity, neural performance.
- [80] Christopher C. Pack and Richard T. Born. Temporal dynamics of a neural solution to the aperture problem in visual area MT of macaque brain. *Nature*, 409, 22 February 2001.
- [81] Giedrius T. Buracas, Anthony M. Zador, Michael R. DeWeese, and Thomas D. Albright. Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron*, 20:959–969, 1998.
- [82] Christopher C. Pack, Richard T. Born, and Margaret S. Livingstone. Two-dimensional substructure of stereo and motion interactions in macaque visual cortex. *Neuron*, 37:525–535, 2003.
- [83] Ning Qian, Richard A Andersen, and Edward H Adelson. Transparent motion perception as detection of unbalanced motion signals: I. Psychophysics. *Journal of Neuroscience*, 14(12):7357–7366, December 1994.
- [84] Jing Liu and William T. Newsome. Correlation between speed perception and neural activity in the middle temporal visual area. *Journal of Neuroscience*, 25(3):711–722, 2005.
- [85] Gopathy Purushothaman and David C. Bradley. Neural population code for fine perceptual decisions in area MT. *Nature Neuroscience*, 8(1), 2005.
- [86] Michael N. Shadlen and William T. Newsome. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86:1916–1936, 2001.
- [87] Xiao-Jing Wang. Decision making in recurrent neuronal circuits. *Neuron*, 60, 2008.